Forthcoming in *Acta Analytica*

Color constancy reconsidered[1]

Wayne Wright

Abstract: This paper proposes an account of color constancy based on an examination of the relevant scientific literature. Differences in experimental settings and task instructions that lead to variation in subject performance are given particular attention. Based on the evidence discussed, the core of the proposal made is that there are two different forms of color constancy, one phenomenal and the other projective. This follows the hypothesis of Reeves et al (2008). Unlike Reeves et al (2008), it is argued that projective color constancy is crucially dependent upon phenomenal color constancy and certain aspects of scene perception. Additionally, it is hypothesized that capacities that support projective color constancy have an important role to play in facilitating our ability to quickly recognize scenes with diagnostic chromatic properties independently of assignments of colors to object surfaces.

Color constancy is the stability of perceived object color across changes in viewing conditions. This is a fascinating achievement. Any given parcel of light reaching the eye from a scene (the "color signal") may have been caused by a vast range of states of the world. There is no way of directly determining the contributions of various physical factors to that light; viz., surface reflectance and illumination. For instance, light of almost identical composition reaches the eye from dandelion flowers in deep shade and dandelion leaves in sunlight, but the flowers look yellow and the leaves green in both conditions (Brown 2003, p.268). Philosophers have shown great interest in color constancy and its empirical study, chiefly due to color's special place in long-abiding debates about the relationship between the mind and the world.

[1] I am deeply grateful to Adam Reeves for providing much helpful feedback that has sharpened my thinking on the matters addressed in this paper. A debt of thanks is also owed to Charlie Chubb, David Hilbert, Kimberly Jameson, Kent Johnson, Louis Narens, A. Kimball Romney, and Jack Yellott.

Contrary to what might initially be supposed, there are unresolved fundamental issues facing empirical research on color constancy. This paper examines some of these issues, focusing mainly on the debate between phenomenal and projective views of color constancy. This exercise is valuable for at least three reasons. First, philosophers have made a number of appeals to the empirical literature on color constancy. A critical examination of that literature would thus benefit the philosophical debates in which such appeals have been made. Second, philosophers can contribute to scientific inquiry by studying bedrock issues and surveying a broad range of available evidence. This is especially so for color science, as throughout its history the field has been dominated by a technological and engineering emphasis that encourages insulated, narrowly focused research. Third, by the end of the paper, it will be possible to sketch a proposal about how color constancy should be understood, based on the evidence on hand. In short, I contend that there are two forms of color constancy, one phenomenal and the other projective, and that projective color constancy is dependent on phenomenal color constancy and mechanisms related to aspects of scene perception.

1. Measuring color constancy

Subjects in many color constancy experiments perform asymmetric matching. They are presented pairs of patterns of surfaces under different illuminants, either simultaneously or successively. Typical experimental stimuli are flat images on computer monitors, but more realistic displays are also employed; see Brainard et al (1997). Subjects use a control device to adjust the chromatic properties of a patch in one pattern so that its color matches that of a standard patch at the same location in the other pattern. The prevailing measure of the degree of color constancy achieved is a color constancy index (CCI), introduced in Arend and Reeves

(1986). This index is so entrenched that data from color constancy experiments using tasks other than matching are routinely transformed so that CCIs can be calculated for them and compared to matching performance. The basic formula for the CCI is $1 - a/b$, where a is the "subjective shift" (i.e., the difference between the subject's match setting and what would be a perfectly color constant match setting) and b is the "physical shift" (i.e., the difference between the standard patch and what would be a perfectly color constant match setting). These shifts are Euclidean distances between points in chromaticity diagrams or color spaces. Perfect constancy and a complete failure of constancy are indicated by CCIs of 1.0 and zero, respectively. The CCI's rationale is straightforward. The standard patch's reflectance properties and those of its counterpart are the same across the different illuminants; this is often true of all corresponding patches in the pattern pairs. Thus the chromaticity of a subject's match settings reveals his or her visual system's ability to compensate for the change in illumination and meet the stability of surface reflectance with a corresponding stability of perceived color.

Despite its intuitive appeal, questions can be asked about the CCI that are useful to a more general examination of color constancy; see Foster (2003). Like any index, the CCI must be properly linked to the phenomenon it is supposed to summarize with a single number. That task turns out to be complicated by difficulties regarding how the target phenomenon should be characterized. Related issues emerge about experimental methods and the data used in calculating CCIs. There is slippage between the matching task often used in studying color constancy and perceived color, and different experimental settings and instructions can produce widely varying results. Thus care needs to be exercised when dealing with the empirical literature on color constancy.

2. Phenomenal, projective, or both?

The description of color constancy given at the outset is insufficient as a guide for inquiry. More needs to be said about the sense(s) in which objects' perceived colors remain stable and what phenomenon is tapped in particular experimental settings. Color constancy might be a matter of phenomenal color undergoing little or no change across different viewing circumstances; to be clear, 'phenomenal color' is used here to designate the distinctively qualitative way in which we are visually aware of color, as opposed to a non-qualitative mode of conscious awareness.[2] On the other hand, color constancy might involve non-phenomenal, visuocognitive elements that provide a stable output regarding object color while phenomenal color varies as viewing conditions change. Such information might be thought of as added to the representation of an object (or its parts) entered in working memory in much the same way that semantic and functional information about objects is hypothesized to be encoded on some theories of vision. While this information about object color would be consciously present or consciously accessible, it might not agree with the qualitative color phenomenology (color appearance, etc.) of the subject's occurrent visual episode.

One example of a non-phenomenal view with deep historical roots posits an ability to infer an object's "true" or canonical color paired with a reflexive attribution of that color to the experienced object despite variation in phenomenal color; see Alhazen (1030/1989), Cohen

[2] To further elaborate, while some things of which I am consciously aware are phenomenally present (e.g., apparent color of the sort tapped in appearance matching tasks), there may be other things of which I am consciously aware that do not have a qualitative nature. As an example of the latter, some researchers hold that a conscious belief does not, in itself, have a phenomenal aspect. Of course, one may feel certain things as a result of entertaining that belief, but that is separate from the belief itself having a phenomenal character.

