Chapter XXX

Drawing the Right Inferences About Color Constancy

1. What is color constancy?

In simple terms, color constancy is the stability of perceived color across changing viewing circumstances. For instance, a ripe orange sitting half in direct incandescent light and half in shadow looks to be uniformly colored all over its surface.[1] When the same orange is viewed entirely in direct sunlight and a passing cloud introduces a shadow over the whole scene, the color of the orange does not seem to change despite the shadow causing things to look different in some color-relevant way. For yet another example, after one has determined that the orange is ripe for picking because of its color appearance when viewed on the tree outside, when one views it a while later indoors under artificial lighting, the orange looks to have pretty much the same color it did when it was picked. The first case illustrates <u>simultaneous</u> constancy and the latter two are instances of <u>successive</u> constancy. The two specimens of successive constancy are distinguished by the degree of adaptation to the prevailing illumination achieved; little or none when the illumination changes quickly and a great deal when one has spent a few hours indoors with the lights on. I trust that the reader's everyday color experience is replete with examples of all three kinds.

While color constancy's ubiquity might make it seem unexceptional, it is a significant feat. The light reaching the eye from an object (the "color signal") depends on <u>inter alia</u> the

---

[1] I am, of course, idealizing about the uniformity of surface color (reflectance properties, etc) for natural objects and materials like oranges, skin, and lawns.

spectral composition of the illuminant and the object's surface reflectance properties; see Arend (2001) for a discussion of other factors. Thus color signals of very different physical character emanate from the orange under the different illuminants. Despite the visual system having direct access only to the varying and confounded light stimulus, the perceiver has the strong impression that the orange does not change color across the shift in illumination in all three cases described above. It is natural to suppose that this is due to the visual system somehow discerning from the color signals that the orange's reflectance properties hold steady while the illuminant does not, although perhaps a cognitive process is (also) involved. Color constancy is of great practical benefit. Without it, the instability brought about by illuminant variation would limit or complicate the use of color for purposes such as object recognition or assessing object properties linked to color appearance, such as ripeness of fruit.

Surprisingly, while the phenomenon has long been recognized, color constancy received relatively little attention from scientists throughout much of the twentieth century. This was largely due to the dominant approaches to vision research being ill-equipped to deal with it in any sort of sophisticated manner (Mausfeld 2003, p.396). Wyszecki & Stiles (1982, p.173) despaired over the then-current state of understanding of the phenomenon and passed over it quickly. Fortunately, a great deal of empirical research has been done during the last four decades, producing a rich body of experimental results and different theoretical perspectives. Philosophers have taken note of this work and their discussions of color and color experience now regularly cite the scientific literature on color constancy, with some philosophers examining it in detail; e.g., Cohen (2008), Davies (2016), and Hilbert (2005, 2012).

A basic issue that has plagued the study of color constancy is deciding how the phenomenon should be understood. While the characterization of color constancy given above is fine as a starting point, it is short on specifics. The missing details matter greatly to sorting out how best to direct experimental work and construct a theory of color constancy. In particular, more must be said about the format or mode in which the sameness of the orange's color is registered by the perceiver. Is it a matter of a sameness of phenomenal color across the changing viewing conditions? If so, an account of what aspects of phenomenal color remain constant is required, as there clearly is some difference in color appearance as the illumination on the orange varies. Several researchers have claimed that color constancy is an inferential achievement. The basic idea shared by these views is that although phenomenal color is not constant, an inferential process takes varying chromatic information as input and generates stable, non-phenomenal object color attributions as outputs; 'inference' is being used somewhat loosely here to gather together a range of different possible processes, without commitment to a full suite of logical apparatus being involved (Arend & Reeves 1986, p.1749; Hilbert 2005, p.149). Of course, one would want to know about the nature of the inputs to and outputs from that activity, as well as the principles that govern it. Perhaps the visual system itself performs the relevant inferences on "low level" sensory representations and includes the conclusions in later-stage perceptual representations. Such an arrangement suggests a "doubling up" of color that is not easily squared with how things seem in everyday experience. Another possibility is that the inferences are a matter of perceivers making explicit judgments based on inconstant phenomenal color experiences. Besides also inviting the suspicion that it involves two color attributions where we appear to only have one, this option faces the problem that color constancy seems so effortless and immediate. Introspection does not reveal such judgments being performed; certainly not on a scale comparable to the prevalence of color constancy effects. Thus we would be owed an explanation of how perceivers remain unaware of engaging in such cognitive acts.

This collection of options and concerns is not meant to be exhaustive; e.g., Foster (2003) questions whether color constancy exists at all. However, it is representative of real positions that have been staked out by both philosophers and scientists. It also is suggestive of the depth of the challenges facing any attempt to develop a rigorous understanding of color constancy.

In an earlier paper (Wright 2013), I argued that color constancy is a collection of different phenomena. A central feature of my account was a “dual representation” of object color. The idea was that object representations constructed by the visual system can include both phenomenal and non-phenomenal surface color attributions. These might be thought of as separate records in an “object file” stored in visual working memory (Kahneman et al 1992), with different tasks (perceptual sets, etc) leading subjects to tap one or the other entry. The entries might not agree about the color of the object. That possible disagreement, however, would account for the “sameness” and “difference” reactions noted above and is relevant to instruction effects in color constancy experiments discussed shortly. In situations in which a high degree of adaptation is achieved for the illuminant on the object being viewed (usually requiring tens of seconds to several minutes), color constancy is a matter of (approximate) sameness of phenomenal color and the non-phenomenal entry is either not accessed or not even created. This addresses one of the cases of successive constancy presented at the beginning of the chapter. For simultaneous constancy and successive constancy with limited adaptation, I held that phenomenal color can show little constancy but the non-phenomenal color entry indicates stability and guides typical subject responses to the scene before the eyes. The hypothesized non-phenomenal color attribution was said to be grounded in inferences performed by the visual system that are dependent on both (i) low-level sensory processes sensitive to relations between reflectance properties of different surfaces and (ii) stable phenomenal color (or, at least, information that underpins stable phenomenal color) for an object under an illuminant to which the perceiver is adapted.[2]

Subsequent reflection and new empirical findings have led me to further develop and refine my view. In particular, I have changed my mind about the proposed process of visual inference. I now contend that the immediate sense that we are engaging with a world of stable colors in circumstances of limited adaptation (whether successive or simultaneous) is largely grounded in the output of the just-mentioned low-level sensory activity alone; see Linnell & Foster (1996, p.227). The signals generated by this low-level activity have an associated phenomenology (at least in some cases), but this does not implicate any sort of stability of perceived color. The inferentially-achieved perceptual representation of object color has been dropped. To the extent that perceptually represented color makes a contribution when adaptation is limited, it is a matter of either variations in phenomenal color still enabling consistent assignment of objects to broader color categories (e.g., green, orange) or explicit reasoning based on the color appearance of differently illuminated objects; regarding the latter, this might be a uniquely human accomplishment, as there is scant evidence to suggest that non-human animal color constancy might involve anything other than low-level sensory adaptation (Kelber & Osorio 2010, p.1620). This chapter sets out the case for my account and compares it to others on offer. I will wait until chapter XXX to address at length the consequences of my position on color constancy for accounts of the nature of color.

---

[2] For simultaneous constancy, this would be one of the illuminants present in the scene at that moment; e.g., one is adapted to daylight and the orange rolls partway into a shadow. For successive constancy, this would (typically) be the illuminant before the shift takes place.

2. Asymmetric matching and instruction effects

Color constancy studies frequently employ an asymmetric matching task. The stimuli for these experiments are pairs of patterns of surfaces under different illuminants, presented either simultaneously or successively. The stimuli tend to be two-dimensional arrangements of colored patches displayed on a computer monitor (many of which are called "Mondrian patterns," due to their resemblance to the Dutch painter's rectilinear compositions), but more naturalistic stimuli have been used, employing either three-dimensional figures rendered on computer monitors or physical objects. The collections of surfaces comprising the two patterns typically are held fixed, except for one patch. Subjects are given control over the chromatic properties of that patch in one of the patterns and are tasked with adjusting it so that it matches in color with its counterpart in the other pattern. Time limits for making matches are rarely imposed.

The color constancy index (CCI) of Arend & Reeves (1986) is the standard for summarizing subject performance in color constancy experiments. For asymmetric matching, this index involves a comparison of differences between (a) the subject's match setting and a perfectly color constant match setting and (b) the patch being matched against and a perfectly color constant match setting. The values for (a) and (b) are determined using Euclidean distances between locations in a chromaticity diagram or color space, such as those of the CIE. These quantities are adapted in appropriate ways when tasks other than asymmetric matching are used. The formula for the CCI is 1- a/b. A CCI of 1.0 corresponds to perfect color constancy and a CCI of zero to a total failure of color constancy; it is worth noting that values greater than 1.0 and less than zero are also possible. While the CCI's rationale is simple enough to grasp – it uses a subject's match settings to assess the degree to which changes in the proximal light stimulus induced by the illuminant shift have been offset (discounted, etc) – there is not an easy path from experimentally-determined CCIs to substantive claims about color constancy. Like any index, the CCI must be properly linked to the phenomenon it is supposed to summarize with a single number. In particular, questions about experimental tasks and conditions need to be addressed before any real progress can be made; see Foster (2003).

Arend & Reeves (1986) found that subjects show different levels of performance in an asymmetric matching task, depending on the matching instructions they are given. A number of other studies have replicated this finding, such as Arend et al (1991), Bäuml (1999), Cornelissen & Brenner (1995), Radonjic & Brainard (2016), Reeves et al (2008), and Troost & de Weert (1991). When subjects are told to make the patch under their control match its counterpart in hue and saturation ("appearance," "phenomenal," or "hue/saturation" match), CCIs are quite low (routinely ≤ 0.3, although some experiments find CCIs ca. 0.4). However, subjects instructed to set their match so that the two patches look as though they were "cut from the same piece of paper" ("surface," "paper," or "material" match) attain CCIs in the range of 0.6 to 0.9.

