

Perception, color, and realism¹

Wayne Wright

ABSTRACT: One reason philosophers have addressed the metaphysics of color is its apparent relevance to the sciences concerned with color phenomena. In the light of such thinking, this paper examines a pairing of views that has received much attention: color physicalism and externalism about the content of perceptual experience. It is argued that the latter is a dubious conception of the workings of our perceptual systems. Together with flawed appeals to the empirical literature, it has led some philosophers to grant color physicalism a scientific legitimacy it does not merit. This discussion provides a useful entry into broader points pertaining to debates about color realism and the relationship between philosophical theories of color and the relevant empirical literatures. A sketch of a novel form of color realism is offered, as is an example that fills in some details of that sketch.

0. Introduction

An account of the nature of color should tell us what, if anything, we see when we have color experiences. Such an account looks potentially relevant to the sciences concerned with color phenomena. Alex Byrne and David Hilbert (2003a, p.4) give voice to this, stating that “the problem of color realism is primarily a problem in the theory of perception” and that the problem is significant not only to philosophers, but “to anyone working in the field of color science.” Against this backdrop, this paper examines a pairing of views that has received much attention over the last twenty years: color physicalism and externalism about the content of perceptual experience. I will argue that the latter, while intuitively appealing, is a dubious

¹ Parts of this paper have been presented at several venues and I am grateful for the helpful comments I have received from those audiences. A particular debt of thanks is owed to Edward Averill, Kimberly Jameson, Peter Ross, Whit Schonbein, Robert Schwartz, Kyle Stanford, and the anonymous referees of this journal. Kent Johnson and Louis Narens provided suggestions that were especially beneficial to the development of the issues raised in section 6.

conception of the workings of our perceptual systems. This view, together with flawed appeals to the empirical literature, has led some philosophers to grant color physicalism a scientific legitimacy it does not merit. The examination of these two theses also provides an entry into broader points pertaining to debates about color realism and the relationship between philosophical discussions of color and the relevant empirical literatures.

I take as a basic point of departure that the reality or otherwise of color is not a matter to be decided from the armchair. The first four sections that follow are thus concerned with establishing that the leading empirically-grounded arguments for the targeted views are unconvincing. I do not take the points raised in these sections to constitute an outright refutation of color physicalism or perceptual externalism. The aim instead is simply to demonstrate that there are legitimate grounds for doubt about the empirical motivations for both views. Those doubts, especially when combined with standard criticisms of color physicalism, are serious enough to make it worth considering alternative conceptions of the relationship between our perceptual experiences and the world and to reconsider what might be expected from philosophical research on color. This sets the stage for section 5, which offers a novel naturalistic perspective on questions about the metaphysics of color and outlines a modest version of color realism. The paper concludes in section 6 with a brief example of the kind of philosophical work that arises from the naturalistic perspective driving the proposed form of color realism.

1. Perceptual externalism

The understanding of perceptual experience at issue here claims that perceptual experiences have representational content and the individuation of that content is metaphysically dependent

Forthcoming in *Erkenntnis*

on external factors. At the level of perceptual qualities like color, those factors are physical, as opposed to, say, social. Let us call this perceptual externalism (PE). Proponents of PE, especially those with naturalistic aims of the sort expressed in the above quote from Byrne & Hilbert, have tended to favor reductively physical accounts of color, although they can differ over the kind of physical property with which color is reductively identified. A good sense of the relationship between PE and theories of color can be gained from Michael Tye's (2006) treatment of a puzzle raised by differences between the color experiences of two standard color perceivers, Jane and John, when both view the same Munsell chip in the same conditions. It looks pure blue to John, but blue with a greenish cast to Jane. To account for the veridicality or nonveridicality of their different experiences, Tye entertains three options that he takes to be mutually exclusive and jointly exhaustive: eliminativism, dispositionalism, and physicalism. Tye contends that eliminativism and dispositionalism are saddled with such deep problems that they are to be dismissed. He backs physicalism, despite its difficulties; Hardin (2003) provides a useful discussion of such difficulties.

Tye describes our perceptual systems as measuring instruments "designed by Mother Nature" (Tye 2006, pp.174-175) to give an "accurate 'readout'" (*ibid.*, p.177) of objects' physical properties. Color physicalism fits easily with this view of the relationship between our perceptual equipment and the world. The basic idea is that the demands of a species' ecological niche drive the evolution of perceptual systems that represent relevant physical properties at an appropriate fineness of grain; hyper-precision is likely not required to meet those demands (*ibid.*, pp.176-177). Thus color vision evolved because perceptual systems are devices that detect and represent physical properties, colors are physical properties of objects, and detecting and representing those physical properties enables a creature to make

Forthcoming in *Erkenntnis*

biologically important discriminations that it otherwise could not make and that perhaps its competitors cannot make as well or at all. If eliminativism were true and objects were not colored, this account of the evolution of color vision could not get off the ground.

Since PE has it that color experiences represent objective properties, it clashes with theories of color in which color experiences essentially figure, such as dispositionalism; see Lycan (2001, p.20). Dispositionalism also seems to entail that the Munsell chip in Tye's example is both pure blue and blue with a greenish tinge, since it causes veridical experiences of those qualities in John and Jane, respectively. Tye (2006, pp.173-174) takes this as a reductio of the view, as uniformly colored objects cannot have two (or more) colors. A standard reply to that objection, not mentioned by Tye, is to relativize dispositional colors to perceivers and circumstances; e.g., redness for subject S in circumstances C. While that maneuver might block Tye's criticism, a perceptual externalist might doubt its compatibility with the evolutionary story just canvassed. Drawing on Byrne & Hilbert's (2003b, p.58) handling of an intra-personal case of relativization, one possible reply to this proposal is that it entails that the color properties my experiences track are not the same ones that were tracked by the experiences of my evolutionary antecedents, which were relativized to them and their circumstances. Thus the forces of selection could not have favored creatures whose visual systems better detected the color properties my visual system detects.

Other proponents of PE have also been forthright about their need to argue for color physicalism; e.g., Churchland (2007, pp.121-125). In the light of their claims about the relevance of the color realism debate to color science, it is interesting to note Byrne & Hilbert's (2003b, p.52) reply to Davida Teller's (2003, pp.48-49) complaint that, to a vision scientist, philosophical disputes over color appear merely terminological:

We conjecture that the reason Teller sees only a tedious squabble about words is that she fails to recognize fully the intentionality, or representational nature, of visual experience.

They go on to explain that they endorse the view that experiences are distinguished by what they represent and they link this with many scientists' claims about the visual system estimating distal properties. Additionally, they object to a relativist account of color on the grounds that it conflicts with their representationalism about phenomenal character (Byrne & Hilbert 2003c, p.792). So, Byrne & Hilbert's attempt to motivate color physicalism is crucial to satisfying the requirements of their theory of perceptual experience; see also Jakab & McLaughlin (2003, p.35). Issues to be discussed subsequently are whether we should accept that theory and what options might become available upon rejecting it.

2. The absence of physical color

As a first step in critically assessing PE and associated color physicalisms, it is instructive to examine how perceptual externalists have handled eliminativism. This discussion is not an argument for eliminativism, but only a demonstration that a commitment to PE leads one to underestimate the merits of views that compete with physicalism. Some points raised in this section are also relevant to the alternative realist account of color sketched later.