(2008), Helmholtz (1924), Joost et al (2002, p.302), and Reeves et al (2008, p.226). Talk of inference suggests a commitment to a mechanism that deals in premises, rules, and conclusions. However, one might endorse a non-phenomenal account of color constancy while holding that the relevant mechanism operates according to principles other than those of inference. Following the remarks of Arend and Reeves (1986, p.1749), the key distinction between phenomenal color constancy and the sort of constancy that might be achieved while color appearance varies is not in the mechanism involved, but "in the perceptual representation of the information" about surface color. What matters here for grouping together non-phenomenal accounts of color constancy is that they accept that stable verdicts regarding surface colors, which are perceived (or perceivable) as properties of external objects, can be reached while phenomenal color varies.

Adam Reeves and his colleagues have used the term 'projection' when discussing non-phenomenal color constancy. This term is neutral about the relevant mechanism and captures the external attribution component of the views I have in mind; see Reeves et al (2008, p.219). While handy, it should be stressed that this usage of 'projection' differs from its standard philosophical understanding. The philosophical use applies to sensations (phenomenal appearances) and is typically 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.[3] Reeves et al (2008,

[3] To illustrate this using anti-realism about colors, it seems open to say that our color experiences represent objects as being objectively colored while nothing actually bears color properties. On such a view, colors exist only in the representational contents of our experiences. That is importantly different than saying colors are properties of experiences. After all, to take a

pp.225-226) are 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 use 'projection' in accordance with its employment by Reeves et al.

The greatest support for projectivism about color constancy comes from marked differences in CCIs depending on the instructions subjects receive; see Arend et al (1991), Arend and Reeves (1986), Bäuml (1999), Cornelissen and Brenner (1995), Reeves et al (2008), Troost and de Weert (1991). When subjects are instructed to adjust a patch so that it "looks like it was cut from the same piece of paper" as another patch under a different illuminant ("material" or "surface" match), CCIs usually range from 0.6 to 0.9, depending on experimental settings. When subjects are instructed to set one patch so that it "matches the other patch in hue and saturation" ("appearance" or "phenomenal" match) their performance shows little disentangling of illuminant and surface reflectance; CCIs drop to something on the order of 0.0 to 0.3. Reeves et al (2008, p.225) observe that subjects in their hue/saturation task behave more like colorimeters for the light reflected from the patch to be matched than detectors of invariant surface reflectance properties. That subjects can separate their judgments about phenomenal color and surface properties, with (at times) the former leading to poor constancy and the latter to good constancy, challenges the idea that color constancy is rooted in stable phenomenal color. On the basis of such instruction effects, projective understandings of color constancy come out

---

property about which objectivism seems safe, whether or not my experience of something looking square is veridical, my experience itself does not have the property of being square; similarly, my thoughts about unicorns do not themselves possess the property of having a horn. The projectivist about color sensations denies this and contends that redness (greenness, etc.) is a property of my experience itself that an external object such as a ripe strawberry is systematically and mistakenly seen as having. See also the remarks about phenomenal transparency later in this section.

as the main alternative to an appearance-based conception. The non-phenomenal attribution of color to objects may be so effortless and ubiquitous that it swamps awareness of variations in phenomenal color, thus explaining the intuitive sense that things look to have the same color when viewing conditions shift.

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). As will be explained shortly, whether the two tasks differ in a meaningful way depends on experimental conditions. Additionally, the distinction between the hue/saturation and surface match tasks is likely to become familiar only with instruction. Findings from Cornelissen and Brenner (1995) are relevant to whether subjects are able to appreciate a 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 and 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 and Brainard (2004, pp.72-74) found 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 not just for the successive presentations of Reeves et al (2008), but also in studies using 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. Ignoring basic differences between asymmetric matching and achromatic settings, several points of contrast are relevant to the dissimilarity of these findings. Scientists often acknowledge these differences. However, some important consequences of these differences for how color constancy should be understood and the conclusions that can be safely drawn from extant color constancy research have not been fully recognized.

The first issue concerns illuminant adaptation. Delahunt and 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) appearance-based settings with full adaptation achieved good constancy. This is tied to their conjecture that there are <u>two</u> constancy processes, one which supports phenomenal constancy in situations of considerable adaptation of the visual system to an illuminant and the other, which supports projective constancy, coming into play when such adaptation is minimal (Reeves et al 2008, p.220).[4] The former is phenomenal and the latter projective. There is no real conflict between the results of Delahunt and Brainard and Reeves et al; see also Delahunt and Brainard (2004, p.74), Kuriki and Uchikawa (1996, p.1634), and Thompson (2006, pp.85-86, fn.15). More will be said soon about adaptation.

[4] David Foster, one of the co-authors of Reeves et al (2008), expressed skepticism about the existence of a constancy of perceived color in his (2003). I will argue later that Foster's (2003) concerns apply to the projectivist component proposed by Reeves et al (2008).

The second point deals with illumination conditions. 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 gradually from one side to the other. They observe (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 and Reeves (1986) involve a gradual illuminant change. A gradual illuminant change allows for greater adaptation than is possible with the sharp illuminant changes typical of other asymmetric matching experiments; see Bäuml (1999, p.1532). It is noteworthy that with their simultaneous displays, Arend and Reeves (1986) employed a condition in which the test and standard Mondrian patterns were each surrounded by a thin strip that displayed the 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, subjects making appearance matches showed low levels of constancy on a par with those made with Mondrians lacking surrounds. This suggests that adaptation has a central role to play in supporting greater appearance-based constancy. However, it does not entail that phenomenal constancy is not also influenced by estimates made by the visual system of the chromatic and intensity properties of the different illuminants on the test and match patches. That subject performance in Brainard et al (1997) was affected by an illuminant estimate is reasonable, 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 further considered in section 5.