Even if attention is limited to the higher end of subject performance, color constancy is only approximate and falls quite short of the mark one is likely to expect based on introspection. The variation in subject performance with different instructions provides a serious challenge to the introspective sense that our color experience is stable and raises questions about what color constancy really amounts to. One natural interpretation of this pattern of results is that it shows that color constancy is not at all a matter of stable color appearance, as phenomenal color can exhibit little or no constancy while responses to surface reflectance properties demonstrate reasonably good constancy. In that case, we would need an account that both captures what is driving performance when good constancy is achieved and makes sense of the seeming stability of the visual world despite the inconstancy of phenomenal color.

Inferentialist accounts look to gain support from the instruction effects for asymmetric matching. According to these views, color constancy involves inferences to stable object colors despite potentially significant variation in phenomenal color. Since the colors of things typically do not seem to change as the lighting varies, it is tempting to conclude that this inferential component – and not phenomenal color, however dazzling it is at times – is at the core of our usual mode of perceptually engaging with the world. One might suppose that the dominance of the inferential constancy response is not noticed due to its phenomenal silence, effortlessness (perhaps automaticity), and pervasiveness. It is only in special circumstances – such as color constancy experiments or when attempting to paint realistic landscapes – that one becomes aware of the inconstancy of phenomenal color. In fact, it can be somewhat difficult to come to terms with the inconstant phenomenal aspect of color experience, as novice painters struggle to "paint what they see, not what they know" and researchers report that subjects in training protocols for color constancy experiments "[need] several iterations of the displays to convince themselves that the same hue and saturation could imply a different material" (Reeves et al 2008, p.222). While Berkeley claimed that the inferences were consciously made by perceivers, it has been much more common for contemporary theorists to hold that the inferences are performed unconsciously, perhaps within the visual system. Cohen (2008), Helmholtz (1924), and Reeves et al (2008) are representative of "unconscious inference" views. Henceforth, I will refer to inferentialist views that rely on conscious inferences as "cognitive" and those that appeal to unconscious inferences as "visuocognitive" or "projectivist."[3]

[3] 'Projectivism' has been used by some researchers to designate unconscious inference views; see Reeves et al (2008, p.226). 'Projection' has the advantage of not immediately suggesting a

Fundamental problems for visuocognitivist interpretations of asymmetric matching performance will emerge shortly. For now, it is useful to address the difference between the matching tasks. As hinted at a moment ago, some researchers have remarked that they or subjects find the hue/saturation task “unusual” or difficult to distinguish from surface matching; see Brainard et al (1997, pp.2093-2094), Hilbert (2005, pp.151-152,157n.7), Reeves et al (2008, p.222). I will soon explain why whether the two tasks differ in a meaningful way depends on experimental conditions and, to a likely lesser extent, subjects’ levels of training and instruction. Findings from Cornelissen & Brenner (1995) are relevant to whether subjects are able to

---

mechanism that deals in premises, rules, and conclusions. As noted earlier, my use of ‘inference’ is intended to encompass a broader range of positions than those limited to strictly logical resources. ‘Projection’ also captures the external attribution component of the views I have in mind. One drawback of the term is its philosophical baggage. The traditional philosophical use applies to sensations (phenomenal appearances) and is tied to a view on which visual qualities (e.g., color, shape, texture) are non-representational, intrinsic features of our experiences themselves that objects are erroneously seen as bearing; see Boghossian & Velleman (1989/1997). Such properties are not (putatively) objective properties of surfaces (etc.) that are represented in our experience, but are instead properties of a subjective object of awareness such as a visual field. While some of Reeves et al’s (2008) remarks about projection give the impression that they have the standard philosophical usage in mind, the broader context of their paper (and their work in general) makes clear that their understanding of projection concerns material properties of objects and is distinct from color sensations (appearances, qualia, etc.). In what follows, I will follow their use of the term.

appreciate a genuine difference between the tasks in the appropriate experimental conditions. In their study, which showed a clear instruction effect on CCIs, subjects' eye movement patterns differed greatly between the two tasks. While the question of whether the change in scan paths is necessary for better color constancy is unanswered by their work (ibid., p.2447), key for now is that the difference in viewing strategies can be readily explained by subjects' grasp of a difference in the two tasks. It should also be pointed out that Cornelissen & Brenner (ibid., pp.2438-2439, 2443-2445) found that the differences in scan paths do not explain differences in CCIs across the two tasks in terms of greater chromatic adaptation being facilitated by more time spent looking at the target patch in the surface matching condition than in the appearance matching condition.

Delahunt & Brainard (2004, pp.72-74) reported only a slight effect of different task instructions on achromatic adjustments; the achromatic adjustment method requires subjects to set the chromaticity of a test patch so that it appears achromatic. For asymmetric matches, a substantial instruction effect has turned up for both successive and simultaneous presentations. However, Brainard et al (1997, p.2105) found fairly good constancy (mean CCI of 0.61) for simultaneous asymmetric matching using "nearly natural" stimuli with an appearance-based task. I will ignore basic differences between asymmetric matching and achromatic settings, instead focusing on other matters that are relevant to subject performance with appearance and surface instructions.

The first issue concerns illuminant adaptation. Delahunt & Brainard (2004) had subjects adapt to their experimental images for one minute before adjusting chromaticities, whereas Reeves et al presented the patterns in their successive pairs for one second each, with no interval between. Reeves at al (2008, p.228n.1) also note that in a previous study (Arend 1993, p.2141) appearance-based settings with full adaptation achieved good constancy (CCIs ca. 0.6 - 0.7); see also Arend & Reeves (1986, p.1743). The visual system includes fast and slow mechanisms that adapt its responses to the color signal, based on the mean and variance of the intensity and chromaticity properties of stimuli; see Schultz et al (2006, p.1103) and Webster & Mollon (1995, p.694). Such mechanisms re-scale receptoral and post-receptoral signals in order to compensate for the effects of different illuminants. One might wonder whether such low-level adjustments alone underpin color constancy. Crucially, a complete normalization of visual signals by adapting low-level responses so that the signals from a surface under one illuminant are equated with those from the same surface under a standard illuminant, is undesirable. It would eliminate at the front end of the visual system ecologically valuable information about illuminant properties; see Smithson (2005, p.1341).

The second point also deals with illumination, in this case the structure of illumination within a scene. In Brainard et al (1997), subjects viewed target and match stimuli mounted on different ends of a wall that was illuminated in a way that changed incrementally from one side to the other. They observe (ibid., p.2096) that this illuminant gradient is akin to what one might find in an outdoor scene in which the relative contributions of diffuse blue skylight and directional yellow sunlight vary gradually across the scene. It is unlike the abrupt illuminant changes characteristic of going from sunlight to shadow or turning on tungsten lighting in a room that had been lit only by daylight seeping in under lowered shades. Neither the successive displays of Reeves et al (2008) nor the simultaneous displays in experiments such as those of Arend & Reeves (1986) involve a gradual illuminant change. A gradual illuminant change allows greater adaptation than is possible with the sharp illuminant shifts typical of other asymmetric matching experiments; see Bäuml (1999, p.1532).

It is noteworthy that with their simultaneous displays, Arend & Reeves (1986) employed a condition in which the test and standard Mondrian patterns were each surrounded by a thin strip composed of their illuminant color; the annuli did not overlap one another. Subjects were told that each annulus was indicative of the illumination falling on the Mondrian it surrounded. Even with this information about the illuminants explicitly available both cognitively and for visual processing to exploit, subjects making appearance matches showed low levels of constancy on a par with those made with Mondrians lacking surrounds. This supports the idea that adaptation has a central role to play in supporting greater appearance-based constancy. However, it does not entail that subjects do not register illuminant properties and that phenomenal constancy performance is never influenced by illuminant estimates made by the visual system. It is reasonable to think that subject performance in Brainard et al (1997) was affected by an illuminant estimate, given subjects' reports about their subjective impressions of the test and match stimuli under the illuminant gradient; see Brainard et al (1997, p.2098) and Maloney and Yang (2003). Related issues will be discussed later in the chapter.

The final point has to do with subject instruction. Brainard et al (1997, pp.2093-2094) simply asked subjects to make a color match. Delahunt & Brainard (2004) gave subjects instructions and a demonstration specific to the one task they were assigned. Reeves et al (2008) ran all their subjects through instructions and training that addressed both tasks. Subjects received multiple demonstrations of hue/saturation matching, due to the earlier-noted difficulty subjects had in accepting that different material properties could cause experiences that are identical in apparent color. Only when those demonstrations were understood did the training regimen continue. While illumination-related factors likely account for most of the divergence between these studies, differences in subjects' level of familiarity with the tasks should be considered, too.

Reeves et al (2008, p.220) tie illuminant adaptation to their conjecture that there are two constancy processes. When the perceiver's visual system has achieved a high degree of adaptation to the illuminant falling on the scene they are viewing, phenomenal color is approximately constant. This might happen as one takes a long walk outdoors and the daylight spectrum slowly changes over time or gradually across the scene, or when one has come inside and spent some time under the artificial lighting in one's home. When adaptation is minimal – as when the sun suddenly breaks through the clouds on a grey day or when one is viewing simultaneously presented Mondrians in an asymmetric matching experiment – the idea is that phenomenal color varies but an inferential mechanism achieves constancy. Thus the higher CCIs for the appearance-based tasks of Brainard and his colleagues, as well as other matching experiments with full adaptation, are likely explained by the operations of mechanisms that play little or no role in other matching experiments. This also applies to the study of Schultz et al (2006) that found good constancy in a hue-scaling task, as they used presentations blocked by illuminant condition and facilitated adaptation with a training phase before each illuminant block. There is no real conflict between results such as those of Delahunt & Brainard and Reeves et al; see also Delahunt & Brainard (2004, p.74), Kuriki & Uchikawa (1996, p.1634), and Thompson (2006, pp.85-86, fn.15).

The role of illuminant adaptation should be kept in mind when mulling over how the instruction effects debate has played out, especially amongst philosophers; this point is echoed by Davies (2016, p.546). Instruction effects in asymmetric matching experiments do not make for a decisive showdown between inferential and phenomenal conceptions of color constancy.

