

Color Constancy and the Complexity of Color

David Hilbert
University of Illinois, Chicago

[W]e are not in the habit of observing our sensations accurately, except as they are useful in enabling us to recognize external objects. On the contrary, we are wont to disregard all those parts of the sensations that are of no importance so far as external objects are concerned. —*H. V. Helmholtz, 1924*

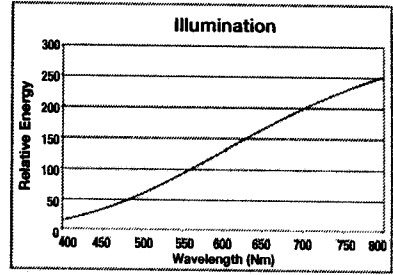
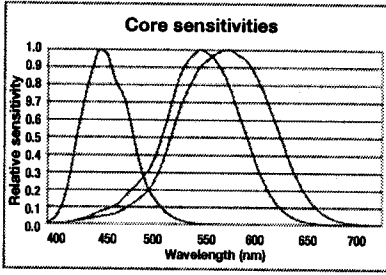
1. THE PROBLEM

1.1 A BRIEF OVERVIEW OF COLOR VISION AND COLOR CONSTANCY

We can start with a definition. “[C]olour constancy is the constancy of the perceived colours of surfaces under changes in the intensity and spectral composition of the illumination” (Foster et al. 1997). Given the definition we can now ask a question: Does human color vision exhibit color constancy?¹ The answer to the question depends in part on how we interpret it. If the question is understood as asking whether human color vision displays constancy for every possible scene across every possible illumination then the answer is no.² If the question is understood as asking whether human color vision displays some degree of constancy for some scenes across some range of illuminants then the answer is yes. The more interesting questions involve characterizing the degree of constancy human vision displays, the types of scenes and ranges of illuminants for which approximate constancy can be achieved and the mechanisms that produce constancy. These questions are difficult ones and have been the subject of vigorous debate within the

color science community. In order to understand some of what makes these questions so difficult it will be useful to review some elementary facts about the causal process that underlies the perception of color. In a simple case the process starts with a source (or sources) of light which has a specific distribution of energy over the wavelengths of the visible spectrum. The light is reflected off the objects in the surrounding environment, each of which reflects a fixed percentage of the energy at each wavelength (the surface spectral reflectance or reflectance), and some of it enters the eye of the observer where it is (selectively) absorbed by the cone pigments. We'll call the light reflected from the object *the color signal*, since it is what carries information about the reflecting properties of the object in the form of how energy is distributed over wavelength. The cone output results from the response of the three human cone types to the color signal and is subject to further processing in both the retina and various cortical areas, and the end result is that tomatoes (usually) look red and their leaves green. The early stages of this process are summarized in Figure 1. Notice, in particular, that the character of the light reaching the eye (the color signal) from the various objects in the environment is the joint product of the character of the illuminant and the spectral reflectance of the surface. One consequence of this fact is that any change in the illuminant will result in a corresponding change in the color signal generated by each object in the scene. Since natural and artificial lights vary substantially in both their intensity and spectral characteristics (distribution of energy over wavelength), the color signal from a surface with a fixed reflectance can vary substantially from one lighting condition to another. Since all of the information about the characteristics of objects in the scene is carried by the color signal, which varies with the illuminant, it may seem that color constancy is impossible. The only illumination-invariant factor involved in the process is the surface reflectance but the color signal completely confounds the reflectance and the illuminant.

Nevertheless human color perception does display some degree of constancy. To take a very simple example involving only lightness, consider a page of black print on white paper viewed first under an indoor reading light and then under direct sunlight. The intensity of the light reaching the eye from the white area of the page in indoor illumination is roughly equal to the intensity of the light reaching the eye from the black print in sunlight (Kaiser and Boynton 1996, 199). In spite of this rough equality, the page looks white under the indoor illumination and the print looks black under sunlight. Although this example concerns lightness only, similar effects are present for the other dimensions of color as well. If perceived color were a function only of the momentary character of the color signal at each point in the scene then no sensible form of color constancy would be possible. What our ability to read both indoors and out demonstrates is that perceived color is not solely a function of the local signal. Some of the factors that underlie this aspect of visual performance are known. Various types of adaptation play an important role and adaptation amounts, in essence, to comparing the current signal to the average signal over some period of time and/or some region of space.



