Comments on Wayne Wright's 'Opponent processing and the physical basis of color'

Presented at the 2008 Annual Meeting of the Southern Society for Philosophy and Psychology, March 2008.

Whit Schonbein
Department of Philosophy / College of Charleston
schonbeinw@cofc.edu

Two reflectances are metameric with respect to the human visual system if they elicit the same perceived color. Metamers pose a problem for color realism because metameric reflectances have nothing in common other than the fact they elicit the same perceived color, thus defeating attempts to identify color with reflectance profiles. In his (2007), Paul Churchland claims to have found a solution to this problem: All metameric reflectances share the same objective 'canonical approximation'. These approximations are calculated by taking standard 2D representations of reflectance profiles, rolling them to form cylinders, and slicing them with a plane (see Churchland (2007) for details). The results are ellipses, and these form his canonical approximations. His proposal, then, is that (i) all metamers share the same ellipse, and (ii) ellipses are objective facts about reflectance profiles. In this way Churchland aims to circumvent the problem of metamers.

In his commentary, Wayne Wright argues that Churchland's ellipses are not objective, in the sense that, when considered independently of the human visual system, the canonical ellipses are arbitrary mathematical constructions – they have no objective status from the perspective of physics.

Wright's analysis is quite convincing, and further reflection on Churchland's project only leads to more critical questions. Consequently, my goal in these comments is to articulate these additional concerns about Churchland's method of canonical ellipses. In section one, I extend Wright's criticism to show that the ellipses can only be understood as relative to a species, and this in turn has important consequences for Churchland's argument. In section two, I scrutinize Churchland's ellipse methodology in more detail, arguing that the proposal can be tested by comparing it with other established methods for calculating what all metamers have in common. Finally, in section three, I consider whether Churchland can sacrifice the objectivity of ellipses, concluding that his domain portrayal semantics renders such a move unavailable.

1. The Species-Relativity of CA Ellipses

As Wright notes, Churchland takes CA ellipses to be objective features of the world. Churchland is very clear about this claim, writing, "the [CA] for a given [reflectance] profile is an objective fact about that profile and about the material object that possess that profile. Its specification makes no reference to the human visual system nor to the nature of its phenomenological responses to anything." (p. 130)

In reply, Wright argues that CA ellipses are *not* objective from the point of view of physics. He nicely summarizes his argument as follows:

> Physics alone does not offer a reason for thinking that any one of the ellipses cut from the cast number of cylinders that could be created is more "canonical" or "natural" than the others. Different choices of wavebands, any of which is physically appropriate, used to construct CA ellipses will induce different equivalence classes on reflectances. Thus CA ellipses are irrelevant to physical theory" (p. 7)

The issue is that there is simply no purely physical justification for preferring one range of wavelengths to another when constructing CA ellipses. What Wright does not investigate (in this short paper, at least) is what why Churchland chooses the range that he does. The answer, I think, is that Churchland tacitly appeals to the response characteristics of the *human* visual system. Consequently, his claim that CA ellipses are not specified relative to the human visual system also is false.

The tacit species-relativity manifests itself in at least two ways. First, Churchland considers only reflectance profiles for wavelengths visible to humans.[1] Given that other types of visual system are sensitive to higher or lower wavelengths[2], and since CA ellipses are relative to wavelength range, it follows that Churchland's method is broadly species-relative.

Second (and more importantly), metamers are *defined* relative to the response of a given type of visual system: A set of reflectances is metameric just in case they elicit the same perceived color in a given type of visual system. Churchland considers *only* reflectances that are metameric relative to the *human* visual system. So the tacit species-relativity of CA ellipses resides not only in the choice of wavelength range, but also in the fact that, within that range, the choice of reflectance profiles is determined by how the human visual system responds to those profiles. So I think we can expand Wright's criticism a bit: Not only is there is no objective physical reality to the ellipses, but also what justification there is derives from facts about the human visual system.

This raises a problem. Churchland claims that the human visual system tracks CA ellipses, and offers as support the observation that his method generates similar ellipses for metameric pairs of reflectances, thereby suggesting that the pairs elicit the same perceived color *because* they instantiate the same ellipse. But even if the ellipses have an objective physical reality, Churchland must also show that there are no profiles with the same ellipse yet perceived as being a *different* color, for in such a case the visual system would fail to track. That is, you cannot simply say the visual system tracks ellipses because all metamers have the same ellipse; you also have to show that there are no reflectances with the same ellipse that are *not* tracked, and by focusing only on

---

[1] I do not recall Wright noting this explicitly, but it seems to be a 'corollary' of his point that we may have reason to take into account reflectances across wider ranges of wavelengths.