Some philosophers who have engaged with the empirical literature on color constancy seem to have overlooked that these different forms of constancy can coexist and hunker down in support of one kind of account or the other. This is evident in Hilbert (2005) and Cohen (2008). Hilbert goes to considerable length to defend a phenomenal conception of color constancy in the face of instruction effects (even suggesting that subjects tasked with making hue/saturation matches are instead matching to other dimensions of phenomenal color, despite the training regimens used in several studies showing instruction effects) while also seeking to undermine the metaphysical and phenomenological coherence of inferentialism; see Hilbert (2005, pp.148,152,156). Cohen expresses unqualified doubt about phenomenal stability in all instances of color constancy and presents his inferentialist account as applying to color constancy in general; this is evident in the remarks of Cohen (2008, pp.79-85). The debate over color constancy looks as though it often proceeds from the assumption that color constancy is a unified phenomenon for which inferentialism and phenomenalism are binary options. Like Davies (2016), I believe the situation is much more complex than is suggested by either pure inferentialism or pure phenomenalism, although we differ over important details of in what that complexity consists.

3. Categorical constancy and the sense of stability

Let us focus on subject performance in conditions that limit adaptation. Consider the earlier examples featuring an orange or, my favorite example to use when talking about color constancy, a lawn that is part in sunlight and part in tree shade. When viewing the lawn, one does not have the impression that the grass changes color or material properties as the illumination quickly shifts from yellowish and bright to bluish and dark. The orange does not look to take on a new color when a cloud passes by and obscures the sun. What explains our strong sense that the colors of things remain stable in such cases, despite the variation in phenomenal color? My account has three elements: categorical constancy, stable cone ratios, and explicit reasoning. Categorical constancy is the subject of this section and the latter two are addressed in the next two sections.

Although appearance-based asymmetric matching performance is routinely low, phenomenal color can still contribute to the everyday sense that surface colors do not change with the illumination. Troost & de Weert (1991, p.596) observe that from the standpoint of understanding the contribution it makes to vision in naturalistic settings, much of color constancy's value plausibly stems from it enabling successful interaction with the world by avoiding erroneous color attributions. In that case, the modest constancy found for appearance matches turns out to be more helpful than might initially be expected, as it facilitates quickly achieving a good deal of categorical color constancy. Categorical color constancy is the stability of the color category to which an object is assigned across different viewing conditions. This sort of constancy could obtain while an object's more fine-grained color appearance varies considerably.

As part of a study in which they also replicated Arend & Reeves' (1986) instruction effects for simultaneous asymmetric matching, Troost & de Weert (1991) had subjects perform a color naming task for 144 different targets viewed under four different chromatic illuminants and a white illuminant used as the standard against which naming responses in the chromatic illumination conditions would be evaluated. Each stimulus consisted of a target disc presented against a background that had the chromaticity of the illuminant. The target patches had their chromaticities modeled for the five different illuminants. For each brief stimulus presentation, subjects were to select the most appropriate term from a list of twelve monolexemic (Dutch) color names. Subjects were not told anything about illumination, material properties, and so forth. While a time limit was not placed on responses nor were responses timed, accuracy and speed of response were emphasized in the instructions given.

For the standard illuminant and each chromatic illuminant, subject responses for each color category were averaged to determine the categories' positions in a chromaticity diagram; e.g., the chromaticity coordinates of all stimuli selected as orange under the yellow illuminant were averaged to determine the location of the orange category under the yellow illuminant. The category locations under the different illuminants were used to assess the effect of the illuminant change on the categories. In the first version of this experiment, Troost & de Weert presented the stimuli blockwise by illuminant, which facilitated adaptation. In this condition, the measurement of color constancy (given by a Brunswik ratio, not the CCI) was midway between the measurements they found for appearance and surface asymmetric matching with simultaneous presentations; compare the average of the means for color categories reported in their table 4 (ibid., p.598) with the averages of the matching means listed in their table 2 (ibid., p.594).

Troost & de Weert also reported the number of categorization violations for their stimuli under the chromatic illuminants; a categorization violation is a difference in color categorization for a target between the standard illuminant and a chromatic illuminant. These results are especially revealing about the potential value of limited phenomenal constancy. Approximately 40% of their targets had zero violations (i.e,, color categorizations across all four chromatic illuminants agreed with that of the white illuminant) and roughly 25% of their targets had just one violation, with the average number of violations per target being in the neighborhood of 1.25; see their figure 6 (ibid., p.598). Categorization held up quite well despite the poor constancy of appearance matches. Relevant to the present concern with situations in which adaptation is limited, they found only a slight difference in the number of violations when illuminant conditions were randomized, as compared to the blockwise presentations; see their figure 8 (ibid., p.600).[4]

It is worth considering the basis for subject responses in the color naming task. Specifically, is color category membership controlled entirely by processes within the visual system or does it involve cognitive resources external to the visual system itself? If the latter, it would have to be determined whether the cognitive contribution comes about through labelling percepts in conscious judgment or an unconscious top-down interaction that shapes the outputs of visual processing. How these matters are resolved bears on the interpretation of categorical constancy performance (in experimental and everyday settings) and theorizing about the mechanisms supporting color constancy.

Troost & de Weert take color category membership to be part of the visual representation of an object (ibid., pp.591,596,600). In broad outline, this fits nicely with the account of "full-blown" perceptual content offered by Raftopoulos (2009), in which abstract categories figure. This late-stage representation, which is said to supply the content of our usual mode of visual experience, is contrasted with iconic, lower-level phenomenal content that is temporally volatile and not recorded in working memory. For example, Raftopoulos (p.117) contends that the higher-level visual representations that are entered into visual working memory "do not contain information about, say, the specific hue or shade of a color, only information about the category

[4] The current focus on color constancy with limited illuminant adaptation is the reason why certain other studies of categorical color constancy, such as Olkkonen et al (2009), are not discussed here.

of the color (say, bright red).” The later stage contents “reflect but do not record” the contents of earlier stages (ibid., p.164). However, there looks to be a crucial difference between Troost & de Weert and Raftopoulos when it comes to the source of the categories involved. Troost & de Weert (1991, p.591) have it that (some) color categories are “structural properties of the visual system.” The categories used in their color naming experiment are examples of colors that have been claimed to be “pan-human perceptual fundamentals,” each of which perhaps having a dedicated neural process (Boynton 1997, p.148; Uchikawa et al 1989, p.882). Raftopoulos instead holds that perceptual content involves top-down contributions from cognitive centers and is dependent on the conceptual resources of the perceiver. His position tallies with that of Troost & de Weert when it comes to categorical color perception not requiring conscious judgment on the part of the perceiver.

Longstanding debates over linguistic relativity and the cognitive penetrability of perception bear on this issue. Many researchers have pointed to cross-cultural color naming patterns as evidence that certain color categories are tightly linked to visual processing. For instance, Malkoc et al (2005, p.2154) take similarities in the clustering of basic color terms across the languages studied in the World Color Survey (WCS) to “[suggest] that the special and shared status of basic color terms may reflect special and shared properties of the human visual system or the visual environment.” Others have focused on the color categorization abilities of human infants and non-human animals (e.g., Bornstein et al 1976), which would seem to implicate an innate, visual (not learned and linguistic) source for the relevant categories. On the other hand, the next chapter will include some discussion of a serious challenge that has been raised against the use of WCS data to ground color categorization in perceptual universals (Jameson & D’Andrade 1997; Jameson 2010). There is also evidence of top-down effects on categorical color perception, such as a lack of categorical perception across the blue/green boundary for perceivers who speak languages that do not distinguish those colors (Roberson, Davies, & Davidoff 2000) and color category learning affecting perceptual discrimination (Ögzen & Davies 2002). I am largely in agreement with Ögzen & Davies (2002, p.479) in granting that some aspects of categorical color perception are rooted in hardwired features of visual processing while others are shaped by social or pragmatic factors; Jameson (2010, pp.195-196) expresses similar sentiments. The latter plausibly include top-down influences that affect perceptual processing and the content of experience, although in other cases the effect looks to be limited to perceivers' judgments about their percepts. To the extent that such top-down factors are in play, categorical constancy may vary across different perceivers.

Troost & de Weert (1991, p.596) note that their findings agree with Jameson & Hurvich's (1989, p.7) observation that variations in phenomenal color appearance induced by illuminant changes usually allow objects to still be recognized by their color, due to the shifts in phenomenal color not crossing category boundaries. Achieving consistent color categorization by means of instantaneous adaptational processes is beneficial to creatures who must act before the visual system has had enough time to determine the fine-grained colors of things; this "quick and efficient" advantage also attaches to the phenomenon discussed in the next section (see Foster 2011, pp.680-681). While categorical color constancy helps explain the sense of stability, it can be only part of the story. This is reflected in the fact that Troost & de Weert's subjects showed lower constancy performance for color naming, even with blocked illuminant presentations, than for surface matching judgments. After all, it is not that my experience of the differently illuminated parts of the lawn leads me to believe that the grass's color in the sunlight is largely similar to (is naturally grouped together with, etc) its color in the shade. Rather, nothing in my experience of the differently illuminated regions suggests that its determinate color and material properties change at all; the grass plainly has one color, not two (or more). Despite its limitations, good categorical constancy could work together with a non-phenomenal mechanism tied to a more fine-grained invariance across illumination shifts to create the overwhelming immediate impression that our experience reveals to us a world of objects with unchanging colors.

4. Surface matching, cone ratios, and relational constancy

A constancy of perceived object color – phenomenal or inferred – depends on the ability to compensate for (factor out, etc.) the effects of the illuminant in order to arrive at information about the contribution of a surface's reflectance properties to the color signal. As noted before, there is ample evidence that the visual system adapts its responses in ways that help offset the effects of the illuminant. A representation of the illuminant might also be used to extract surface reflectance information from the confounded retinal stimulus. However, such processes may prove unnecessary for successful performance in experiments used to study color constancy or for achievements outside the laboratory typically associated with color constancy. In fact, Foster (2003, p.441) argues that perceivers can make good surface matches and accurate discriminations between reflectance and illuminant changes, quickly and reliably, in circumstances in which any kind of stable perceived color is impossible; see Amano et al (2005) and Craven & Foster (1992). Here, I present some of the evidence Foster relies on, extend his argument with further evidence, and show how Foster's own claim (made with Reeves and Amano) about the existence of a mechanism of projective color constancy is threatened by his discussion of what subject performance in surface matching experiments relies on.