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 and 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, as they “needed several iterations of the displays to convince themselves that the same hue and saturation could imply a different material” (p.222). Only when those demonstrations were understood did the training regimen continue. While adaptation and illuminant cues likely account for most of the divergence between these studies, differences in subjects’ level of familiarity with the tasks should be considered, too.

Turning to the details of projectivist views that might be developed (at least in part) based on instruction effects, Hilbert (2005) alleges a lack of clarity in such views. Hilbert’s primary target is inferentialist views in the Helmholtzian tradition, but his objection could be adapted to other projectivist accounts. His charge is founded largely on two worries he has about the premises and conclusions of the hypothesized inferences. The first is that the similarities and differences in color appearance observed under changing illumination conditions are perceived as external to our minds and not as due to something internal to the perceiver, such as properties of an experience itself; see footnote 3. This turns on the much-discussed “transparency of experience,” according to which, in having an experience, it seems to the perceiver as though what she is aware of is what her experience is an experience of, and not the experience itself or any of its properties. Transparency of this sort sits uneasily with the inferentialist account, which has it that a conclusion regarding a stable property of objects is drawn while phenomenal colors can vary. The inferentialist would seem forced to conceive of the latter "internally," as sensational properties of experiences themselves. The second claim is that if 'color' is used in a univocal, external fashion in describing the premises (color appearance) and conclusion (surface color), then any given patch of surface would be experienced as having more than one color at any moment. One color would be stable over time (the conclusion) and the other subject to variation over time as viewing conditions change (the premises). Hilbert's suggestion is that this seems phenomenologically inapt and is likely metaphysically confused.

Inferentialists are not without reasonable replies to Hilbert's concerns. Moreover, the appearance of such problems for inferentialists might be primarily due to unfortunate word choices on the part of some vision scientists and their lack of familiarity with (or concern for) certain philosophical debates about the nature of perceptual experience. As an example of an inferentialist response, it certainly seems open for an inferentialist to grant that experience does not involve awareness of intrinsic experiential properties, but to insist that the perceived stability of surface color is not due to a stable feature of qualitative experience. The inferentialist could claim that we have a mechanism for determining the colors of surfaces which issues outputs that are consciously accessible but qualitatively absent. Since there are visual information stores with contents that affect perceivers' reports without being phenomenally experienced, such as for visual long-term memory (Henderson and Hollingworth 1999) and scene gist (Oliva 2005), the idea that there might be such a mechanism is not fantastical. The outputs of this hypothesized mechanism would shape our judgments about the colors of things and may do so in such a pervasive, routine fashion that we are not aware of the interpretive nature of our reports about perceived color.

It is a commonplace in psychology that judgments about stimuli are made relative to a context (or level of reference), rather than being absolute in nature; see Helson (1964), Smithson and Zaidi (2004, p.707). Thus there is no difficulty in understanding how subjects' judgments about the colors of objects, formed relative to one (perhaps default or familiar) context, can come apart from their judgments about the phenomenal color attributes encountered in their experience, formed relative to another context. The difference in contexts is presumably due to the different instructions participants are given. In terms familiar from the work of Daniel Dennett, how things seem to seem depends on what judgment one is making. None of this requires thinking that subjects are confronting or drawing inferences based on intrinsic properties of their experiences. Also, if the two judgments are formed relative to different levels of reference, there is no issue of simultaneously and inconsistently attributing multiple colors to the same patch of surface.

After sufficient exposure to the illumination conditions, the phenomenal appearances of external things and the outputs of the hypothesized mechanism might largely agree. They might differ substantially in other viewing conditions, but that difference would not seem remarkable (and might usually not be noticed) because of how natural it is for our judgments about the colors of things to rely on the output of the hypothesized mechanism; see also the remarks about categorical constancy that soon follow. Another possibility is that the hypothesized inferential mechanism only kicks in when the chromatic properties of the phenomenal percept are outside an acceptable range. A potential benefit of the arrangements sketched here is that while we are usually concerned with the surface properties of objects, ecologically important illuminant properties can be determined by a comparison of the two different judgments. For example, while trekking through a forest, I observe that the region ahead is green (inferred) and its yellowish cast (phenomenal) tells me that I am nearing the forest's edge, where sunlight is more abundant. All this is speculative, but it does offer a workable inferentialist reply to Hilbert's worries. See Cohen (2008, pp.84-85) for remarks in a similar spirit.

Once the effects of adaptation and an illuminant estimate are recognized, it is clear that the debate about the nature of color constancy is not necessarily a matter of a binary choice between phenomenal and projective accounts. 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 and 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 suffice for color constancy or higher-level processes are also needed to compensate for the effects of the illuminant on the color signal. 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 visual system seems to also use a weighted set of cues to form an estimate of the chromaticity and intensity of the illuminant, in attempting to disentangle (and likely separately represent) illuminant and surface reflectance properties.

Following Troost and De Weert (1991, p.596), a compelling way of interpreting hue/saturation task results with limited illuminant adaptation is in terms of fast adaptation mechanisms accounting for at most 40% of the change in sensitivity needed to achieve perfect color constancy. Such adaptation is sufficient to facilitate a fairly high degree of categorical color constancy in some circumstances, but it cannot preserve an approximate sameness of fine-grained phenomenal color (ibid., pp.596-600). Categorical color constancy is the ability to assign a surface to the same general color category (e.g., green, brown, blue) based on its appearance, under different illuminants. From the perspective of projectivism, categorical constancy might help explain why inferences (etc.) to objects' surface colors typically go unnoticed: at a coarse-grained level phenomenal color appearance may often agree with the output of the process.

