The Physical Unnaturalness of Churchland's Ellipses[1]

Wayne Wright
UCI – Dept. of Logic & Philosophy of Science * waynew@uci.edu

ABSTRACT: This paper addresses Paul Churchland's (2007) attempt to identify colors with surface reflectance spectra. Of particular concern is Churchland's novel method of approximating surface reflectance spectra. While those approximations are generated by objective means and yield a striking match with human phenomenological color space, they are not physically meaningful. The reason for this is that the method used to produce the approximations induces equivalence classes on surface reflectances that are not invariant under physically appropriate changes of measurement convention. This result undermines Churchland's proposed response to the objection from metamers commonly raised against color physicalism, as his surface reflectance approximations are supposed to provide an objective, physical unifying basis for metamers.

1. Introduction

Many philosophers and scientists have argued that, whatever else is to be said about color, we should not identify colors with physical properties; i.e., we should not be color physicalists. Nonetheless, color physicalism has supporters, particularly those with an externalist representationalist account of perceptual experience; see Byrne & Hilbert (2003b, p.52) and Lycan (2001, p.20). Color physicalists have shown great ingenuity in replying to their critics and have made considerable use of the empirical literature in attempting to carve out a physical structure with which to identify color. However, the leading color

---

[1] Parts of this paper have been presented in various forums and I have learned much from the feedback provided by those audiences. A particular debt of thanks is owed to Edward Averill, C.L. Hardin, Kimberly Jameson, Laurence Maloney, Louis Narens, A. Kimball Romney, Whit Schonbein, Michael Webster, and John I. Yellott.

physicalisms developed from such efforts might leave something to be desired; see the commentaries on Byrne & Hilbert (2003a) for some relevant criticisms. Most notably, they end up committed to the claim that the precise colors of things are unknowable to us; see Tye (2006, p.177) and Byrne & Hilbert (2003a, p.21n50).

This result is potentially problematic for at least two reasons. First, there is a long-standing case against color physicalism based on both philosophical and empirical considerations. An opponent of color physicalism could plausibly insist that any candidate physicalist theory of color has a high standard of proof to meet in order to be accepted. It is hard to see how such a high standard could be met by – or why we should have much interest in – a theory that tells us that colors are physical properties, but we cannot know which physical property a given color is. A second, related point has to do with the stated interdisciplinary aims of some advocates of color physicalism. For example, Byrne & Hilbert (2003a, p.4) write that the problem of color realism is significant not only to philosophers, but "to anyone working in the field of color science." A critic might contend that if one adopts a form of color physicalism that has it that we cannot know which physical property is which color, one is led to do so because a favored metaphysical ("physicalizing") theory of mental representation or consciousness depends on the truth of color physicalism, not because of any benefit expected to accrue to empirical inquiry; e.g., providing an account of physical color that can be put to use in inter- and intra-species assessments of veridicality and variation.

The need to embrace unknowable color facts stems largely from the choice, made by theorists such as Tye and Byrne & Hilbert, to identify colors with physical properties that are arbitrary from the standpoint of physics. The basic motivation for this choice is that it is commonly thought that a color physicalism grounded in a kind that is natural from the perspective of physics is ill-equipped to meet common objections to realist theories of color. In order to isolate the properties with which colors should be identified, “non-physics physicalists” have appealed to the idea that certain perceivers experience the true colors of things. For example, a given surface color turns out to be a set of surface reflectance spectra that have no unifying characteristic from the inventory of physics but that are perceptually equivalent to one another for a certain class of perceivers in certain viewing conditions. However, the widespread variations amongst color perceivers classified as normal by the usual methods (e.g., performance on tasks with the Ishihara color plates or the Farnsworth-Munsell colored disks) prevent these physicalists from specifying whose color experiences are genuinely veridical and thus which set of reflectances should be identified with a given color. There is also the difficult question of in what circumstances veridical color experience would take place. See Hardin (2003) for further discussion of these and similar difficulties for views like those of Tye and Byrne & Hilbert.

Paul Churchland (2007) steps into this breach, offering an account of color that is supposed to tell us exactly which physical property is which color and to provide a physical property that unifies the diverse set of surface reflectance spectra that are the same in apparent color. Churchland (2007, pp.147-148) makes plain that he cannot brook the epistemic limitations attached to views like Byrne & Hilbert’s. His basic strategy is to demonstrate that the classes of spectral reflectances treated as equivalent by the human visual system (i.e., metamers) correspond to classes of surface reflectances that are physically equivalent. If successful, this would give Churchland a powerful reply to a prominent objection to realist theories of color, a reply that avoids the epistemic limitations just canvassed. The following argues that the structure Churchland claims is supposed to unify a diverse set of reflectances that are same in apparent color is not, contrary to what he asserts, "of interest to a physicist" (ibid.).

