

How Things Look (And What Things Look That Way)

Abstract What colour does a white wall look in the pinkish light of the late afternoon? What shape does a circular table look when you are standing next to it? These questions seem simple enough, but philosophers disagree sharply about them. In this paper, I attempt to provide a new approach to these questions, based on the idea that perception modifies our epistemic dispositions regarding specific environmental objects. I shall argue that by determining *which* object is involved in this way, we can determine the subject of visual predication. This enables us to parcel out visual features to different visual objects in a way that enables us to reconcile conflicting philosophical intuitions.

Keywords: perceptual appearance, perceptual constancy, visual objects, predication, colour.

Index terms: appearance, constancy, scene-parsing, perceptual norm, conditioning, habituation, gestalt psychology, illumination, shadows, shape perception, colour perception, darkness, visual objects, discounting the illuminant, reflectance, appearance properties, subjects, predicates.

(I hope all personal names will be indexed.)

This paper owes a lot to a lot of people, including Imogen Dickie, Todd Ganson, Jennifer Nagel, Bence Nanay, Calvin Normore, Susanna Schellenberg, and Dustin Stokes.

HOW THINGS LOOK

What colour does a white wall look in the pinkish light of the late afternoon? – Well before sunset, that is, when the light turns lurid. Does it look different in colour than at noon, when the light is harsh and bluish-white? What shape does a circular table look when you are standing next to it? – Viewing it from a moderately oblique angle, that is, not as if you were squatting down by it. Does it look different in shape than if you were suspended above its centre, looking directly down?

These questions seem simple enough, but philosophers disagree sharply about them.

(1: Variable Appearance) Some imply that the white wall looks pinkish in the late afternoon, and that the table viewed obliquely looks elliptical. This used to be the majority view.

(2: Constant Appearance) Others insist, on account of the phenomenon known as “colour constancy”, that the wall looks white. Shape constancy would lead them to say that the table looks circular. (I discuss perceptual constancy at length in sections II – IV below.)

(3a: Both) Sydney Shoemaker (1994, 2006) takes a more complex view: he acknowledges constancy, and holds that in pink light, the wall looks white, just as it does at noon, but insists that it nonetheless looks different in some “unnamed” colour-related way; this “unnamed” property is supposed to account for the pinkish tinge in the scene. Shoemaker would presumably hold that the table looks circular, but nonetheless different than when viewed straight on in some unnamed shape-related way. (These “unnamed” properties are sometimes known as “appearance properties”: the wall is supposed to have an appearance property that makes it look different in the late afternoon. I shall return to a brief consideration of appearance properties in section VI below.

HOW THINGS LOOK

Susanna Schellenberg 2008 defends an objective version of appearance properties – she calls them “situation-dependent properties.”)

(3b: Both) Alva Noë (2005) says that the table looks *both* circular and elliptical, these being, as he says, two “moments” of visual perception. He takes a similar position on colour.

(4: Neither) Finally, some are persuaded that the wall looks *no* colour, the table no shape.

An intractable dispute, because each party seems motivated by considerations quite different from the others, and seem not even to agree about how to go about settling their differences. Introspection has not helped: if it rendered an easy verdict, there would presumably be no disagreement.

In this paper, I attempt to provide a new approach to these questions. After a brief overview of this complex dispute (section I), I embark on a re-examination of perceptual constancy (sections II-III). This leads to a reinterpretation of this phenomenon, which I will call the Scene-Parsing Thesis (section IV). In the second half of the paper, I apply the Scene-Parsing Thesis to the task of showing that, although it seems most correct to say that the wall looks white in pinkish light, and the table looks circular when viewed at a moderately oblique angle – in line with the Constant Appearance option above – none of the options mentioned above gives a full account of visual phenomenology. The questions posed above assume that colour, shape, and other visual features are always ascribed *to the wall* or *to the table*. On the account I propose this is not correct: some features are visually ascribed to other things in the scene. For it might be that though it is true that the wall looks white, it is nevertheless true that *the light* looks pink, and that, as a consequence, the wall looks as white things do in pink light. This introduces a complexity to the discussion of the issue that helps sort out some of the intuitions that lead to the conflicting positions described above.

HOW THINGS LOOK

My main goal is theoretical, of course, but I also hope that my discussion will actually enrich introspection and help readers become more acutely aware of certain aspects of their own visual experiences.

I. The Disagreement and Its Context

1. *Locality* Much empiricist theorizing about perception begins from the vague intuition that sensory experience reflects the state of the sensory receptors. Hermann von Helmholtz, the great physicist and pioneer of sensory psychology, propounded a precise version of this doctrine.¹ He thought that things produce sensory experiences that reflect the energy patterns by which they excite sensory receptors. He held, first of all, that “the excitation of any given nerve attached to a sense-organ produces its own experience.” And further, he held that the functioning of each such nerve is “discrete”, in the sense that it cannot influence the others. He concluded that:

Locality Thesis The visual experience occasioned by an object depends solely on the excitation of the retinal neurons onto which that object projects.

2. *Immediacy* Some empiricists draw from the Locality Thesis (or some vaguer precursor) the rather startling notion that we do not *immediately* see things as having *any* properties. Here is the first step of their argument:

By definition, we immediately see *x* only if seeing *x* does not have to be explained in terms of seeing something *else*.