When subjects are given longer exposure to the illuminant or when the illuminant shift is gradual, constancy performance in appearance-based matching tasks improves considerably, due to contributions of the slower adaptation mechanisms and (plausibly) a reliable estimate of illuminant properties being formed. Thus the higher CCIs for the appearance-based tasks of Brainard and his colleagues, as well as in Arend's (1993) experiment 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. Importantly, the stimulus configuration may impair the reliability of the illuminant estimate, thereby limiting color constancy even with full adaptation; besides the disparity between shadows and illuminant gradients, differences between computer-rendered and natural displays may be relevant.

There are obvious benefits to being able to consistently deal with objects' surface properties in an appropriate manner across abrupt and gradual, spatial and temporal variation in illumination. The results of a wide range of studies indicate that humans and other animals have such abilities to a considerable extent. However, the fact that the end result is in a general sense the same – behavior that is in some way sensitive to the reflectance properties of objects – does not entail the same psychological phenomenon is involved across the different kinds of cases. For example, the two color constancy mechanisms posited by Reeves et al could be linked to very different information-acquisition processes and specific behaviors. The varied results of the studies discussed suggest that a possibility worth taking seriously is that the impression that color constancy research targets a single phenomenon, or a unified set of phenomena, might be an artifact of an overreliance on commonsense ideas about perception. Perhaps several quite distinct phenomena have been unwittingly lumped together as color constancy; note that this is separate from whether multiple mechanisms can provide the basis for a single phenomenon across varying circumstances. Following a suggestion by Donald Hoffman (2003, p.274), these phenomena that might be better studied separately as part of research on, inter alia, object perception, scene perception, motion perception, and visual memory. This worry becomes particularly acute once certain aspects of common experimental methods are taken into account.

3. Cone ratios and surface matching

A constancy of perceived object color – phenomenal or projected – depends on the ability to compensate for the effects of the illuminant in order to arrive at information about the contribution of surface reflectance 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 and that it constructs a representation of the illuminant that could be used to extract surface reflectance information from the confounded retinal stimulus. However, such processes might not be required for successful performance in some experiments used to study color constancy or for achievements outside the laboratory typically associated with color constancy. In fact, David Foster (2003, p.441) argues that good surface matches can be made in circumstances in which stable perceived color is impossible. He also questions what basis there is for thinking color constancy exists. 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.

Foster and his colleagues have demonstrated that good surface matching performance can be attained without useful illuminant cues or adaptation; see Amano et al (2005), Craven and Foster (1992), Foster (2003). Consider the study by Amano et al (2005) of performance when 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) makes it unlikely that adaptation effects alone would facilitate 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) affords useful cues to the illuminant, but (ii) is so structurally simple that it prevents 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. Surprisingly, the availability of illuminant cues provides no benefit to surface matching performance. 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. While subjects might base their surface match settings on perceived color in situations that allow for substantial compensation for the effects of the illuminant, important for the issue of what conclusions can be safely drawn from experimental results is that a stability of perceived color is not necessary (at least under conditions relevantly similar to those of this experiment) for making good asymmetric surface matches (Amano et al 2005, p.1012).

The abilities of cerebral achromatopes 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 ratios of cone excitation levels caused by light reflected from scene elements. 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 the entire 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 limited circumstances also indicates that cortical mechanisms contribute to some aspects of surface matching performance.

The ability to make good surface matches independent 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 judgments about whether chromaticity changes in a scene are due to an alteration of illuminant or material properties. Relational color 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 perceived relations between scene elements, not the specific colors assigned to them. The sense in which those relations are perceived 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 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. While it is an open question that a common mechanism accounts for both MS's preserved abilities and normal subjects' performance with minimal stimuli, the extent to which relational color constancy might require any kind of perceived color is equally unsettled.

Leaving this aside, 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 any 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.” Despite the somewhat unfortunate wording, it is clear that they are advocating the sort of projectivism described earlier. On the same page, they observe that subjects could have performed all 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. A consequence of this is that surface match results alone do not reveal anything about an inference (etc.) to the stable colors of stimuli. Thus it is unclear what experimental evidence there is for projectivism. Projectivism is both consistent with the effects of different instructions on subject performance and intuitively appealing. However, more needs to be done to join together the projectivist hypothesis and experimental results before the hypothesis can be properly evaluated.

4. Relational color constancy and scene recognition

In addition to the relevance of relational color constancy to scene perception pointed out by Foster (2003), there is a potential connection between relational color constancy and the ability to quickly recognize scenes with diagnostic colors. The proposal made here is that cone excitation ratios between different scene elements facilitate quick recognition of the global context of some scenes, in as little as a single glance, independently of any recognition of those scenes' objects and their properties (including their color). This global context, or "gist", guides visual search and critically shapes how perceivers extract information from a scene; see Henderson and Hollingworth (1999), Rensink (2000). While much interest in color constancy stems from the essential role a stability of perceived object color is thought to have in constructing a stable visual world, the possibility fleshed out here would be an example of contributing to a stable visual world in a way that does not implicate a constancy of perceived object color.

Scene recognition is challenging, as scenes of the same kind can vary in many ways; e.g., illumination, configuration of component objects. A standard understanding of scene recognition is that it depends on the reconstruction of objects and their properties from the retinal input. For example, I can tell that I am in my kitchen and not at the beach from my perception of objects that are fitting for the former and not the latter. Color constancy could play a role in this process, as some objects can be quickly and reliably identified on the basis of their typical color. However, there is research demonstrating that scene recognition need not hang on the recognition of the identity of objects within the scene and that we can get the "gist" of some scenes before we get their objects; see Schyns and Oliva (1994, 1997). The visual cues mediating fast scene recognition would have to be low-level. One potential cue is chromatic composition: scenes having a diagnostic chromatic composition could be identified from early visual processing of chromatic information. The basic arrangement of "chromatic blobs" in the scene might facilitate scene recognition for perceivers with prior exposure to scenes of the same kind. Also required are the ability to form memory templates for scenes and a comparison operation for stored templates and the current scene. One class of scenes with diagnostic chromatic properties is natural scenes such as forests, deserts, and seacoasts.