Tye (2000) remarks that "massive error" theories of color experience, which follow from eliminativism, are "totally implausible" (p.46) or "incoherent" (p.166). Byrne & Hilbert (2003b, p.59) think it "is something of a mystery" that evolution could produce systems that represent a range of properties that no thing actually has. So, physicalism might have as

Forthcoming in *Erkenntnis*

advantages over eliminativism both rescuing us from implausible or incoherent views and affording a compelling explanation of the evolution of color vision.

Setting aside PE dampens the feeling that the evolution of color vision in an achromatic world is fantastic. Suppose perceived colors are not measurements of a physical structure, even a high-level one like Byrne & Hilbert's reflectance-types. Color vision could still provide survival and reproductive advantages, if there are relations holding with robust statistical regularity between color experiences and ecologically relevant properties of certain objects that cause those experiences. The emergence of such relations likely hinges on there being stable links between those objects' "interesting" properties and the bits of their physical structure that determine how they interact with light. Such regularities could be exploited in ways that improve a creature's success in foraging, breaking camouflage, social maneuvering, mating, etc. Importantly, what regularities a species might become sensitive to depends not only on the demands of its ecological niche, but also on what members of the species can do with information about those regularities. It will not do the creature much good to have its perceptual systems sensitized to an environmental feature, if its cognitive and motor systems are unable to use information about that feature in appropriate ways.

Adam Reeves' reply to Byrne & Hilbert (2003a) provides further useful details. Reeves outlines a "partial color physicalism," but he explicitly does not agree that objects are colored (2003, p.46). He is trying to develop an understanding, in terms of physical features of the environment, of why we might have the kind of color experiences we do. That is importantly different from claiming that perceptual content is to be individuated along the lines of physical properties or that the evolutionary benefits of perception should be understood in terms of a general-purpose goal of reconstructing such properties. Of the three hypothesized opponent

Forthcoming in *Erkenntnis*

processing channels, two have plausible interpretations in terms of physical variables: the achromatic channel – which sums the output of long (L) and middle (M) wavelength sensitive cones – corresponds to luminance and the blue/yellow channel – which compares the activity of short (S) wavelength sensitive cones with that of L and/or M cones – to the phase of daylight. The red/green channel – which compares L and M cone responses – has no natural counterpart, but its orthogonality to the other two dimensions “[permits] more of the natural variability of lights and surface reflectances to be picked up” (*ibid*). Also noteworthy are that (i) there is much more red/green than blue/yellow output from the primate retina and (ii) red/green changes are uncommon in natural settings; see Lovell et al (2005, p.2060) and the findings of Chiao et al (2000) discussed in section 6. If, as a contingent fact about the creature’s way of life and what its environment affords, it is important for some things to be represented in a highly salient way, that red/green channel can be put to effective use in informing the creature about rare but biologically significant states of its physical or social environment.

Consider also the variant and invariant color-opponent responses under illuminant changes in natural scenes and the roles they can play at different times of day; Lovell et al (2005) is followed closely here. Due to its close tie to the phase of daylight, the blue/yellow system has considerable variability in its response as illumination conditions change throughout the day. The bluish character of shadows also poses a problem for the blue/yellow channel. Under a relatively flat illuminant spectrum, however, the blue/yellow system does a reasonable job of segmenting the visual scene on the basis of reflectance, although with not nearly the contrast between many fruits and their backgrounds as is found in the red/green system. The latter is set up in a way that makes for substantial color differences between many

Forthcoming in *Erkenntnis*

(but not all) ripe fruits and their backgrounds, while minimizing the color differences between background elements. The red/green system is capable of ignoring the effects of illumination change and shadows on the color signal from objects throughout much of the day. It does stumble around dawn and sunset, when long wavelengths dominate. At such times, the blue/yellow channel proves valuable.

A further benefit of this arrangement is that, if paired with a sufficiently rich memory and mechanisms that enable comparisons with such memories, outputs from the two chromatic opponent channels could be combined in a way that exploits variant and invariant responses in order to determine whether the object currently being viewed has the same material properties as one viewed before in different circumstances. This scheme could also facilitate discerning telling illuminant properties, such as direction, when the source is not visible; see Zaidi (1998, 2001). Also relevant is the evidence provided by Hansen et al (2005) that memory can have a top-down effect on the perceived color of objects with typical colors, such as bananas, lettuce, and oranges. Consistent with Reeves' (2003, p.46) contention that the basic elements of his partial color physicalism do not entail that objects are colored, the exploitation of reflectance and illuminant spectra in color experience need not be understood in terms of our visual system's measurement of physical properties. It is plausible that such exploitation is instead based on an internally determined set of perceptual categories and structures that have been shaped to capitalize on variant and invariant responses to ecologically relevant stimuli, and which fit together with cognitive and motor factors that are also sensitive to the significance of certain stimuli; e.g., the red berry against green surroundings, the pinkish flush on the skin of a receptive potential mating partner, the yellowish-green appearance of a sunlit tree canopy, the bluish cast, quick luminance change, and low color contrast of a shadowy hiding place.

Byrne & Hilbert (2003b, p.56) note that “Reeves’ claim is compatible with the visual system having the overall function of acquiring reflectance information.” That perhaps is true, but what matters in this context is that Reeves’ claim is also compatible with, and provides a useful basis for explaining, the evolution of color vision in an achromatic world. The considerations raised thus far bear on the connection color physicalists draw between (i) the typically close tie between surface reflectance and color appearance and (ii) the idea that the function of color vision is to represent the reflectance properties of objects. One way of understanding the current point is that rather than having the biological function of recovering or reconstructing the reflectance properties of objects (Hilbert 1992, pp.365-366), color vision allows creatures to take advantage of the reflectance properties of objects (and other relevant features of a creature’s environment, such as atmospheric and illuminant regularities) for various purposes suitable to their ecological niche. The reflectance-dependent comparisons and discriminations that have earned their evolutionary keep need not involve a goal of having perceived color be a measurement (estimate, etc.) of surface reflectance spectra or sets of them, and a separate argument is required to motivate the idea that the visual system has such a goal. In the following, these points will receive further elaboration and defense.

3. More on evolution

Perception is ordinarily thought to help us navigate the world by organizing, or extracting meaning from, the constantly changing sensory input from external sources; see Zimbardo & Gerrig (1998/2002, p.133). It is an error to directly infer from that simple characterization that perception has the function of representing a privileged set of physical properties of the external world; for similar observations, see Matthen (2005) and Mausfeld (2002, 2003).

Forthcoming in *Erkenntnis*

While our perceptual systems might have the purpose of (to some degree) representing the state of the world in terms of physical categories, whether that is actually so is an empirical matter, no matter how intuitively implausible it might seem that they do not. PE is surely an alluring way to think of our perceptual systems and something at least very much like it has been a useful source of ideas that have inspired empirical inquiry. That is not sufficient for making it a cornerstone of perception theory.

Biological considerations in addition to those of the previous section underscore that PE is not on obviously secure empirical ground. The key point is that maximizing fitness and churning out perceptual systems that produce veridical experiences need not go hand-in-hand (Hoffman 2009). For one thing, the heavy metabolic, computational, and architectural burdens that would go along with constructing representations true to the physical details of the scene before the eyes speak against PE; see Dominy et al (2001), Findlay & Gilchrist (2003), Tweed (2003). When it comes to promoting adaptive behavior, those burdens make an approach that accommodates evolution producing perceptual systems that rely on shortcuts and gross simplifications – and, in some cases, built-in deceptions (Trivers 2006, p.xx) – much more plausible than one that heavily emphasizes accurate readouts of physical properties.