2. Churchland's reflectance physicalism

Churchland identifies colors with surface reflectance profiles in the visible band of wavelengths; i.e., the 400 – 700 nm range.[2] A surface's reflectance <u>profile</u> or <u>spectrum</u> is its disposition to reflect a certain percentage of incident light at each wavelength across a range of wavelengths (Palmer 1999, p.123). C.L. Hardin (1988, p.64) argues that <u>metamers</u> provide a reason to reject the identification of colors with reflectance spectra. Metamers are physically different stimuli that are perceptually indistinguishable under a given illuminant (Palmer 1999, p.102). Some metamers – "strong" metamers – can physically differ from each other quite substantially; see Thornton (1998). Two surfaces metameric under one illuminant might not be metameric under another illuminant; note that Churchland's treatment of metamers does not address in any detail the illumination dependence of metamerism. One way of putting the problem metamers pose is that there is no objective, physical property that unifies a set of metameric reflectances. A set of metameric reflectances is "held together" only by the structure of our visual system; i.e., the only thing the metamers have in common is that we find them to be indistinguishable in certain viewing conditions. Thus while

---

[2] Churchland's treatment of "reflectance colors" is the sole concern of this paper, so I will not address his distinction between them and "self-luminous colors" (pp.144-147). I will correspondingly limit my remarks to the surface reflectance properties of objects. Nothing about these simplifications affects my argument.

perceived color typically depends greatly on surface reflectance, it cannot be identified with surface reflectance.

Churchland's first act in trying to overcome this objection is to change the manner in which reflectance profiles are formally displayed. Reflectance is usually plotted as a function of wavelength in a two-dimensional graph: wavelength runs along the x-axis and reflectance along the y-axis. Some examples of such representations of reflectance spectra are presented in section 4. Churchland picks up that flat figure and wraps it around to form a cylinder, with the left- and right-most ends (400 and 700 nm, respectively) of the original x-axis touching and the height of the cylinder indicating the percentage of incident light reflected that was labeled on the original y-axis. The reflectance curve is now projected on the interior surface of the cylinder. The key step is to determine a planar cut through the cylinder that satisfies two requirements:

(i) Equate the total area above the cut but below the upper reaches of the reflectance curve with the total area below the cut but above the lower reaches of the reflectance curve.

(ii) Minimize the total area of those two areas, relative to all cuts that would satisfy (i).

The cut will produce an elliptical cylindrical section. The ellipse that is the top cap of that section gives the "canonical approximation" (CA) of the original reflectance profile (p.130). Figures 4 and 5 of Churchland's article (pp.127 and 129, respectively) illustrate the procedure just described.

More precisely, the CA includes the rotational position of the highest point of the ellipse on the cylinder's vertical surface, the angle between the lowest and highest points of the ellipse on the cylinder’s vertical surface, and the altitude within the cylinder of the ellipse’s center point. Of course, once the prescribed cut is made through the cylinder, the cylinder can be unrolled to reveal an approximating curve in the traditional two-dimensional frame. Importantly, no reference is made to particular perceivers, human or otherwise, in the procedure for determining a CA ellipse; that a reflectance profile has the CA it does is an objective fact about it. Additionally, for every empirical reflectance curve, there is a unique CA and the same CA can correspond to different reflectance curves. Note that Churchland’s account is similar to the “Gaussian World” approach to modeling surface reflectances of Donald MacLeod and Jürgen Golz (2003), which is an alternative to the more established linear models approach; on the latter, see Maloney (1999) and Wandell (1995). For MacLeod & Golz, the position and height of a Gaussian function’s centroid and the function’s curvature play similar roles to the properties of CA ellipses highlighted by Churchland; see also Mizokami et al (2006).

With the objective CA ellipses in place, Churchland goes on to connect them with human color vision in order to secure his color physicalism. He projects all possible CA ellipses into a space with three dimensions, corresponding to the special properties of CA ellipses. Churchland observes that the resulting structure bears a stunning resemblance to a three-dimensional rendering of human phenomenological color. That match “[topographically] speaking … is the answer to a color realist’s prayer” (p.133). There are deformations present in human phenomenological color space that aren’t present in CA ellipse space, but this is no refutation of the proposal. All that shows is that our instrument for measuring CA ellipses has the sort of reduced performance at extreme levels and limits of resolution that invariably confront the design of all complex biological systems and measuring instruments more generally (pp.133-134, 141).