Oliva and Schyns observed that subjects more quickly recognize natural scenes having diagnostic colors when they are presented in their normal colors than when presented in either grey-scale or with their normal colors inverted or swapped along chromatic axes in CIELAB space; see Goffaux et al (2005) and Oliva and Schyns (2000). There was no difference in the speed at which scenes lacking a diagnostic color were recognized across all conditions. One of their experiments was intended to probe whether the diagnostic color advantage is due to object colors or more coarse-grained chromatic information (Oliva and Schyns 2000, pp. 194-199). They used low-pass filtered versions of normally colored and grey-scale images of scenes with diagnostic and nondiagnostic colors. Low resolution images, especially below 2 cycles per degree of visual angle, consist of little more than colored (or grey-scale) blobs, with no useful information about object identities. For scenes with diagnostic colors at low resolutions, subjects showed clearly better performance in scene categorization for normally colored images over their grey-scale counterparts. The differences in categorization performance between normally colored and grey-scale images begin to fall off at 2 cycles per degree of visual angle and stabilize at a small difference from 4 cycles per degree of visual angle and up. This suggests that the contribution of chromatic information to express scene recognition is at a level of spatial resolution that excludes objects and their properties. Recordings of event-related potentials indicate that categorization for scenes with diagnostic colors can occur "within the first 150 ms of visual processing" (Goffaux et al 2005, p.890). It is unknown what aspects of visual processing enable diagnostic colors to be exploited for express scene recognition.

The invariance of cone excitation ratios under illuminant changes might fill in some missing details of this account. Bear in mind the brief time course just noted, that scene gist can be extracted from a single fixation, and that chromatic diagnosticity occurs at resolutions too coarse for the representation of objects and their properties. Additionally, as Smithson (2005, p.1335) notes, "there is evidence that spatial cone-excitation ratios might be an elementary feature extracted from the visual scene." The invariance of cone excitation ratios would allow quick, low-level computation of information about chromatic properties of scene regions. Using cone excitation ratios to exploit diagnostic chromatic properties of scenes would obviate the need to go through whatever processes might be required to compensate for the illuminant and construct and assign colors to objects. This is valuable, as scenes are often differently illuminated from one viewing instance to another. Many variations in the configuration of scene elements across different scenes of the same kind, which could substantially alter fine-grained chromatic properties, would be usefully omitted at low resolutions. As for the representation of scene colors in memory, ratios of excitations in different cone classes between scene regions corresponding to coarse-grained "blobs" could be stored as a template to match against in future experience.

This section is intended to motivate the idea that some aspects of surface matching performance are grounded in capacities for scene perception not concerned with objects and their properties, including their reflectance properties. To be clear, my claim that invariant cone excitation ratios mediate express scene recognition is speculative and needs further development.[5] However, to the extent that the proposal gains further credence from future empirical work, it is evidence that some phenomena may have been improperly lumped under

[5] One prediction of this proposal is that, for any arbitrary scene, if it shares the same relevant cone ratio properties with a scene having diagnostic colors, it should be classified as the same sort of scene as the latter.

the single heading “color constancy.” Even as speculation, this proposal highlights the need to consider a more nuanced understanding of color constancy and its role in our interactions with the world.

5. The dimensions of color

Based on the discussion thus far, the one case of color constancy with clear empirical support involves apparent color in situations that allow for considerable illuminant adaptation. However, even this presents complications in understanding the sense in which perceived color is stable. The structure of psychological color might rule out the identity of color appearance across different viewing conditions. If so, asymmetric appearance matches really are no such thing, but instead involve a different sort of judgment by subjects. Reports of subjects’ dissatisfaction with their match settings bring this out:

> The observers were able to set reliably what they regarded as the best match. At this point, however, the test and the match surfaces looked different, and the observers felt as if further adjustments of the match surface should produce a better correspondence. Yet turning any of the knobs or combinations of knobs only increased the perceptual difference. (Brainard et al 1997, p.2098)

Despite this, and perhaps suggestive of the power of commonsense conceptions of perceptual phenomena, Brainard et al go on to describe their subjects’ matches as “establish[ing] pairs of stimuli that look identical” (1997, p.2100). That aside, similar observations about the impossibility of identical color appearance under different illuminants have fueled skepticism about familiar characterizations of color constancy (see Mausfeld 2003, p.397). The issue again is just what it means to say that perceived color is stable across changes in viewing conditions.

To explain subjects' dissatisfaction with their match settings, Brainard et al consider the possibility that color experience should be described by more than just three variables of hue, saturation, and lightness. A dimensionality of greater than three could make sense of why subjects succeed in setting matches for hue, saturation, and lightness (which Brainard et al took their subjects to be doing) while being unable to set perfect appearance matches. The differing illuminants on two stimuli might impose differences on some dimension(s) that cannot be counterbalanced by adjusting the chromaticity of a patch.