Consider the study by Amano et al (2005), in which simultaneous surface matches were made for:

i. targets embedded in the center of 7 x 7 Mondrian checkerboards; each of the 49 surfaces in both displayed patterns were 1° x 1°

ii. the right-hand surface in 1 x 2 arrays of surfaces; the two patches in both displayed patterns were either each 1° x 1° or 3.5° x 7°

The nature of the displays (simultaneous presentations under different illuminants without a gradient between them) prevents adaptation effects from facilitating good color constancy. Thus if surface matching depends on a constancy of perceived color, performance should be better with stimuli that enable better illuminant estimates. Configuration (i), which is representative of the stimuli often used in asymmetric matching experiments, affords useful cues to the illuminant. In contrast, (ii) is so structurally simple that it does not allow for reliable estimation of the illuminant. The mean CCIs for these conditions were: two 1° x 1° surfaces: 0.72; two 3.5° x 7° surfaces: 0.78; 49 1° x 1° surfaces: 0.73. The availability of illuminant cues provides no benefit to surface matching performance; recall that Arend & Reeves (1986) found that subjects making appearance matches did not take advantage of explicit illumination information that was pointed out to them. Subjects could not have relied on a sameness of perceived color when setting surface matches for the simplistic stimuli – limited adaptation and the absence of useful illuminant cues prevent that – but their settings were just as good as in a condition that offered information about the illuminant. In conditions relevantly similar to those of this experiment, a stability of perceived color is <u>not necessary</u> for making good asymmetric surface matches (Amano et al 2005, p.1012).

The abilities of patients with cerebral achromatopsia are interesting to consider at this juncture. These patients have lesions in higher visual areas that leave them unaware of phenomenal color, but they can perform well in some surface matching tasks (Hurlbert et al 1998). Cerebral achromatopsia does not involve damage to retinal and other low-level mechanisms, but it interferes with the ability to make use of global chromatic information from across the entire scene; computations on global information are required on all leading proposals for estimating the illuminant using image statistics (Lennie 1999, p.246). Hurlbert et al suggest that their patient, MS, is likely able to set matches for simple patterns based on an ability to exploit local (e.g., between neighboring patches) ratios of cone excitation levels caused by light reflected from scene elements. These ratios very well could be preserved in perceivers with damage limited to higher visual areas, as “there is evidence that spatial cone-excitation ratios might be an elementary feature extracted from the visual scene” (Smithson 2005, p.1335). Ratios of cone excitations between scene elements are nearly invariant under many natural and artificial illuminants (Amano et al 2005, p.1012). These ratios are not between cone classes for a given surface (e.g., L:M:S for the test patch), but rather within a given cone class between one patch and another in the display or between a patch and (say) the mean of some portion of the pattern in which it is embedded. For example, the ratio of L cone excitations between two differently reflecting surfaces should remain approximately the same under illuminant changes. MS’s success in only a limited range of circumstances as compared to normal subjects also indicates that cortical mechanisms contribute to some aspects of surface matching performance.

The ability to make good surface matches independently of information about or adaptation to the illuminant is relational color constancy. Foster (2003, p.439, 441-442) observes that this ability is relevant to scene perception, as it enables fast judgments about whether chromaticity changes in a scene are due to an alteration of illuminant or material properties. A more speculative possibility is that it would allow for the nearly immediate recognition of the global context ("gist") of a scene, prior to the visual construction of the objects in the scene; see Wright (2013, pp.446-447). Relational constancy has to do with "the constancy of perceived colour relations between surfaces under different illuminants" (Amano et al 2005, p.1011-1012) and can be facilitated by the aforementioned invariance of cone excitation ratios. All that matters are the relations between scene elements, not the specific colors assigned to them. The means by which those relations are registered is not obvious, given patient MS's performance. Focusing only on the normal perceivers presented with the minimal stimuli of Amano et al (2005), one might conclude that they were sensitive to stable relations between the apparent colors assigned to patches in their experience. The patches' experienced colors would not be constant, but some relations holding between them would be, given the approximate invariance of cone excitation ratios. However, MS is unaware of phenomenal color and also fails in tasks that might be mediated by a projective mode of color perception.

The results of Nascimento & Foster (1997) support the claim that subjects can and do exploit a signal related to the stability of cone ratios to discriminate illuminant changes from reflectance changes. They manipulated the cone ratios induced by a temporal (i.e., successive) illuminant change to a Mondrian pattern. In one presentation, the Mondrian was simply modeled as undergoing an illuminant change, without any further alteration. This resulted in slight natural variations from perfect invariance of cone ratios. In another presentation, the same Mondrian was modeled as undergoing the same illuminant change, but with a correction applied so that there were no deviations from perfect invariance of cone ratios. Subjects were tasked with judging which of the presentations represented a "natural illuminant change." Despite the "perfected" changes corresponding to "highly improbable natural events" (ibid., p.1399), subjects consistently mistook the sequences with corrected images for natural illuminant changes. The rate of such misidentifications increased with the degree to which the cone ratios induced by the actual illuminant change deviated from perfect invariance (Nascimento & Foster 1997, p.1397, fig.1). This indicates a built-in preference or expectation for stable cone ratios. Subjects reported that a phenomenological signature distinguished modified changes from imperfect natural ones. For the former, the illuminant shift was accompanied by a "wash" over the entire scene, whereas for the latter the effect was spatially uneven across different patches in the scene and led to some patches standing out from the others (ibid., p.1397). Craven & Foster (1992, p.1364) also relate that subjects experienced similar distinctive phenomenologies that distinguished abrupt successive illuminant shifts and material changes. This phenomenological cue enabled fast decisions about whether a material or illuminant change had taken place.

Further details regarding the processes that support relational constancy can be adduced. Subjects in Linnell & Foster (1996) could reliably determine whether a reflectance change had occurred along with an illumimant change in successive presentations, when the illuminant change occurred over the course of ca. 200 ms or less. Once the illumination shift took much longer 200 ms to complete, performance fell off significantly. Foster, Amano, & Nascimento (2001) reported that surface matching performance was ca. 10% better when Mondrian stimuli were presented sequentially (with an abrupt transition) in the same location, as compared to simultaneous, side-by-side presentations that required refixations as subjects looked back and forth between the Mondrians. Of course, even with performance reduced due to refixations, their subjects' simultaneous match settings showed good constancy. This is likely a consequence of the fact that if one's gaze is alternated quickly enough between two stimuli, nearly the same pattern of retinal stimulation results as when the two stimuli are seen in the same location in immediate succession (Foster 2011, p.682, fn.13). In line with the finding that performance declines with more gradually occurring successive changes, slower scanning across targets with simultaneous presentations would impair performance that relies on relational constancy.

Reflecting on such findings, Foster, Nascimento, et al (2001, p.8155) write that “[the] transient nature of the cue [to surface reflectance changes] and its failure to persist over refixations suggests that visual working memory contributes little, if at all, to performance.” Foster (2011, p.680) observed that one is thus naturally led to consider

> the suggestion that the ability to judge whether color relations were preserved or violated was the result of fast, relatively low-level, spatially parallel visual processing (Foster et al 1992). This notion was supported by subsequent measurements in which during successive illuminant changes, material changes in one or more surfaces in an array of other surfaces were shown to be readily detected almost independently of the numbers of surfaces (Foster, Nascimento, et al 2001).

The “pop out” of material changes observed in Foster, Nascimento, et al (2001) dovetails with the phenomenological reports from subjects in Craven & Foster (1992) and Nascimento & Foster (1997). Recall that for those subjects, a “wash” over the entire scene was diagnostic of an illuminant shift while an inhomogeneous effect that gave some scene elements a prominent appearance indicated a material change. Foster, Nascimento, et al (2001, p.8155) also found that the spatial window for the parallel processing of color constancy violations is just a bit more than 4° and coincides with the anatomical fovea, which is densely packed with chromatically-sensitive cone cells; cone cell density decreases markedly with eccentricity from the fovea. The correspondence between the fovea and the spatial window for detecting the (in-) stability of reflectance properties makes a good deal of sense. It is the fovea that is moved about and fixated on objects (and other scene elements) as one actively explores one's environment in search of information, resources, and other relevant items. The fovea also receives a disproportionate share of processing capacity in the visual brain ("cortical magnification"); see Findlay & Gilchrist (2003). Peripheral vision is much more suited to motion detection (and other changes that produce flicker) than detailed analysis of illuminant properties and object features such as color; see Wright (2006).

An upshot of relational color constancy is that there is no guarantee of an approximate sameness of perceived object color across circumstances in which good surface matches can be made. This is clearly troubling for attempts to link matching performance with a constancy of phenomenal color. However, it seems to have gone unrecognized in philosophical and empirical discussions of color constancy that things are no better for projectivist approaches. They, too, involve a stable assignment of colors to objects across viewing conditions. Compensation for contributions of the illuminant to the color signal is plausibly required for <u>any</u> consistent assignment of colors to surfaces. This should be kept in mind when weighing the proposal by Reeves et al (2008, p.226) that performance in their surface property tasks "may depend on a mechanism by which we unconsciously 'project' a subjective experience, such as color, back onto the physical world as an object property." On the same page, they observe that subjects could have performed <u>all</u> the tasks in their experiments without compensating for the illuminant. Performance in surface matching tasks shows that subjects are capable of reliably determining whether a surface's reflectance properties remain the same across changes in the chromatic and intensity properties of the scene. However, that can be accomplished without any determination of the reflectance properties of the surface in question and thus need not have anything to do with a mechanism like that suggested by Reeves et al.