[2] For example, some raptors are sensitive to wavelengths in the ultraviolet range, such as are reflected by vole urine.

reflectances already known to be metameric to humans, Churchland's method ignores this requirement.[3]

2. How should we understand Churchland's proposal?

The concern just raised has to do with his method for calculating ellipses: Does it ever treat a non-metameric reflectance as metameric? Other issues emerge on further examination. As Churchland and Wright point out, the features of the ellipses – rotation, tilt, and vertical location – correspond to hue, saturation, and brightness, respectively (Wright, p. 4, Churchland, pp. 132-3). When the cylinder is unrolled, we get a curve whose features – phase, slope, and amplitude[4] – correspond to those properties. Notice that this method for representing hue, saturation, and brightness is a standard one (e.g., figure 1)[5].

Now, take all metamers in an equivalence class. They all have the same perceived color – the same hue, brightness, and saturation. So it stands to reason that there is some curve on the 2D plot, common to all members of the class, representing this information. This means that the fact that there *is* an ellipse (which is just a curve) common to all members of the class should not surprise us – it is to be fully expected given the definition of a metamer and how color information is represented on a standard 2D plot. What *is* surprising, I suggest, is that Churchland's method succeeds in finding (an approximation of?) this curve.

Why? Consider dominant wavelength. The dominant wavelength of a reflectance is that wavelength corresponding to the perceived hue of the reflectance. Therefore, all metamers in the class have the same dominant wavelength. This suggests (*prima facie*) that, since Churchland's method solves for hue, it also solves (at least in part) for dominant wavelength. We can then compare Churchland's method with other known methods for calculating dominant wavelength to assess the success of his method.

---

[3] One possible reply is to argue that, if there are such *prima facie* failures in tracking, they occur for principled reasons. This seems to be Churchland's strategy in section 4 of his paper, where he argues that the differences between the space of possible ellipses and the human phenomenological color space can be accounted for by appealing to features of the human visual system (e.g., his discussion of the limits of transducer sensitivity on p. 134). Regardless, my concern is directed at his *method* for calculating ellipses given reflectances – what guarantee do we have that, by using his method for calculating ellipses, some reflectances do not result in an ellipse that should be in a different region of ellipse space if it is going to correspond correctly with phenomenological space?

[4] Typically, brightness is defined as either the amplitude of the curve or the area under the curve. It may be that the latter is a better interpretation of Churchland's method, but nothing turns on this point, as far as I can tell.

[5] In figure 1, all three curves have the same hue (the peak of each curve is in the same location, i.e., the phase is the same) and the same brightness (the amplitudes are the same), but their saturation differs (red curve = less saturated, blue curve = most saturated). In ellipse terms, they have the same rotation, amplitude, and different tilts. So both simple 2D charts of wavelength distributions and Churchland's ellipses represent the same information using the same parameters. Now, Churchland's method for curve construction cannot construct the curves in figure 1 – among other things we would have to allow the center of the ellipse to go below the 0% reflectivity threshold. But the general point still stands – Churchland's ellipses use the same basic representational features as standard 2D plots. (Figure 1 taken from http://whatis.techtarget.com/definition/0,,sid9_gci212262,00.html.)

On first glance, that it successfully solves for dominant wavelength is doubtful: Dominant wavelength is calculated relative to visual system type (e.g., by using the CIE color space, which is calibrated to the human visual system), and a single reflectance curve on one of these 2D plots does not carry information about visual system type (aside from assumptions about visible wavelengths). It would therefore be quite surprising if Churchland's method successfully solves for dominant wavelength.[6]

One might think that Churchland can claim he's not out to calculate dominant wavelength. But here's the rub: *It is essential for the success of Churchland's proposal that his method calculate dominant wavelength*. On his account, the human visual system tracks ellipses. So, as he argues in his paper, the space of possible ellipses corresponds to human phenomenological color space, i.e., the space of possible perceived colors. Consequently, the *hue* of an ellipse corresponds to *perceived hue*. Furthermore, the perceived hue of a reflectance corresponds to the dominant wavelength. Consequently, if Churchland's method does not calculate dominant wavelength, then the hue of an ellipse does not correspond to perceived color, and the visual system does not track ellipses.