Churchland identifies colors with empirical reflectance profiles, not with CA ellipses. However, he contends that the visual system tracks CA ellipses and the objective, physical feature that unifies a class of metameric reflectances is the CA ellipse they all share (p.135). Contrary to what is widely thought, a reflectance's inclusion in a set of metamers is not dependent on the idiosyncratic functioning of the human visual system. If true, this would block the objection from metamers. Colors can be identified with surface reflectance spectra and we see objective, physical colors by means of our visual system's ability to measure their objective, physical CA ellipses. As for the traditional three dimensions of color experience, Churchland takes the rotational position of a CA ellipse's highest point to correspond to a surface's hue, the ellipse's tilt its saturation, and the ellipse's center's altitude its lightness.

3. The physical status of CA ellipses

Despite the impression that Churchland's neat, tidy manipulation of reflectance curves delivers color physicalism from the metamerism objection, the original problem still remains. CA ellipses are not physical, at least not in a way that is helpful to Churchland's desired brand of physicalism. While CA ellipses are generated by what can be granted are "objective means" – all that happens is a mathematical transformation on a representation of something that is clearly physical – that is no guarantee that the resulting construction has any interpretation that is natural from the standpoint of physics.

Churchland takes seriously Hardin's objection from metameric stimuli and has set for himself the task of identifying a physical property that will unify a family of metamers in a way that does not implicate the particular workings of our visual system. The following passage makes clear that Churchland imposes on himself – in order to avoid the epistemic limitations of views like Byrne & Hilbert's – the requirement that CA ellipses be "of interest to a physicist":

> Byrne and Hilbert acquiesce in the received wisdom that the family of metamers for any commonsense color displays no unifying intrinsic feature specifiable in purely physical terms. They should not have acquiesced to this claim, because we can indeed specify, in terms "of interest to a physicist," the feature that unites the family of metamers for a given commonsense color: they all share the identical reflectance-space ellipse as their canonical approximation. Moreover, this shared objective feature is precisely what gets mapped within the human subjective or phenomenological color space. Accordingly, we can see how the space of human color sensations counts as a structurally accurate map of an objective domain of properties, a real achievement by our visual system that remains either denied or unrecognized in their view. (pp.147-148)

The difficulty here is finding a physically natural structure that corresponds to whatever mathematical structure is generated by manipulating reflectance curves. All manner of operations could be performed on a standard surface reflectance curve to produce new mathematical objects. For example, one could re-draw the two-dimensional plot of reflectance between 400 and 700 nm by taking, at each wavelength, the cube-root of reflectance and multiplying that by the square of the surface's reflectance at the wavelength that is the same absolute distance from 550 nm, but in the direction opposite along the x-axis from the wavelength in question. The results of any such operations would surely be objective in some sense, but it is likely that few, if any, of them would find a home in physical theory.

Lurking in the background is the matter of just what it means for something to be physically interesting. In describing their own view, Byrne & Hilbert (2003a, p.11) write that the reflectance types with which they identify 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." In the extended passage quoted just above, Churchland emphasizes that the property he uses to unify a set of metameric reflectances – their shared CA ellipse – is supposed to be "of interest to a physicist." Unfortunately, in neither case is anything further said to unpack just how such talk should be understood.

While perhaps neither Byrne & Hilbert nor Churchland have this in mind, it would pay to attend to the notion of "physical meaningfulness" at work in physics and psychophysics. Central to this idea is that quantitative relationships that are invariant under admissible transformations of scale are physically meaningful, and those that are not invariant are physically meaningless; see Narens & Mausfeld (1992, pp.467-468), which is followed closely here. As Narens & Mausfeld observe, this criterion is commonly used in physics when deciding whether to include a proposed quantitative concept and it is often thought to dictate which psychophysical relationships should be considered. Simple examples of quantitative concepts that fail to be physically meaningful in this sense are temperature ratios based on Celsius or Fahrenheit scales. The ratio for both (i) 10 °C and 20 °C and (ii) 20 °C and 40 °C is 1/2. That ratio will not obtain – and (i) and (ii) will differ in their ratios – when the same temperatures are measured in Fahrenheit; (i) will be 50 °F and 68 °F while (ii) will be 68 °F and 104 °F. The basic idea is that the equivalence classes induced by a physically meaningful concept will not break if one adopts different conventions of measurement. In the light of both the well-established usage of this understanding of physical meaningfulness and the absence of alternative proposals from Churchland or other color physicalists, I will employ this understanding in the following. The key concern thus becomes whether the equivalence classes induced by Churchland's CA ellipses remain invariant when one changes – in physically permissible ways – the conventions by which they are determined.