Cone Signal (Relative)

L	M	S
1.00000	0.85890	0.00011

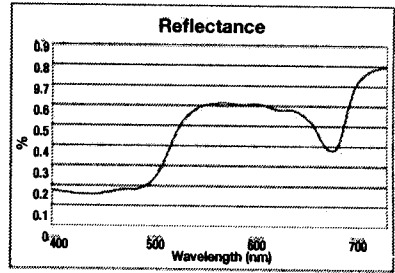
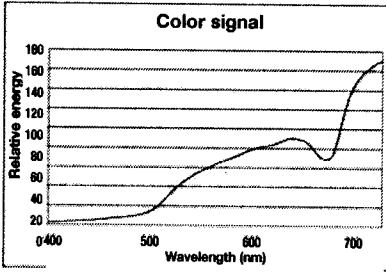



FIGURE 1: The causal process underlying color perception

Adaptation does not fully explain human visual performance and there continues to be lively debate over approaches to explaining color constancy.

Many theories of color constancy take the form of explaining how it is that the visual system manages to extract information about the reflectance of the objects in a scene from the color signal from those objects.³ Since this involves separating the contributions of the reflectance and the illuminant to the color signal these theories are often characterized as “discounting the illuminant.” Perfect color constancy in these terms would involve accurate recovery of reflectance for any scene under any lighting conditions. The perceived color of objects would be perfectly correlated with their reflecting characteristics and not vary at all with changes in the illuminant or the composition and arrangement of objects in view. This type of perfect color constancy is not possible. There is no method that will work across all

scenes and all illuminants. The color constancy of human color vision is necessarily imperfect and it does not require sophisticated experiment to display this fact. We are all familiar with the substantial excursions in perceived color produced by many common types of artificial lighting. What is much more difficult to answer is the question of exactly how good is the constancy of human color vision. This question is difficult for two reasons. First, how much constancy we display depends on the range of environments and illuminants that are being considered. We have pretty good constancy for some scenes viewed under some range of illuminants and poor constancy for other scenes viewed under the same range of illuminants or the same scene viewed under different illuminants. There is no sensible general answer to how good our color constancy is. Second, and less well recognized, there is a problem with the characterization of color constancy with which we began. To help motivate discussion of this problem it will be useful to be reminded of some familiar phenomena of color vision.

1.2 PHENOMENOLOGY

Our experience of color is not one of absolutely invariant color appearance as we view objects under changing conditions. A well-known passage from Russell (1912) displays one kind of appearance change quite nicely.

To make our difficulties plain, consider the table. To the eye it is oblong, brown, and shiny . . . Although I believe that the table is “really” of the same colour all over, the parts that reflect the light look much brighter than the other parts, and some parts look white because of the reflected light. I know that, if I move, the parts that reflect the light will be different, so that the apparent distribution of colours on the table will change. It follows that if several people are looking at the table at the same moment, no two of them will see exactly the same distribution of colours, because no two can see it from exactly the same point of view, and any change in the point of view makes some change in the way the light is reflected. (8–9)⁴

As Russell observes, the appearance of many objects changes depending on the angle between the viewer, the surface of the object, and the light source. In addition to the highlights Russell describes it is also common for the surfaces of objects to be visibly shaded with some areas appearing brighter and others dimmer. It also common for there to be visible effects of inter-reflection between objects in a scene. A white sheet of paper next to a tomato in bright light will take on a distinct pinkish cast in those areas closest to the tomato. All of these effects can be manipulated by changing the spatial relationships among objects, light sources, and perceivers.

In addition to these relationships between color appearance and the spatial layout of the environment there are also, equally familiar, effects on color appearance of the character of the illuminant. Although the white (black) areas of the printed page from our earlier example look white (black) under both the indoor and outdoor illuminations they don't look the same. Outside the white seems whiter and

the black blacker than when viewed under the dimmer indoor light. As Russell goes on to observe:

And we know that even from a given point of view the colour will seem different by artificial light, or to a colour-blind man, or to a man wearing blue spectacles, while in the dark there will be no colour at all, though to touch and hearing the table will be unchanged. (9)⁵

Although there may be a sense in which the perceived color of objects is independent of the illuminant under which they are viewed, it is not a sense in which the appearance of the object is unchanging. It's also important to emphasize that it's not a sense in which objects look to have two colors simultaneously. Rather, the printed page looks both similar and different when viewed under the different illuminations but not because there are multiple apparent colors associated with it. It's important to keep in mind that it is the printed page that appears both similar and different. There is nothing in the phenomenology of these cases that points to the change being internal and the constancy external.