Suppose that seeming to see *x* does not guarantee that one is *really* seeing *x*.

Then, seeming to see *x* has to be explained by seeing something other than *x*.

¹ See D. W. Hamlyn (1977) for a concise and informative account, from which the following quotations are drawn. See also Helmholtz (1977), Hatfield (1990), and Turner (1994).

HOW THINGS LOOK

Here is an example that illustrates the move. It might seem that we immediately see three-dimensional objects such as tables and chairs. However, if we were presented with only the hollowed out facing surfaces of the tables and chairs that we see, we would still seem to see the three-dimensional tables and chairs. Now, there is something in common between these situations – in both cases, we see the hollowed out surfaces. The philosophers we are talking about claim that in such a case, we see three-dimensional tables and chairs by seeing the common element – that seeing the facing surfaces partially explains how we see the tables and chairs. Generalizing from this, they claim that whenever error is possible, there is something that we see both in the non-erroneous situation as well as in the erroneous one. This common element is explanatorily what we see *immediately* in both situations.²

Reasoning along the lines outlined above, some classic empiricists constructed immediate seeing along the following lines:

Definition 1 S immediately sees x – where x may be a material object, fact, or event – if and only if the occurrence of S's visual experience *guarantees*, necessarily and *a priori*, the presence or occurrence of x – i.e., if no error is possible with respect to seeing x. (I'll often drop the 'immediately' in what follows.)

Berkeley, for example, argued (apparently following William Molyneux) that we cannot immediately see distance. For given Locality, our visual experience changes with conditions of viewing. As a consequence, there is nothing about the patch that a thing projects on the retina which guarantees that the thing is any particular distance away. By Definition 1, this means that we do not see things as at any definite distance from ourselves (*New Theory of Vision*,

² For a discussion of this kind of argument, see Martin (2005) – who uses the term 'directly' where I have written 'immediately'.

HOW THINGS LOOK

paragraphs 2-20).³ In fact, since no visual state guarantees the truth of any external state of affairs, we don't immediately see anything external.

3. *Isolation* How do these precepts apply to the white wall viewed in the late afternoon, and the circular table viewed from a moderately oblique angle? (I will be returning repeatedly to these examples, and will often refer to them simply as "the wall" and "the table".) The wall reflects a pinkish packet of light to the eye, and the relevant retinal neurons respond in a way that correlates with this colour. But the same pinkish packet of light could have been reflected by a pinkish wall in white light. By the Locality Thesis, it follows that this too would produce the same kind of colour experience. This experience is compatible, then, both with the wall being pink and with its being white. So it is not an immediate visual apprehension of either colour in the wall, or, for that matter, of any other. The same holds true of shape-perception. A thing looks a certain shape because of the topography of the retinal points to which it projects. Given different angles of viewing, the same projection could have been thrown by an object of a different shape. Hence, we do not (immediately) see any particular shape.

Generalizing from this, and assuming that what we immediately see corresponds to visual experience, some conclude that:

Isolation Thesis Visual experience does not constitute cognition of (is not "as of") anything possessing an objective property (where an "objective property" is one that an object has independently of anybody experiencing or thinking about it).

Anil Gupta (2006) offers what is perhaps the best recent development of this view. He says: "An experience taken in isolation does not pronounce on how things are" (cf. Charles Travis

³ Berkeley's argument is somewhat more complicated than my representation here. He allows that in binocular vision, the inward turning of the eyes betokens distance. However, he takes distance to be a tactual quality, and insists that this betokening is by learned association.

HOW THINGS LOOK

2004).⁴ This, of course, corresponds to position (4) above. It is a thesis about sensation, about the kind of cognition that it is or contains – about the cognitive significance of how things look. Gupta holds that experiences lead to perceptual judgements only when conjoined with “views” about the world, but that taken by themselves, they imply nothing about how the world is.

This is the argument for position (4: Neither) above.

4. *Externality* Let us simply grant that there is no sensory state *S* such that my being in *S* implies that the external world is a certain way. The *Isolation Thesis* is obviously not forced on us, since we may reject *Definition 1*. Some philosophers maintain, on the contrary, that:

Externality Thesis A visual experience is *in and of itself* (by virtue of its phenomenal character, for instance) as of an *objective* visual sense-feature.⁵

The idea behind *Externality* is that in order to know the message conveyed by a perceptual state, a perceiver needs nothing more than to undergo the experience itself. When I look at the table, I know in virtue of the experience itself, and nothing else, that it looks a certain colour. In order to know this, I do not need to know whether my being in this state increases the likelihood that the table is that colour. Moreover, these philosophers claim, the message

⁴ Gupta and Travis both attribute the argument to Descartes, but Descartes holds in the *Sixth Meditation* that: “The proper purpose of the sensory perceptions given me by nature is simply to inform the mind of what is beneficial or harmful for the composite of which the mind is a part.” It seems from this that Descartes holds that sensation pronounces on something outside the perceiver. The *Third Meditation* could be read as testing the minimum extent to which external objects are implicated in perceptual content.

⁵ Tyler Burge (1986) argues that certain scientific approaches to sensory processing are committed to something like *Externality*.

HOW THINGS LOOK

conveyed by perception is about external things: when we look at the wall or at the table, we receive a message that is about purportedly objective properties of these things.

5. *Normative Construals* Now, a considerable complication.