In the case of the wraparound configuration Churchland employs, from the standpoint of physics alone, there is nothing about the 400 – 700 nm range in the electromagnetic spectrum or surface reflectance properties across that range that suggests that a structure with a physically natural interpretation can be extracted from the mathematical object Churchland creates. Note that this is not an attempt to reject Churchland's color physicalism merely on the grounds that it – and human color vision – privileges the 400 to 700 nm range over others. That our visual system has a limited range of wavelength sensitivity is by itself no problem for color physicalism. The point is instead that there is no physically natural relation that might obtain between either light at the extreme ends of the visible band or the dispositions to reflect a certain percentage of incident light at those wavelengths that is captured by depicting a surface's reflectance continuously (likely, piecewise) across 699, 700, 400, 401 … nm; i.e., there is nothing physically canonical about wrapping around the ends of the visible spectrum. Thus it is a matter of convention that the cylinder used to construct an approximating ellipse includes only the 400-700 nm interval. This is important, given the characterization of physical meaningfulness at work here.

For any surface, a reflectance profile can be specified that extends in both directions beyond the 400 – 700 nm interval; of course, one could also focus on a subinterval of 400-700 nm, intervals that include one of the extrema of the visible band but not the other, or intervals entirely outside the visible band. For example, a "reflectance cylinder" could be constructed by taking the range of electromagnetic radiation that reaches the earth's surface (the optical window; ca. 300 – 1000 nm) and wrapping it around so that 300 and 1000 nm abut one another. If the reflectances from that range were included in the cylinder used to construct an ellipse to approximate a surface's reflectance spectrum, that could considerably affect the altitude, tilt, and rotational position of the ellipse that is constructed. Relevant examples include the ultraviolet reflectance of some red flowers (Nassau 2001, p.350), the marked rise in reflectance beyond 700 nm for some fruits that are staples of some primates' diets (Sumner & Mollon 2000, p.1991), and the spike around 300 nm in the otherwise low and flat reflectance spectrum of short-tailed weasel pelts (Reynolds & Lavigne 1981, p.302); see also section 4. Thus a surface's reflectance in the visible band could be represented differently – in terms of the behavior of the portion of the ellipse's curve that passes through that region – depending on what range of wavelengths is used to construct a reflectance cylinder from which an approximating ellipse is cut. Two reflectance spectra that share the same approximating curve in the visible band (because they share the same approximating ellipse) when their ellipses are cut from cylinders using only 400-700 nm could very well have different approximating curves in the 400-700 nm region (because they have different approximating ellipses) when their ellipses are cut from cylinders using, for example, 350-900 nm.

In dealing with the metamerism objection and generating the striking figure for CA ellipse space (fig.3b on p.131), the values Churchland uses for the three highlighted attributes of CA ellipses – tilt, angle, and rotation – depend crucially on the fact that the base of the cylinder he employs is created by the physically arbitrary act of joining together the two extreme ends of the 400 – 700 nm band. Physics alone does not offer a reason for thinking that any one of the ellipses cut from the vast number of cylinders that could be created is more “natural” or worth singling out than the others. While the curve in a 2-D plot of surface reflectance depends only on the surface’s reflectance at each wavelength, a CA ellipse’s properties depend on relations that obtain between reflectances in the range of wavelengths considered; note that the relational character of CA ellipses runs counter to what is suggested by Churchland’s talk of a “unifying intrinsic feature” in the block quote above. Such relational properties may be interesting to us for various reasons (e.g., for our visual system to exploit in certain ways), but they result in different equivalence classes being induced on surface reflectances by different choices of wavelength ranges. Therefore, CA ellipse representations of surface reflectance are meaningless from the standpoint of physical theory.