It is equally a mistake to overemphasize the degree and kind of change that we actually observe. The character of the illumination varies frequently and substantially as we move about the world and within a single setting. Nevertheless there is a definite perception that objects do not appear to be frequently and substantially changing in color. As we will see shortly it is not easy to account for the complex mixture of stability and change that characterizes our experience of color in ordinary circumstances. It is this problem, how to account for the complex phenomenology of color constancy that was referred to at the end of the preceding section and that will make necessary a change in the characterization of color constancy.⁶

2. ACCOUNTING FOR COLOR CONSTANCY

2.1 CONSTANCY AND JUDGMENT

One possible response to this combination of change and constancy is to claim that the changing aspect is the purely sensory one and the constancy is the result of inference or judgment. The visual system responds to the color signal and although the response may be modified by the effects of adaptation at the level of the receptors the constant element is not itself a product of the senses but rather of a (more or less) cognitive interpretation of the sensory data. The apparent color is constantly changing but we are able to infer what it would be under some standard viewing condition and it is this inferred color that is stable across changes in illuminant and other viewing conditions. This type of view is, of course, the familiar approach to perception that has dominated modern philosophy since the empiricists and still has a powerful grip on philosophical thinking about perception. This approach is usefully illustrated by Berkeley's account of visual size constancy. According to Berkeley the immediately perceived visual magnitude of an object is

constantly changing. The very same object looks small at a distance but large when it is close. On the other hand, objects usually appear the same size to touch (tangible magnitude). We learn correlations between various aspects of the visual appearance, the circumstances of perception and the tangible magnitude of the object. The changing visual appearance serves as a premise in an inference to the (usually) unchanging tangible magnitude. If we focus on what we immediately perceive we will be aware of the constantly changing character of the visual idea. If we focus, as we usually do, on the conclusion of the inference we will be aware of the unchanging tangible magnitude. The perception of a constant size is a cognitive or intellectual accomplishment rather than a visual one (Berkeley 1979, 52–66).

Inferential theories must involve either two distinct determinable properties or the same determinable applied to different objects: one to figure in the premises of the inference and the other to figure in the conclusion. It is important in general and for the particular problem of accounting for the pattern of change and constancy in color perception that these properties be of distinct types or instantiated in different objects. In general, inferences from a premise to a conclusion logically incompatible with the premise are to be avoided. Berkeley avoids the possibility of conflict in size by making the premise and conclusion involve distinct properties. In the specific case at hand, the thesis under consideration involves simultaneous awareness of both premise and conclusion and without two distinct properties to be aware of our experience would simply be incoherent. Perceiving a red ball with one side brightly lit and the other shaded would then be like perceiving an impossible triangle or an Escher drawing. But this kind of paradox, although possible with color, is not a feature of our experience of common cases of color constancy. In the case of color, there are several possibilities for the two relevant property types or the two distinct objects to which color is attributed. For the representative theory, the inference is from properties of sense-data or sensational properties to external color properties. In color science, a fundamentally similar proposal is often put in terms of the distinction between color sensations and object color. The distinction between the reflectance (standing color) and the color signal (transitory color) is sometimes understood in this way, where the perceived transitory color is used as a basis of the inference to the standing color (Campbell 1969). Jonathan Cohen has proposed that the inference is from experienced color to a counterfactual property: the property of being such that it would appear to have some specified color under some range of illuminants (Cohen 2003).

The inference theory can take several different forms. For Berkeley, the inferences are potentially conscious and lead from consciously accessible premises to consciously acceptable conclusions. Cohen, however, treats the inference as unconscious although both premise and conclusion are consciously accessible. Most current theories in color science involve what are, in essence, inferences from premises that are not consciously accessible to conclusions that are. One problem all versions of the inference theory face, with the possible exception of the transitory color/standing

color pair, is some tension with the phenomenology. The change in appearance of the printed page when carried outdoors strikes us as a change in something related to the stable whiteness and blackness, perhaps an intensifying. Analyzing the change in appearance in terms of the difference between a property of a sensation and an object property or a color appearance and a counterfactual concerning color appearance does not get this aspect of the phenomenology correct. Since standing colors and transitory colors are related properties of the same objects an inference theory formulated in these terms avoids this problem.

2.2 CONSCIOUS INFERENCE

There are several problems with the inference theory in its Berkeleyan form. First, it runs the risk of over-intellectualizing a process that is in fact automatic and largely independent of cognitive abilities. Fish, small children, and chickens have all been demonstrated to have some degree of color constancy and there is some implausibility in supposing that they make complicated inferences based on experience. The role of learning in these inferences is particularly suspect since there is evidence that color constancy emerges in human infants in the first few months of life (Dannemiller 1998).

Most important, however, the necessary inferences don't seem to happen. There are two lines of reasoning in support of this claim. First, we just seem to be presented with colors, not to engage in reasoning to arrive at conclusions concerning them. Berkeley attempted to evade this problem by claiming that the inferences happen so quickly we hardly notice making them but surely they would occasionally obtrude themselves on consciousness. Second, the inferences don't seem to interact with the rest of our cognitive apparatus.⁷ The color constant appearance is not strongly affected by beliefs about the illumination (Rutherford and Brainard 2002). It is not necessary for color constancy that the objects be familiar or have any standard color (Land 1977). Finally, even if the conclusion of the inference is believed to be incorrect the appearance persists. The appearance of sameness in color across changing viewing conditions seems paradigmatically sensory rather than cognitive. Although none of these considerations is conclusive on its own, the whole package strongly suggests that it is worth searching for an alternative to the conscious inference theory.

