Perceptual Constancy

Jonathan Cohen[*]

Our eyes deceive us when we look down railway tracks, but our brains do not. The rails appear to converge in the distance, but we know that the rails are parallel. We know that they are the same distance apart a mile down the track as they are where we are standing, so the brain says, "The tracks only appear *to converge because they are distant." But how does the brian know that the tracks are distant? The brain answers, "They* must *be distant because they appear to converge." (The flow of this logic must shock computer programmers, but they are accustomed to the limitations of inferior hardware.)* (Hunter *et al.*, 2007, 82)

1 Introduction

Students of perception have long known that perceptual constancy is an important aspect of our perceptual interaction with the world. Here is a simple example of the phenomenon concerning color perception: there is some ordinary sense in which an unpainted ceramic coffee cup made from a uniform material looks a uniform color when it is viewed under uneven illumination, even though the light reflected by the shaded regions to our eyes is quite different from the light reflected by the unshaded regions to our eyes (see figure 1). Or consider this example concerning size perception: there is some ordinary sense in which two telephone poles look the same size when the first is viewed from 100 meters and when the second is viewed from 1 meter, even though the visual angle subtended by the two poles on our retinae is very different (see figure 2). Or consider this example concerning shape perception: there is some ordinary sense in which a penny looks round both when viewed head on and when viewed from an acute angle, even though the area projected by the penny onto our retinae under these two conditions is very different (see figure 3). Or, finally, consider this example concerning auditory volume perception (which I cannot depict graphically): there is some ordinary sense in which a speaker's voice sounds the same volume when heard from across the room and when heard from a distance of 1 meter, even though the energy striking our ears under these two conditions is very different.

[PUT FIGURE 1 ABOUT HERE.]

---

[*] Department of Philosophy, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0119, joncohen@aardvark.ucsd.edu

[PUT FIGURE 2 ABOUT HERE.]

[PUT FIGURE 3 ABOUT HERE.]

The kind of perceptual constancy exemplified in these cases, and others like them, is ubiquitous, ordinary, and central to the way perception tells us about the world in which we live. Without this kind of constancy, we would experience the world as a Jamesian blooming, buzzing confusion — a constant flux of colors, shapes, and sounds with no apparent organization. For, unavoidably, the perceptual signals incident on our transducers are the results of not only the kinds of distal individuals there are and properties they exemplify, but also the constantly changing details of the circumstances under which we perceive (the angle and distance from the perceived object, the lighting conditions, the ambient noise, our own cognitive and perceptual histories and futures, our expectations, and so on). If perception were incapable of representing the world as in some ways constant despite various changes in our perceptual circumstances, it would radically misrepresent the distal world: it would fail to reveal ways in which the world is stable. And since these ways underpin our engagement with that world, this would (disastrously) undermine the possibility of effective action and empirical knowledge.

However, despite its recognized ubiquity and importance, there are several respects in which the phenomenon of perceptual constancy is poorly understood. Aside from the independent interest in getting clear on these matters, perceptual constancy has figured prominently in recent debates about the ontology of colors and other sensible qualities, knowledge, attention, mental modularity, the contents of mental representation, and the objectivity of our representations of the world.[1] Therefore, in this essay I'll review some of what is and is not known about perceptual constancy with an eye to drawing connections with ongoing controversies in the philosophy of perception and elsewhere.[2]

2 Perceptual Constancy as Perceptual Stability

As both its name and the initial examples used to introduce the phenomenon above suggest, perceptual constancy is, in some sense yet to be explained, about the absence of change. Indeed, the textbook characterization has it that perceptual constancy is nothing more or less than a stability in perceptual

---

[1] Recently a number of philosophers have returned to issues about constancy anew; for example, see Hilbert (2005); Thompson (2006); Cohen (2008); Bradley (2008); Hatfield (2009); Gert (2010); Matthen (2010); Wright (2013). Also see Burge (2010), for whom perceptual constancy is used as a touchstone for the objectivity of intentional representation quite generally.

[2] Because there is vastly more research, by both philosophers and psychologists, on perceptual constancy in vision than in other modalities (and, even more particularly, on color constancy), this entry is, regrettably, unavoidably visuocentric in its choice of examples and theories discussed. There remains much work to be done in this area.

response across a range of varying perceptual conditions.[3] Thus, in the case of the unevenly illuminated coffeecup (figure 1), the idea is that the perceptual system represents the distinct regions of the cup as bearing the same color even though there is variation in the illumination incident on them (and, therefore, in the total amount of light energy they reflect to our retinal transducers). Or, again, in the case of volume perception, the thought is that perception represents the speaker's voice as having the same volume even though there is significant variation in the distance from which it is heard (and, therefore, in the total amount of auditory energy absorbed by our aural transducers).

While I will want to qualify the above characterization in what follows, one of the ways in which it is useful and interesting is that it presents perception as an active process of engagement with the world. It suggests that perception is not just a matter of passively registering the impinging energy array, but of somehow articulating or decomposing that array to arrive at a representation of a subset of the distal features that contribute to the configuration of the array.

Unfortunately, the textbook characterization of perceptual constancy just presented can't be quite right by itself. (Or, alternatively, we can retain that characterization by itself, but only at the cost of emptying the phenomenon of all of its instances). For it is not true that our perceptual responses are entirely constant in the kinds of cases at issue. Returning once again to the unevenly illuminated coffeecup, we know there must be a difference in a subject's perceptual response to the shaded and unshaded regions of the cup, or else she would be unable to discriminate the luminance boundary between them. Likewise in canonical cases of size constancy (subjects' perceptual responses can clearly distinguish in some size-related way between the perception of the telephone pole at 100m and the perception of the telephone pole at 1m), shape constancy (there is clearly a discriminable difference between the subject's perception of the penny seen head on and her perception of the penny seen at an acute angle), auditory volume constancy (there is clearly a discriminable difference between the subject's perception of the speaker's voice from across the room and her perception of the speaker's voice from a distance of 1 meter), and all of the other canonical instances of perceptual constancy.