While it is possible that relational color constancy performance depends on perceivers' ability to register and exploit relations amongst varying phenomenal colors, it is certainly left open that perceivers' sensitivity to illuminant changes and material changes is instead grounded in a sensory signal generated at early stages of visual processing. The signal itself might directly affect visual behavior independently of any phenomenological trace it might leave (Moore & Brown 2001, p.193). It is plausible, though, that subjects rely on the "pop out" and "wash" phenomenology closely associated with such signals. Neither of the alternatives just entertained involves inferences based on phenomenal relations. Moreover, even if subjects do make inferences regarding phenomenal relations in performing some tasks connected to color constancy, that may very well be a matter of basing explicit judgments on them, rather than a visuocognitive mechanism taking them as inputs. A consequence of this is that surface match results alone do not reveal anything about an unconscious inference to the stable colors of stimuli. Thus one is left to wonder what evidence there is for projectivism. Projectivism is consistent with the effects of different instructions on subject performance and is intuitively appealing (for at least some theorists, my former self included). However, more needs to be done to join together the projectivist hypothesis and experimental results. The next section examines empirical considerations that further complicate matters for projectivism.

5. A direct challenge to projectivism

In a study using simultaneous presentations, Radonjic & Brainard (2016) replicated the familiar instruction effects on subject performance. Certain patterns in their data bear on the prospects for projectivism, as they suggest that, depending on the instructions received, subjects (exclusively) either adopted a perceptual strategy based on matching phenomenal color or engaged in conscious reasoning about their phenomenal color percepts. Subject reports gathered through a debriefing questionnaire support this interpretation of their data. Radonjic & Brainard (ibid., p.863) are quite clear that they take their findings to undermine the case for projectivism. While I would not say that Radonjic & Brainard deliver a knockout blow to projectivism, their results advance the case against it and influenced my re-thinking of the projectivist view I held in Wright (2013).

Radonjic & Brainard (2016) employed matching and selection tasks with simultaneous presentations. Each subject performed both tasks and unlimited time was provided for making responses. Trials with the two illuminant changes used were randomly intermixed and thus adaptation was limited.[5] In the matching task, subjects adjusted a patch so that it matched in color with the target patch. For the selection task, subjects chose the patch most similar in color to the target from a competitor set displayed under a different illuminant. The selection task was intended to be more representative of everyday uses of color vision. To illustrate, one does not go out to the garden and manipulate the color of the strawberries to make them appear just right for picking. Rather, one inspects the color of the strawberries hanging on the vines and determines which ones look ready to be harvested.

Each subject was assigned to one of four instruction groups: neutral, physical spectrum, and two concerning reflectance that ultimately proved to not be significantly different from one

[5] I am ignoring Radonjic & Brainard's (2016) condition in which the entire stimulus was under the same illuminant; i.e., illuminant constant condition.

another (ibid., p.863). These instructions, aside from the neutral set, specified the criteria subjects should use for matches/selections. The neutral instructions made no reference to anything besides the unexplicated notion of color. Subjects in the neutral instructions group did not receive any training, while the other instruction groups were familiarized with the task they would be performing and relevant concepts were explained (e.g., the difference between surface reflectance and reflected light). The physical spectrum instructions stated that subjects should engage with the stimuli as illuminated physical surfaces and aim to equate the light reaching the eye from the patch they selected or adjusted with that from the target. These subjects were specifically directed to focus entirely on the light from the patches in question and to disregard as much as possible the effects of the illuminant on the background against which the patches were displayed. The reflectance instructions also told subjects to treat the stimuli as illuminated physical surfaces, with the goal of making settings/selections that matched the target's reflectance properties. Two separate experiments were run with these instructions and tasks, one that used quite simple two-dimensional stimuli on a computer monitor and another that used computer renderings of a three-dimensional object (a cube).

With simple stimuli, both settings and selections showed a pronounced separation between, on the one hand, the two kinds of reflectance instructions and, on the other, the neutral and physical spectrum instructions. The former resulted in modest levels of constancy with matching and rather strong constancy with selections, whereas selection and matching CCIs were close to zero for the latter. It is noteworthy that the highest matching CCI with reflectance instructions was only ca. 0.60, whereas some reflectance subjects' selection CCIs approached 1.0. For both tasks there was a good deal of performance variation across the reflectance instruction subjects. Despite subjects not receiving feedback, data analysis revealed that three of the eight reflectance subjects showed a notable increase in performance in the selection task across experimental sessions; there were not enough matching sessions to determine whether there was a similar effect for that task. No such trend emerged for the other five reflectance subjects, nor for the entire group of physical spectrum and neutral subjects.

In the experiment using more naturalistic stimuli, average performance with reflectance instructions continued to be higher than that for neutral and physical spectrum instructions, although the difference was not nearly as great as with the simplified stimuli. This is largely attributable to physical spectrum and neutral CCIs being much higher in the second experiment than in the first. The neutral and physical spectrum groups attained mean CCIs somewhat higher than the typical high-end of performance for appearance matches in other studies, with only a very small performance advantage for selections over matching. Reflectance instructions again led to a great deal of variation between subjects and performance was considerably better for selections than for matching. There was no sign of the sort of improvement of performance over sessions demonstrated by some subjects in the first experiment.

Before discussing Radonjic & Brainard's interpretation of their findings and subject responses to a post-experiment questionnaire, some remarks are in order regarding the stimuli used in both experiments. With respect to the simple stimuli, the structure of the background likely hindered attempts to exploit relational constancy in reflectance tasks. The colored textures making up the background were quite small (0.17° x 0.19°; the target and competitor squares were 2.6° x 2.6°) and were randomized across the display; see Radonjic & Brainard (2016, p.849, fig.1). Thus subjects could not have relied on cone ratios between the patches of interest ("squares") and any particular patch(es) in their immediate surroundings. Perhaps an average over some part of the region surrounding or adjacent to each square could have been used instead of a single local patch (Foster 2011, p.692). However, there does not appear to be anything about the set of texture elements used that would constrain local space-averaged cone responses in a way that would make them well-suited for such a purpose; e.g., that the average reflectance of the texture elements in the 0.7° annulus around each square would be equal (an annulus of that size would bring each patch plus its surround to the 4° size of the window for parallel processing of constancy violations). A similar issue arises with Radonjic & Brainard's more naturalistic stimuli. The faces of the cubes are essentially Mondrian patterns, each composed of a 7 x7 random array of colored square patches; see Radonjic & Brainard (2016, p.856, fig.6). The patches of interest ("buttons") were circular and depicted as attached to the cube's surface. While the size of the background elements is not an issue, the randomization of the elements across the faces of the cube is. This would thwart using cone ratios between the buttons and a neighboring patch for reflectance tasks. As with the simplified stimuli, a space-average of the region surrounding each button might be calculated, but there again looks to be nothing that would ensure that such an average could be reliably put to good use for relational constancy; e.g. that the Moore-8 neighborhood of each button had the same average reflectance. None of this is intended to suggest that Radonjic & Brainard's stimuli are un-ecological or that their methodology in suspect. There certainly are natural viewing conditions of which these stimuli are representative. However, the aspects of their stimuli just highlighted should be kept in mind when assessing what their findings show about color constancy and what drives subject performance in color constancy experiments.

Radonjic & Brainard (2016, p.863) interpret their results as showing that neutral and physical spectrum instructions lead to subjects using a strategy based on color appearance alone, while reflectance instructions cause subjects to rely on a process of explicit reasoning about their percepts. In support of this interpretation, they point to certain features of performance with reflectance instructions that were not present with neutral and physical spectrum instructions (in addition to the typical result of higher CCIs with reflectance instructions):

i. the large degree of intersubject variability – in both experiments, the range of individual subject CCIs ran from nearly zero to just shy of 1.0;

ii. the consistently better performance in selection over matching in both experiments; and

iii. the significant improvement over sessions shown by some subjects in experiment 1.

The variability between reflectance subjects suggests that individual subjects adopted idiosyncratic response strategies, strategies that involved reasoning about how best to satisfy the match/selection criteria provided in the instructions given the content of their color percepts. Compare this with the neutral and physical spectrum results, in which constancy was lower, but so was variability across subjects. It is plausible that subjects made their responses with these instructions directly on the basis of what was present in their percepts, without having to engage in any deliberate thinking about how to fit their responses with the instructions they received, thus avoiding one possible source of variation. Given the similar performance with neutral and physical spectrum instructions, it would seem that the "normal" mode of engaging with these stimuli is driven by the phenomenal color percept and that the reflectance instructions forced subjects into an unfamiliar situation that required them to devise strategies for making use of the phenomenal percept in order to satisfy the instructions.

That selection results in higher CCIs than matching with reflectance instructions, while there was a minimal difference between the tasks with the other two kinds of instructions, strengthens the case for the claim that responses with reflectance instructions were not driven purely by the subjects' color percepts. It is reasonable to expect that if reflectance instructions tapped perceived color alone, the different tasks should produce basically the same CCIs. The different tasks would simply amount to two different ways of accessing the same information for effectively the same ultimate purpose; viz., equating the reflectance properties of the target and another patch under a different illuminant. As Radonjic & Brainard (2016, p.861) note, the higher selection CCIs could be explained by reflectance subjects hitting upon a fairly straightforward strategy for making selections: inspect the competitors and choose the one that, were it the target, appears to have undergone the largest shift in the direction of the illuminant change. Given the set of competitors used, such a strategy could lead to overconstancy (i.e., excessive correction for the illuminant change), which was found in the selection responses for seven of the eight reflectance subjects in experiment 1. The following response to the post-experiment questionnaire given by reflectance subject mdd (experiment 2) lends credibility to the idea that such a policy was implemented by at least some subjects:

> If the illumination was different (as in general impression that one side was more "blue"), I would pick the button that seemed more "blue" for example.