Cumulative selection also invites differences between the world model the visual system creates and the world itself. As trial-and-error processes acted on the perceptual systems of our antecedents, maintaining some elements while discarding others, the solution space for new engineering problems was restricted by the materials and structures already on-hand. This limits how perceptual systems can go about the business of making information available for guiding behavior. As already noted, perceptual systems are not independent of the processes that operate on the representations they generate. Nature is unlikely to maintain systems that

Forthcoming in *Erkenntnis*

waste valuable resources on acquiring information (or making acquired information available in a way) that is unusable at subsequent levels of processing. This strongly suggests that our experience of the world results from satisficing and that it is mistaken to insist that it aims to render the objective state of the world; see Brenner & Smeets (2001), Gigerenzer et al (1999, p.363), Ramachandran (1990). What matters is that, when conjoined with our other systems (which are also built with the same priority on getting the job done well enough), our perceptual systems operate in a way that contributes to behavior that promotes our survival and reproductive success.

Further sapping the appeal of an evolutionary rationale for PE, more sophisticated perceptual systems – what a proponent of PE might say are more powerful detection devices – can perform worse than less sophisticated ones at critical tasks, depending on details of the creature’s internal functioning and ecological factors. For example, despite the advantages of trichromacy over dichromacy, dichromats can outperform trichromats in some situations. Morgan et al (1992) offer an example in which dichromats (protanopes and deuteranopes) were easily able to penetrate color camouflage that posed great difficulties for trichromats. In their experiments, dichromats and trichromats were presented briefly flashed patterns and were tasked with identifying a subregion in which the elements differed in orientation and/or size from the rest of the pattern. In conditions in which all the elements of the display were the same color and in settings in which the elements were either randomly red and blue or randomly green and blue, there was no significant difference in performance between dichromats and trichromats. There was a significant difference, though, when the randomly colored elements were red and green: trichromats’ ability to discern the target was compromised whereas dichromats were insensitive to the camouflage.

Forthcoming in *Erkenntnis*

In their Discussion, Morgan et al (*ibid*, p.294) explain the color camouflage effect in terms of principles employed by the visual system to segment the visual scene. The dichromats' visual systems could not use color to segment the display and they were left to rely on a shape-based segmentation, which revealed the target pattern. A preference in the visual system for color-based over shape-based segmentation would explain why trichromats did so poorly in detecting the target pattern. There is a strong tendency for them to see the display as "segmented into differently coloured regions, even though in reality the coloration was random" (*ibid*). This dovetails with empirical research into the propensity of our minds to find patterns, even when they are not actually present. Since our visual system evolved so that only one interpretation is available to us at a time, trichromats would have difficulty seeing through the misleading-but-preferred color-based segmentation in order to detect the target pattern. As Morgan et al (*ibid*, pp.294-295) note, if there is a food source that dichromats can identify but trichromats cannot and the food sources that trichromats can more easily identify become scarce, then (up to the limits imposed by the scarcity of their food supply) dichromats will be at a selective advantage; see also Saito et al (2005), Verhulst & Maes (1998).

Contrary to the evolutionary reasoning frequently offered in support of color physicalism, not only is a richer or higher-dimensional quality space not always needed or beneficial (e.g., over-engineered relative to the demands of one's ecological niche), it can be disadvantageous. Whether it is or is not depends on factors that go beyond whether the perceptual system in question provides accurate readouts of physical properties. PE's narrow focus on mappings between perceptual categories and physically-characterizable properties of the environment dangerously obscures the complicated nature of how such categories contribute to our navigation through the world. The idiosyncratic features of how a species'

visual system functions, how that visual system interacts with other faculties in members of the species, and the structure of the environment can combine to make ineffective or harmful something that would be quite valuable in a different setting.

To clarify, the argument is not that we could get by with perceptual systems that are radically disconnected from the state of the world in which our struggle for existence takes place. Our perceptual systems do put us in contact with significant aspects of our environment, but that is accomplished through a highly complex interplay of psychological, physiological, and external factors. Examples of the latter include terrain, the capabilities of and ecological demands on predators and prey, and the social setting; see Sterelny (2003). Even if attention is limited to external factors that shaped the evolution of our perceptual systems, single-minded concern for how things are characterized in terms of physical categories illegitimately ignores or trivializes biological and social factors, which need not admit of any useful description in terms of physical properties. The physical story is important, but there is no reason to grant it a priori privilege when developing an account of the structure and content of perception.

While none of this refutes PE, the real world factors just discussed cast PE in an unfavorable light. One cannot use the typically close dependence of perceived color on surface reflectance to make a straightforward case on evolutionary (or more broadly biological) grounds for the claim that perceived color is reconstructed surface reflectance. At this point, the appropriate moral to be drawn is that scientists should not be antecedently committed to – moreover, there are grounds for being leery of – a way of theorizing about perception that employs kinds having a categorical mapping onto objective, physical properties of the environment. Hence perceptual externalists who wish to defend color physicalism and claim a relevance to empirical research for their theories will have to look for support in the details of

actual scientific practice or benefits expected to accrue to scientific research from the adoption of PE. That is just what they have done.

4. Perception and the world

Perception researchers distinguish between properties of the physical stimulus and properties encountered in the subjective experience caused by the stimulus. Examples of the former are luminance, sound pressure, and weight and of the latter are brightness, loudness, and heaviness (Laming 1997, p.1). There is an ongoing controversy over just how strong the distinction should be between the two sets of qualities. The history and current state of perception research might be interpreted as suggesting that there is little distinction to be made. A widely shared view in Bayesian and computational approaches to perception is that the visual system has the purpose of recovering properties of the physical scene before the eyes from the confounded sensory input. For example, Tomaso Poggio (1990, p.147) writes that “the goal of color vision is to recover the invariant spectral reflectance of objects (surfaces)” and other scientists speak of the color vision system “estimating” the spectral reflectance characteristics of distal objects. Byrne & Hilbert (2003b, pp.52, 55-56) explicitly take such claims about recovery and estimation of distal properties as a scientific basis for their response to Teller’s complaint about the substance of the color realism debate. Given the language often used by, and perhaps also the actual practices of, many vision scientists, it is no wonder that some would come to believe that “[asking] exactly which properties are being estimated” (Byrne & Hilbert 2003b, p.52) is essential to one’s theory of perceptual experience.

While estimation and recovery play a prominent role in perception research, some perception scientists have accused others of taking such talk too literally. Picking up on points

raised before, some have decried the “passive” approach to vision research – which understands vision as purely a process of interpreting static images, with the goal of reconstructing the distal scene – that is encouraged by an exclusive emphasis on inverse optics; see Churchland et al (1994), Findlay & Gilchrist (2003). Their main complaint is that passive vision research, while it has provided insights into the workings of the visual system, is incapable of accounting for crucial features of how we use our visual equipment to successfully interact with the world. In a different vein, Rainer Mausfeld (2002, 2003) criticizes the strong interpretations of recovery and estimation that give rise to externalist models of perception. He argues that the easy adoption of physical categories into accounts of perceptual representations is an abdication (or empty performance) of an essential task of constructing a theory of perception: developing a taxonomy that properly characterizes the content and structure of perceptual experience. So, despite the initial appearance that scientific practice provides a basis for individuating perceptual content in terms that either match or are close counterparts to physical categories, perhaps greater care is needed in handling the claims from scientists which give that impression.