While it may be doubtful that Churchland's method calculates dominant wavelength, my real point is more general: Churchland's proposal looks surprising on the surface, since it finds something in common to all metamers, but a proper understanding of the method indicates we should not be surprised. Furthermore, once we recognize his method as solving for the hue, brightness, and saturation shared by all metamers, we can compare it with other methods for calculating the same information. Doing so will reveal important information about his proposed method (e.g., how accurate is it?), and also allows us to *test* his thesis – if the ellipses do not correctly characterize those features known to be tracked by the human visual system as computed through alternative methods, then the ellipse hypothesis fails.

3. What happens if Churchland gives up his ellipse proposal?

Regardless of everything I've said thus far, Wright's original point still stands: Churchland's ellipses have no basis in physical reality.

One possible response for Churchland is to give up hope that there is some physically interesting sense in which all metamers are the same. This seems to be Byrne & Hilbert's strategy in their *BBS* target article (2003): They propose to identify perceived colors with *sets* of reflectances, where the criterion for set inclusion is how a reflectance interacts with a given visual system (B&H, 2003, pp. 10-11). They then concede that such sets are not formed by 'carving nature at its joints', saying, "the reflectance types that we identify with the colors will be quite uninteresting from the point of view of physics or any other branch of science unconcerned with the reactions of human perceivers." (p. 11). So, perhaps Churchland could sacrifice the claim that the ellipses

---

[6] For better or for worse, I picked up this info on dominant wavelength from Wikipedia. The details are a bit gorier, because the actual perceived color is a mixture of the dominant wavelength and the complementary wavelength, so Churchland's method is actually solving for this mixture. However, there are methods for calculating both of these wavelengths, and both types of wavelength are calculated relative to a visual system, so my general point still holds.

are objective, reserving them as an account of how the human visual system groups otherwise unrelated reflectances, while retaining his color realism (which, recall, holds that colors are to be identified with individual reflectances).

Unfortunately, this proposal runs afoul of Churchland's 'domain portrayal semantics' (Churchland, 2001, 2007).[7] According to Churchland, the appropriate way to conceive of the *format* of occurrent mental representations is not in terms of sentences constructed from a language-like system of symbols, but rather as points in a high-dimensional state space (such as one gets in the hidden layer of a connectionist network). These points of activation occur against a richly structured 'background map' of possible activations, and "[such] activations … code or index where, in the space of background possibilities comprehended by the map, the creature's current objective situation is located." (2007, p. 123). For example (and roughly speaking), the well-known classical color space is a three-dimensional background map of possible color experiences, and an occurrent color experience is a point of activation in that space (Churchland, 2007, p. 123). Other background maps describe the space of possible faces and possible limb positions (2001, p. 158).

The *content* of a mental representation is determined by the structure of the background map and its relation to the world.[8] To illustrate, Churchland offers an analogy with a GPS navigation system:

> what makes [the] stored map [in the GPS navigation system] a *portrayal* or *representation* of the entire urban area is the usual *relation-preserving, abstract, projective mapping* that makes any map a map. … It is a map because of its own internal structure, and because there is an abstract, relation-preserving mapping from the global street system surround to that stored internal structure. (2001, p. 155)

This 'abstract, relation-preserving mapping' is a homomorphism, as Churchland makes clear in his (2007):

> The basic virtue of such background maps – as with any map – is a structural homomorphism between the map-as-a-whole, on the one hand, and the entire feature-domain that it attempts to portray, on the other. The family of proximity relations that configure the many map elements of the brain's internal map must have a relevant homomorphism with the family of similarity relations that configure the many landmark features within the domain-to-be-portrayed." (2007, p. 123).

---

[7] Formerly known as 'state space semantics'.

[8] Actually, teasing out what determines the content of mental states on Churchland's theory is difficult. On the one hand, Churchland claims that his semantics is "strictly internalist" (2001, p. 127), but on the other hand, (i) the structured state spaces upon which his internalist account is built are the result of sustained causal interactions with the world (2001, pp. 138-9), (ii) the accuracy (truth?) of a mental representation depends on its relation to the world (2001, p. 156), and furthermore, (iii) there must be a homomorphism between features of the world and the structure of state space (2007, p. 123), each of which independently suggests an externalist component. Here I bracket these complications by focusing just on the homomorphism requirement.