2.3 UNCONSCIOUS INFERENCE

Many of the problems the inference theory faces could be avoided if the inferences were not conscious. That the perceptual constancies are the result of unconscious processes that are relatively independent of higher cognition is the dominant view among contemporary perceptual theorists. The most prominent historical figure holding this type of view is Helmholtz, who viewed perceptual constancies as the result of a species of unconscious causal inference.

The psychic activities that lead us to infer that there in front of us at a certain place there is a certain object of a certain character, are generally not conscious activities, but unconscious ones. In their result they are equivalent to a *conclusion*, to the extent that the observed action on our senses enables us to form an idea as to the possible cause of this action; although, as a matter of fact, it is invariably simply the nervous stimulations that are perceived directly, that is, the actions, but never the external objects themselves (Helmholtz 1924, Vol. III, 4).

For Helmholtz, the action of light on the eye produces a reaction in the nervous system that is directly perceived resulting in a sensation. This nervous reaction, and its accompanying sensation, varies with the proximal stimulus and thus, in the case of color, is not illumination independent. This sensation then serves as a starting point for a process of unconscious inference that results in a conclusion as to the color of the object. The processes of unconscious inference, although modulated by perceptual experience, are not under voluntary control and do not interact with explicit belief and other cognitive information. Although memory is involved, it is memory as a store of past sensations rather than explicit belief. The general principle governing such inferences according to Helmholtz is that:

whenever an impression is made on the eye . . . such objects are always imagined as being present in the field of vision as would have to be there in order to produce the same impression on the nervous mechanism, the eyes being used under ordinary normal conditions. (Helmholtz 1924, Vol. III, 2)

Helmholtz's theory avoids some of the difficulty of the conscious inference theory but other problems remain. The general independence of perceived color from belief about color, the inaccessibility of the inferential process and the lack of control over the conclusion of the process are all nicely explained. Nonetheless, a fundamental mystery remains. Helmholtz has two factors that can vary independently, the sensation and the perceptual conclusion, to explain the pattern of stability and change that we find with changing illumination. In spite of this he has the difficulty noted in section 2.2 in finding appropriately related candidates to serve as the two objects of awareness. Helmholtz's explicit descriptions would lead to the conclusion that the process starts with the illumination-dependent sensation of color and its conclusion concerns a property of external objects. But the sameness and difference we are concerned with are all perceived as external. It's the properties of the printed page that are seen as changing in some ways and not in others and properties of sensations are nowhere involved. Helmholtz himself uses the word "color" to characterize both premise and conclusion, which suggests that just a single property is involved. Here a different problem would arise in that objects would appear to have multiple colors at any given point in time, some of which change with the illumination and others which generally don't. This also seems unfaithful to the phenomenology (although below I shall rescue a grain of truth from this description).

Helmholtz's fundamental insight, that perceptual processing is devoted to recovering information about the distal causes of proximal stimuli, is embodied in the most influential current approach to explaining color constancy. These theories, however, depart from Helmholtz in two significant ways. First, the unconscious process is described in computational terms rather than as a process of logical inference.⁸ Second, and for my purposes more important, the inputs for these computations are typically regarded as being inaccessible to consciousness and cognition and can concern external objects rather than sensations. For Helmholtz, although it is difficult to become aware of our sensations, it is, at least in principle, always possible. What is inaccessible to conscious awareness is the process that generates perceptual conclusions from sensory inputs. Computational theories of color constancy generally take the inputs as well as the processes that produce the outputs to be unavailable to conscious awareness.

Computational theories of color constancy typically take as their starting point the color signal at each point on the retina. The problem to be solved becomes that of computing an illumination-independent quantity on the basis of this input (plus possibly other information). The specific illumination-independent quantity is usually taken to be some measure of surface reflectance.⁹ Since, as was noted above, the general problem of ascertaining an arbitrary reflectance under an arbitrary illuminant in an arbitrary scene is insoluble, all such theories take advantage of constraints on the range of reflectances and/or range of illuminants and/or scene composition. Given that the environment meets the constraints then the proposed algorithm accurately recovers the reflectances of objects in the scene. Although the computational problem posed by color constancy has no solution in the general case there are a variety of solutions that take advantage of special features of the scene or of restricting the range of possibilities that have to be considered. Whatever the details of the theory, and a very wide range of possibilities have been explored, the end product is an estimate of the reflectances of objects in the field of view.

Although computational theories of color constancy are very interesting they cannot, as usually formulated, give an adequate account of the aspect of the phenomenology of color vision on which we have focused. By generating an estimate of surface reflectance they have the potential to account for the fact that some aspect of color appearance remains unchanged across some changes of illumination. In their usual formulation, however, this is the sole output of the computation and thus they cannot account for that aspect of color appearance that depends on the illuminant. Because the inputs to the computation are held to be inaccessible to consciousness, computational theories seem to be even less phenomenologically plausible than either form of the inference theory.