PE is clearly committed to the idea that isolating physical structures that are reconstructed (mapped, etc) is necessary in developing a theory of, in this case, color experience. Should vision scientists’ talk of estimation and recovery be construed as supporting an understanding of perceptual experience according to which distal stimuli and perceptual achievements are to be characterized using the same sorts of categories? Or is it of more heuristic value? While it is likely that perceptual externalists’ reading of such talk agrees with what many vision scientists have in mind (see Hoffman 2009), the latter option offers a more promising way of proceeding. It is clearly a goal of cognitive science to explain

perception in terms of stimuli, but from that it does not immediately follow that the categories of perception theory can, should or need to be put into one-to-one correspondence with physical attributes of stimuli.

Indeed, a strong interpretation of “recovery” can lead one to take a too simplistic view of the relationship between distal properties and the perceptual experiences they typically cause. As Donald MacLeod (2003, p.433) puts it:

When surface reflectance, for instance, is estimated by heuristics such as taking ratios of retinal stimulus intensities, we may talk of accurate or inaccurate “recovery” or estimation of a physical characteristic of the perceived object The term “recovery” might encourage non-scientists to commit a Rylean category mistake; the “estimated” quantity may have no simple and well-defined physical referent; and the perceptual “estimate” need not have a simple or well-defined functional dependence on its physical counterpart, beyond approximate monotonicity.

The treatment of “recovery” by perceptual externalists commits the mistake MacLeod describes. The importance given to finding physical counterparts – whether carved out along the lines of physics (e.g., Churchland’s 2007) or higher-level physical properties that are “uninteresting from the point of view of physics” (Byrne & Hilbert 2003a, p.11) – for perceptual qualities echoes the attitude that ultimately doomed behaviorism to either concocting excessively complex and unwieldy theories or altogether foregoing theorizing: a reluctance to theorize about causally efficacious internal structures that cannot be characterized in “external” terms; see Heijden (1992, pp.7-9). In the current situation, there are two main ways in which this mindset threatens to harm inquiry: it can lead to mischaracterizations of perceptual phenomena by trying to force upon them a vocabulary

derived from physical theory or it might encourage mistaken attributions of features that are only present in experience to the stimulus. The former is Mausfeld's (2002) "physicalistic trap" and the latter is Köhler's (1929/1947) "experience error." Relatedly, this outlook might set off a quest to discover (or stitch together) "natural kinds" that correspond to perceptual categories, but which turn out to be so fractured, vacuous, or ad hoc that they are of no aid to inquiries into the nature of perception.

One possible reply is that a monotonic dependence of a perceptual feature on a physically-characterizable quantity suffices for the former being an estimate of the latter.² Basically, the suggestion is that so long as color vision is systematic with respect to surface reflectance properties – which surely must be so for it to be beneficial – we should expect color physicalism to come out true. However, as Hardin (1988, 2003) has argued, surface reflectance does not suffice for explaining all the phenomena that should figure in a theory of color; e.g., structural features of phenomenal color, metamerism, contrast effects. For this reason, we should resist the identification of perceived object colors with surface reflectance properties. Various physicalist attempts to deal with those phenomena run into problems that illustrate the dangers just noted.

For example, Byrne & Hilbert (2003a, pp.10-11) identify surface colors with reflectance-types, rather than token reflectances. This is supposed to enable them to deal with the familiar objection from metamers (Hardin 1988, p.64). However, those reflectance-types seem ill-suited to playing a causal-explanatory role when it comes to the kinds of color experiences we have. Given how they are characterized, unless one has an antecedent commitment to PE, there is no reason to understand a given reflectance-type as a physical

² I owe this point to an anonymous referee.

Forthcoming in *Erkenntnis*

property to be type-identified with a given perceived color (with the reflectance-type being reconstructed in color experience), rather than as an equivalence class that is unified solely by the shared potential effects on our visual system (rather than shared nature) of certain heterogeneous token spectral reflectance profiles; see Botterell (2003) and van Gulick (2003). However, not only is PE susceptible to the doubts set out herein, but, as previously noted, the truth of PE depends on the truth of color physicalism. Thus PE could not be fairly cited as a reason to endorse a color physicalist stance with respect to reflectance-types. Additionally, Byrne & Hilbert must admit that the actual mapping of token reflectances to reflectance-types is “unknowable” (2003a, p.21n.50) and their proposal for linking reflectance-types and hue magnitudes runs into an empirical snag (2003b, p.55; 2003c, p.792). Regarding the latter, which is critical to their response to objections from phenomenal structure, Byrne & Hilbert (2003b, p.55) say that “this is no embarrassment; we were simply trying to illustrate our view using a very simple model, and to show how there is no obvious barrier to supposing that individual hue magnitudes are physical properties.” The sort of complications Byrne & Hilbert sweep into the bins of unknowable color facts and simplified examples might really be unproblematic. However, it is not obvious why someone who does not already share their commitments would accept that. Such a party may very well instead find those complications to illustrate flaws inherent to any attempt to make color physicalism relevant to color science.

Regarding Churchland’s (2007) color physicalism, he identifies colors with token surface reflectance profiles. Churchland attempts to handle the metamerism objection by carving out an objective, physical property – a canonical approximation (CA) ellipse – that unifies a set of metameric reflectances. This strategy fails on both empirical and measurement-theoretic grounds; see Kuehni & Hardin (forthcoming) and Wright (2009). Churchland’s

Forthcoming in *Erkenntnis*

treatment of the aspects of phenomenal structure that are supposed to be problematic for color physicalism also merits attention. Using a comparison to maps, Churchland (2007, p.140) aptly points out that it is mistaken to demand that representations encode only information about what is represented. Some structure might be added. As Churchland (*ibid.*) puts it, the troublesome structural relations of phenomenal colors are “an incidental feature” of the way our visual system was set up to detect CA ellipses. The dismissal of, for example, the similarities and differences amongst the hues as merely incidental to what colors “really are” is unmotivated and detrimental to perception theory. In the light of the discussion of sections 2 and 3, the survival and reproductive challenges that faced our ancestors give good reason for thinking that those relations were shaped by selective pressures and are more than a side-effect of mechanisms dedicated to the alleged true function of color vision, despite not neatly mapping onto a physical structure. They look to be a crucial part of the format of visual experience, enabling us to make ecologically important discriminations, identifications, and comparisons. Churchland’s maneuver can be plausibly viewed as drawing whatever appeal it has from an a priori commitment to the idea that only those features that can be characterized in purely physical terms count as proper to the function of perception. It altogether bypasses direct engagement with the difficult issues pertaining to the actual structure and content of visual experience and how that influences our interactions with the world.