The difficulties with appearance matches highlighted by Brainard et al prompted Logvinenko and Maloney (2006) to have subjects rate the apparent dissimilarity of stimuli in a lightness constancy study, rather than make matching adjustments to the stimuli; this also figured in the experimental design of Reeves et al (2008). Based on their findings, Logvinenko and Maloney argue "that there are two distinct perceptual dimensions of achromatic [surface] color" (2006, p.82); see also Izmailov and Sokolov (1991, p.257). One of these dimensions, surface lightness, is a surface's perceived degree of total reflectivity (albedo). The other, surface brightness, is determined by the interaction of illuminant intensity and reflectivity. Each illuminant intensity induces its own brightness scale, with idiosyncratic distances between points, on which surface lightness values fall. Since surface brightness depends on the illuminant, changes along one achromatic dimension under an illuminant shift cannot be offset by changes on the other. Identical appearance matches for surfaces under different illuminants are impossible and subjects can do no better than choose pairs of differently illuminated stimuli that appear least dissimilar to one another. Thus appearance-based constancy is best understood as a matter of a high degree of similarity, rather than a sameness, of phenomenal color across different viewing conditions. Such interactions need not be limited to albedo and illuminant intensity, but could extend to chromatic properties (see Mausfeld 2003, Tokunaga and Logvinenko 2010). The further consequences of such interactions for the analysis of appearance matching data is an important topic that deserves extended discussion elsewhere.

Returning to an issue raised in section 2, the existence of aspects of color appearance determined by contributions from both the illuminant and surface reflectance implies that perceived color hinges on processes that separate the retinal stimulus into objects and illumination. Thus the illuminant is not simply (as is often stated) discarded altogether so that a sameness of object color can be achieved across varying circumstances. Mausfeld (2003, p.417) notes that the dependence of perceived color on representations of both surfaces and illuminants would "not necessitate that the illumination is also phenomenally present as a distinct separate representation." Rather, we might phenomenally experience surfaces, their colors, and the way they are illuminated. Hilbert (2005, pp.150-151) endorses this plausible idea.

Hilbert (2005, p.152) finds in the existence of more than three dimensions of color experience a possible defense of phenomenal color constancy from the challenge posed by instruction effects in matching tasks:

> Suppose … that color experience has [five dimensions] … [Even] if color constancy is perfect … there may be no exact match on all five dimensions. Because the illuminants are different it can be the case that even if there is a perfect match for hue, lightness, and chroma (the dimensions related to reflectance), there may still be differences in brightness and colorfulness. The instructions could have the effect of changing the relative weight of the various attributes of appearance determining the best match.

Hilbert (2005, pp.150-151) points out that theories which yoke phenomenal color to an estimate of a surface's spectral reflectance properties can admit that stimuli with identical reflectance properties may differ phenomenally along dimensions that include effects of the illuminant. The idea is to deny that color constancy is a matter of sameness along all dimensions of color appearance and to construe it as having to do only with sameness of color appearance along dimensions that map onto surface reflectance properties alone. As he parenthetically states, Hilbert takes hue and saturation to be of this kind.

As is evident from the earlier discussion, instruction effects on matching performance do not threaten the idea that there can be a constancy of phenomenal color; leading supporters of projectivist views (such as Amano, Arend, Foster, and Reeves) acknowledge this. It has also been pointed out that there is no direct route from such effects to the existence of a projective form of color constancy. Thus Hilbert's reply might be directed toward a problem that does not exist. It is plausible that there could not be any sort of constancy of perceived color in situations in which pronounced instruction effects are found, because of the difficulties facing the visual system's attempts to compensate for the effects of the illuminant and form a representation of surface reflectance properties. In that case, it would be more appropriate for Hilbert to treat the instruction effects as no concern at all for a purely phenomenal conception of color constancy, instead of appealing to illumination-dependent dimensions of color appearance to explain away a seeming problem for such a conception.

Nonetheless, something is amiss in Hilbert's reply to the (supposed) challenge from instruction effects to a purely phenomenal understanding of color constancy. It is reasonable to expect that subjects place the greatest weight on hue and saturation in the hue/saturation task. In addition to the studies that find an instruction effect on matching performance when the hue/saturation and surface tasks are given to different subject groups, there is a clear separation of CCIs for individual subjects who both rated hue and saturation and made judgments about surface properties; see Reeves et al (2008, p.225, fig.4). The implication is that, in certain experimental conditions, subjects' judgments about hue and saturation do not reliably track surface reflectance properties. The existence of residual differences along dimensions of color appearance involving the interplay of illumination and surface reflectance when hue and saturation are matched does not change that. Hilbert could counter by insisting that subjects attend most to hue and saturation in the surface matching task and most to illumination-dependent attributes in the hue/saturation task. However, there appears to be no principled basis for that move and it sits uneasily with the instruction and training regimen of Reeves et al (2008). See Cohen (2008, pp.67-69) for a similar discussion.

6. A proposal

The preceding examination suggests a proposal about how color constancy should be understood. The proposal has two central claims related to phenomenal and projectivist conceptions of color constancy. The first is that there is a constancy of phenomenal color in conditions that allow for the contributions of slower adaptation mechanisms and reliable estimates of illuminant properties. This constancy will rarely be perfect and might often be considerably imperfect, due to the retina's limited wavelength sampling of the confounded color signal and factors that can limit the accuracy of the illuminant estimate. In situations in which a representation of the illuminant figures in the phenomenally conscious color percept through dimensions of interaction with reflectance properties, phenomenal color constancy is a matter of a high degree of similarity, rather than sameness, of color appearance for a given surface across different illumination conditions. The main sources of support for this claim were discussed in sections 2 and 5.

The second claim is that there can be a projective constancy of color, despite the points made earlier about the existence of projective color constancy not following directly from instruction effects. The sense that surface colors are stable across different viewing conditions, even ones that do not allow for significant compensation for the shift, is undeniable. Projectivism presents a highly plausible way of helping to account for that impression. The key is to restrict the contributions of such projections to specific circumstances, ones that make projective color constancy dependent upon phenomenal color constancy and relational color constancy.