Despite all that has been said, there is something physically salient about the 400 – 700 nm region of the electromagnetic spectrum. This might seem to justify Churchland’s approach. Specifically, it is in that “narrow band where the interaction of radiation with electrons first becomes important” and only in that range “is the energy of light well attuned to the electronic structure of matter” (Nassau 2002, p.31). Sunlight is also abundant in that band, although care needs to be taken in understanding this claim; see Lynch & Soffer (1999). Of course, some animals have infrared sensitivity or, much more commonly, vision that extends into the ultraviolet, but that can be set aside for present purposes. Key to recognize at this juncture is that the existence of a physically grounded rationale for why the human visual system samples a certain portion of the electromagnetic spectrum does not automatically confer a physical interpretation on the reflectance cylinders Churchland employs. To see this, one need only recall the earlier point about the open-ended number of mathematical operations that could be performed to re-draw a 2-D reflectance curve in the visible band.

In fact, what makes the 400 – 700 nm range physically interesting speaks against the idea that Churchland's wraparound configuration taps a physically natural property. As Kurt Nassau (ibid) puts it:

> Radiation at lower energies [i.e., infra-red] induces relatively small motions of atoms and molecules, which we sense as heat, if at all. Radiation at higher energies [i.e., ultra-violet] has a destructive effect, as it can ionize atoms, that is, completely remove one or more electrons, and can damage molecules permanently.[3]

So, just beyond one end of the visible spectrum, electromagnetic radiation has a relatively insignificant effect on electrons, whereas just beyond the other end it has dramatic effects. Each makes it useful for our visual sensitivity to extend only so far in a particular direction, but together they offer no insight into there being a physical basis for focusing exclusively on a continuous (piecewise) representation of surface reflectance across those two extremes. Churchland's choice of 400-700 nm as the base of his cylinder still lacks a physical rationale for being taken as canonical – as opposed to conventional – and thus the case presented against the physical meaningfulness of his approximating ellipses stands.

---

[3] Energy is inversely related to wavelength.

4. On the applicability of CA ellipses to empirical reflectances

Aside from the question of the physical status of CA ellipses that is main the concern of this paper, Churchland's framework looks to share a problem with the Gaussian World approach of MacLeod & Golz (2003); similarities between the two views were noted in section 2. Many empirical reflectance spectra fall into a pattern in which greens, blues, and purples (if one closes the visible spectrum back on itself as Churchland does) have a characteristic element that is roughly bell-shaped, like Churchland's figure 1.d (p.122). Such spectra are good candidates for useful approximation by Gaussian curves or sine waves; Churchland's ellipses amount to the latter for chromatic colors. Reds, oranges, and yellows, however, often have a sigmoidal shape, in which there is a low, basically flat component throughout the short and middle wavelength regions of the spectrum, with a fairly rapid rise beginning somewhere in the middle or middle-long region that eventually peaks and subsequently remains more-or-less flat out to (and to some extent beyond) 700 nm. These patterns can be found in the reflectance spectra for Munsell painted chips and natural reflectances (e.g., flowers, leaves, berries) measured by the University of Joensuu Color Group, some examples of which are displayed in figure 1.[4]

---

[4] These spectra are available on-line at http://spectral.joensuu.fi/index.php?page=database. The Munsell color system is based on the concepts of hue, value, and chroma. These attributes provide for a three-dimensional, sphere-like color solid:

- Value (lightness) is the vertical dimension. Steps along the axis run from 0 (black) to 10 (white).
- Radii from the achromatic value axis define different hues. Hues are denominated by terms such as 'red', 'blue', and 'green yellow', with numbered gradations within each hue category. There are ten total hues, each having 10 steps. Hue step 5 is routinely taken to be "most representative" of a hue category.