Returning to the main thread of discussion, scientists, including vision researchers, seek useful or interesting patterns in their data to exploit in, for example, constructing mechanistic models of the systems they study. Their efforts to explain the patterns in their data will often lead them to posit structures that do not have any direct mapping onto a physical correlate. To bind oneself to characterizing the products of perceptual processing in terms that

straightforwardly correspond to physical properties of the external world is to don a theoretical straitjacket. Teller (2003, pp.49) is making just this point when she writes that the likely reason so many vision science textbooks stake out positions on color in strongly non-physicalist terms is that they “are trying primarily to insist on the distinction between physics and perception.” This is not a shallow attempt to defend one’s intellectual turf. Rather, Teller is making explicit what is widely recognized in actual scientific practice: a vast amount of useful and productive work is done in a wide variety of fields, independently of any concern about whether the kinds of those fields have a “justification” in terms of a reduction to physical properties, entities, relations, etc. This indifference to physical reduction is key to successful research in fields such as linguistics, biology, and cognitive psychology.

Pace Byrne & Hilbert, researchers like MacLeod, Mausfeld, and Teller have not failed to recognize the intentional nature of perception; i.e., that it represents the world as being such-and-such way. Instead, they realize that a great preoccupation with whether perceptual states represent the structures in the environment that typically cause them is a methodological hazard that only invites confusions and errors. It poses a grave risk of distorting, ignoring, or over-simplifying the perceptual phenomena that actually need to be explained. Mausfeld’s (2002, 2003) recommendation that the “internal conceptual structure” of perception be studied entirely in isolation from consideration of distal features is an extreme response to such concerns. However, his basic point about the importance of not allowing the characterization of perceptual phenomena to be dictated by how features of the environment might be categorized (and vice versa) can be accommodated within a perspective that allows distal causes to be taken into account when explaining (at least) some aspects of color experience.

5. A modest color realism

What is to be made of the ontological status of color, given this discussion? One result is that color physicalism does not have the backing from empirical research that some have claimed for it. The model of perception that makes color physicalism so appealing turns out to have little to recommend it, once one considers the biological nature of our perceptual systems and methodological practices that are vital to productive scientific inquiry. Nonetheless, there are numerous instances in which scientists do invoke color properties to explain the phenomena they study; e.g., camouflage, foraging, sexual display, markers of health. There are scientifically interesting questions that can be formed around those phenomena regarding the veridicality of a creature's color experiences and variation in color experiences, both inter- and intra-species. Moreover, even if one is careful to distinguish between perception and physical properties in constructing a theory of perception, there is an important level of analysis at which it is appropriate to address issues of veridicality and variation. Do such considerations bring us right back to the points raised by proponents of PE and demand that we take a position in the color realism debate – specifically realism, likely of the physicalist variety – despite what has been argued so far?

The answer to that question is a qualified “yes”. Yes, in that to the extent that color properties are invoked in successful (promising, etc.) scientific theories, we should be realists about color. The basic reasoning endorsed here is that epistemology drives metaphysics and our best epistemology is embodied in our best current scientific theories. The qualification is that we need not commit ourselves to physicalism, dispositionalism, primitivism, or any other such account of color in so doing. To be clear, I am not offering an argument that any of those theories are false. The argument thus far has only been that color physicalism, in connection

with a theory of perceptual experience to which it has been closely tied, should not be accepted a priori and is lacking in empirical motivation. However, there is a broader moral for philosophical discussions of color that should now be brought to the fore. For the purpose of building a realist theory of color based on the roles colors play in scientific theories, metaphysical theorizing of the familiar sort is unlikely to prove fruitful. A major reason for this has to do with the current state of scientific knowledge. A careful review of the empirical literature reveals that few substantive metaphysical questions about color are appropriate at our current stage of scientific understanding of the relevant phenomena; some relevant examples will be noted shortly. Thus for present purposes, colors are best treated as high-level statistical constructs that explain some of the variance or correlations present in the data gathered in studies that systematically explore the connections between, inter alia, color experiences (inferred from behavioral data), task instructions, physical properties of distal stimuli, and viewing conditions; see Johnson & Wright (2006). A consequence of this is that the particular details of one's account of color are always provisional, subject to revision or rejection in the light of new empirical findings.

While some may recoil from the provisional nature of this approach to color, it is actually a virtuous result, as it conforms to how the discovery of hidden structure is often handled in the sciences. It also accords with the fact that there are unresolved basic issues for color science, such as the structure of psychological color, the nature of color constancy, and the reliability of color spaces and chromaticity diagrams based on the CIE Standard Observers and Grassman's laws; see Burns & Shepp (1988), Foster (2003), and Thornton (1998, 1999). These foundational matters, which must be addressed in order to accumulate the sort of data

needed for the account sketched here, provide a clear opportunity for philosophers to directly contribute to scientific research on color.

Nothing that has been said so far is meant to imply that colors are constructs that can be extracted from statistical regularities. Rather, that is all we can infer about them now and perhaps ever. Nonetheless, associating colors with statistical constructs is enough of a theory of color to suit the needs of working scientists. In connection with the concerns raised a moment ago, individuating colors along the lines of constructs built out of correlations and variance present in the relevant scientific data enables scientists to handle cases of inter- and intra- species variation in color experience (e.g., in terms of the degree of vagueness revealed by analysis of the data) and to make judgments about the veridicality of a creature's color experiences (e.g., in terms of outliers). If necessary, it can do service in accounts that take color to be a heterogeneous kind, even when considering only the roles it plays in the experiences of a single type of perceiver; e.g., it can provide a basis for the strong distinction Mausfeld (2003) makes between color as a property of surfaces and color as a property of illuminants. This account is also capable of encompassing aspects of color vision that go beyond surface reflectance spectra, illuminant spectral power distribution, and the like. For example, there is surely a place within this approach for the use birds (and other creatures) make of specialized photoreceptors (e.g., oil droplets containing particles of magnetite) for magnetoreception and magnetic orientation; see Deutschlander et al. (1999). The same goes for honeybees' use of rhodopsin molecules for sensitivity to the polarization of light in the ultraviolet range. For related remarks about specialized biological uses of color vision, see Thompson (2000).

This approach surely does not directly address many questions philosophers have asked about color. It may even strike some as a naïve scientism that is, at bottom, empty of genuine philosophical content and incapable of making any real contribution to the sciences, despite how it has been billed. In order to quiet such concerns, I turn now to a brief example of the kind of philosophical work one might be led to pursue from this perspective. Noteworthy about the example presented is that it bears on the prospects for computational accounts of color constancy that are thought to favor PE; e.g., compare PE with the “linking proposition” of Brainard et al (2003, pp.313-314), which posits a (presumably bijective) function between a three-parameter estimation of surface spectral reflectance and experienced color.

6. Evaluating low-dimensional linear models

Considerable effort has gone into developing low-dimensional linear models of surface reflectance (and illuminant) spectra. Such models might have applications in, for example, color reproduction systems, color order systems, and theories of color vision. Regarding the latter, linear models may aid attempts to understand the relationship between distal stimuli, experienced color, and various stages of processing from the light incident at the retina to color experience; consistent with the preceding discussion, an interest in these matters should not be read as an endorsement of PE. The basic rationale of low-dimensional linear models is that they provide a means for approximating a set of empirical reflectance curves using a small number of basis curves across the entire set and assigning appropriate weights for the basis curves to each surface. For example, in a model with three basis functions (derived, perhaps, by a data reduction technique like principal component analysis; PCA), the reflectances \underline{R} of a set of surfaces can all be approximated as $\underline{R}(\lambda) \approx C_1V_1(\lambda) + C_2V_2(\lambda) + C_3V_3(\lambda)$ at wavelength

intervals from 400 – 700 nm, where the $V_{i\lambda}$ s are basis functions and the $C_{i\lambda}$ s are an individual surface's weights. A key issue for such models is how few basis functions are needed to produce an adequate approximation of surface reflectance. Given the constraints imposed by normal human perceivers' possession of three cone classes, proponents of certain computational theories of color constancy (and color physicalists who find support for their views in such theories), in particular, would be quite happy if three basis functions would suffice; see Maloney (1999).³ A major normative concern involves how such approximations should be developed and evaluated. From the perspective described in section 5, such normative issues are just the sort of thing that philosophers are encouraged to take up.