Other subjects (e.g., bce from experiment 2) made remarks that indicated a conscious attempt to base their responses on the illuminant's effects in much the same vein. While subjects could also adopt a somewhat similar approach for reflectance matches, they would not have an example of an extreme shift in the direction of the illuminant already presented to them as a possible response. This would leave them much less likely to systematically overcorrect for the effects of the illuminant. No overconstancy was found for reflectance matches, which is consistent with the results of other experiments showing instruction effects for asymmetric matching.

Radonjic & Brainard (2016, p.861) hypothesize that some form of learning took place for the three reflectance subjects in experiment 1 who showed improvement in selection performance over sessions. Presumably this would involve adjustments of the conscious strategies employed by these subjects that took them from CCIs on a par with the subjects who received neutral or physical spectrum instructions, to CCIs much more in line with typical performance in reflectance-based tasks; i.e., this does not look like a case of perceptual learning. This sort of improvement is interesting and comports with the notion that these subjects were not relying on their percepts alone in making reflectance selections. However, it should be kept in mind that four of the eight reflectance subjects in experiment 1 showed no change over sessions and one subject showed a slight worsening, while none of the reflectance subjects in experiment 2 showed a significant improvement over sessions. Additionally, a review of the replies to the post-experiment questionnaire for the three subjects whose performance improved does not reveal anything about a conscious adjustment of their response strategy over sessions. Interestingly, though, reflectance subject hsc from experiment 2 made remarks that certainly seem to indicate such an adjustment:

> At first I would use the squares around the target and tests to compare the colors to one another on either side. I began to be able to notice more subtle differences in the test buttons regarding lighting over time which I believe helped me make more correct choices.

While hsc's performance did not improve across sessions, it may well have been the case that this subject is describing a process that took place during the first session. In any event, Radonjic & Brainard's interpretation is worth taking seriously, but more investigation is needed to fully evaluate it.

In their post-experiment questionnaire, Radonjic & Brainard asked the subjects about their strategies for the tasks they performed. As already indicated by the responses quoted above, the answers from reflectance subjects in both experiments contain many remarks that support the interpretation that they relied on explicit reasoning and imagination about what the target should look like under a different illuminant. Here are some further examples from reflectance subjects:

> [With] the [matching] experiment, it was very difficult to determine what the center square would look like under the changed illumination, so I had to think more about it. (subject eoh, experiment 1)
>
> The [matching] task was very counter-intuitive. It seemed like I had to think backwards constantly, which was frustrating. (subject nke, experiment 1)
>
> [Once] I identified the color of illumination I applied this to what I thought the target surface color would show under this illumination. (subject fvh, experiment 2)

The neutral and physical spectrum subjects had a strong tendency to focus on similarities and differences in appearance between the patches and the steps they went through in performing their task. The following give a sense of what these subjects had to say:

> When choosing the test square that matched my eye went straight to one then I would double check the middle target and glance at the other boxes to confirm none were close to the target also. (subject iul, experiment 1)
>
> I often pondered if I should select a color or adjust a color to compensate for different illumination, but I thought it would be best if I made decisions based on my observations

> instead of what I believed would compensate for different illumination situations. (subject jfa, experiment 2)

> I would first get the brightness as close as possible while adding color to the cube (moving the stick to the right). Then I would worry about the hue. Once I was close I would make finer adjustments with the minimum step size. (subject vlz, experiment 2)

Taking their subjects' comments together with the data from their experiments, Radonjic & Brainard (2016, p.863) conclude that instruction effects should not be read as supporting unconscious inference views that posit a color representation separate from phenomenal color. Surface matching turns out to not be drawing on a perceptual phenomenon, but instead depends on subjects' reasoning about reflectance and illuminant properties based on the phenomenal colors they encounter.

I find Radonjic & Brainard's interpretation of their findings persuasive when it comes to circumstances relevantly akin to those of their experiments. I am also partial to their rejection of projectivism. However, Radonjic & Brainard are not in a position to make a general statement about the basis for subject performance in reflectance tasks, such as their claim that "our view is that instructional effects are telling us about [how subjects can reason from their percepts]" (ibid., p.863). Many viewing situations that are importantly different from those of Radonjic & Brainard's experiments enable subjects to achieve good reflectance matches by means of relational constancy. While those situations look like they do not implicate a mechanism of unconscious inference to stable object colors, there is also no reason to suppose that they involve conscious inferences. The low-level signal connected to cone ratios discussed by Foster and his colleagues is sufficient. It is not surprising that Radonjic & Brainard's reflectance subjects would resort to conscious reasoning in order to complete their tasks. As noted before, the structure of their stimuli compromised or eliminated the opportunity to use the visual system's sensitivity to cone ratios across the illumination change. Responses such as that from hsc quoted above make it evident that some of their reflectance subjects made a deliberate effort to notice and reason about relations of apparent color connected to the effects of the illuminant on patches in the display. That, though, is markedly different than the "quick, accurate, and effortless" performance of subjects in situations that facilitate taking advantage of relational constancy (Foster 2011, p.688). Compare the transient nature of the signal that is relevant to relational constancy to the memory demands of the explicit judgments Radonjic & Brainard's subjects report engaging in. Note also that Radonjic & Brainard's reflectance subjects' reports do not mention the "wash" reported by Nascimento & Foster's (1997) subjects nor any other phenomenological signature of stable reflectance properties.

In the light of Radonjic & Brainard's findings and (especially) their subjects' reports, it is interesting to consider Cohen's (2008) visuocognitivist counterfactual theory of color constancy. The following passage summarizes Cohen's view:

> I propose … that, in cases of colour constancy, one of the responses of visual systems amounts to an answer to a question about the counterfactual properties of the regions under comparison. Namely, these visual systems answer the question: would region R_1 (presented under illuminant I_1) share an apparent colour with region R_2 (presented under illumination I_2) if, contrary to fact, both regions were presented under the same illumination – namely, both under I_1 or both under I_2. (Cohen 2008, p.80)

Cohen claims that the visual system's answer to the question about counterfactual color appearance properties "drives … the invariance/surface match reaction" to stimuli presented under different illuminants (ibid). Cohen's account is hard to square with the answers Radonjic & Brainard's reflectance subjects gave to the post-experiment questionnaire. Their reports make it clear that they were consciously engaging in an imaginative process to make reflectance-based selections or matches, not relying on a perceptual verdict about shared counterfactual color appearance. Having identified the nature of the illuminant shift (e.g., in the bluish direction), the subjects tried to picture in their minds what the target patch would look like under the other illuminant. If Cohen's account were correct, the right counterfactual appearances should be readily available for subjects to exploit in making matches or selections. At a minimum, the subjects' percepts should allow them to recognize fairly straightforwardly which competitor is the right one to select or when their match setting is satisfactory; recall again the characteristics of the reports from neutral and physical spectrum subjects.[6]

In response to this challenge, Cohen might appeal to the structural features of Radonjic & Brainard's stimuli examined earlier. Just as those features interfered with the process that supports relational color constancy, they might also explain why subjects had to engage in conscious reasoning and imagining rather than being able to rely on a determination by the visual system of whether two patches shared a counterfactual appearance property. Cohen is cautious

[6] It is important to stress that Cohen's entire account is based on a visuocognitive mechanism for inferring counterfactual sameness of color appearance and he intends for it to address color constancy in both humans and non-human animals. For reasons that should be clear by now, I think such a one-size-fits-all account is implausible. I will also note again the earlier quoted remark from Kelber & Osorio (2010, p.1620) about the lack of evidence for color constancy mechanisms beyond low-level gain control in non-human animals.

about committing to any specific details of the operation of his hypothesized mechanism for inferring counterfactual color properties and I am even more cautious about speculating too much about how he might fill in the details of his account. One possibility, though, is that he would connect his proposed mechanism to Foster's account of relational constancy in terms of a transient signal regarding the stability of cone ratios. The visucognitivist view that I previously held, which has some similarities to Cohen's, included such a link with relational constancy.[7]

[7] One important difference between our accounts is that Cohen's features a visuocognitive mechanism that renders a verdict about counterfactual color appearance properties for differently illuminated surfaces, while my account included a visuocognitive mechanism that took as inputs both (i) the phenomenal color of a surface under an illuminant to which the perceiver is adapted and (ii) the sensory signal based on cone ratios that indicates that two differently illuminated surfaces share the same material properties, in order to generate a non-phenomenal color attribution for the surface under the illuminant to which the perceiver is adapted. In short, my hypothesized mechanism is not the least bit concerned with counterfactual appearance properties and instead only targets the sameness of material color across the illuminant shift. My sense is that it is rather odd to suppose that there is a visual mechanism dedicated to determining whether two surfaces would match in phenomenal appearance if viewed under the same illuminant. To be clear, my skepticism about Cohen's view is not due to a wholesale rejection of the idea that perceptual systems might represent counterfactual properties (Cohen 2008, p.85), but instead is concerned with the kind of counterfactual properties Cohen appeals to. The visual system is in the business of helping us find our way around the environment, thus we should expect its outputs to inform us about our surroundings, not the nature of visual states that we are not

However, I doubt that this move would succeed for Cohen. For one thing, cone excitation ratios may be relevant only to responses to stability or instability of reflectance or illuminant properties in particular circumstances and have nothing at all to do with determining color appearance. In that case, cone ratios would be ill-suited for the sort of work they would have to do as part of Cohen's account. Another issue is that it is not clear that cone ratios should be necessary for the visual counterfactual mechanism to do its job. Take Radonjic & Brainard's reflectance subjects as an example. Their visual systems have determined the color appearance of the target and the illuminant shift seems to be perceptually present. Both the simple and the complex stimuli are structurally rich enough to allow their visual systems to form an estimate of the illuminants in the scene. That should be enough information for the hypothesized mechanism to work with in order to make a determination of shared counterfactual color appearance, yet the subjects had to resort to explicit reasoning and imagination. Lastly, Cohen's view faces the same difficulty confronting all visuocognitivist accounts: once we recognize what drives subject performance across different tasks in which good constancy is demonstrated (viz. a sensory signal based on cone excitation ratios), there is no obvious need to posit a visuocognitive mechanism that generates non-phenomenal color attributions.