Indeed, the *non-constancy* of our perceptual responses across variations in the perceptual circumstances is not only immediately apparent, but underlies another much-observed and much-discussed aspect of perception — the phenomenon of perceptual contrast.[4] It is easy to find instances of perceptual contrast once one begins to look for them. For example, figure 4 illustrates an instance of simultaneous lightness contrast: although the two central patches depicted here are qualitatively intrinsically identical, the perceptual system represents them as different in color because of the different ways in which they contrast in lightness with surrounding items. Simultaneous lightness contrast plays a role in many classic visual illusions, such as the appearance

---

[3] See, for example, Byrne and Hilbert (1997, 445), Zaidi (1999, 339), Palmer (1999, 312–314, 723), Goldstein (1999, 567), Brainard *et al.* (2003, 308–309).

[4] Whittle (2003) provides an excellent overview of the importance of perceptual contrast for color vision.

of grey dots at the intersections of an achromatic grid (the Hermann grid illusion, figure 5), the interpretation of a pair of opposed lightness gradients as two constant lightness regions separated by an edge (the Cornsweet illusion, figure 6), and the appearance of light or dark bands next to the boundary between two different lightness gradients, even when the lightness on both sides of the boundary is the same (Mach bands, figure 7).[5]

[PUT FIGURE 4 ABOUT HERE.]

[PUT FIGURE 5 ABOUT HERE.]

[PUT FIGURE 6 ABOUT HERE.]

[PUT FIGURE 7 ABOUT HERE.]

Perceptual contrast is by no means restricted to the perception of lightness/brightness; within vision there are also simultaneous contrast effects for chromatic color, size, spatial frequency, orientation, motion, and speed, *inter alia*. For example, figure 8 illustrates an instance of simultaneous size contrast: although the central circles are the same geometric size, the perceptual system represents them as different in size because of the contrast with the different elements surrounding them. Moreover, in addition to *simultaneous* contrast — contrast between simultaneously perceived items, there are also ubiquitous instances of *successive* contrast — effects of contrast between successively perceived items for each of these dimensions. And, of course, contrast occurs in non-visual modalities as well (although there is *much* less systematic investigation of contrast outside vision). Thus, in gustation, we commonly observe that sweet wines strike us as markedly less sweet when consumed with dessert items (which contain much more sugar than the wines) than on their own. In audition, we find that it is much easier to detect variations in pitch (say, while tuning a guitar string) by contrasting the target against other (simultaneously or successively perceived) tones. Or, again, in kinaesthesia, Gibson (1933) reports that after blindfolded subjects run their fingers over a curved surface for three minutes, straight edges seem to them to be curved in the opposite direction.

[PUT FIGURE 8 ABOUT HERE.]

In each of these cases, the perceptual system reacts differently to objects depending on how they contrast with other perceived items. Perceptual contrast occurs because perceptual systems tend to be responsive to magnitude differences, as opposed to magnitudes themselves.[6] For our purposes, the phenomenon of contrast is important because it makes for a vivid demonstration of the observation made above: contrary to the textbook characterization, our perceptual responses to an object/property are *not* constant, but instead

---

[5] For a discussion of the role of contrast in many lightness illusions, see Adelson (2000).

[6] The standard physiological explanation of this generalization turns on lateral inhibition between neurons carrying perceptual information (e.g., retinal ganglion cells, in the case of lightness perception). Lateral inhibition results in the suppression of all but the most stimulated/least inhibited neurons; consequently, the overall firing pattern is highest in cells corresponding to parts of the stimulus where there is a steep spatial/temporal gradient — where a small population of most active cells is left relatively uninhibited by the firing of their neighbors.

change in interesting and systematic ways across variations in the perceptual circumstances.[7]

3 Psychophysics and Measurement

So far our discussion has been framed by questions of which qualitative discriminations are made by perceivers. However, for many purposes it is useful to have quantitative measures of similarity/dissimilarity in cases of perceptual constancy. The standard technique used for this purpose is to measure the dissimilarity between a subject's reaction to two stimuli by measuring how much of a change she must make to one of them, holding the other fixed, before she regards the two as a perceptual match.[8]

Thus, for example, the main quantitative measure by which contemporary psychophysicists assess color constancy, known as asymmetric color matching (Wyszecki and Stiles, 1982, 281–293), involves asking subjects to change the chromaticity (or lightness, in lightness constancy experiments) of a test patch under one illuminant until it perceptually matches a standard patch under a different illuminant. The size of the chromaticity (/lightness) difference between the test and the standard patches required to achieve a perceptual match, then, is a quantitative measure of the effect of the illumination difference between test and standard patches on the subject's total perceptual response to them — it is an operational measure of the extent to which perceptual responses are unchanging across variations in perceptual conditions.

Such quantitative measures reinforce the assessment made above on the strength of qualitative reactions: in canonical instances of color constancy, subjects' perceptual responses are not simply unchanging — rather, they are in some respects similar or unchanging and in some other respects dissimilar or changing. Moreover, interestingly, (most) subjects can be made to switch between attending to the respects of similarity and the respects of dissimilarity in many canonical instances simply by changing the experimental instructions. For example, Arend and Reeves (1986) found that subjects in an asymmetric color matching paradigm responded to instructions to "adjust the test patch to match its hue and saturation to those of the standard patch" (1744) by making large chromaticity changes (suggesting that their perceptual systems initially represented the test and the standard patch as quite different), although the

---

[7] Objection: The cases I have used to highlight contrast (the Hermann grid illusion, the Ebbinghaus illusion, etc.) are often put forward as textbook cases of perceptual illusion. They give no reason to suppose there is substantial non-constancy in veridical cases of perception.

Response: Contrast is pretty clearly at work in ordinary perception; I have relied on textbook cases of perceptual illusion only because they make the results of perceptual contrast so vividly apparent. However, a theory of perception that set aside cases involving the operation of perceptual constancy would have little to say about the kinds of perceptual systems we happen to enjoy.