At first sight, *Externality* directly contradicts the position that Travis and Gupta take: it asserts, contrary to them, that visual experience *does* pronounce on how things are. But it may be an oversimplification to counterpose *Externality* and *Isolation* in this way. For the *Isolation Thesis* may be intended, by at least some of its proponents, *normatively*, not “naturalistically”.⁶ For one might make something like the following a condition of adequacy on the description of visual states:

Epistemic Norm of Correlation Whereas *seeing x* is a source of knowledge about *x*, so it is stipulated that:

In order for visual state *V* to count as a seeing of *x*, *V* must be correlated with *x* thus: (. . .).

Definition 1 can be understood as one way of filling in the dots in the above epistemic norm.⁷ It sets a very high standard, of course, but this is a product of a certain kind of attitude towards *knowledge*, not perception. One could get a more accommodating norm of correlation by lowering one’s demands on knowledge – however that might be, one’s conception of perception would derive from these demands.

⁶ See Gary Hatfield 1990, especially chapter 1, for a stimulating discussion of this distinction. Hatfield takes Descartes and Locke to be concerned with norms: for Descartes in particular, “to call something a ‘thought’ is to imply that it meets certain standards” (17).

⁷ Gupta insists that his theory is not empirical, but logical (182); it is about the justificatory relationship that perceptual experience bears to judgement. Travis is harder to read on this point.

HOW THINGS LOOK

On a normative construal, the *Isolation Thesis* is a conceptual consequence of the fact that no visual state satisfies the chosen epistemic norm with respect to any external object. By contrast, *Externality* tells us only how visual states instinctively seem to us – as determined by introspection, or by the kind of investigation of this question that takes place in a psychological laboratory. On this construal, there is no conflict between the two perspectives. Externality is advanced as a fact about phenomenology; understood normatively, Isolation seeks to undermine the epistemological pretensions of visual phenomenology.

This complication about normativity need not concern us right away; the debate that will occupy our attention for most of this essay occurs within the province of the *Externality Thesis*, though in due course, I shall sketch an alternative to the *Epistemic Norm of Correlation*.

6. *The Disagreement* Bracketing the *Epistemic Norm*, then, here is the debate. There are some Externalists, firstly, who claim, on the basis of introspection or on other psychological grounds, that *the wall* looks pinkish in the late afternoon, and that *the table* looks elliptical when one is standing next to it – or more generally, that these external things look different as conditions of observation change.⁸ Change the light and the sensed colour-properties of the wall also change; change the angle of viewing and (generally speaking) the sensed shape-properties change.

Thus:

⁸ Noë (2005, 89) borrows from painters' techniques a way to *prove* these claims. He says that the table is "looks elliptical to me because, to indicate its shape, I can (and indeed, in some sense, *must*) move my hand in a characteristic manner". Similarly, he says, the white wall looks pink because you can tell that it matches a pink chip (*ibid*, 140). While his approach is suggestive, I do not agree that it proves that the table looks elliptical or the wall pink. (I'll return to this in section V.) However this may be, it is worth noting that Noë is proposing a principled way of deciding a question that others attempt to adjudicate on intuition alone.

HOW THINGS LOOK

Variation Thesis When (relevant) viewing conditions change, things look different (provided that these things themselves stay the same in the relevant respect).

Because *Variation* is meant to be a *phenomenological* fact, it is not hostage to the *Locality Thesis* or to any other psychophysical theory, even though it is often inspired by *Locality* and ends up with similar results. Supporters of the *Locality Thesis* (such as Helmholtz) are committed to *Variation*, but supporters of *Variation* (such as Berkeley, Shoemaker, or Noë) need not take any position on *Locality*.

Other Externalists deny *Variation*, insisting instead that:

Constancy Thesis In good conditions of viewing, things look the way they are, even if these conditions change (provided that they remain “good”).⁹

Proponents of this way of thinking point out that, contrary to what one might expect given *Variation*, normally endowed perceivers have no tendency at all to mistake the colour of things in polychromatic light (cf. Bradley and Tye, 2001), or their shape at moderate angles of viewing. Indeed, if asked to describe things in such circumstances, or to draw a picture, they tend, spontaneously and immediately, to represent them as they really are, and not as the *Variation Thesis* supposes they look. Of course, there are situations in which things look otherwise than they are. In a particularly vivid sunset, white things might look violet. Or there may be

⁹ Note a certain parallelism between Constancy and Definition 1. The latter sets a normative standard of correlation that a visual state must meet in order to count as a perception of something. Constancy suggests that in good conditions a similar standard is satisfied – though not so strong a standard as Definition 1 assumes. This opens up the possibility of plugging good conditions into the Epistemic Norm of Correlation to produce something less demanding than Definition 1. I’ll ignore this: I am interested in formulating a rather different kind of Epistemic Norm.

HOW THINGS LOOK

situations in which our visual states don't present things very clearly: it may not always be visually evident whether something is circular or elliptical. (This is why I have been emphasizing good viewing conditions in relation to *Constancy*.)

Finally, there are hybrid positions. Shoemaker endorses both *Constancy* and *Variation* but with reference to two different sets of properties: he suggests that *Constancy* is true of color and shape, while *Variation* is true of a different set of unnamed color-related and shape-related properties (the appearance properties). Noë suggests that *Constancy* and *Variation* mark two different "moments" of visual experience – in one way, or relative to one attentional attitude, the wall looks white, relative to another, it looks pink.