As noted in sections 2 and 3, in many situations which limit illuminant compensation, it is still possible to determine whether two surfaces have the same reflectance properties due to the invariance of cone excitation ratios. Without a determination of individual surface reflectance properties, though, there can be no stability of perceived object color. It bears repeating that the cone excitation ratios in question are not between different cone types for a single surface, but rather for each given cone type between different surfaces (or other scene elements). The minimalistic stimuli of Amano et al (2005) are one clear instance of a situation in which subjects can make relational but not absolute judgments about reflectance properties. This would also seem to be the case for some richer stimuli, such as Mondrian checkerboards, that are presented without adaptation to either the standard or the match stimulus. However, in situations in which one is already highly adapted to an illuminant and thus able to (approximately) determine the particular reflectance properties of a given surface, if an abrupt (spatial or temporal) illuminant change is introduced, it is possible to assign the same surface color across that change, despite variation in phenomenal color across the change. Examples of such changes would be a lawn half in sunlight and half in shadow or a room in which the sole source of illumination suddenly changes from fluorescent tubes to a halogen lamp by the flip of a switch. The shadowed and sunlit portions of the lawn do not have the same phenomenal appearance, but the grass in shadow could be determined to have the same material properties, and hence the same surface color, as the grass that appears some shade of green in the sunlight. It is in dealing with such situations that projectivism is most compelling. In other cases, projective color constancy is either otiose, as phenomenal color is approximately constant, or impossible, as surface reflectance properties cannot be extracted from the color signal.

This proposal conflicts with strong forms of projectivism, such as the Helmholtzian idea that color constancy is a matter of an unconscious inference to the actual colors of objects. Nonetheless, its projective component does not diverge all that far from the main ideas expressed in more careful statements by other projectivists. For example, Reeves et al (2008, p.226) conclude with the observation that "even when the quality of a particular chromatic change alters perceived hue and saturation, observers can reliably infer the cause – namely, the constancy of the underlying reflecting surface." In that sense, perceivers can separate their judgments about surface properties from their judgments about color appearances and there is no disagreement between this claim and the current proposal. However, surface property judgments can involve a stable projection of color "onto the physical world as an object property" (ibid.) only in cases where phenomenal color supplies a reliable basis for attributing a color to at least one of the regions involved. As Foster (2003, pp.441-442) stresses, the relational judgments that enable telling whether two surfaces share the same reflectance properties do not reveal the absolute reflectance properties of either surface. Setting aside potential contributions of color memory (see below), when two surfaces can be determined to share the same material properties but both are presented in circumstances that prevent extraction of surface reflectance information, a perceiver cannot accurately assign a surface color to either. In such cases, one could infer that were those two surfaces presented under the same illuminant, they would appear the same in color (Cohen 2008), but one would be unable to say what that color is. Only when the surface color of one of the stimuli is determined – which requires conditions that facilitate phenomenal color constancy – can a stable projection of surface color onto the other take place.

To further elaborate, consider the above example of a lawn that is half in direct sunlight and half in shadow. Assume a perceiver adapted to the sunlight. That perceiver should achieve a high degree of phenomenal color constancy between her occurrent visual episode and past viewings of the lawn when she had been adapted to the (potentially different) illumination. However, given the quick luminance change, low color contrast, and bluish cast characteristic of shadows, her phenomenal color experience across the change of illuminants currently falling on the two patches of grass is likely to vary. Nonetheless, due to relational constancy, the perceiver sees the entire lawn as being (approximately) the same with respect to its material properties. This registered sameness of material properties induces an automatic attribution of surface color to the shaded area based on the greenish color appearance of the sunlit portion of the grass; statistical properties of the current scene and the perceiver's current state of adaptation would likely be relevant to which of the two regions is used as the basis of the attribution. Once the surface color of some object in the non-adapted region is determined in this way, the visual system might also be able to determine (and project) the surface colors of other objects in that region, even ones that have no counterpart in the adapted region; e.g., a purple ball that rests only in the shaded region of the lawn while there are no purple objects in the sunlit portion.

Compare the above example of the lawn with a situation in which the perceiver is adapted to sunlight and is then presented with a scene in an unfamiliar room in which two regions are separately illuminated by strongly deep-bluish and orangish lights, with no gradient between them. Not only will color appearance vary across the illuminant shift, the perceiver's state of adaptation is such that were she ever to view this same scene in sunlight, phenomenal constancy between that episode and her current experience of either of the two regions now before her would be weak or non-existent. Relational constancy enables her to recognize the sameness of material properties across the illuminant change; e.g., the floor and the walls do not appear to change their reflectance properties where the lighting abruptly turns from deep-blue to orange. However, since the perceiver cannot reliably determine the absolute reflectance properties of (say) the floor under either illuminant, she is in no position to accurately attribute a surface color to it. Additionally, considering the poor fit between statistical properties of the two regions and her current state of adaptation, any attribution of surface color that takes place might be tentative or unstable.

It is not being suggested that all cases in which perceivers are able to assign a (roughly) stable color to a surface in conditions that prevent phenomenal color constancy are instances of projective color constancy. In some such conditions, perceivers seem able to use color memory to assign colors to familiar objects or objects with typical colors; as an aside regarding the plausibility of color memory effects, there is evidence that color memory can affect phenomenal color (Hansen et al 2006). A memory-based attribution of surface color could be achieved independently of whether the perceiver has the sort of information about absolute and relational reflectance that is needed for projective constancy. Of course, like projective constancy, this use of color memory relies on phenomenal color experience to assign surface color, only now it is past experience that is relevant. Surface color attributions of this sort are much more cognitive than visual or visuocognitive, as they do not emerge from the visual system compensating for the effects of the current illumination and making use of relational information about reflectance.

Thus this memory-based attribution of surface color is a different phenomenon than the one currently considered.

The understanding of projective color constancy recommended here should be of scientific interest. It combines well-established empirical findings in a way that explains a phenomenon that is evident in introspection and that has been appealed to, in some form, by a number of vision researchers going back one thousand years. Given what is known about the abilities to determine absolute and relational reflectance information, there is good reason to believe that perceivers are capable of the sort of projective color constancy proposed here. The information needed to achieve projective constancy is available to perceivers and already used for other purposes. Admittedly, though, empirical work is needed to establish that perceivers actually have and make use of a mechanism that produces the hypothesized projections based on the information acquired about absolute and relational reflectance. For now, phenomenological observations and empirical plausibility are the main sources of support for the proposal. A virtue of this conception of projective color constancy is that it is clear enough about the sorts of information and processes needed to achieve it, to provide useful guidance regarding how it should be experimentally investigated.