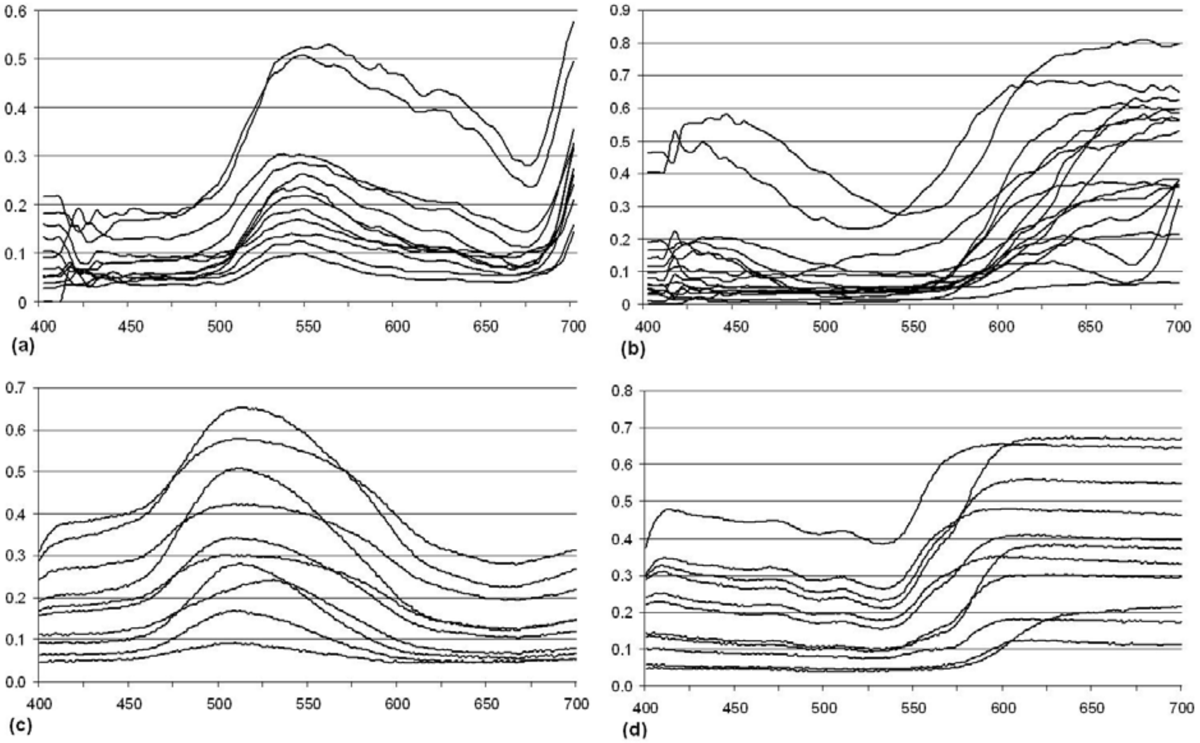

Figure 1: Sample reflectance spectra for (a) green and (b) red natural reflectances (colors indicated by the labels provided by the University of Joensuu Color Group for these samples) and (c) 5G and (d) 5R Munsell chips of varying value and chroma. Note that the two U-shaped spectra in (b) plot as red-purple, rather than simply red, in the Romney & Fulton (2006) "prime color" color appearance model.

- Steps along each radial spoke indicate chroma (roughly, saturation, level of purity). Stimuli with chroma zero are achromatic and different combinations of hue and value have their own maximum chroma.

The standard format for referring to locations in Munsell space is Hue Value/Chroma. For example, chip 2.5RP 6/4 has level 2.5 of hue RP (red-purple), value 6, and chroma 4.

Larry Maloney (2003, p.243) observes that “[it] is difficult to see how [MacLeod & Golz’s] model could fit the approximately step-function spectral reflectance functions that typically correspond to reddish or yellowish biological colourants;” Maloney cites Lythgoe (1979) and Chittka et al (1994) in support of his claim about reddish and yellowish reflectance profiles.

Essentially the same point applies to Churchland’s “canonical approximations” method. When 400 to 700 nm is used, the different lengths of the elements flanking the rising portion of strongly reddish and yellowish curves will play havoc with CA ellipses, given the requirement that the cut through the reflectance cylinder is supposed to equate the area above the cut and below the upper reaches of the reflectance curve with the area below the cut and above the lower reaches of the reflectance curve. Orangish spectra tend to be in the midst of their climb in the vicinity of 550 nm, resulting in lower and upper pieces of basically the same length, which makes it likely that a cut satisfying the equal area requirement will peak at an appropriate point. See Kuehni & Hardin (forthcoming) for specific examples of discrepancies between mappings into CA ellipse space and locations in perceptual color space related to this issue.

As can be seen in figure 1, some natural reddish reflectances have a “humped” shape in the long wavelength region, instead of the flatness characteristic of Munsell reds, yellows, and oranges; this is also true of some yellowish and orangish natural reflectances. A primary reason for this has to do with the balance between chlorophyll and other pigments in plants. For example, leaves with high concentrations of anthocyanins (e.g., cyanidin) or carotenoids (e.g., beta-carotene) and low concentrations of chlorophyll will show little if any modulation in the portion of their reflectance curve after their steep rise. As chlorophyll content increases, the reflectance spectra for leaves containing anthocyanins or carotenoids develop a "dip" near 678 nm, which is one of the absorption peaks of chlorophyll (the other is ca. 450 nm). See Lee (2007, chpt.3), Gitelson et al (2002), and Merzlyak et al (2003) for further relevant discussion of plant pigments and their reflectance and absorption properties.