Now, there are two issues that emerge from the above sketch of Churchland's theory of mental representation. First, as suggested in a quotation given above, maps are supposed to map *objective* features of the world (2007, p. 123; 2007, p. 158).[9] Specifically, the background map involved in perceived color represents objective features of the world: "the structured activation space of our own color-opponent neurons in the LGN or V4 is a (somewhat problematic) map of the range of possible *objective electromagnetic reflectance profiles*." (2001, p. 158, original emphasis). Even more specifically, this map is a portrayal of possible CA ellipses. (2007, pp. 130-3) But if Churchland drops the claim that CA ellipses are objective features of reflectances, then he concedes that the background map fails to map objective features of the world, and the human visual system doesn't track objective properties. So, relinquishing the objectivity of ellipses conflicts with Churchland's domain portrayal semantics.

Second, a similar issue arises from the homomorphism requirement. A homomorphic mapping from world to background-map is operation-preserving.[10] So, if we perform an 'operation' on the world, that 'operation' should be mirrored in the color map. For example, take an object O with a reflectance R, characterized by ellipse E, and with perceived color C. Now, perform an 'operation' on O. Since we are assuming that ellipses are not objective properties of reflectances, there is no such thing as an objective 'ellipse-operation' either; instead, the objective manipulation will be directly on the reflectance. If the visual system map is homomorphic with respect to reflectances, and we change the reflectance, then – by virtue of the operation-preserving nature of the homomorphism – the perceived color of O should change as well. But this does not always happen – metamers are a case in point. The issue is that the actual 'operations' that do exist objectively, such as switching reflectances, are not preserved in the color map – instead, the map represents different reflectances with the same perceived color. In short, if Churchland gives up the objectivity of ellipses, then his theory of mental representation conspires to resurrect the problem of metamers.

If this is correct, then Churchland is committed to there being some objective property of all metamers such that they are mapped by the visual system in the way they are, and he cannot adopt Byrne & Hilbert's solution without infringing on this commitment. To wrap up, I'll put two possible replies on the table. First, one could retain the color realism (in the form of identifying colors with individual reflectances), and also drop the homomorphism requirement, at least for phenomenological color space, thereby embracing the pervasive misrepresentation of color by the human visual system (most obviously in the case of metamers). One issue with this strategy is that we loose some of the motivation for color realism in the first place, since the 'solution' raises the question of why we should take our color experiences to be veridical at all.[11]

[9] So, for example, in his (2001) he writes, "And so, I suggest, is *every* structured neuronal activation-space within the brain an abstract map of *some* **objective** domain of features, structures, or processes." (p. 158). The bold highlight is mine.

[10] This follows from the definition of 'homomorphism'. Homomorphisms are by definition operation preserving.

[11] Byrne & Hilbert (2003, p. 17) seem content with pervasive misrepresentation *within* a color category within a species, but they also deny that misrepresentation occurs *across* species, where the color categories may be different (p. 16). That is, while I may misrepresent something as a shade of green when compared to other humans, I do not misrepresent it when compared to the color experiences of, say, bees, which represent it as, say, a shade of red. I may be misinterpreting their position – I found it hard to reconcile

A second option is to keep the homomorphism requirement while trying to find some other objective property of the world tracked by the visual system. One possibility is that the human color map portrays the space of *possible objective relationships between reflectances and the human visual system*. As an analogy, consider the taller/shorter-than relationship (or relative hight). First, such a relation is arguably *objective*. Second, we can easily imagine a 'map' of the space of possible taller/shorter-than relationships that may hold between two individuals; a 2-d space where each axis is the height of a person would do the trick, and each point in this space represents the heights of two individuals. Finally, a person could use this map to represent his or her *own* height relative to another individual. Similarly, perhaps the color map portrays the possible objective relationships that may hold between the human visual system and individual reflectances. An activation point in that space represents the objective fact that the individual's visual system is currently standing in a particular relationship with a reflectance. One issue with this proposal (and there are many) is that it threatens color realism insofar as perceived colors are not really of the colors of objects, but rather of the 'colors' of agent/reflectance relationships, which is a peculiar form of realism.