[8] Note that perceptual matching is a statistical notion: two stimuli count as a perceptual match for a subject if the subject is unable to discriminate one from the other over several presentations at a rate higher than that attributable to chance.

same subjects responded to instructions to "adjust the test patch to look as if it were 'cut from the same piece of paper' as the standard, i.e., to match its surface color" (1744) by making very small chromaticity changes (suggesting that their perceptual systems initially represented the test and the standard patch as quite similar).[9]

4 Stability and Instability

It seems, then, that the right thing to say is not, or not just, that the perceptual system responds in a constant or unchanging way in the face of variations in the perceptual conditions — either as a general matter or even in the cases that have been put forward as parade instances of perceptual constancy. On the other hand, neither does it seem that the perceptual system responds by treating objects as merely *approximately* the same in different perceptual conditions — the similarities and dissimilarities that perception recognizes are not collapsed into a single scalar value somewhere between the extremes of perfect qualitative match and perfect qualitative mismatch. Rather, what we should say is that perception represents both some aspects of similarity and some aspects of dissimilarity in its responses to objects across changing perceptual circumstances. Moreover, we should recognize that both the respects of similarity and the respects of dissimilarity are in many cases available to the perceiving subject for the purpose of making perceptual discriminations.[10]

---

[9] That the perceptual system displays this sort of bimodal behavior has been understood for a long time; see Evans (1948, 163–164); Beck (1972, 66–67) for an overview of some of the earlier work. For more recent work (mostly on cases of simultaneous color constancy), see Blackwell and Buchsbaum (1988); Valberg and Lange-Malecki (1990); Arend *et al.* (1991); Troost and deWeert (1991); Cornelissen and Brenner (1995); Bäuml (1999). While there has been far less systematic investigation of this effect with respect to cases of successive color constancy, investigators have found the same sort of bimodal pattern of results here too (Delahunt (2001, 114–117); Delahunt and Brainard (2004, 71–74)).

[10] Many philosophers and psychologists working in this area have tended to be so impressed by the constant aspects of our perceptual responses that they have played down, dismissed, or, more frequently, just ignored the inconstant aspects of our perceptual responses to the same scenarios. Thus, one sometimes sees assertions to the effect that the inconstant aspects of perception are "unnatural and sophisticated ... [and] difficult to attain" (Smith, 2002, 182; cf. 178). Whatever else we think of such claims, I suggest that an adequate theory of perception must account for all of the ways in which perceptual systems respond to the world rather than only some of them — whether these responses are natural or unnatural, naive or sophisticated, and easily attained or not.

Emphasis on constant aspects of our perceptual responses at the expense of inconstant aspects also shows up in a prominent line of argument for the view that colors are illumination-independent features of objects (I discuss these arguments critically in Cohen, 2008). For example, Tye (2000, 147–148), Hilbert (1987, 65), and Byrne and Hilbert (2003, 9) explicitly appeal to constancy reactions in color perception as cases where the very same feature can be extracted despite variation in the ambient illumination, and infer from this claim that color (which they reasonably assume is indeed represented by color perception) is itself illumination-independent. However, if it is reasonable to take constancy reactions to show that perception represents constant features, it is no less (and no more) reasonable to take inconstancy reactions to show that perception represents inconstant features. But if color perception represents both constant and inconstant features, there is no sound inference from the premise that color is represented

This raises an important puzzle for the understanding of perceptual constancy. Given that there is clearly substantial variation in our perceptual responses to objects across changes in perceptual circumstances even in canonical cases of constancy (such as those used to introduce the topic in §1), it won't do to think of constancy simply in terms of stability of perceptual response. Rather, if we want to be able to say that there is perceptual constancy in such canonical cases, then we owe a characterization of just which kinds of perceptual similarity, in the context of just which kinds of variation in perceptual circumstances, are necessary for the exemplification of perceptual constancy. Moreover, we need a characterization that is applicable across the broad range of cases to which we want to apply the notion. Unfortunately, there is at present no adequate and fully general characterization of this sort, and therefore no general understanding of what perceptual constancy amounts to.

5 Computation and Constancy

While the problems just discussed should not be underplayed, neither should they make us lose sight of the initial observation that makes perceptual constancy so interesting: in canonical cases there *is* some interesting respect in which perception is unchanging in its treatment of an object despite differences in the conditions under which it is perceived, and despite the attendant differences in the total signals impinging on our sensory transducers, even if these must be characterized in a case by case way.

This observation naturally invites the important question about how perception pulls off the feats of constant representation in the face of inconstant perceptual circumstances that it does. That is, given the complex total signal striking the transducers — a signal that is determined jointly by the features of perceived objects and perceptual circumstances, and therefore that changes as circumstances vary — how does the perceptual system arrive at a verdict about whether the perceived objects change? How, for example, does the perceptual system start with the varying array of light intensities reflected by the cup in figure 1 and end with the information that the entire cup is uniform in color (or, more cautiously, in some color-related respect)?

A burgeoning subfield of perceptual psychology has attempted to build empirically adequate computational models that would answer this question. Perhaps the dominant approach within this tradition is to think about perception as computing a solution to an "inverse problem": the job is to find ways of factoring apart the complex resultant that is the impinging energy array to arrive at a representation of the distal features that contribute to the

---

by perception to the conclusion that color is a constant (here, illumination-independent) feature. (Nor, for that matter, is there a sound inference from that premise to the conclusion that color is an inconstant/illumination-dependent feature.) Consequently, the sort of appeal to perceptual constancy made by these authors does not successfully motivate the claim that colors are illumination-independent object features.

resultant. Thus, for example, consider color constancy once again, since that is the area in which the most intense research on computational methods has been carried out.[11] In color perception the perceptual system begins with an array of light intensities on the retina which is the joint product of two factors — the features of the illumination incident on surfaces and those of the surfaces that reflect light to our eyes. The leading approach to computational color constancy has involved finding methods of estimating the properties of the illuminant so that the system can, as it were, subtract off this factor from the total signal (in Helmholtz's phrase, "discounting the illuminant"), leaving an illumination-independent characterization of the reflecting surface (Maloney and Wandell, 1986; Brainard *et al.*, 1997; Brainard, 1998). Crucially, since this characterization is illumination-independent, the thought is that it will be shared by distinct regions of a uniform surface that happen to be illuminated differently (e.g., the regions of the cup in figure 1). Therefore, a perceptual system that performed this sort of computation would be able to treat such regions as (in this one respect) perceptually similar, even though they are clearly discriminably different.