Recognition of the role played by the absorption spectrum of chlorophyll in a vast range of natural reflectance patterns, allows further detail to be added to the earlier noted examples of reflectance spectra that might share the same CA ellipse when the 400-700 nm range is used but have different CA ellipses when another range of wavelengths is employed. Consider the two spectra in figure 2, one of which is a for a maple leaf with roughly balanced chlorophyll and carotenoid content (approximated from #4 in figure 2 of Gitelson et al 2002) and the other of which is for chip 5GY 5/6 from the Munsell set (taken from the University of Joensuu Color Group data set cited earlier).

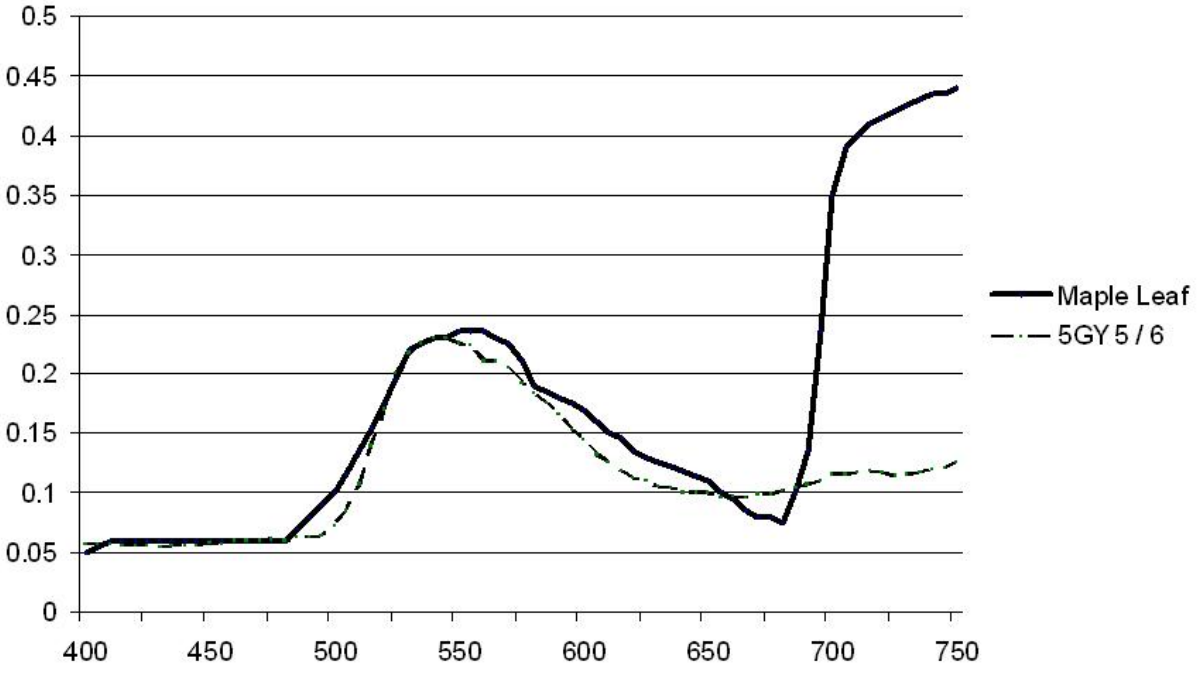

Figure 2: Reflectance spectra for Munsell chip 5GY 5/6 (dashed line; source: University of Joensuu Color Group) and a maple leaf (solid line; source: Gitelson et al (2002, p.275, fig.2).

These two spectra greatly overlap in the 400-700 nm interval and would have nearly identical CA ellipses in that range. They also plot very close to one another in the Romney & Fulton (2006) "prime color" color appearance model: in the three-dimensional color space they have a Euclidean distance less than that between chips 5GY 5/6 and 5GY 5/8, which differ only in two steps of chroma.[5] Despite these similarities, the significant differences in their reflectance spectra beyond 700 nm (and also out past the 750 nm cut-off shown here) would lead to very different CA ellipses being cut for them when one includes wavelengths longer than those of the visible band. Many examples of the same sort – involving Munsell chip and plant spectra that match (or nearly so) perceptually and have essentially the same CA ellipse in the 400-700 nm range while having very different reflectance spectra (and thus CA ellipses) outside that range – can be generated quite easily from the sources cited herein. The existence of such examples makes evident that the problems raised for Churchland's approximation method developed in section 3 are pervasive and rooted in empirical issues regarding the physical causes of color experiences.