The fact that human color constancy is not perfect provides a possible response to this objection. Take a circumstance under which, for a given computational theory, the reflectance-estimate will change slightly. If that theory correctly describes

human vision there will be a slight change in color appearance. Perhaps our experience of both similarity and change across change in illuminant can be accounted for in terms of small shifts in color appearance as the illuminant changes. Because the shift is small we are aware of a similarity in appearance. Because the shift takes place we are aware of a difference in appearance. Setting aside the question of whether or not this is an adequate account of the phenomenology of the case of the illuminant changing over time, it is clearly not an adequate account of the illuminant changing across space.¹⁰ The keyboard on which I am typing has white keys and is lit from above. As a result of the position of the light source relative to the keyboard the tops of the keys are more brightly lit than the sides. This shading is clearly visible but at the same time there is no appearance of a change in material or color. If there were errors in color constancy the difference between the top and sides would appear to be a change in the surface itself. Computational theories appear to lack any account of the perception of shading.¹¹

3. SOLVING THE PROBLEM

3.1 A SIMPLE SOLUTION

All of the issues with computational theories can be resolved by supposing that in addition to delivering information about the reflecting properties of objects the visual system also delivers information about the way in which those objects are illuminated. When we look at the printed page under indoor illumination we see not only that some parts of it are white and others are black but that the whole of it is dimly lit. When we look at the same page outdoors we see the same distribution of colors but also see that it is brightly lit. When I look at the keyboard of my computer I see both that the keys are a particular shade of white all over and that the illumination on the sides is dimmer than the illumination on the tops. The intuitive description of the pattern of stability and change in our perception of color is, according to this proposal, the right description. Any computational theory of vision that estimates reflectance is in a position to also generate an estimate of illumination. Recall that the light reaching the eye from an object, the color signal, is the joint product of the reflectance and the illuminant. Since, we will assume for now, the color signal is known, an estimate of either the illuminant character or the reflectance will allow an estimate of the other. Many color constancy theories work, in fact, by first deriving an estimate of the illuminant and then using that, in conjunction with the color signal, to generate an estimate of the reflectance.¹² Rather than discarding the illuminant estimate, however, the current proposal is that it is retained and is part of the visual representation of the scene.

Above I described the character of the illuminant as being visually represented. This is ambiguous as to whether what is represented is a property of the object or a property of the light source. Do we see how an object is illuminated or do we see the

illumination itself? On phenomenological grounds the first option seems better to me. What we see as changing with the illumination is an aspect of the object itself, not the light source or the space surrounding the object. One consequence of this choice is that the color appearance of an object must have more than the traditional three dimensions of variation. Although this consequence may seem to contradict well-established empirical facts, it is well known that the three-dimensionality of color is an abstraction from a rather messy and more complicated set of phenomena. Those color scientists interested in the details of color appearance have been aware for some time that three dimensions are not enough to fully characterize color appearance. The list of independent dimensions from one standard survey includes five dimensions of color appearance: brightness, lightness, colorfulness, chroma, hue (Fairchild 1998, 107–9). The point here is not to endorse any specific proposal but rather to show that the claim that there are more dimensions to color appearance than the usual three is not in conflict with the empirical literature. The exact number of extra dimensions and their correct characterization is an empirical issue and, thus, not to be decided *a priori*. In particular, we need to not assume that any precise representation of the way an object is illuminated is visually represented.

3.2 SOME EXTENSIONS AND APPLICATIONS

In discussing a similar response to the phenomenological complexity of color vision David Chalmers observes that there is an apparent similarity between a shaded area of a white object and an unshaded area of a gray object (Chalmers 2006, 88). Chalmers own account of this similarity is in terms of an additional property, superficial grayness, which is possessed by both objects that are shaded and are gray. A related but somewhat different account can be motivated by considering the physical properties underlying surface gray and shading respectively. A gray surface is one whose reflectance is nonselective, i.e., fairly constant across the visible spectrum, and intermediate, i.e., lower than white but higher than black. Shading is a nonselective, partial decrease in the amount of light falling on a surface. If the visual system were to represent surface grays in terms of proportions of the best white and to represent illuminants in terms of proportions of the average illumination for the scene then there would be a formal analogy between the representation of gray surfaces and shaded surfaces. In addition, there is a similarity in the content of the representations in that both are concerned with the object's relation to the light in its environment. In other words, a gray surface is one that reflects significantly less than all the available light, while a shaded surface is one that is illuminated by significantly less than all the available light. This similarity in content may be enough to account for the similarity in phenomenology.