David Foster and his colleagues have challenged traditional approaches to evaluating the adequacy of low-dimensional models of surface reflectance; see Nascimento et al (2005) and Oxtoby & Foster (2005), which in what follows will be referred to together as “Foster et al.” They contend that too much emphasis has been placed on theoretical considerations when determining how few dimensions a linear model needs to satisfactorily approximate reflectance spectra. For example, variance accounted for (VAF) has “overwhelmingly” been the measure of choice for evaluating linear models (Maloney 2003, p.289), with the usual criterion of adequate approximation being to account for >99% of the total variance. A robust result across many studies is that three (or four) basis functions are sufficient to account for >99% of the variance in the data sets studied. Instead of only three basis functions, Foster et al's psychophysical results indicate that at least eight are needed to create approximations of natural scenes that are perceptually indistinguishable from their originals; they find that five or

³ Given how the theories are constructed, if the number of required basis functions exceeds three (i.e., the number of cone cell classes), there is no unique solution for the “intrinsic color” specified by a surface's basis function weights; see Maloney (1999, pp.399-401).

Forthcoming in *Erkenntnis*

more basis functions are needed for Munsell chips. While more work is needed to fully assess these claims, their approach raises a host of interesting questions about the development and evaluation of linear models, as well as how they might be used in color research.

One study that is suggestive in this context, despite its reliance on VAF, is Chiao et al (2000). They studied reflectance spectra of scenes that creatures are likely to encounter in their natural habitats. Using VAF as their criterion, they were satisfied with having found that three principal components (PCs) account for ca. 98% of the variance in their data. They observe:

Although the percentage of the total variance of the third PC in forest scenes is less than that of the first two PCs, the information present in this component might be important biologically. (*ibid.*, p.223)

The VAF of the PCs they derived for forest scenes was, from first to fourth: 89.2042, 5.4875, 3.8834, 0.5158 (*ibid.*, p.221). They go on to note that their third PC suggests red-green (RG) variation and link this with the hypothesis that trichromacy evolved in our lineage as an aid to visual search for ripe fruits against green foliage. Recall also the earlier point about the dominance of red-green output from the primate retina. Another important observation comes from Nascimento et al (2005, p.1021):

The salience or otherwise of poorly approximated spectral reflectances in natural scenes can only be determined empirically. Numerically small errors may assume a much larger visual significance when present in real scenes.

The upshot of all this is that, at least far as the visual system is concerned, not all variance is created equal. Given our current state of knowledge of the workings of the visual system and the presence of good reasons for thinking that some reflectance patterns might have (or have had) greater biological significance than others, a proper assessment of the perceptual

significance of any bit of variance requires experimental study. Of particular importance is how different choices in data analysis might affect the interpretation of results of studies like those of Foster et al.

In their studies, Foster et al used indistinguishability of originals and approximations as the criterion for evaluating the adequacy of n-dimensional linear models. The harder it is to discriminate the two, the better the approximation. They experimentally investigated, using human participants, how many basis functions are required to:

1. Produce approximations of natural scenes that subjects cannot discriminate from the originals; subjects viewed both simultaneous and successive presentations of originals and approximations that were generated using one to eight basis functions.
2. Produce pairs of Mondrian patterns such that subjects cannot distinguish patterns constructed using the Munsell set from their approximations; subjects viewed simultaneous presentations of originals and approximations that were generated using one to eight basis functions.

Performance at near-chance level (~55%) was used as the standard for indistinguishability.

The following focuses on the study of natural scenes in order to streamline the presentation.

Nascimento et al (2005) subjected reflectance spectra of twenty outdoor (urban and rural) scenes to a mean-centered PCA, deriving eight PCs. Important to note is that PCA was performed on each scene separately and thus there is no common set of basis functions used across all approximations. For this reason, Nascimento et al take their conclusion that a minimum of eight PCs are needed to produce adequate approximations to be conservative. Eight PCs accounted for >99% of variance in each image. Nascimento et al worked with the original reflectance spectra, rather than a transformation of them, prior to doing PCA. Other

researchers, though, have achieved impressive results using manipulations of the original reflectance data. Laurence Maloney (1986, p.1679) found that weighting the differences between empirical and modeled reflectances by the inverse of the luminous efficiency function (V_λ) results in VAF over 99.6% using just three or four basis functions. A.K. Romney (2008) applied a cube-root transformation to Munsell reflectance spectra before performing a singular value decomposition on them, which gave outstanding results with just three basis functions.

It is natural to ask whether Nascimento et al might have achieved better approximations, resulting in a faster decline in discriminability as more components are used, by employing some such technique. That, of course, provokes questions about what is to be made of their finding that (at least) eight PCs are needed to produce approximations of natural scenes that are indistinguishable from the originals, given their choices in data analysis. On the other hand, there are issues regarding the significance of (maybe) getting a better fit by using those techniques, when it comes to matters such as the prospects for computational accounts of color constancy. For example, some have criticized the CIE 1931 color matching functions, one of which is V_λ ; see Stockman & Sharpe (1998). It is also appropriate to ask about the empirical plausibility of the idea that the visual system has (or is able to behave as though it has) a representation of its luminous efficiency function.⁴ Additionally, while there are nonlinearities in the visual system, it is not obvious what empirical interpretation might be

⁴ A similar question arises about the idea that the visual system achieves color constancy by performing – or behaving as though it performs – a matrix inversion to solve for a surface’s basis function weights; see Brown (2003). However, Usui et al (1992) show that a five-layer neural network can perform an operation equivalent to singular value decomposition.

Forthcoming in *Erkenntnis*

given to a cube-rooting of surface reflectance spectra prior to deriving basis functions.⁵ In other words, addressing the normative issue of what criteria are appropriate for determining the adequacy of low-dimensional models of reflectance spectra requires coming to grips with what bearing a candidate model might have on our understanding of color phenomena. For present purposes, I will limit my attention to one specific issue turned up by reflection on the work of Foster et al.

While Foster et al emphasize the absolute values of their discriminability measures, Dannemiller (1992, pp.507-508) previously made a compelling case that relative rates of change are a safer measure to work with. Thus compare in figure 1 (a) the rate of change in the average colorimetric difference of pixels (across the entire image) in CIELAB space ($\overline{\Delta E^*_{ab}}$) between originals and approximations as more basis functions are used to construct approximations and (b) the rate of change in discrimination performance as basis functions are added; these results are from Nascimento et al (2005).

⁵ Wright (submitted-a) offers a proposal that effectively moves Romney's cube-rooting of reflectance spectra inside the visual system.