Modellers have pursued a wide variety of strategies for estimating the separate contributions to the retinal array made by illuminants and surfaces. For example, Maloney (1986); Maloney and Wandell (1986) show how a system with more classes of receptors than there are degrees of freedom in (the system's linear models of) surface reflection profiles can exploit its multiple receptoral signals to recover representations of surfaces. Other approaches solve the inverse problem by adding as constraints assumptions about the kinds of scenes perceptual systems will encounter. Thus, Buchsbaum (1980) proposes a model that rests on the assumption that the median lightness value in a scene corresponds to a middle grey surface, and computes from this assumption what the incident illumination would have to be to result in the observed intensity array. A related but distinct strategy proceeds from the assumption that anchors some part of the visual image (rather than a mean) to an extremal lightness value — for example, by treating the lightest visible surface as white (Land and McCann, 1971; Gilchrist *et al.*, 1999). Others have proposed estimating illuminants from information about mutual reflections in the scene (Funt *et al.*, 1991), the boundaries of regions known to be specular reflections (D'Zmura and Lennie, 1986; Lee, 1986), and shadows (D'Zmura, 1992). Still others propose to solve the inverse problem by appeal to higher-order scene statistics, such as the correlation between redness and luminance within the scene (Golz and MacLeod, 2002) or the statistical distribution of colors within the scene (MacLeod, 2003; Brainard *et al.*, 2006). In recent years, many theorists have advocated "Bayesian" probabilistic models as solutions to the illuminant estimation problem. According to Bayesians, the visual system first selects as its estimate that hypothesis about the illuminant with the highest

---

[11] Much of the work in this tradition is restricted to the perception of surface colors (as opposed to the colors of lights, volumes, films, and so on). Moreover, many (but not all) of the models depend on the simplifying assumptions that surfaces are illuminated by constant or smoothly varying, and exclusively diffuse, illumination.

probability conditional on the data received by the transducers, constrained by the prior probability of that illuminant hypothesis; then it goes on to select as its estimate about distal surfaces that hypothesis with the highest probability conditional on the transducer data and the illuminant estimate obtained at the first step, again constrained by prior probabilities assigned to the various hypotheses about surfaces (Brainard and Freeman, 1997).[12] It is possible, of course, that human color constancy involves a combination of these methods, or others.

However, there is a different class of computational models for perceptual constancy — one that has received much less attention from philosophers — that rejects the assumption that constancy requires factoring out of the perceptual signal a representation of the distinctive contribution made by the perceived object and its features. Thus, Craven and Foster (1992); Foster and Nascimento (1994); Dannemiller (1993); Zaidi (1998, 2001); Amano *et al.* (2005) suggest that perceptual systems compute color constancy not by deriving an illumination-independent representation of object surfaces, but by comparing total perceptual signals in light of what is known about the illumination or other properties of the total scene. Crudely, the idea is that the system can ask whether the difference between the two perceptual signals it gets from two perceptual episodes (simultaneous or not) can be accounted for by the behavior of the illumination (rather than by a difference in the surfaces perceived on the two occasions). If, say, the system represents that the illumination profile includes a shadow cast over the scene (say, by a partially occluded light source) then this would have predictable effects on the perceptual signal: there would be higher intensities in the (portion of the) signal corresponding to the directly illuminated regions and lower intensities corresponding to the (portion of the) signal corresponding to the region in shadow. Therefore, the system can treat the image regions as being relevantly alike although they cause different perceptual signals (i.e., it can display perceptual constancy) if it can conclude that the two different perceptual signals lie in the graph of a transformation consistent with illumination variations.

Here, as in more traditional computational models, the computation of color constancy depends on deriving from the perceptual signal an estimate of the illumination. But unlike more traditional models, the suggestion is that the system can compute constancy directly from the perceptual signal and the illumination estimate, without going to the trouble of separately deriving a closed-form representation of object surfaces. Also unlike more traditional models, here there is no suggestion that the perceptual system discounts or discards the illuminant — on the contrary, the claim is that the system's continuing to represent the illuminant is absolutely vital to the computation

---

[12] In such models, the kinds of substantive assumptions about the distal world that ground the deterministic models described above — e.g., about the way illuminants vary smoothly in ecological settings, about where the mean lightness values can be expected, and so on — show up as well, but here in the form of the prior probabilities about both illuminant and surfaces used to constrain the assignment of posterior probabilities.

of constancy.[13] And, though these are proposals about color constancy in particular, the general lessons they teach may well be applicable for other visual and non-visual instances of perceptual constancy as well.

6 Is Perceptual Constancy Perceptual?

Perceptual constancy shows that perceivers are not passive receivers of the array of energy falling on their receptors — for if they were, they could not react in similar ways (in some respects), as they sometimes do, when there are large differences in that array. Something more must be going on. But is that something more a *perceptual* process? Or is it a post-perceptual process that gets its start at the point where perception ends? It is clear that, for example, subjects will (under some experimental instructions) judge that the penny in figure 3 is relevantly alike in shape when presented from two distinct viewpoints. But what is not clear is whether that judgment is informed by the output of perceptual systems by themselves, or by the integration of perceptual systems together with certain kinds of cognitive corrective factors (e.g., memories about the canonical colors, shapes, etc. of similar objects).[14]

An early instance of a post-perceptual/cognitive view about perceptual constancy is the proposal, defended by von Helmholtz (1962) and Hering (1964), that color constancy is (at least in part) driven by our memory/knowledge about the colors of familiar objects.[15] This "memory theory"

---

[13] There are several further pieces of evidence that confirm the prediction of such models that perceptual systems maintain representations of the illumination rather than simply discarding them. Perhaps the most direct is just that subjects can, when asked, make matches of ambient illumination as opposed to surface lightness (Katz, 1935; Gilchrist, 1988; Hurlbert, 1989; Jameson and Hurvich, 1989; Zaidi, 1998).