There has been some debate as to whether color constancy is a genuinely visual phenomenon as opposed to the effect of some kind of reasoning or inference. One kind of evidence that has been taken to be significant for this debate is whether or not the color matches subjects make across varying illumination are sensitive to the

effects of the verbal instructions they are given. Although I am not sure that I understand the issue precisely, the outline of the argument is clear enough. If the constant element in color experience is the result of specifically visual processing then the matches subjects make will be largely independent of the verbal instructions they are given. Visual processing, being unconscious, is automatic and thus not responsive to conscious goals and information. Since there is some experimental support for the existence of instruction effects in human color vision there is evidence that color constancy is the result of conscious inference rather than unconscious processing.¹³ If, however, the content of color experience is complicated in the way suggested here then another possibility for explaining instruction effects is available. Suppose for definiteness that the color experience has the five dimensions given above. In this circumstance, even if color constancy is perfect, when subjects are asked to make a color match between objects viewed under different lighting conditions 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 given could have the effect of changing the relative weight of the various attributes of appearance in determining the best match. This kind of dependence of subject performance on what they are told is wholly uninformative as to the nature of the kind of processing that produces the data on which the subjects are relying. Whether this kind of analysis can be given for any specific set of experimental results will depend on the details of the experiment and the details of the results obtained. This possibility does, however, show that it is not possible to conclude solely from the existence of instruction effects that the process underlying color constancy is conscious and cognitive as opposed to unconscious and sensory.

3.3 OTHER VISUALLY DETECTABLE PROPERTIES RELATED TO COLOR

There is good deal of variation in how objects interact with light and not all of the variation is usefully captured by talking only of reflectance and illuminant. The discussion so far has implicitly followed the convention found in many discussions of color constancy of focusing exclusively on the diffuse (nondirection dependent) component of the light reflected by an object. Most objects do reflect much of the light that falls on them diffusely, but many common objects also have a significant specular (direction-dependent) component. This component is visually detectable in the shifting highlights that Russell describes on his table and also underlies the perception of gloss. The characteristic speckle that some metallic surfaces possess is the result of complicated reflection resulting from a mixture of specular and diffuse reflection at different scales and different depths within the surface (Germer and Nadal 2001). Transparency and translucence provide yet more examples of visually detectable aspects of the interaction between objects and light from the environment. The tendency of philosophers has been to suppose that all these fea-

tures of the visual world are somehow derived from the perception of shape and color, but there is no reason to think this is true. What is true is that perception of all these properties derives from the same set of photoreceptor outputs, but unless we confuse what we see with the photoreceptor outputs there is no reason to suppose that these features are not perceptually primitive. The main moral of the discussion so far has been that it is time to complicate our picture of the content of visual experience and that this more complicated content allows us to provide a better understanding of visual phenomenology. In the next section, I will turn to some of the consequences that befall philosophers who fail to appreciate the complexity of visual content.

4. PHILOSOPHICAL CONSEQUENCES

4.1 PEACOCKE ON SENSATIONAL PROPERTIES

One of Peacocke's arguments for the distinction between the representational and sensational properties of experience depends on what he calls, "the problem of the additional characterization." (Peacocke 1983, 12) The problem is supposed to be that there are phenomenal features of visual experience that cannot be explained in terms of the representational content of that experience. Peacocke offers several examples of the problem only one of which is relevant to our concerns. We are asked to:

Imagine you are in a room looking at a corner formed by two of its walls. The walls are covered with paper of a uniform hue, brightness and saturation. But one wall is more brightly illuminated than the other. In these circumstances, your experience can represent both walls as being the same colour; it does not look as if one of the walls is painted with brighter paint than the other. Yet it is equally an aspect of your visual experience itself that the region of the visual field in which one wall is presented is brighter than that in which the other is presented. (Peacocke 1983, 12–13)

As we have seen, however, the difference in brightness that separates the two walls can be a feature of the walls themselves and part of what is represented about the walls. There is no need to postulate an internal visual field to bear the brightness properties, and there is no characterization of the experience that cannot be accounted for in terms of its content. This reply is not fatal to his arguments since his other examples do not involve color, but nevertheless it shows the way in which the failure to consider illumination properties has figured in philosophical argumentation (for a similar point, see Byrne 2001, 221–22).