Figure 1: (a) Average Euclidean distance in CIELAB space between originals and approximations for rural and urban scenes; (b) Discriminability of originals and approximations of rural scenes for one observer. Images adapted from Nascimento et al (2005), with connecting lines used to highlight contrasting patterns in rates of change.

$\overline{\Delta E^*_{ab}}$ changes rapidly as approximations go from using one component to five and slows considerably as five to eight components are employed. Percent-correct discrimination is mostly flat and stays high through four components, dropping quickly to ca. 75% as a fifth component is added. The subsequent decline is fairly steep from five to eight components.

The $\overline{\Delta E^*_{ab}}$ curve behaves as one might expect from a model in which most of the variance in the system is quickly accounted for by the first few basis functions (hence the initially fast changes in $\overline{\Delta E^*_{ab}}$) and the remaining variance is extremely small and broken out in small pieces across the remaining basis functions (hence the latter part of the curve being low and flat). The discrimination curve looks nothing like that. In fact, it resembles the $\overline{\Delta E^*_{ab}}$ curve reflected about both the vertical and horizontal axes. Going from one to four components makes little or no difference to discriminability, but discrimination drops precipitously over the latter half of the curve. The perceptual adequacy of approximations increases rapidly as components are added that account for very little and progressively less of the total variance. The addition of the later components provides only minimal improvements in the colorimetric quality of the approximations. It is also at these later components where typical linear models have exceeded or are close to exceeding 99% of total variance.

This suggests a need for more careful study of the variance structure of surface reflectances. PCAs run on several different, broad data sets, while they vary in their details, tend to have a first PC that accounts for ca. 90% of the total variance and which is relatively flat. Thus there is a prima facie case for leaving that component intact and treating it as linked to lightness. However, it would be worthwhile to investigate the results of rotating the other PCs while keeping the first in place; it may even prove beneficial to also rotate the first PC.

Such (orthogonal or oblique) rotation would be to a new set of coordinate axes in the subspace spanned by the original PCs. A review of the literature reveals that researchers, including Foster et al, are strongly disposed to simply take the components as they come out of their mathematical software, which fits with the traditional approach's emphasis on VAF. Since PCs are ordered according to the maximum amount of remaining variance projectable on them, a goal of determining how much variance can be accounted for by as few basis functions as possible would not naturally lead one to consider rotating the PCs. In the light of both Foster et al's findings and developments leading up to their experiments, this practice should be questioned. A more promising approach would respect the fact that the visual system is much more sensitive to some chromatic variations than others and would regard VAF as altogether too blunt an instrument for evaluating or putting to use linear models.

Given its greater perceptual significance with respect to discrimination, one idea is to analyze the ~10% of the variance left after the first PC as though it were 100% of the variance to be explained. A goal that naturally emerges is to explore ways of modeling that variance so, for instance, the resulting new $\overline{\Delta E^*_{ab}}$ and (experimentally determined) discriminability curves come into some sort of alignment. It is plausible that satisfying this goal would provide insight into the relationship between natural surface reflectance patterns and color experience. On a more practical note, it could also suggest more effective means of color reproduction. Key to this project is a thorough inspection of the "locations" of ΔE^*_{ab} differences for approximations generated using different choices of coordinate axes. CIELAB color space is not perceptually uniform and ΔE^*_{ab} is only a "rough guide to ... discriminability" (Brainard 2003, p.203). The location in CIELAB space of a particular value of ΔE^*_{ab} matters to its perceptual significance. Such an examination of the distribution of colorimetric differences

between approximations and originals should prove beneficial in determining what rotations of the PCs might result in approximations that are more indistinguishable from their originals with a small number of basis functions. Of course, if better approximations can be achieved through this sort of work, there remain the earlier questions of what to make of the empirical significance of the matches found, given the choices made in handling data and constructing the model.

This section likely raises more questions than it answers, but it serves the purpose of illustrating the sort of work that would go into constructing a theory of color of the sort sketched in section 5 and its significance for color scientists; different examples are developed at greater length in Wright (submitted-a,b). This discussion amounts to more than just a recitation of empirical details or an appeal to empirical results to advance or criticize a theory of the metaphysics of perception or color. Rather, it directly engages with philosophical issues at the root of an important area of scientific work on color, work that promises to inform the “statistical construct” approach to color realism favored here. Further activity along these lines will not only be useful for assembling an account of color based on our best current scientific theories, it can also aid the empirical research that goes into forming those theories.

7. References

- Botterell, A. 2003. “Colors as explainers?” *Behavioral and Brain Sciences*, 26, 785-786.
- Brainard, D. 2003. “Color appearance and color difference specification.” In S. Shevell (ed) *The Science of Color*. Oxford: Elsevier.
- Brainard, D., J. Kraft, and P. Longère. 2003. “Color constancy: developing empirical tests of computational models.” In Mausfeld & Heyer (2003).
- Brenner, E. & J. Smeets. 2001. “We are better off without perfect perception.” *Behavioral and Brain Sciences*, 24, 215-216.

Forthcoming in *Erkenntnis*

- Brown, R. 2003. "Backgrounds and illuminants: the Yin and Yang of color constancy." In Mausfeld & Heyer (2003).
- Burns, B. & B. Shepp. 1988. "Dimensional interactions and the structure of psychological space: the representation of hue, saturation, and brightness." Perception & Psychophysics, *43*, 494-507.
- Byrne, A. & Hilbert, D. 2003a. "Color realism and color science." Behavioral and Brain Sciences, *26*, 3-21.
- _____. 2003b. "Color realism redux." Behavioral and Brain Sciences, *26*, 52-63.
- _____. 2003c. "Color realism revisited." Behavioral and Brain Sciences, *26*, 791-794.
- Chiao, C, T. Cronin, & D. Osorio. 2000. "Color signals in natural scenes: characteristics of reflectance spectra and effects of natural illuminants." Journal of the Optical Society of America A, *17*, 218-224.
- Churchland, P.M. 2007. "On the reality (and diversity) of objective colors: How color-qualia space is a map of reflectance-profile space." Philosophy of Science, *74*, 119-149.
- Churchland, P.S., V.S. Ramachandran, & T.J. Sejnowski. 1994. "A critique of pure vision." In C. Koch & J. Davis (eds.) Large scale neuronal theories of the brain. Cambridge, MA: MIT Press.
- Dannemiller, J. 1992. "Spectral reflectance of natural objects: how many basis functions are necessary?" Journal of the Optical Society of America, A, 507-515.
- Deutschlander, M., J. Phillips, J, & C. Borland. 1999. "The Case for Light Dependent Magnetic Orientation in Animals." Journal of Experimental Biology, *202*, 891-908.
- Dominy, N., P. Lucas, D. Osorio, & N. Yamashita. 2001. "The sensory ecology of primate food perception." Evolutionary Anthropology, *10*, 171-186.
- Findlay, J. & I. Gilchrist. 2003. Active Vision. Oxford: Oxford University Press.
- Gigerenzer, G., P. Todd, & the ABC Research Group. 1999. Simple Heuristics that Make Us Smart. New York: Oxford University Press.
- Hansen, T., M. Olkkonen, S. Walter, & K. Gegenfurtner, K. 2006. "Memory modulates color appearance." Nature Neuroscience, *9*, 1367-1368.
- Hardin, C.L. 1988. Color for Philosophers: Unweaving the Rainbow. Indianapolis: Hackett.
- _____. 2003. "A spectral reflectance doth not a color make." Journal of Philosophy, *100*, 191-202.
- Heijden, A.H.C. van der. 1992. Selective Attention in Vision. New York: Routledge.
- Hilbert, D. 1992. What is color vision? Philosophical Studies, *68*, 351-370.
- Hoffman, D. 2009. "The interface theory of perception." In S. Dickinson, M. Tarr, A. Leonardis, &