It is worth noting that the possibility of computing constancy without deriving specific object/object-property representations undercuts the (oft-made) claim that object tracking and reidentification depend on representing condition-independent object properties.

[14] Obviously, one's approach to this last question will be shaped, in part, by how one understands the cognition/perception distinction. I won't attempt to settle this vexed issue here, but will simply take for granted that e.g., memory for the colors/shapes/sizes/etc. of objects and other apparent instances of concept deployment fall on the cognitive side of the divide, and that, e.g., receptoral adaptation effects are perceptual. What is at stake is (of course) not the labels, but instead what kinds of causal explanatory resources are invoked to explain observed instances of perceptual constancy.

[15] An even earlier post-perceptual view of constancy emerges from Locke's discussion of the role of judgment in sensation:

> When we set before our eyes a round globe of any uniform colour . . . it is certain that the *Idea* thereby imprinted in our Mind, is of a flat Circle variously shadow'd, with several degrees of Light and Brightness coming to our Eyes. But we having by use been accustomed to perceive, what kind of appearance convex Bodies are wont to make in us; what alterations are made in the reflections of Light, by the difference of the sensible Figures of Bodies, the Judgment presently, by an habitual custom, alters the Appearances into their Causes: So that from that, which truly is variety of shadow or colour, collecting the Figure, it makes it pass for a mark of Figure, and frames to it self the perception of a convex Figure, and an uniform Colour; when the Idea we receive from thence, is only a Plain variously colour'd, as is evident in Painting (Locke, 1975, II.ix.8).

of color constancy faces several difficulties. First, Katz (1911) showed that there is color constancy for random and presumably newly encountered objects (for which there could not be color memory), and thereby demonstrated that the sort of memory/knowledge enlisted by the memory theory is not necessary for successful color constancy. Second, it is doubtful that our memory for color is sufficiently accurate to underwrite observed levels of constancy (Hurvich (1981, 2); Halsey and Chapanis (1951, 1058)). A third line of concern for memory (and, more generally, cognitive) explanations of color constancy is that one can dissociate the capacity for color constancy from (what are generally taken to be) cognitive capacities in both directions. In one direction, there appears to be robust color constancy in goldfish, honeybees, and several other non-human animals (see the review in Neumeyer (1998)) and human infants somewhere between 9 and 20 weeks old (Dannemiller and Hanko, 1987; Dannemiller, 1989), whose cognitive/conceptual resources are usually assumed to be pretty limited. In the other direction, there is (admittedly more limited) evidence from lesion studies where color constancy is impaired but memory and other conceptual capacities are spared (Rüttiger *et al.*, 1999).

These reasons, among others, have led investigators to search for less obviously cognitive explanations of color constancy. For example, contemporary explanations of color constancy often cite several kinds of retinal adaptation (changes in the sensitivity of retinal receptors as a response to incident light) including adaptation over temporally and spatially local regions (so-called von Kries adaptation), adaptation to the spatial mean of the whole scene, and adaptation to the region of highest intensity in the scene (McCann, 2004). However, there is evidence suggesting that these factors are not always sufficient for color constancy by themselves (Kraft and Brainard, 1999). Moreover, even if they are not by themselves sufficient for constancy, it appears that cognitive factors may make an important contribution to constancy after all: several investigators have found that familiarity for types of objects perceived (e.g., common fruits and vegetables) enhances color constancy (Hurlbert and Ling, 2005; Olkkonen *et al.*, 2008, 2012).

A similarly complicated mix of findings seems to be the pattern for shape and size constancy. On the one hand, there is evidence that the visual system can in some conditions (e.g., at short distance ranges) compute constant size and shape from relatively low-level perceptual cues such as vergence (information about the relative ocular positions of the two eyes in their sockets) and disparities in the retinal projections from the two eyes. And, once again, there is double dissociation between constancy for size/shape and cognitive sophistication. Thus, for example, there appears to be size constancy (at least at short distance ranges) in comparatively cognitively unsophisticated creatures such as newborn human beings (Granrud, 2006, 2012; Slater *et al.*, 1990), non-human primates (Fujita, 1997; Barbet and Fagot, 2002), goldfish (Douglas *et al.*, 1988), and amphibians (Ingle, 1998). And, in the other direction, Cohen *et al.* (1994) give evidence of the selective impairment of certain kinds of

---

size constancy that spare general cognitive abilities. All that said, it is also true that higher-level, cognitive cues — e.g., memory for the canonical shape and size of recognized objects, comparison to other perceived items whose shape and form are established independently, the smoothed appearance of texture from a distance — enhance shape and size constancy substantially (for a useful overview, see Palmer, 1999, ch.5, ch. 7).

Cumulatively, these results suggest strongly that perceptual constancy is neither exclusively perceptual nor exclusively cognitive. Instead, it appears that "the" phenomenon of perceptual constancy, even considered as constancy for a single dimension of a single quality within a single modality (e.g., just for lightness), is an interaction effect produced by several different mechanisms operating across different spatial and temporal scales — some possibly more and some possibly less cognitive than others, depending on how one chooses to mark the cognitive/non-cognitive distinction.[16] Whether any one of these mechanisms contributes to perceptual constancy on any particular occasion will depend on the details of many features of the perceptual circumstance.

7 Conclusion

While I have argued that the perceptual stabilities emphasized by traditional characterizations of perceptual constancy can only be part of the story, it remains true, indisputable, and important that some aspects of our perceptual responses are stable even through changes in perceptual circumstances that result in changes in transduced perceptual signals. It is no less indisputable that there are important lessons to be learned from the phenomenon of perceptual constancy, although many unresolved questions remain.