II. Perceptual Constancy

The philosophical disagreement among those who subscribe to the *Externality Thesis* is about how things look phenomenally. As noted in the introductory paragraphs of this essay, the persistence of this disagreement seems to imply that casual introspection will not decide the issue. We will need, cautiously, and with due regard for proper philosophical methodology, to consult empirical studies of perceptual experience to help and enrich introspection. In the bulk of what follows, I shall be concerned with *Variation* and *Constancy* – I shall try to show that neither tells the whole truth, and propose a third alternative – the *Scene-Parsing Thesis* – that also presupposes *Externality*. I shall make some remarks about the *Isolation Thesis* as we go along, and conclude with a brief assessment in section VIII.

Let's begin with a well-known phenomenon, or rather, with the well-known aspects of a phenomenon that, as we shall see, turns out to be surprisingly complicated. It has been known for a very long time that in a wide range of circumstances, things have a tendency to look the way they are. Here is Robert Woodworth's (Woodworth and Scholsberg 1954) classic description:

It is a fact of common observation that coal looks black even in sunlight, and chalk white even in shadow. Yet under these conditions the eye receives much stronger light from the coal than from

HOW THINGS LOOK

the chalk. The retinal image of the coal is much the brighter, but the observer's impressions seem to conform to the objects rather than to the stimuli. He sees no paradox here, for if coal is really black, it certainly ought to look black in a good light. Not a bad answer, perhaps, but one that poses a problem for the psychologist. (*ibid.*, 428)

Woodworth's observation is one of many that seems to support Constancy. As I observed earlier, it is also commonly observed that things look the colour they are in a variety of viewing conditions. Similar observations have been made with regard to the visual perception of shape. Pierre Jolicoeur and Keith Humphrey (1998) review evidence that given certain "canonical" angles of viewing – for example, views that reveal three sides of an object without excessive foreshortening – even relatively irregular objects (wrenches, saucepans, and even odd-looking contrivances with arm or nose like attachments) can reliably be recognized after rotation.

These observations are bolstered by experiments in comparative vision. Train a bee or a bird to expect pure water in green dishes and sugar water (which they generally prefer) in yellow ones. Now place the yellow dishes in conditions where they reflect greenish light – under heavy leaf cover, for instance. The bees will still go to the yellow dishes in search of the sweet water they prefer. This shows that they recognize the yellow dishes for the colour they are, despite the colour of light they reflect.

Notice the three elements of this behavioural complex:

First, an *unlearned* or *unconditioned* input, perceived colour,

Second, an *innate operation* – for example, classical or (as in this case) operant conditioning – that modifies the organism's responses to the environment in response to discovered correlations, and

Third, a *learned* or *conditioned* response to yellow dishes.

Note that bees that have been trained to expect sugar-water in yellow dishes will go to these dishes the very first time they see them in non-standard illumination. Learning is involved with

HOW THINGS LOOK

regard to where the bees search for sugar water – in yellow dishes, not green – but not in how they perceive colour in non-standard illumination.

In short, animals react to the same *colour* or the same *shape* the same way, even when these features are presented in unfamiliar conditions of viewing. And this is true very widely across the animal kingdom. As Christa Neumeyer (1998) has written,

In all animals investigated so far, color constancy has been found. Even animals with color vision systems probably developed during evolution independently of each other, such as the honeybee, the goldfish, and the rhesus monkey, show this property (347).

Colour constancy in animals is tested by learning or conditioning experiments such as the one just described. Now, it is possible that nothing *looks* any way to honeybees; perhaps honeybees do not possess consciousness at all (cf. Hardin 1988, 151-2). Still, it is not unreasonable to suppose that as consciousness emerges in evolution, it sits, however precariously, on top of the kind of visual processing that is present in more primitive visual systems. In other words, the unlearned input to conditioning in these more primitive systems is preserved in more evolved systems as a determinant of conscious visual experience. This is confirmed by the fact that in humans, who can directly report how things look, conditioned responses to things correlate with how they look subjectively.

The Gestalt psychologists emphasized a shape-related phenomenon of the same kind. Kurt Koffka (1935) describes it in this way:

If the after-image of a circle is projected on a plane which is not frontal-parallel, it will appear as an ellipse . . . (*ibid.*, 211) What is true of shape is equally true of size, the size of an after-image being a direct function of the distance at which it is projected . . . Mrs Frank (1923). . . had her students project after-images on a plane surface on which a perspective drawing of a deep tunnel had been made. Then the size of the after-image varied with the place on the sheet of paper on which it was projected; if it fell on a part corresponding to a near part of the tunnel it was considerably smaller than when it was thrown on a part corresponding to a remote part of the tunnel . . . (212)

HOW THINGS LOOK

Koffka cites this phenomenon in support of the view that the perception of wholes is not simply composed of the perceptions of their parts. Rather, as Max Wertheimer (1924) put it, “what takes place in each single part already depends upon what the whole is.”

All of this vindicates the *Constancy Thesis*. *Variation* cannot be a proper account of how things look – at least not on the simplistic reading on which it is taken to imply that things always look a different colour in different coloured light, or a different shape when viewed from different angles. (Note that Shoemaker and Noë, who endorse *Variation* do not take it this way: both think that things look the same colour in different conditions – but also a different colour.) Organisms clearly react to colours and shapes as such, even though they encounter them in different conditions of viewing.