Were it to turn out that this form of color constancy exists, it would bear on our understanding of the ways in which perception informs us about and enables us to act in the world. Among other things (and picking up on remarks from section 2), it would be suggestive regarding the nature of perceptual representations of objects and the processes that exploit them. Loosely using the object-files framework for purposes of illustration (Kahneman et al 1992), it would seem that some entries in an object's representation code color in a way that is phenomenally present while other entries for color are phenomenally silent. In some circumstances, the former are accessed while in others the latter are drawn on. This, in turn, would have consequences for what is to be made of the role of phenomenal experience in guiding us around our surroundings. For example, even when it comes to basic perceptual qualities like color, we ought not suppose (as is intuitively appealing) that the phenomenal appearances of things are the sole source of the perceptual information that feeds into decision-making and action-guiding processes.

The proposal just sketched in terms of two forms of color constancy offers an intuitive way of accounting for empirical findings and phenomenological observations. However, before concluding, it is worth entertaining an alternative reading, one consistent with concerns voiced earlier. This alternative does not suggest a revision to core details of the preceding proposal, but it is potentially relevant to broader considerations regarding how the kinds of perception theory are to be individuated and studied. One might argue that instead of two forms of color constancy, there is really only one form of color constancy – phenomenal color constancy – and that the other supposed form of color constancy is an ability that results from the interaction of phenomenal color constancy and certain mechanisms of scene perception. This ability often enables us to compensate for limitations and trade-offs inherent in phenomenal color constancy. Since that further ability is crucially dependent on both phenomenal color constancy and capacities related to scene perception, and the relevant scene perception capacities are unconcerned with the specific surface properties of objects, perhaps it should be treated (at least for some purposes) as a wholly separate phenomenon. The idea is that projective constancy might be simply a byproduct of the organization of the rest of our visual system and cognitive machinery, and that it is similar to phenomenal color constancy only in very broad terms but not in a way that justifies gathering them together under the same umbrella from the standpoint of perception theory. It is plausible that projective constancy resulted from evolutionary processes seizing on (or stumbling across) certain fortuitous consequences of the "informational leakiness" of minds/brains. However, another possibility is that the considerable use made of projective constancy stems from cultural or technological factors rather than a specific function of the visual system itself. Whether any ideas in this neighborhood get traction depends on what might be learned in future research about the specific contributions of phenomenal and projective constancy to how we and other animals get around in the world. Following Mausfeld (2003), success in that pursuit depends on first answering questions about the structure and content of perceptual experience that have thus far not received sufficient attention and achieving better understanding of the interdependencies between perception, cognition, and action.

7. References

Alhazen, A. 1030/1989. The Optics of Ibn al-Haytham, Books I–III. Edited by A.I. Sabra. London: Warburg Institute.

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. and A. Reeves. 1986. "Simultaneous color constancy." Journal of the Optical Society of America A, 3, 1743-51.

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: MIT Press.

Brainard, D. 1998. “Color constancy in the nearly natural image. 2. Achromatic loci.” Journal of the Optical Society of America A, 15, 307-325.

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

Brainard, D., J. Kraft, and P. Longère. 2003. “Color constancy: developing empirical tests of computational models.” In Mausfeld & Heyer (2003, pp.307-328).

Brown, R. 2003. “The Yin and Yang of colour constancy.” In Mausfeld & Heyer (2003, pp.247-272).

Cohen, J. 2008. “Colour constancy as a 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.

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

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

Goffaux, V., C. Jacques, A. Moraux, A. Oliva, P. Schyns, & B. Rossion. 2005. “Diagnostic colors contribute to the early stages of scene categorization: behavioural and neurophysiological evidence.” Visual Cognition, 12, 878-892.

Hansen, T., M. Olkkonen, S. Walter, & K. Gegenfurtner, K. 2006. “Memory modulates color appearance.” Nature Neuroscience, 9, 1367-1368.

Helson, H. 1964. Adaptation Level Theory. Oxford: Harper & Row.

Henderson, J. & A. Hollingworth. 1999. “High-level scene perception.” Annual Review of Psychology, 50, 243-71.

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

Hoffman, D. 2003. “Colour construction.” In Mausfeld & Heyer (2003, pp.273-274).

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.

Izmailov, C. & E. Sokolov. 1991. "Spherical model of color and brightness discrimination." Psychological Science, 2, 249-259.

Joost, U., B. Lee, & Q. Zaidi. 2002. "Lichtenberg's letter to Goethe on 'Färbige Schatten'." Color Research and Application, 27, 300–303.

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

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.

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

Maloney, L. & J. Yang. 2003. "The illumination estimation hypothesis and surface color perception. In Mausfeld & Heyer (2003, pp.335-360).

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.

Oliva, A. & P. Schyns. 2000. "Diagnostic colors mediate scene recognition." Cognitive Psychology, 41, 176-210.

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

Rensink, R. 2000. "Seeing, sensing, and scrutinizing," Vision Research, 40, 1469-1487.

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

Schyns, P. & A. Oliva. 1994. "From blobs to boundary edges: evidence for time and spatial scale dependent scene recognition." Psychological Science, 5, 195-200.

_____. 1997. "Flexible, diagnosticity-driven, rather than fixed, perceptually determined scale selection in scene and face recognition." Perception, 26, 1027–1038.

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

Smithson, H. & Q. Zaidi. 2004. “Color constancy in context.” Journal of Vision, 4, 693-710.

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.

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