As we have seen, there is no completely general account of which dimensions of perceptual response must remain fixed, and which may vary, across which kinds of variation in perceptual conditions, for a perceptual episode to count as an instance of perceptual constancy. Moreover, there is no general understanding of the relation between perceptual constancy and perceptual contrast. And, partly because so much less is understood about both constancy and contrast in non-visual modalities, it is so far unclear what (if any) systematic cross-modal generalizations hold for each. Finally, the range of computational strategies that perception uses to extract stabilities, of the mechanisms underlying their implementation, and of the ways these distinct strategies and mechanisms are combined with one another in real-time perception remain incompletely understood.

Notwithstanding these substantial gaps in our knowledge, it seems clear that constancy is an absolutely fundamental aspect of perception, and therefore

---

[16] Cf. Foster (2003), who points to the heterogeneity of the factors in operation as a reason to be skeptical about the very existence of color constancy.

that it will figure centrally in our ultimate understanding of mind-world interaction.[17]

References

Adelson, E. H. (2000). Lightness perception and lightness illusions. In M. Gazzaniga, editor, *The New Cognitive Neurosciences, 2nd Ed.*, pages 339–351. MIT Press, Cambridge, Massachusetts.

Amano, K. D., Foster, D. H., and Nascimento, S. M. C. (2005). Minimalist surface-colour matching. *Perception*, **34**, 1007–1011.

Arend, L. and Reeves, A. (1986). Simultaneous color constancy. *Journal of the Optical Society of America A*, **3**(10), 1743–1751.

Arend, L., Reeves, A., Schirillo, J., and Goldstein, R. (1991). Simultaneous color constancy: patterns with diverse Munsell values. *Journal of the Optical Society of America A*, **8**, 661–672.

Barbet, I. and Fagot, J. (2002). Perception of the corridor illusion by baboons. *Behavioural Brain Research*, **132**(1), 111–115.

Bäuml, K.-H. (1999). Simultaneous colour constancy: how surface color perception varies with the illuminant. *Vision Research*, **39**(8), 1531–1550.

Beck, J., editor (1972). *Surface Color Perception*. Cornell University Press, Ithaca.

Blackwell, K. T. and Buchsbaum, G. (1988). Quantitative studies in color constancy. *Journal of the Optical Society of America A*, **5**, 1772–1780.

Bradley, P. (2008). Constancy, categories and Bayes: A new approach to representational theories of color constancy. *Philosophical Psychology*, **21**(5), 601 – 627.

Brainard, D. H. (1998). Color constancy in the nearly natural image 2. achromatic loci. *Journal of the Optical Society of America A*, **15**(2), 307–325.

Brainard, D. H. and Freeman, W. T. (1997). Bayesian color constancy. *Journal of the Optical Society of America A*, **14**(7), 1393–1411.

Brainard, D. H., Brunt, W. A., and Speigle, J. M. (1997). Color constancy in the nearly natural image I. asymmetric matches. *Journal of the Optical Society of America, A*, **14**(9), 2091–2110.

Brainard, D. H., Kraft, J. M., and Longere, P. (2003). Color constancy: Developing empirical tests of computational models. In R. Mausfeld and D. Heyer, editors, *Colour Perception: Mind and the Physical World*, pages 307–328. Oxford University Press, New York.

---

[17] Many thanks to Damon Crockett, Joshua Gert, Gary Hatfield, Don MacLeod, Mohan Matthen, and Sam Rickless for discussion and comments on earlier drafts.

Brainard, D. H., Delahunt, P. B., Freeman, W. T., Kraft, J. M., and Xiao, B. (2006). Bayesian model of human color constancy. *Journal of Vision*, **6**, 1267–1281.

Buchsbaum, G. A. (1980). A spatial processor model for object colour perception. *Journal of the Franklin Institute*, **310**, 1–26.

Burge, T. (2010). *Origins of Objectivity*. Oxford University Press, Oxford.

Byrne, A. and Hilbert, D. R., editors (1997). *Readings on Color, Volume 2: The Science of Color*. MIT Press, Cambridge, Massachusetts.

Byrne, A. and Hilbert, D. R. (2003). Color realism and color science. *Behavioral and Brain Sciences*, **26**(1), 3–64.

Cohen, J. (2008). Color constancy as counterfactual. *Australasian Journal of Philosophy*, **86**(1), 61–92.

Cohen, L., Gray, F., Meyrignac, C., Dehaene, S., and Degos, J.-D. (1994). Selective deficit of visual size perception: two cases of hemimicropsia. *Journal of Neurology, Neurosurgery, and Psychiatry*, **57**, 73–78.

Cornelissen, F. W. and Brenner, E. (1995). Simultaneous colour constancy revisited: an analysis of viewing strategies. *Vision Research*, **35**, 2431–2448.

Craven, B. J. and Foster, D. H. (1992). An operational approach to colour constancy. *Vision Research*, **32**(7), 1359–1366.

Dannemiller, J. L. (1989). A test of color constancy in 9- and 20-week-old human infants following simulated illuminant changes. *Developmental Psychology*, **25**, 171–184.

Dannemiller, J. L. (1993). Rank orderings of photoreceptor photon catches from natural objects are nearly illuminant-invariant. *Vision Research*, **33**, 131–140.

Dannemiller, J. L. and Hanko, S. A. (1987). A test of color constancy in 4-month-old human infants. *Journal of Experimental Child Psychology*, **44**, 255–267.

Delahunt, P. B. (2001). *An Evaluation of Color Constancy Across Illumination and Mutual Reflection Changes*. Ph.D. thesis, University of California, Santa Barbara, Santa Barbara.

Delahunt, P. B. and Brainard, D. H. (2004). Does human color constancy incorporate the statistical regularity of natural daylight? *Journal of Vision*, **4**, 57–81.

Douglas, R. H., Eva, J., and Guttridge, N. (1988). Size constancy in goldfish (carassius auratus). *Behavioural Brain Research*, **30**, 37–42.

D'Zmura, M. (1992). Color constancy: surface color from changing illumination. *Journal of the Optical Society of America A*, **9**, 490–493.

D'Zmura, M. and Lennie, P. (1986). Mechanisms of color constancy. *Journal of the Optical Society of America A*, **3**, 1662–1672.