Now, Woodworth’s “problem for the psychologist” is this. How does the visual system accomplish the feat of converting variable retinal signals received from objects into invariable features as they really inhere in objects? Historically, the key to solving this puzzle lay in the discovery that, contrary to the *Locality Thesis*, the *same* local retinal response can produce very different observer-impressions in different contexts. One particularly revealing demonstration was provided in 1789 by Gaspard Monge. Here it is, in the words of J. D. Mollon (2003):

To illustrate his lecture, Monge had hung a red cloth on the wall of a house opposite the west facing windows of the meeting room of the Academy. He invited his fellow *académiciens* to view the red cloth through a red glass. The appearance of the cloth was counter-intuitive. Seen through a filter that transmitted predominantly red light, it might have been expected to continue to look a saturated red. But no, it looked pale, even whitish. The same was true when the assembled company inspected one of their fellows who happened that day to be wearing a red outfit. A yellow-tinted paper examined through a yellow glass looked absolutely white. Monge was aware that his illusion (we may call it the Paradox of Monge) was strongest when the scene was brightly lit and when there was an array of variously colored objects present in the scene, including objects that one knew to be naturally white. When all that was visible through the red glass was a red surface, the effect was abolished. (*ibid.*, 20)

HOW THINGS LOOK

So much for the *Locality Thesis*¹⁰: the apparent colour and shape of an object does not co-vary with the image *it* projects to the retina taken in isolation. For as the Gestalt psychologists realized, how an object looks depends on the array of light received by a much larger area of the retina, or perhaps the whole of it. The visual system computes the true colour of a small region using characteristics of the whole scene. (Early theorists, such as Ewald Hering, were inclined to think that constancy is entirely a retinal effect, involving an inter-neuron damping process now known as lateral inhibition. This is almost certainly wrong, as the computational approaches discussed later on imply.)

These experiments and demonstrations (and hundreds of others like it) led investigators like Monge to suppose that the context-sensitive process that leads to evident mistakes of perception – the red cloth looking white – leads in more normal circumstances to things looking the way they really are. The mistakes are, in other words, diagnostically valuable errors of normal sensory processing. This would show that the Variation Thesis outlined in the previous section must be wrong if it is taken in a certain way – a rather simplistic way, as we shall see in the following section. Woodworth’s lump of coal looks *black* in good light, even though it is reflects less light than the chalk in shadow right beside it. In *this* respect, at least – with respect to looking black, that is – the coal looks the same in bright light as it does in shadow.¹¹

¹⁰ William James was already well aware of this. In a section of *The Principles of Psychology* written by E. B. Delabarre, we find the following “Law of Contrast”: “The color and brightness of one object always apparently affects the color and brightness of any other object seen simultaneously with it or immediately after” (662). James noted that Helmholtz had attributed the Law of Contrast to a “deception of judgment”, but he himself endorsed Ewald Hering’s contrary view that the effect depends “purely” on the “physiological action of the terminal apparatus of vision” (667). (This is probably wrong, as I note in the main text.)

¹¹ Interestingly, Woodworth qualifies this: “The regression or shift is from the stimulus color, size, or shape *toward* the object color . . . the coal looks black in any light but not quite so

HOW THINGS LOOK

Similarly, the very same after-image looks a different size and shape when projected on different backgrounds.

Of course, this is not to claim that things will *always* look *exactly* the colour they really are. The claim is just that *in a wide range of circumstances* – in “good” conditions of viewing – they will look the colour they are. Even this has to be qualified. It has been demonstrated that constancy is only approximate. Observers cannot match colour samples shown by side by side in different illumination. They are usually successful in identifying what broad colour name (e.g., blue, yellow, etc.) should be applied to a sample. But they make plenty of mistakes within such categories. Nevertheless, there is some level of description under which things look the same colour in different conditions of viewing.

III. *Extracting the Illuminant*

Let these observations about *Constancy* rest for a while – we are not quite ready to deal with them systematically. Here, however, is a question prompted by the material just covered. Does *Constancy* imply that things look in all respects the same in good viewing conditions?

Even disregarding the approximate nature of colour and shape constancy, surely not. For the simple introspective observation that proponents of *Variation* offer does seem in some sense correct. It is impossible to deny that a lump of coal looks different in bright light than it does in shadow; a wall looks different in pinkish light than it does at noon. How exactly do they look different? Many – among them the proponents of *Variation* – find it reasonable and straightforward to answer that the coal looks lighter; the wall pinker. But on the face of it, this contradicts the constancy phenomena we just visited. So we have a problem. If things look the

black under very bright as under a relatively dim light.” This seems wrong to me. Assuming the presence of an “array of variously colored objects” in the scene (cf. Mollon), it seems to me that the coal will look blackest in bright light.

HOW THINGS LOOK

same in colour and shape despite changes of viewing conditions, how are we to account for the widely acknowledged – undeniable, really – difference of appearance?