6. Fitting together the pieces

From the preceding discussion the following account of color constancy can be assembled:

---

actually undergoing. It seems much more likely that the visual system would "settle" for something less extravagant and more useful in the immediate moment; viz., that the two surfaces share the same material color property.

1. Phenomenal color shows good, but imperfect, constancy in situations that either facilitate a considerable amount of adaptation to the illuminant or involve only a small change in illuminant properties. The main factor here is the timescale at which adaptation processes work, although I am not ruling out either other possible effects on the level of adaptation achieved or non-adaptational factors that might help facilitate phenomenal color constancy.
2. When adaptation is limited, phenomenal color is quite inconstant, but some viewing conditions allow perceivers to reliably register (distinguish between, etc) illuminant changes and surface reflectance changes by means of signals generated by visual processing. The strongest effects along these lines occur with quick successive illuminant shifts, although rapid eye movements would allow certain kinds of spatial variation in illumination to yield the same kind of result. There are also circumstances in which perceivers are able to make such discriminations with extremely brief (1 ms), simultaneously presented stimuli that do not allow eye movements to play a role (Foster, Craven, & Sale 1992). One source of the relevant signals is low-level visual processing concerned with local spatial cone ratios. It is left open that later-stage visual processes that operate on global properties of the scene (e.g., mean and variance of cone excitation levels or combinations of them) might also generate or modulate such a signal. A distinctive phenomenology (viz., the "wash" and "pop out" noted earlier) occurs in some of these cases and perceivers might rely on that phenomenology when engaging in various tasks; note that the phenomenology is not found in, for example, the "immediate" color constancy studied in Foster, Craven, & Sale (1992). It is possible, though, that even when those phenomenological characteristics are present, visual behavior (e.g., direction

and capture of attention) is instead driven by the visual signal itself, with subjects having access only to its "surface" phenomenology and lacking any insight into its typical effects on their ways of visually engaging with the world; see Moore & Brown (2001, p.193). Besides the absence of the phenomenological signature when immediate constancy is achieved, another reason for considering this possibility comes from the good constancy performance in surface matching tasks shown by cerebral achromatopsia patient MS. It is not at all obvious that his colorless visual experience would support the phenomenological signatures reported by normal subjects. A comparison of MS's preserved abilities to blindsight comes to mind, although such a comparison should not be pushed too far, given that MS is consciously aware of the items he is viewing (Cowey & Heywood 1997, p.137).

3. When adaptation is limited and viewing conditions do not support the operation of the visual mechanism discussed in (2), phenomenal constancy is poor and perceivers lack any other visual response that alone would support good constancy performance. In such circumstances, perceivers have to resort to explicit reasoning in order to determine whether two differently illuminated surfaces have the same reflectance properties, whether a difference in color appearance is caused by an illuminant change alone or is also due to a difference in reflectance properties, and so forth. This process can resemble (2) by involving consideration of relations between different surfaces/regions in the scene, but in this case it is a matter of reflecting on aspects of phenomenal color and (for example) imagining how things would look if the phenomenal relations between two surfaces under one illuminant were preserved under a different illuminant. In the course of routine engagement with our surroundings, it is likely that perceivers rarely stop to

consider such things, because either nothing prompts them to do so or there is no time for such wondering. Furthermore, when the occasion arises, if one cannot make an immediate determination that two differently illuminated surfaces are the same color, one can often move the objects of interest (e.g., baseballs, lemons, socks, paint chips) so that their color appearance can be compared under the same illuminant. Thus the sort of reasoning called for here may very well be unusual, leaving subjects to invent strategies for how to do it on an as-needed basis.

4. Even with the variable phenomenal appearance that comes with limited adaptation, a useful degree of categorical constancy can often be achieved through adaptation processes that are fast-acting (although incomplete). This can occur in connection with any of (1)-(3). This constancy of color category assignments helps maintain the impression of a stable visual world, despite variations in phenomenal color. This impression would be especially compelling if, as proposed by Raftopoulos (2009), our routine mode of engaging with the world has contents that are constructed out of abstract categories. The combination of categorical constancy with the visual process(es) of (2) is beneficial to perceivers making their way through real-world environments without the luxury of unlimited response time that is often granted to subjects in studies of surface color perception (Moore & Brown 2001, p.193). Together they would quickly provide (for example) (a) a visual indication that two differently illuminated surfaces have the same reflectance properties and (b) a classification of those two surfaces that results in their similarities being emphasized (and differences minimized) for cognitive tasks pertaining to color.

While this account portrays color constancy as a quite varied phenomenon, there is no reason to think that perceivers have much, if any, awareness of its diverse nature. Subjects might skillfully shuttle back and forth between these processes, as the situation demands, without reflecting on the details of how they go about doing things. This sort of skillful activity may even include exploitation of different results generated by these processes to accomplish some tasks. Consider, for example, someone trekking through a dense forest and looking for a way out. She might observe that the region ahead is green and composed of the same flora as her immediate vicinity (relying on relational and categorical constancy), while the yellowish cast on the foliage ahead (revealed by a comparison between the phenomenal colors she is aware of when looking ahead and looking at things nearby in the illumination conditions she has adapted to) allows her to infer that she is approaching the forest's edge, where sunlight is more abundant. This picture of color constancy comports with the "patches in a façade" (or "maps in an atlas") account of our practically-oriented conceptual behavior offered by Wilson (2008).

As heterogeneous as this account is, it does not include a projective mechanism like that proposed by Cohen (2008) or Reeves et al (2008). Simply put, I do not see any circumstance in which such a mechanism might be required to bridge a gap left by (1)-(4) above and I know of no evidence of such a mechanism contributing to constancy performance. I previously believed that there was work for such a mechanism to do in accounting for our immediate perceptual sense that objects have the same color across illuminant changes, as well as our thought about and behavior in such scenes. However, after thinking further about the processes that support relational constancy (viz., the fact that the information about cone ratios does not get encoded in working memory) and reading the responses of Radonjic & Brainard's (2016) subjects (in conjunction with the results of their experiments), I have reconsidered that position. Of course, I do not intend for my account to rule out the possibility of a projective form of constancy. Should subsequent empirical work show that, for example, non-phenomenal surface color attributions figure in perceptual representations of objects, I would have no problem (once again) including a projective mechanism in my account.

Also not included in my account are the complex dimensions of color appearance that figure in Davies' (2016) pluralist theory of color constancy. Davies argues that appearance matching tasks and surface matching tasks both engage perceived color, but different aspects of it. Specifically, he appeals to a distinction between material and illumination dimensions of color, with hue, saturation, and intensity attributes for each. Mausfeld (2003) discusses the history of such "duplicate dimensions" views and proposes an account along such lines. Logvinenko (Logvinenko & Maloney 2008; Tokunaga & Logvinenko 2010) has experimentally investigated the number of dimensions needed to model subjects' dissimilarity judgments for differently illuminated stimuli. Using multidimensional scaling to explore the structure of the dissimilarity ratings subjects gave to pairs of stimuli, Logvinenko concluded that six dimensions are called for, three each for illumination and reflectance; the standard three dimensions suffice when a single illuminant is involved. These dimensions are not independent. Rather, the lighting dimensions emerge "because the object colour manifold varies with illumination ….whilst remaining three-dimensional, the object colour manifold is different for different illumination" (Tokunaga & Logvinenko 2010, p.1746). Focusing specifically on the intensity attributes for simplicity's sake, the idea is that different levels of surface lightness (albedo) maintain their order across different illuminants, but each illuminant intensity induces its own idiosyncratic lightness continuum that is incommensurable with other such continua; this is nicely illustrated by Logvinenko & Maloney (2008, p.78, fig.2). Due to the incommensurability of the lighting continua, it is strictly speaking impossible to establish an exact appearance match between any two surfaces – including two identical surfaces – under different illuminants. The best that one can do is to find the two surfaces that are least dissimilar in appearance across the illuminant shift. Given the preservation of albedo ordering along each continuum and certain other properties of the family of continua, a given surface is always least dissimilar to itself under different illuminants.

The incommensurable nature of the lighting dimensions is taken to explain the dissatisfaction some subjects have reported with the best matches they are able to achieve; see Brainard et al (1997, p.2098) for an example of this kind of dissatisfaction. Tokunaga & Logvinenko (2010, p.1746) claim that the ease with which perceivers can distinguish material changes from illuminant changes (as in Foster's research on relational constancy) is due to "the immediate phenomenological difference between material and lighting dimensions." According to Davies (2016, p.552), surface matching responses are driven by material dimensions of color while appearance matching responses are based on lighting dimensions. Davies (ibid., p.557) also endorses Tokunaga & Logvinenko's take on the basis for relational constancy. All of this is intended to allow Davies to grant that there is a kind of significant variation in phenomenal color across different illuminants while preserving stability for some aspect of phenomenal color. Thus he would avoid being driven by instruction effects to hold that color constancy has a non-phenomenal basis.

While I am intrigued by the idea that the object color manifold depends on the nature of the illuminant, I do not find Davies' account of color constancy based on it compelling. My main concern is with the claim that surface matching subjects respond on the basis of perceptually accessed material dimensions of color. As Radonjic & Brainard (2016, p.863) observe, it is entirely possible that the subjects in Logvinenko's multidimensional scaling studies are relying on explicit reasoning in making their judgments about the dissimilarity of differently illuminated surfaces. Moreover, as was discussed earlier, it seems very much as though Radonjic & Brainard's subjects were not able to rely on perceptual representations of material dimensions of color for the differently illuminated stimuli. The remarks from their reflectance subjects regarding having to think backwards and to imagine what the target should look like under a different illuminant, as well as the seeming adoption of an overconstancy strategy by subjects in the selection task, are not at all what one would expect from subjects who based their responses on perceptual representations of material properties of the differently illuminated surfaces.