4.2 NOË AND ENACTIVISM

Alva Noë, as part of his enactivist theory of perception, defends a view of color constancy that in outline is very similar to that of Helmholtz (Noë 2004, ch. 4). Noë

first describes how the apparent color of an object varies with changes in the illumination and other conditions of perception. He then claims that we have implicit knowledge of the patterns found in the changes of apparent color with viewing condition. Finally we are able to use this implicit knowledge in order to “experience the actual color of the object as, so to speak, that condition which governs or regulates the way these changes unfold” (128). Although Noë does not use the language of inference the overall structure of the explanation is very similar. We start with something experienced, the apparent color. We end with a conclusion about a condition outside experience, the condition which governs the way the changes unfold. Noë also insists, following the tradition of the inference theory, that the output of the process is not sensory. According to Noë, “We experience the presence of a uniform color that, strictly speaking, we do not see” (28). (In context it is clear that Noë intends “uniform color” here as a synonym for “actual color” in the first quote.) There are important differences between Noë and Helmholtz as to the nature of perceptual experience and the conclusions drawn from it. Noë’s theory of the nature of perceptual experience and perceptual knowledge is the radical part of his theory and the one that justifies describing it as enactivist. Structurally, however, the theory is exactly similar to the unconscious inference theory of Helmholtz.

The similarities are strong enough that the main difficulties that are faced by the unconscious inference account are also faced by the enactivist theory. First, both theories maintain that the constant element in color experience is not itself sensory. Noë seems convinced that this is phenomenologically correct. His view entails that experience of a shaded area of a white surface and the experience of a gray surface will be identical in their sensory properties. This claim is, at the very least, not obviously true. Second, for Noë as for Helmholtz, the two species of color are entirely different properties that characterize entirely different types of objects. The starting point in both cases are characteristics of the perceiver and the ending point is a characteristic of something that affects the perceiver. But, as noted above, our experience does not reveal this division into two sorts of properties and all of the properties we experience seem directed onto the external causes rather than onto ourselves.

If Noë provided some independent motivation for the Helmholtzian structure of this theory then there would be a reason to consider whether the objection mischaracterizes the phenomenology. But when we look for the motivation for Noë’s claim that color constancy is not a sensory phenomenon we find the usual examples of white walls that are differentially illuminated and the like. The main motivation is exactly the pattern of stability and change in color experience that we have been discussing. As we have seen, however, these kinds of examples cannot by themselves motivate a theory with the structure of the inference theory since there is an alternative account with a different structure. Noë does provide one additional piece of evidence that purports to support the structural features of his theory. He claims that the correlations that we implicitly know and that support our conclusions about the “actual” color are learned from experience (127). That the visual

system changes over time in response to features of the environment is certainly true. That this process is best described as a form of learning is much less obvious. Many processes important for color vision and color constancy in particular are best described as forms of adaptation. Adaptation happens on several different time scales ranging from seconds to weeks and is best understood in terms of sensitivity adjustment. The visual system adjusts its sensitivity so that the response to the average input is roughly in the middle of the output range. Different kinds of adaptation affect different sensory channels and different stages of sensory processing. This kind of “learning” provides little support for the claim that constancy effects are not sensory.¹⁴

4.3 INTROSPECTION AND PHILOSOPHY

As Helmholtz reminds us in the epigraph, gaining an accurate picture of the phenomenology of perception is difficult. Our normal approach to our sensory experience is dominated by various kinds of practical concerns. These concerns can vary from our interest in finding peaches against a cluttered background of foliage to the painter’s interest in exactly what pattern of pigments will produce a particular effect on the viewer. Because our focus is on using our senses to do things, we focus on particular aspects of the information about the world that they deliver. When we look inward to gather data to support our theories these practical interests and habits of attention can produce an overemphasis on some aspects of our rich perceptual experience and a neglect of others. When practical concerns don’t dominate there is still the danger of “seeing” our experience through the filter of our theoretical expectations. Philosophers, including myself, rarely concede any phenomenological awkwardness to the position they defend. In fact, it is an unusual philosophical paper in which it is not claimed that a careful consideration of the phenomenology wholly supports the position taken in the paper. As readers will have noticed this paper is no exception.

Scientists studying perception almost entirely avoid introspection as a direct source of data on perceptual processes and for the most part rely on nonverbal, behavioral data. The bread and butter of vision science is data on discrimination and matching, supplemented with physiological data.¹⁵ When this type of data is available and relevant it would be better for philosophers to use it rather than display the naïve faith in our introspective powers that characterizes so much of the literature in philosophy. Unfortunately for our particular topic, there is relatively little hard data that bears directly on the issues. It is clear that subjects make matches under many circumstances that are more nearly correlated with the reflecting properties of surfaces than with the illumination-dependent color signal. It is also clear that these processes are capable of operating even in circumstances in which there is no relevant explicit belief that could support the particular matches that are made. The fact that it has proven difficult to converge on a single theory concerning the mechanisms of color constancy suggests that individual perceivers, even the

motivated and reflective scientists and philosophers studying the subject, lack any explicit knowledge of how they make the matches that they do. These facts, and others like them, fleshed out with the appropriate detail are the constraints that theories of color constancy in science and philosophy have to negotiate. Phenomenological insights can help motivate and orient but are too unreliable in the end to be relied on for reaching definite conclusions. Introspection is what we look to when we have no other data, but it is a tool to be avoided otherwise.