Evans, R. M. (1948). *An Introduction to Color*. Wiley, New York.

Foster, D. H. (2003). Does colour constancy exist? *Trends in Cognitive Science*, **7**(10), 439–443.

Foster, D. H. and Nascimento, S. M. C. (1994). Relational colour constancy from invariant cone-excitation ratios. *Proceedings of the Royal Society of London B*, **257**, 115–121.

Fujita, K. (1997). Perception of the Ponzo illusion by rhesus monkeys, chimpanzees and humans: similarity and difference in the three primate species. *Perception & Psychophysics*, **59**, 284–292.

Funt, B., Drew, M., and Ho, J. (1991). Color constancy from mutual reflection. *International Journal of Computer Vision*, **6**, 5–24.

Gert, J. (2010). Color constancy, complexity, and counterfactual. *Noûs*.

Gibson, J. J. (1933). Adaptation, after-effect and contrast in the perception of curved lines. *Journal of Experimental Psychology*, **XVI**(1), 1–31.

Gilchrist, A. L. (1988). Lightness contrast and failures of constancy: A common explanation. *Perception & Psychophysics*, **43**, 415–424.

Gilchrist, A. L., Kossifydis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., Annan, V., and Economou, E. (1999). An anchoring theory of lightness perception. *Psychological Review*, **106**(4), 795834.

Goldstein, E. B. (1999). *Sensation & Perception (5th Edition)*. Brooks/Cole Publishing, Pacific Grove, California.

Golz, J. and MacLeod, D. I. A. (2002). Influence of scene statistics on colour constancy. *Nature*, **415**, 637–640.

Granrud, C. E. (2006). Size constancy in infants: 4-month-olds' responses to physical versus retinal image size. *Experimental Psychology: Human Perception and Performance*, **32**(6), 1398–1404.

Granrud, C. E. (2012). Judging the size of a distant object: Strategy use by children and adults. In G. Hatfield and S. Allred, editors, *Visual Experience: Sensation, Cognition, and Constancy*, chapter 1, pages 13–34. Oxford University Press, Oxford.

Halsey, R. M. and Chapanis, A. (1951). Number of absolutely identifiable hues. *Journal of the Optical Society of America*, **41**, 1057–1058.

Hatfield, G. (2009). On perceptual constancy. In G. Hatfield, editor, *Perception and Cognition: Essays in the Philosophy of Psychology*, chapter 6, pages 178–211. Clarendon Press, Oxford.

von Helmholtz, H. (1962). *Helmholtz's Treatise on Physiological Optics*. Dover, New York. Originally published in 1867.

Hering, E. (1964). *Outlines of a Theory of the Light Sense*. Harvard University Press, Cambridge, Massachusetts. Originally pulished in 1878; translated by Leo M. Hurvich and Dorothea Jameson.

Hilbert, D. R. (1987). *Color and Color Perception: A Study in Anthropocentric Realism*. CSLI, Stanford.

Hilbert, D. R. (2005). Color constancy and the complexity of color. *Philosophical Topics*, **33**, 141–158.

Hunter, F., Biver, S., and Fuqua, P. (2007). *Light Science and Magic: An Introduction to Photographic Lighting*. Focal Press, Oxford, third edition edition.

Hurlbert, A. C. (1989). *Cues to the color of the illuminant*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.

Hurlbert, A. C. and Ling, Y. (2005). If it's a banana, it must be yellow: The role of memory colors in color constancy. *Journal of Vision*, **5**(8), 787.

Hurvich, L. M. (1981). *Color Vision*. Sinauer Associates, Sunderland, Massachusetts.

Ingle, D. (1998). Perceptual constancies in lower vertebrates. In V. Walsh and J. J. Kulikowski, editors, *Perceptual Constancy: Why Things Look as They Do*, pages 173–191. Cambridge University Press.

Jameson, D. and Hurvich, L. M. (1989). Essay concerning color constancy. *Annual Review of Psychology*, **40**, 1–22. Reprinted in Byrne and Hilbert (1997), 177–198.

Katz, D. (1911). *Die Erscheinungsweisen der Farben und ihre Beeinflussung durch die Individuele Erfahrung*. Barth, Leipzig.

Katz, D. (1935). *The world of Color*. Kegan, Paul, Trench Truber & Co., London. translated by R. B. MacLeod and C. W. Fox.

Kraft, J. M. and Brainard, D. H. (1999). Mechanisms of color constancy under nearly natural viewing. *Proceedings of the National Academy of Sciences, U. S. A.*, **96**, 307–312.

Land, E. H. and McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, **61**.

Lee, H. C. (1986). Method for computing the scene-illuminant chromaticity from specular highlights. *Journal of the Optical Society of America A*, **3**, 1694–1699.

Locke, J. (1975). *An Essay Concerning Human Understanding (1689)*. Oxford University Press, New York.

MacLeod, D. I. A. (2003). Colour discrimination, colour constancy, and natural scene statistics. In J. D. Mollon, J. Pokorny, and K. Knoblauch, editors, *Normal and Defective Colour Vision*. Oxford University Press, London.

Maloney, L. T. (1986). Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *Journal of the Optical Society of America A*, **3**, 1673–1683.

Maloney, L. T. and Wandell, B. A. (1986). Color constancy: a method for recovering surface spectral reflectance. *Journal of the Optical Society of America A*, **3**(1), 29–33.

Matthen, M. (2010). How things look (and what things look that way). In B. Nanay, editor, *Perceiving the World*. Oxford University Press, New York.

McCann, J. J. (2004). Mechanisms of color constancy. In *Proceedings of the IS&T/SID Color Imaging Conference*, volume 12, pages 29–36, Scottsdale, Arizona.

Neumeyer, C. (1998). Comparative aspects of color constancy. In V. Walsh and J. Kulikowski, editors, *Perceptual Constancy: Why Things Look as They Do*, pages 323–351. Cambridge University Press, Cambridge.

Olkkonen, M., Hansen, T., and Gegenfurtner, K. R. (2008). Color appearance of familiar objects: Effects of object shape, texture, and illumination changes. *Journal of Vision*, **8**(5), 1–16.