Common sense suggests a way to resolve this difficulty. Admittedly, sunlit coal looks the same in one way as it does in the dark – it looks black. But it simply *can't* be the case that a lump of coal looks *just* the same in bright light as it does in shadow. If it did, how could one tell by looking at it whether it is in bright light or in shadow? And surely one *can* tell this. It follows that the coal should, at the very least, look, when it is in bright light, or shadow, as if it were in bright light, or shadow. And this suggests a different possibility concerning how it looks – it may look the same in one way, different in another. Perhaps it looks *black* both in sunlight and in shadow, but *differently illuminated* in different conditions of illumination. Similarly, the table: perhaps it looks round viewed perpendicularly, and also from any moderately oblique angle, but when it is viewed obliquely, it looks different in that it looks as if it is being viewed obliquely. Perhaps, then, the cure for the ills of constancy is more constancy – constancy with respect to conditions of viewing as well as constancy as to properties of objects. (Note, however, that this accommodates the Variable Appearance attitude (1) sketched in the Introduction above, but with a subtle amendment: the coal looks different (as does the wall or the table), but it looks different with regard not to colour, but to how it is illuminated or viewed).

This elementary reflection is taken up and supported by a seminal but somewhat neglected reflection on colour constancy. Dorothea Jameson and Leo Hurvich (1989) complain that recent computational approaches to constancy concentrate exclusively on a single question: how constant surface reflectance properties of objects can be derived from reflections under unknown illuminants. But, they protest:

Apart from its use in the computation of surface reflectance characteristics for object recognition, perceptual information about the different conditions of illumination as relevant in its own right [is] largely if not totally ignored in [the computational approach]. Our own judgment is that human visual systems . . . are likely to have evolved a design that provides perceptual information about

HOW THINGS LOOK

change as well as constancy – about light, weather, and time of day, as well as about the relatively constant physical properties of mainly opaque objects within a scene. (*ibid.*, 2-3)

In other words, much of the literature assumes that the visual system's overriding task is to discover constant properties of external objects and simply to discard the interfering variable information that confuses the retinal image. This is the idea implicit in Helmholtz's famous phrase, "discounting the illuminant". But, as Jameson and Hurvich imply, the visual system does not simply discard or discount information about transitory environmental conditions – it provides us with this information also. It would be much more accurate to say that the system *extracts* the illuminant, as well as many other features of the scene. In short, colour vision does not tell us simply that an object is white or pink, but also that the illumination in which it stands is bright or dull, bluish-white or slightly pink. Earlier, I implied that the wall might look as a white thing does in when it stands in pinkish light. If Jameson and Hurvich are right, this would be a consequence of the wall looking white *and* the light looking pink.

It is easy to extend this reasoning to the perception of shape. We recognize shapes from different angles of viewing. Suppose you are shown a long-handled oval saucepan, and then shown it again after it has been rotated by a small amount (say 20°) in three dimensions. The rim might project a different shape; the handle might now be somewhat foreshortened. Still, you would not find it difficult to recognize the shape, or to detect a switch with a pan of different shape or dimensions. Presumably, the visual system is able to register the shape of the pan, and then to rotate this shape (or rather its image) for purposes of comparison. The difficulty of this feat of visual recognition should not lead one to think that in such cases the system is interested simply in shape. Though it is only a transitory property of a viewer in a scene, orientation, is just as important in general: after all, it is key to how we locate ourselves and find our way around. When the visual system identifies shapes presented at different angles, it is simultaneously making a diagnosis of how the perceiver is oriented relative to the thing in question.

HOW THINGS LOOK

V. Constancy vs Scene-Parsing

In the traditional view, there is a river of undifferentiated information that flows from different parts of the visual field to the observer. In response to perceptual “constancies”, theorists have asked how and to what extent the perceiver, or her sub-personal visual system, is able to extract object-information from this flow by discounting changes and variations of viewing conditions. In the previous section, we glimpsed an alternative way of posing the problem. We can take the visual system to be *parsing* the flow of information into its various components. The task is not to assign a definite value to one variable despite the interference of others; it is rather to assign values to all – perhaps it would be better to say ‘many’ – of these variables *simultaneously*. I’ll call this the

Scene-Parsing Perspective Where sensory receptors are excited by signals determined by several environmental objects varying independently in their properties, the function of the sensory system is to extract from the state of the sensory receptors an assignment of properties to all such environmental objects as are of interest to the organism. (Generally, it does so by exploiting normally existing environmental correlations found in *scene-wide* statistics.)

Terminological Note Up until now, I have been using the the term ‘object’ somewhat vaguely, but mostly to refer to material objects. This is in line with the views that I have been discussing, which restrict their focus to material objects, and to certain other things with definite boundaries, such as figures on a ground. I will continue to use ‘object’ this way, i.e., vaguely, but with a focus on material objects. However, I will henceforth use the terms “environmental object” and “stimulus” to refer more broadly to anything in the environment that is of interest to the organism. This would include illumination, for example.

It might be thought that the two problems of solving constancy problems and parsing scenes are not very different. Given that the colour and brightness of the light received by the retina depend on both object-reflectance and ambient illumination (as well as a number of

HOW THINGS LOOK

other variables that need not be mentioned here), the visual system has to solve a single equation in more than one unknown in order to assign a value to reflectance. The problem, as traditionally conceived, is to figure out where in the undifferentiated information flow the system finds enough information to solve for two unknowns simultaneously. Now, it is surely no more difficult to solve a simultaneous equation in order to determine reflectance than to solve it in order to determine both reflectance and illumination. What then is the advantage of adopting the new perspective?