Forthcoming in *Erkenntnis*

- B. Schiele (eds) Object categorization: Computer and human vision perspectives.
Cambridge (UK): Cambridge University Press.
- Jakab, Z. & B. McLaughlin. 2003. "Why not color physicalism without color absolutism?"
Behavioral and Brain Sciences, 26, 33-34.
- Johnson, K. & W. Wright. 2006. "Colors as properties of the special sciences." Erkenntnis, 64,
139-168.
- Köhler, W. 1929/1947. Gestalt Psychology. New York: Liveright.
- Kuehni, R. & C.L. Hardin. "Churchland's metamers." British Journal for the Philosophy of
Science. Forthcoming.
- Laming, D. 1997. The Measurement of Sensation. Oxford: Oxford University Press.
- Lovell, P., D. Tolhurst, C. Parraga, R. Baddeley, U. Leonards, J. Troscianko, T. Troscianko. 2005.
"Stability of the color-opponent signals under changes of illuminant in natural scenes."
Journal of the Optical Society of America A, 22, 2060-2071.
- Lycan, W. 2001. "The Case for Phenomenal Externalism." In J. Tomberlin (ed.) Philosophical
Perspectives, 15. Cambridge: Blackwell Publishers.
- MacLeod, D. 2003. "Phenomenology and mechanism." In Mausfeld & Heyer (2003).
- Maloney, L. 1986. "Evaluation of linear models of surface reflectance spectra with small numbers
of parameters." Journal of the Optical Society of America A, 3, 1673-1683.
- _____. 1999. "Physics-based approaches to modeling surface color perception." In K.
Gegenfurtner & L. Sharpe (eds) Color Vision: From Genes to Perception. New York:
Cambridge University Press.
- _____. 2003. "Surface color perception in constrained environments." In Mausfeld & Heyer
(2003).
- Matthen, M. 2005. Seeing, Doing, and Knowing. Oxford: Oxford University Press.
- Mausfeld, R. 2002. "The Physicalistic Trap in Perception Theory." In D. Heyer & R. Mausfeld
(eds.) Perception and the Physical World. Chichester: Wiley & Sons.
- _____. 2003. "'Colour' as part of the format of two different perceptual primitives: The dual
coding of colour." In Mausfeld & Heyer (2003).
- Mausfeld, R. & D. Heyer (eds.). 2003. Colour Perception. New York: Oxford University Press.
- Morgan, M., A. Adam, & J.D. Mollon. 1992. "Dichromats detect colour-camouflaged objects that
are not detected by trichromats." Proceedings of the Royal Society B, 248, 291-295.
- Nascimento, S. D. Foster, & K. Amano. 2005. "Psychophysical estimates of the number of
spectral-reflectance basis functions needed to reproduce natural scenes." Journal of the

Forthcoming in *Erkenntnis*

- Optical Society of America A, 22, 1017-1022.
- Oxtoby, E. & D. Foster. 2005. "Perceptual limits on low-dimensional models of Munsell reflectance spectra." Perception, 34, 961-966.
- Poggio, T. 1990. "Vision: The 'other' face of AI." In K.A. Mohyeldin Said, W.H. Newton-Smith, R. Viale, & K.V. Wilkes (eds.) Modeling the Mind. Oxford: Clarendon Press.
- Ramachandran, V.S. 1990. "Interactions between motion, depth, color, and form: the utilitarian theory of perception." In C. Blakemore (ed.) Vision: Coding and Efficiency. New York: Cambridge University Press.
- Reeves, A. 2003. "Color as a factor analytic approximation to nature." Behavioral and Brain Sciences, 26, 46.
- Romney, A.K. 2008. "Relating reflectance spectra space to Munsell color appearance space." Journal of the Optical Society of America A, 25, 658-666.
- Saito, A. A. Mikami, S. Kawamura, Y. Ueno, C. Hiramatsu, K. Widayati, B. Suryobroto, M. Teramoto, Y. Mori, K. Nagano, K. Fujita, H. Kuroshima, T. Hasegawa. 2005. "Advantage of dichromats over trichromats in discrimination of color-camouflaged stimuli in nonhuman primates." American Journal of Primatology, 67, 425-436.
- Sterelny, K. 2003. Thought in a Hostile World. Malden, MA: Blackwell Publishing.
- Stockman, A. & L. Sharpe. 1998. "Human cone spectral sensitivities: a progress report." Vision Research, 38, 3193-3206.
- Sumner, P. & J.D. Mollon. 2000. "Chromaticity as a signal of ripeness in fruits taken by primates." Journal of Experimental Biology, 203, 1987-2000.
- Teller, D. 2003. "Color: A vision scientist's perspective." Behavioral and Brain Sciences, 26, 48-49.
- Thompson, E. 2000. "Comparative Color Vision: Quality Space and Visual Ecology." In S. Davis (ed.). Color Perception: Philosophical, Psychological, Artistic and Computational Perspectives. Vancouver Studies in Cognitive Science, Volume 9. Oxford: Oxford University Press.
- Thornton, W. 1998. "How Strong Metamerism Disturbs Color Spaces." Color Research and Application, 23, 402-407..
- _____. 1999. "Spectral Sensitivities of the Normal Human Visual System, Color-Matching Functions and Their Principles, and How and Why the Two Sets Should Coincide." Color Research and Application, 24, 139-156.
- Trivers, R. 2006. Foreword to the first edition of R. Dawkins, The Selfish Gene. New York:

Forthcoming in *Erkenntnis*

Oxford University Press.

Tweed, D. 2003. Microcosms of the brain: What sensorimotor systems reveal about the mind. New York: Oxford University Press.

Tye, M. 2006. "The puzzle of true blue." Analysis, 66, 173-178.

Usui, S., S. Nakauchi, and M. Nakano. 1992. "Reconstruction of Munsell color space by a five-layer neural network." Journal of the Optical Society of America A, 9, 516-520.

Van Gulick, R. 2003. "Beautiful red squares." Behavioral and Brain Sciences, 26, 50-51.

Verhulst, S. & F.W. Maes. 1998. "Scotopic vision in colour-blinds." Vision Research, 38, 3387-3390.

Wright, W. 2009. "The physical unnaturalness of Churchland's ellipses." Philosophy of Science, 76, 391-403.

_____. submitted-a. "Prime colors and the hues."

_____. submitted-b. "Color constancy reconsidered."

Zaidi, Q. 1998. "Identification of illuminant and object colors: heuristic-based algorithms." Journal of the Optical Society of America A, 15, 1767-1776.

_____. 2001. "Color constancy in a rough world." Color Research and Application, 26, S192-S200.

Zimbardo, P. & R. Gerrig. 1998/2002. "Perception." In D. Levitin (ed.) Foundations of Cognitive Psychology. Cambridge, MA: MIT Press.