Olkkonen, M., Hansen, T., and Gegenfurtner, K. R. (2012). High-level perceptual influences on color appearance. In G. Hatfield and S. Allred, editors, *Visual Experience: Sensation, Cognition, and Constancy*, chapter 9, pages 179–198. Oxford University Press, Oxford.

Palmer, S. E. (1999). *Vision Science: Photons to Phenomenology*. MIT Press, Cambridge, Massachusetts.

Rüttiger, L., Braun, D. I., Gegenfurtner, K. R., Petersen, D., Schönle, P., and Sharpe, L. T. (1999). Selective color constancy deficits after circumscribed unilateral brain lesions. *The Journal of Neuroscience*, **19**(8), 3094–3106.

Slater, A., Mattok, A., and Brown, E. (1990). Size constancy at birth: Newborn infants' responses to retinal and real size. *Journal of Experimental Child Psychology*, **49**, 314–322.

Smith, A. D. (2002). *The Problem of Perception*. Harvard University Press, Cambridge, Massachusetts.

Thompson, B. (2006). Colour constancy and Russellian representationalism. *Australasian Journal of Philosophy*, **84**(1), 75–94.

Troost, J. M. and deWeert, C. M. M. (1991). Naming versus matching in color constancy. *Perception & Psychophysics*, **50**, 591–602.

Tye, M. (2000). *Consciousness, Color, and Content*. MIT Press, Cambridge, Massachusetts.

Valberg, A. and Lange-Malecki, B. (1990). 'Colour constancy' in Mondrian patterns: A partial cancellation of physical chromaticity shifts by simultaneous contrast. *Vision Research*, **30**(3), 371–380.

Whittle, P. (2003). Contrast colours. In R. Mausfeld and D. Heyer, editors, *Colour Perception: Mind and the Physical World*, pages 115–138. Oxford University Press, New York.

Wright, W. (2013). Color constancy reconsidered. *Acta Analytica*, **28**(4), 435–455.

Wyszecki, G. and Stiles, W. S. (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae*. Wiley, New York. Second edition.

Zaidi, Q. (1998). Identification of illuminany and object colors: Heuristic-based algorithms. *Journal of the Optical Society of America A*, **15**, 1767–1776.

Zaidi, Q. (1999). Color and brightness induction: from Mach bands to three-dimensional configurations. In K. R. Gegenfurtner and L. T. Sharpe, editors, *Color Vision: From Genes to Perception*, pages 317–343. Cambridge University Press, Cambridge.

Zaidi, Q. (2001). Color constancy in a rough world. *Color Research and Application*, **26**, S192–S200.

Figure 1: There is some good sense in which the regions of the cup in shadow and the regions of the cup in direct sunlight look the same in color. Photograph © Jonathan Cohen.

Figure 2: There is some good sense in which the telephone poles seen from different distances look the same size. Photograph © Jonathan Cohen.

Figure 3: There is some good sense in which the penny looks the same in shape when seen from two different angles. Photograph © Jonathan Cohen.

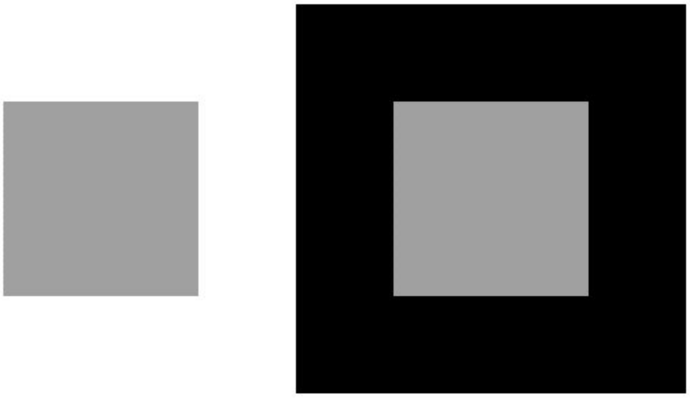

Figure 4: An instance of simultaneous lightness contrast: the central patches are qualitatively identical, but perception represents them differently because of the contrast with surrounding items.

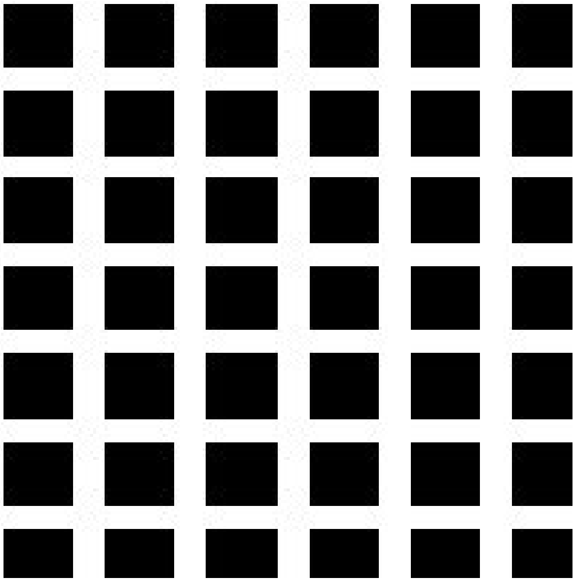

Figure 5: The Hermann grid illusion.

Figure 6: The Cornsweet illusion.

Figure 7: Mach bands.

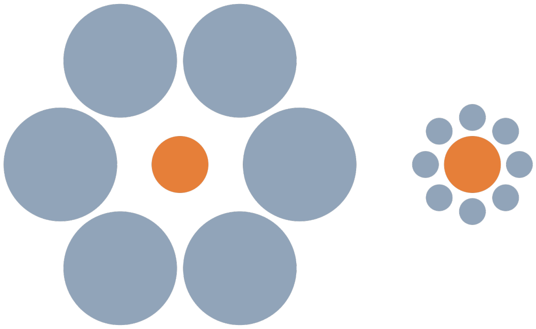

Figure 8: The Ebbinghaus illusion is an instance of perceptual simultaneous size contrast.