The question is correct in its presuppositions, but I am not proposing the change in perspective in order to lighten the computational load on the visual system, or to simplify its task. (And it is surely a good thing that I am *not* suggesting that more than two centuries of work on constancy be thrown away.) The advantage I am claiming is that the Scene-Parsing perspective enables us to reinterpret the *function* and *output* of visual processing, and thus to look differently at a number of longstanding puzzles.

Already, one old puzzle has taken on a new aspect. A number of philosophers have wondered how perceptual constancy is compatible with the obvious fact that in some clearly colour-related (or shape-related) way, things look different when illuminated differently (or when viewed from a different angle). This puzzle is dissolved when we note that things may look different *not only* when they look a different colour (or shape), but because they look differently illuminated, or as if they are being viewed from a different angle.

A second puzzle, or rather challenge, is one that C. L. Hardin (see, e.g., 2003) has been posing forcefully for quite a while. Hardin protests that colour constancy is a myth. It is certainly true, he would concede, that there is less variation in appearance than the Locality Thesis might lead one to expect. Still, it is demonstrable that observers make a lot of mistakes about the colours of things in different illumination, different surrounds, different conditions of adaptation, and so on. (See Jameson and Hurvich 1989 and Foster 2003 for more details.) And similarly, it is demonstrable that people make a lot of mistakes about the shape of things after

HOW THINGS LOOK

they have been rotated (Jolicoeur and Humphrey 1998). Why then should anyone insist that the colour vision system is trying to discount the illuminant, or that the shape perception system is trying to discern object shape? If these systems are so bad at doing these things, why suppose that this is their function? Thus some psychologists and some philosophers propose that constancy is not what the system is after: constancy is merely a side-effect of certain properties of neurons, they say – lateral inhibition, contrast, and such – not a scene-wide constraint.

However, the Scene-Parsing Thesis need not and does not insist on absolute constancy; indeed, given the complexities of the computational problem alluded to above, it is better that it should not do so. The claim is that the organism has an interest in a number of different variables, not just one. The light that reaches the eye contains information about many of these all mixed up together in a way that would impede the visual system's attempt to deal with them separately. The following questions have very different implications for the perceiver. Is that thing round or elliptical? Am I positioned directly above it or off to one side? If information about shape and orientation were inextricably mixed together, the perceiver has no way of dealing appropriately with either. *Any* help with this matter is useful – “more or less circular”, “somewhat oblique” are vague answers, but they improve on the confusion of the retinal image. There is, therefore, no need to insist on *constancy* in a sense that betokens a precise assignment of values. We should rather focus on the separation of relevant aspects of the scene with as much precision as the task and the circumstances allow. On the Scene-Parsing proposal, then, Hardin's challenge misses the mark.

It is entirely plausible, then, that perceptual constancy is, as Hardin has said, at best an approximation. But the visual system's function ensures that things will look different in different conditions of viewing, and this function can be adequately performed even if things do not end up looking exactly the way they are in all good conditions of viewing. The system is trying to gather useful information about a variety of environmental factors. Information can

HOW THINGS LOOK

be useful even if it is not exact. The benefit of the Scene-Parsing Thesis is that it emphasizes separation, not exactness.

VI. The Things We See

As noted before, many philosophers who embrace the Externality Thesis multiply sense features in the attempt to illuminate the complexity of visual experience. For example, Shoemaker correctly argues that there is more to colour experience than the colours attributed to external things. He says:

If the surface of a table is partly in shadow, one may say that the way the shadowed part of it looks is different from the way the unshadowed part of it looks without implying that the two parts look to have different colors (2006, 462)

He concludes that there are “phenomenal ways that things look”, unnamed visual *properties* that account for the different colour-appearance of things that look the same in colour. The idea is that the brown of the shadowed table appears differently to a perceiver than the same brown when unshadowed because it has a different *appearance* property. *Brown* is a colour, but there is also an “appearance property” associated with *brown* in each condition of viewing, and vision attributes appearance properties to things, not just colours. Similarly, Noë argues that perceptions of shape and colour have two “moments”. When a circular table is viewed obliquely, it is perceived as both circular and elliptical, though in different ways. This is an attempt to deal with the fact that the table looks different when viewed from different angles, though it does not look as if it had changed shape.

The Scene-Parsing Thesis offers us a different way of accounting for complexities of visual appearance. Where Shoemaker and Noë account for the phenomenology by supposing that vision attributes new kinds of properties – appearance properties – to the same object, or attributes different familiar properties to this object in multiple “moments”, Scene Parsing appeals to familiar kinds of properties, and attaches them to a plurality of familiar *stimuli*, including material objects, but also illumination, etc. The claim is that the colour vision system

HOW THINGS LOOK

separates information concerning illumination from information concerning colour. Thus, it might determine not only that the surface of the table is brown, but also that the light that falls on it is uneven – bright here, less bright there. If so, the different appearance of different parts of the table lies not just in the sense-features that these parts possess in themselves, but in differences concerning the illumination in which these parts stand.

Here, then, are two accounts of the logical form of the content of the visual state. Shoemaker's analysis runs like this:

Appearance Property analysis Brown (whole table) and Appearance-different-Brown (unshadowed part) and Appearance-yet-different Brown (shadowed part).

I am suggesting an analysis along different lines;

Scene Parsing analysis Brown (whole table) and In-the-light (unshadowed part) and In-shadow (shadowed part).