5. CONCLUSION

There are two conclusions and a moral I would like to draw from this extended discussion of a very narrow topic. First, there is no reason to be found in a consideration of the facts of color constancy and the associated pattern of stability and change to reject standard computational and representational theories of vision. These facts are perfectly compatible with any reasonable version of such theories. Second, some reasons, although not particularly strong ones, favor these theories as accounts of color constancy over alternatives such as the conscious and unconscious forms of the inference theory and their modern descendants. Although as an advocate of a representational approach to perception these conclusions are important to me, it is the moral that I most wish philosophers to take to heart. Since the early modern period philosophers have been committed to a very limited view of the basic properties that are perceptually represented.¹⁶ There may have been reasons for this view originally, but it has continued largely as a legacy and a legacy that is increasingly at odds with the empirical study of perception. I hope to have motivated consideration of the thesis that the primitive properties of visual representation are considerably richer than their usual description in philosophical work on perception. For many philosophical discussions, the impoverished view makes no difference, but as the discussion of Noë and Peacocke demonstrates there are occasions when it leads to error.

NOTES

This paper is the direct descendant of a talk given at a conference organized by Keith Allen and hosted by the Institute of Philosophy, University of London. Both the stimulus to produce a talk and the comments I received were invaluable. Jonathan Cohen revived my interest in the topic with results I am sure he did not intend. Alex Byrne, Walter Edelberg, and Ed Minar gave me very helpful comments on an earlier draft.

1. Many nonhuman animals possess color vision, of course, and consequently many of the same issues arise as in the human case. For definiteness, I will focus on human color vision although many of the conclusions would also apply to some nonhuman animals.
2. This may seem an implausible interpretation but it is sometimes adopted in the color science literature (see, for example, Fairchild 1998, 156).
3. Not all theories of color constancy take this form and even those that do need not be taken as assuming a physicalist theory of color. Approximate color constancy as defined above is consistent with many theories of color including eliminativist ones.
4. For an interesting discussion of this and other, similar examples, see Brown 2006.
5. Russell goes on to conclude, "This colour is not something which is inherent in the table, but something depending upon the table and the spectator and the way the light falls on the table" (9). Although I will be questioning Russell's characterization of the phenomena below it is worth observing that this is a bad argument even if the premises are granted. For discussion of a related argument in which the variation is due to differences among perceivers, see Byrne and Hilbert 2004.
6. This problem has been discussed in two recent papers (Cohen 2003, 2004; Chalmers 2006). Cohen uses it to motivate an attack on computational theories of color constancy while Chalmers ends up favoring a solution similar in some respects to the one given below.
7. The one possible exception to this is the possibility of what are called instruction effects on perceived color. Some research has reported that subjects will make matches that vary with changing illumination or matches that are relatively constant across changing illumination depending on the instructions they are given. One particular paper reporting such an effect (Arend and Reeves 1986) has captured the attention of philosophers. The particular task described in this paper is somewhat unusual and the experimental literature on instruction effects does not seem to me to support any particularly clear conclusion. (For a useful discussion, see Appendix B of Delahunt and Brainard 2004.) I will return to this issue below.
8. I won't argue for this here but the difference between inference and computation doesn't embody any substantive disagreement.
9. There are exceptions to this focus on reflectance and thus algorithms that aim at other relatively illumination-independent features of a scene.
10. This point is somewhat obscured in the empirical literature by the fact that the standard task testing theories of color constancy utilizes an asymmetric matching paradigm. In a typical example of this kind of task a subject will be asked to compare the appearance of a flat, matte, uniformly lit surface presented first under one illuminant and then under another. The data that theories are tested on thus does not include the kind of variation in illumination due to geometry that makes the phenomenological inadequacy most vivid.
11. Although the history of theories of color constancy is complicated and interesting, it is clear that computational theories take their inspiration from a point emphasized by Helmholtz. According to Helmholtz:

Colours have their greatest significance for us in so far as they are properties of bodies and can be used as marks of identification of bodies. Hence in our observations with the sense of vision we always start out by forming a judgment about the colours of bodies, eliminating the differences of illumination by which a body is revealed to us. (Helmholtz 1924, Vol. II, 286–87)

From this point of view, the illuminant is merely a means of revealing the stable surface properties and not a property of interest in its own right. This focus on reflectance is reinforced by the fact that one source of interest in computational theories of vision is their possible use in artificial vision systems whose main application is to industrial tasks. Here too identification of surface properties tends to be of greater importance than identification of the illuminant. There are no technical reasons that I am aware of that stand in the way of constructing theories that deliver information about the illuminant as well as the surface reflectance.
12. I am simplifying pretty dramatically here but not in a way that affects the argument that follows. Among other things, it is the cone outputs that contain all the information available about the color signal and the cone outputs do not fully capture the color signal.