5. References

Byrne, A. & Hilbert, D. 2003a. "Color realism and color science." Behavioral and Brain Sciences, 26, 3-64.

---

[5] Note that the Romney & Fulton model is used only for purposes of convenient illustration; the point made here does not depend on this particular choice. Other color appearance models could be used to determine a matching Munsell chip for the maple leaf in question. While another color appearance model might generate a (in all likelihood, only slightly) different result, nothing about differences in predictions between different color appearance systems has any bearing on the usefulness of this example against Churchland's claim about CA ellipses being "physically interesting".

_____. 2003b. “Color realism redux.” Behavioral and Brain Sciences, 26, 52-63.

Chittka, L., A. Shmida, N. Troje, and R. Menzel. 1994. “Ultraviolet as a component of flower reflections, and the colour perception of hymenoptera.” Vision Research, 34, 1489-1508.

Churchland, P. 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.

Gitelson, A. Y. Zur, O. Chivkunova, & M. Merzlyak. 2002. “Assessing carotenoid content in plant leaves with reflectance spectroscopy.” Photochemistry and Photobiology, 75, 272-281.

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.

Kuehni, R. & C.L. Hardin. “Churchland’s metamers.” Forthcoming in The British Journal for the Philosophy of Science.

Lee, D. 2007. Nature’s Palette: The Science of Plant Color. Chicago: University of Chicago Press.

Lycan, W. 2001. “The case for phenomenal externalism.” In J. Tomberlin (ed.) Philosophical Perspectives, 15. Cambridge: Blackwell Publishers.

Lynch, D. & B. Soffer. 1999. “On the solar spectrum and the color sensitivity of the eye.” Optics & Photonics News, 10, 28-30.

Lythgoe, J. 1979. The Ecology of Vision. London: Oxford University Press.

MacLeod, D. & Golz, J. 2003. “A computational analysis of colour constancy.” In R. Mausfeld & D. Heyer (eds.) Colour Perception: Mind and the Physical World. London: Oxford University Press.

Maloney, L. 1999. “Physics-based approaches to modeling surface color perception.” In K. R. Gegenfurtner & L. T. Sharpe (eds.) Color Vision: From Genes to Perception. Cambridge: Cambridge University Press.

_____. 2003. “The importance of realistic models of surface and light in the study of human colour vision.” In R. Mausfeld & D. Heyer (eds.) Colour Perception: Mind and the Physical World. London: Oxford University Press.

Merzlyak, M., A. Gitelson, O. Chivkunova, A. Solovchenko, & S. Pogosyan. 2003.

"Application of reflectance spectroscopy for analysis of higher plant pigments." Russian Journal of Plant Physiology, 50, 785-792.

Mizokami, Y., J. Werner, M. Crognale, & M. Webster. 2006. "Nonlinearities in color coding: Compensating color appearance for the eye's spectral sensitivity." Journal of Vision, 6, 996-1007.

Narens, L. & R. Mausfeld. 1992. "On the Relationship of the Psychological and the Physical in Psychophysics." Psychological Review, 99, 467-479.

Nassau, K. 2002. The Physics and Chemistry of Color: The Fifteen Causes of Color. New York: John Wiley & Sons.

Palmer, S. 1999. Vision Science: Photons to Phenomenology. Cambridge: MIT Press.

Reynolds, P. & J. Lavigne. 1981. "Visible and ultraviolet reflectance characteristics of arctic homeotherms." International Journal of Biometeorology, 25, 299-308.

Romney, A.K. & J. Fulton. 2006. "Transforming reflectance spectra into Munsell color space by using prime colors." Proceedings of the National Academy of Sciences, 103, 15698-15703.

Sumner, P. & J.D. Mollon. 2000. "Chromaticity as a signal of ripeness in fruits taken by primates." Journal of Experimental Biology, 203, 1987-2000.

Thornton, W. "How strong metamerism disturbs color spaces." Color Research and Application, 23, 402-407.

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

Wandell, B. 1995. Foundations of Vision. Sunderland, MA: Sinauer.