With respect to relational constancy, it seems plausible that the "pop out" and "wash" phenomenologies reported by Foster's subjects are connected to the phenomenologies of the material and lighting dimensions discussed by Logvinenko and his colleagues. In fact, I suspect that the latter kind of phenomenology depends on the former kind. However, I see no reason to think that relational constancy is achieved by means of a sensitivity to changes along the material dimensions of color appearance. Davies (2016, pp.557-558) is quite clear in holding that, contrary to what is claimed by Foster, the ability to discriminate reflectance changes from illuminant changes depends on perceptual representations of the properties of individual surfaces and illuminants. I strongly disagree. Numerous studies from Foster's work on relational constancy make clear that subjects can perform such tasks without any determination of either illuminant or reflectance properties. Cone excitation ratios provide a physical basis for this ability and need not have anything to do with constructing a representation of a surface's reflectance properties or an illuminant's spectral power distribution. I agree with Davies (2016, pp.556-558) that Foster has not always been clear about how subjects are taking advantage of the relevant color relations; e.g., conscious awareness of phenomenal color relations, a "visual sense" that an illuminant or reflectance change has taken place. However, Foster has been very consistent about the physical underpinning for subject performance. At this point it should be obvious that I believe the transient signal based on cone excitation ratios, either alone in some non-phenomenal sense or through an associated phenomenology that is not tied to representations of specific surface reflectance or illuminant properties, is up to the task of enabling "observers to perceive the material world as changing or stable" (Linnell & Foster 1996, p.227). As was the case with my exclusion of a projective mechanism, though, I am open to the possibility of including an invariance in subjects' perceptual representations of material dimensions of color, should future research show that it is needed to account for constancy performance in some circumstances.

References

Amano, K., D. Foster, & S. Nascimento. 2005. Minimalist surface–colour matching. Perception, 34, 1007-1011.

Arend, L. 1993. How much does illuminant color affected unattributed colors? Journal of the Optical Society of America A, 10, 2134-2147.

Arend, L. 2001. Environmental challenges to color constancy. In B. Rogowitz & T. Pappas (eds.). Human Vision and Electronic Imaging VI. Proceedings of SPIE. Bellingham (WA): SPIE.

Arend, L. & A. Reeves. 1986. Simultaneous color constancy. Journal of the Optical Society of America A, 3, 1743-1751.

Arend, L., A. Reeves, J. Schirillo, & R. Goldstein. 1991. Simultaneous color constancy: papers with diverse Munsell values. Journal of the Optical Society of America A, 8, 661-672.

Bäuml, K. 1999. Simultaneous color constancy: how surface color perception varies with the illuminant. Vision Research, 39, 1531-50.

Boghossian, P. & D. Velleman. 1989/1997. Colour as a Secondary Quality. In A. Byrne & D. Hilbert (eds.). Readings on Color, vol. 1: The Philosophy of Color. Cambridge (MA): MIT Press.

Bornstein, M., W. Kessen, & S. Weisskopf. 1976. Color vision and hue categorization in young infants. Journal of Experimental Psychology: Human Perception and Performance, 1, 115–129.

Boynton, R. 1997. Insights gained from naming the OSA colors. In Hardin & Maffi (1997).

Brainard, D., Brunt, W., & J. Speigle. 1997. Color constancy in the nearly natural image: I. Asymmetric matches. Journal of the Optical Society of America A, 14, 2091-2110.

Cohen, J. 2008. Color constancy as counterfactual. Australasian Journal of Philosophy, 86, 61-92.

Cornelissen, F., & E. Brenner. 1995. Simultaneous colour constancy revisited: an analysis of viewing strategies. Vision Research, 35, 2431-2448.

Craven, B. & D. Foster. 1992. An operational approach to colour constancy. Vision Research, 32, 1359-1366.

Davies, W. 2016. Color constancy, illumination, and matching. Philosophy of Science, 83, 540-562.

Delahunt, P. & D. Brainard. 2004. Does human color constancy incorporate the statistical regularity of natural daylight? Journal of Vision, 4, 57-81.

Findlay, J. & I. Gilchrist. 2003. Active Vision. New York: Oxford University Press.

Foster, D. 2003. Does colour constancy exist? Trends in Cognitive Sciences, 7, 439-443.

_____. 2011. Color constancy. Vision Research, 51, 674-700.

Foster, D., K. Amano, & S. Nascimento. 2001. Colour constancy from temporal cues: Better matches with less variability under fast illuminant changes. Vision Research, 41, 285-293.

Foster, D. B. Craven, & E. Sale. 1992. Immediate colour constancy. Ophthaalmic and Physiological Optics, 12, 157-160.

Foster, D., S. Nascimento, K. Amano, L. Arend, K. Linnell, J. Nieves, S. Plet, & J. Foster. 2001. Parallel detection of violations of color constancy. Proceedings of the National Academy of Sciences, 98, 8151-8156.

Hardin, C. & L. Maffi (eds.). 1997. Color Categories in Thought and Language. Cambridge:Cambridge University Press.

Helmholtz, H.V. 1924. Helmholtz's Treatise on Physiological Optics. Rochester (NY): Optical Society of Americe.

Hilbert, D. 2005. Color constancy and the complexity of color. Philosophical Topics, 33, 141-158.

_____. 2012. Constancy, content, and inference. In G. Hatfield & S. Allred (eds.) Visual Experience: Sensation, Cognition, and Constancy. Oxford: Oxford University Press.

Hurlbert, A., D. Bramwell, C. Heywood, & A. Cowey. 1998. Discrimination of cone contrast changes as evidence for colour constancy in cerebral achromatopsia. Experimental Brain Research, 123, 136–144.

Jameson, D. & L. Hurvich. 1989. Essay concerning color constancy. Annual Review of Psychology, 40, 1-22.

Jameson, K. 2010. Where in the World Color Survey is the support for the Hering Primaries as the basis for Color Categorization? In J. Cohen & M. Matthen (eds.) Color Ontology and Color Science. Cambridge (MA): MIT Press.

Jameson, K. & R. D'Andrade. 1997. It's not really red, green, yellow, blue: an inquiry into perceptual color space. In Hardin & Maffi (1997).

Kahneman, D., A. Treisman, & B. Gibbs. 1992. The reviewing of object files: Object specific integration of information. Cognitive Psychology, 24, 175-219.

Kelber, A. & D. Osorio. 2010. From spectral information to animal colour vision: experiments and concepts. Proceedings of the Royal Society B, 277, 1617-1625.

Kuriki, I. & K. Uchikawa. 1996. Limitations of surface-color and apparent-color constancy. Journal of the Optical Society of America A, 13, 1622-1636.

Lennie, P. 1999. Color coding in the cortex. In. K. Gegenfurtner & L. Sharpe (eds), Color Vision: From Genes to Perception. New York: Cambridge University Press.

Linnell, K. & D. Foster. 1996. Dependence of relational colour constancy on the extraction of a transient signal. Perception, 25, 221-228.

Logvinenko, A. & L. Maloney. 2006. The proximity structure of achromatic surface colors and the impossibility of asymmetric lightness matching. Perception & Psychophysics, 68, 76-83.

Malkoc, G., Kay, P., & M. Webster. 2005. Variations in normal color vision. IV. Binary hues and hue scaling. Journal of the Optical Society of America A, 22, 2154-2168.

Mausfeld, R. 2003. ‘Colour’ as part of the format of two different perceptual primitives: the dual coding of colour. In Mausfeld & Heyer (2003, pp.381-430).

Mausfeld, R. & D. Heyer (eds.). 2003. Colour Perception. New York: Oxford University Press.

Moore, C. & L. Brown. 2001. Preconstancy information can influence visual search: The case of lightness constancy. Journal of Experimental Psychology: Human Perception and Performance, 27, 178-194.

Nascimento, S. & D. Foster. 1997. Detecting natural changes of cone-excitation ratios in simple and complex coloured images. Proceedings of the Royal Society of London B: Biological Sciences, 264, 1395-1402.

Ögzen, E. & I. Davies. 2002. Acquisition of categorical color perception: a perceptual learning approach to the linguistic relativity hypothesis. Journal of Experimental Psychology: General, 131, 477-493.

Olkkonen, M., T. Hansen, & K. Gegenfurtner. 2009. Categorical color constancy for simulated surfaces. Journal of Vision, 9, 1-18.

Radonjic, A. & D. Brainard. 2016. The nature of instructional effects in color constancy. Journal of Experimental Psychology, 42, 847-865.

Raftopoulos, A. 2009. Cognition and Perception: How Do Psychology and Neural Science Inform Philosophy? Cambridge (MA): MIT Press.

Reeves, A. K. Amano, & D. Foster. 2008. Color constancy: phenomenal or projective? Perception & Psychophysics, 70, 219-228.

Roberson, D., I. Davies, & J. Davidoff. 2000. Color categories are not universal: replications and new evidence from a stone-age culture. Journal of Experimental Psychology: General, 129, 369-398.

Schultz, S., K. Doerschner, & L. Maloney. 2006. Color constancy and hue scaling. Journal of Vision, 6, 1102-1116.

Smithson, H. 2005. Sensory, computational and cognitive components of human colour constancy. Philosophical Transactions of the Royal Society B, 360, 1329-1346.

Thompson, B. 2006. Colour constancy and Russellian representationalism. Australasian Journal of Philosophy, 84, 75-94.

Tokunaga, R. & A. Logvinenko. 2010. Material and lighting dimensions of object colour. Vision Research, 50, 1740-1747.

Troost, J. & C. de Weert. 1991. Naming versus matching in color constancy. Perception & Psychophysics, 50, 591-602.

Uchikawa, H., K. Uchikawa, & R. Boynton. 1989. Influence of achromatic surrounds on categorical perception of surface colors. Vision Research, 29, 881-890.

Webster, M. & J.D. Mollon. 1995. Colour constancy influenced by contrast adaptation. Nature, 373, 694-698.

Wilson, M. 2008. Wandering Significance. Oxford: Oxford University Press.

Wright, W. 2006. Visual stuff and active vision. Philosophical Psychology, 19, 129-149.

_____. 2013. Color constancy reconsidered. Acta Analytica, 28, 435-455.

Wyszecki, G. & W.S. Stiles. 1982. Color Science: Concepts and Methods, Quantitative Data and Formulae. New York: John Wiley & Sons.