References

Byrne, A. & Hilbert, D. R. (2003). 'Color realism and color science'. *Behavioral and Brain Sciences*, 26, pp. 3-64.

Churchland, P. M. (2001). 'Neurosemantics: On the mapping of minds and the portrayal of worlds'. in K. E. White (ed.), *The Emergence of Mind*, pp. 117-47. Reprinted in some book by Churchland, CUP, 2007 (page numbers refer to the reprint).

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, pp. 119-149.

Wright, W. (2008). 'Opponent processing and the physical basis of color'. Paper presented at the 2008 annual meeting of the Southern Society for Philosophy and Psychology, New Orleans, LA, March 2008.

Figure 1:

their simultaneous denial of misrepresentation (p. 16) with their acceptance of misrepresentation (p. 17), and the above summary is my current best shot.

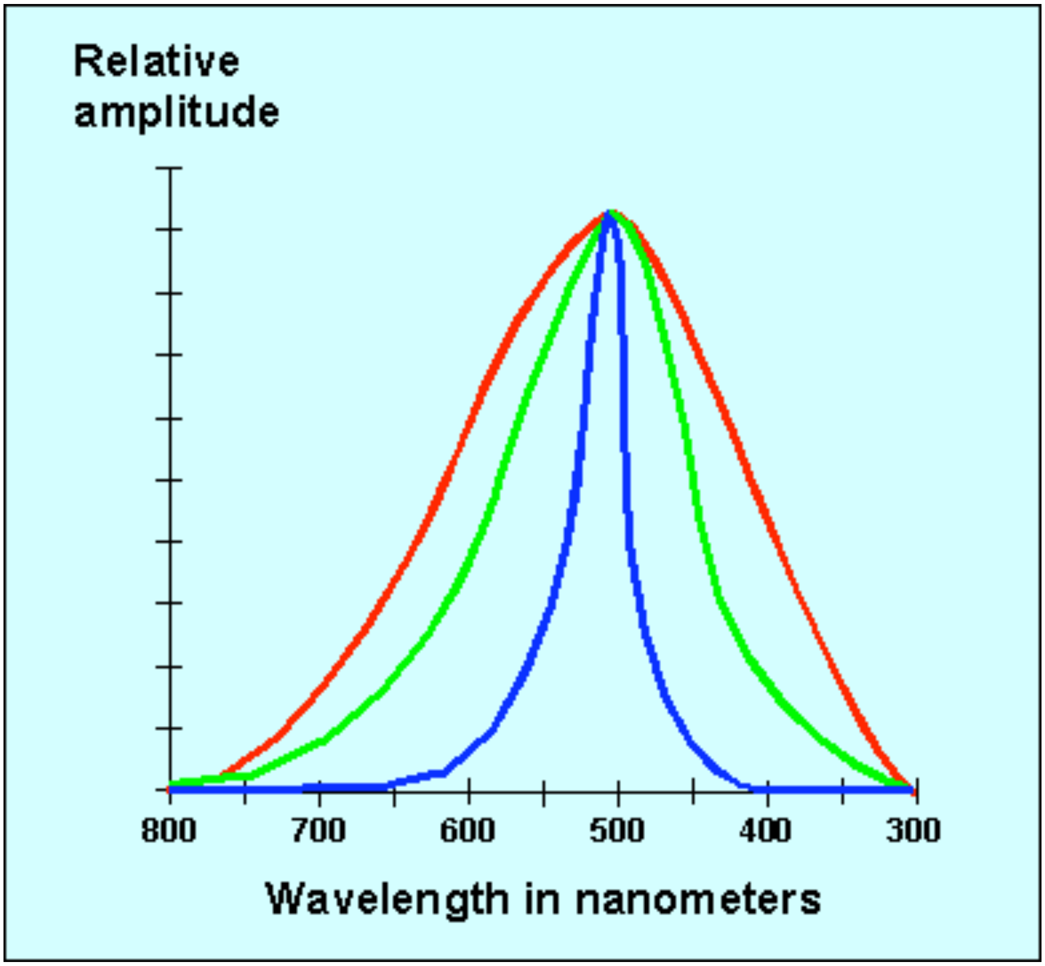
Relative
amplitude
800
700
600
500
400
300
Wavelength in nanometers