Scene-parsing can be a good bit more complex than indicated here. We shall come to that in the next section. The point to note for now is the familiarity of the properties and stimuli to which we appeal. The analysis does not make use of appearance properties, and thus it retains the idea that vision is concerned with properties that environmental things possess independently of observers; it also does not make use of anything like Husserl's "phenomenological reduction" or Noë's "two moments" to account for differences of appearance. I do not want entirely to deny the validity of such approaches, but in my view, they are most at home when one wants deal with *meta*-visual states – attention and reflection. They are out of place in the account of visual content.

To add weight to the Scene Parsing Thesis, consider now how the visual system computes illumination. First, a largely unappreciated fact: there are *direct* indicators of illumination available to vision. The size of the pupillary aperture indicates *average* brightness, and lightness of colour is calibrated across the scene in part by setting surfaces near to this

HOW THINGS LOOK

average as of middling lightness.¹² Moreover, both the brightness and the colour of illumination are evident in shiny surfaces (Lee 1986). These surfaces reflect light in two ways: some of the light incident on them bounces off a shallow layer without much selective absorption (“specular” reflection), and some of it goes deeper and is selectively reflected. The specular reflections are chromatically close to the incident illumination, while the deep ones convey information about the surface reflectance of the object. Thus a leaf looks green because deep reflections absorb and reflect light selectively; nevertheless, its shiny surface gives your visual system a good indication of the quality of ambient illumination. Shiny surfaces are ubiquitous, of course – water, shiny leaves, metallic objects, and human faces. The shiny spot on your friend’s forehead affords your visual system direct access to the quality of illumination in his vicinity. The very fact that colour is inextricably associated with qualities such as *shiny*, *luminous*, *matte* and so on indicates this aspect of the scene-parsing activity of the colour vision system.

Secondly, the visual system employs certain “higher-order scene statistics”, as Jurgen Golz and Don MacLeod have called them (Golz and MacLeod 2002, MacLeod 2003).¹³ For

¹² Lightness is a component of colour sensed as belonging to objects. Pale yellow is lighter than orange; pink is lighter than red. Brightness, on the other hand, is a measure of the amount of light that a surface sends to the eye. Strong light is generally sensed as brighter than dim light, but also light of certain colours (e.g., yellow) is brighter because it excites more cone-cell types than light of other colours (e.g. violet).

¹³ I am going to be deliberately careless in what follows about attributing personal agency to the subpersonal visual system. It would be more accurate, but also annoyingly circumlocutory, to say something like: “Neural activity in the visual system shows certain parallels to a process that computes higher-order scene statistics, and the output of this process correlates with the colour of the illuminant.”

HOW THINGS LOOK

example, it computes the correlation between chromaticity and luminance over the whole scene: if it finds, for instance, that all of the brightest places in the visual field are sending mostly red light to the eye, then it infers that the scene is either being illuminated by red light or being filtered through a red film. (Thus, the Paradox of Monge.)

Now, here is the important point. The computations that have just been mentioned are ways to estimate the brightness and colour of the illuminant at various points of the scene. The assignment of reflectance to surfaces comes (mathematically and computationally) *after* the system has estimated illuminant information. Why? Because the system assumes, as it were, that illumination is a property of the whole scene – not that it is uniform across the scene, but that it obeys certain laws of variation in that scene. If it comes from a certain direction, it will weaken in the opposite direction; if it falls on an angled planes, it will show predictable variations of brightness across the edge; an opaque object in its path will cast a shadow; if there is more than one light-source, there will be complex patterns of brightness and darkness, including illumination fringes. In all these conditions, light spreads across a scene in more or less predictable ways. Visual computation takes such predictability into account. For example, Donald Hoffman (1998, 115) suggests, as a computational rule employed by the visual system: “Interpret gradual changes of hue, saturation, and brightness in an image as changes in illumination.”

Reflectance, by contrast, varies more or less randomly across the scene,¹⁴ and the reflectance – i.e., colour – of something in one place gives no indication of the reflectance of

¹⁴ John Mollon (2000) has argued that primate colour vision is specially adapted to “dappled and brindled” scenes. A brindled scene is presumably one in which there is a high spatial frequency of random reflectance variance, as in an Ishihara plate; a dappled one is where there are equally frequent changes in illumination, as in a forest where there are many pinholes in the leaf cover. Notice that even in dappled scenes, the colour of the illumination is predictable, even though its intensity is highly variable.

HOW THINGS LOOK

other things elsewhere; for example, a thing of one colour offers no indication of the colour of an object next to it. Thus, the system figures out the predictable first, and computes the unpredictable in the light of its conclusions. Once illumination is mapped out, the residual variations in colour of light arriving from different parts of the scene are attributed to differences of surface colour.

All of this means that from very early in visual data-processing, the various causes of receptor activity at a point are separated out and differently treated. These separated flows of information result eventually in our awareness of different *kinds* of environmental stimulus – material objects and illumination have been mentioned so far; I shall discuss some other kinds of stimuli in the following section. Suppose that red light is emanating from a certain place *P* in the visual field. How the visual system utilizes this information depends in part on the kind of stimulus the red light is attributed to. Suppose that *P* is diagnosed as occupied by a shiny surface specularly reflecting a spot of red light. The visual system will construe the red light coming from *P* this as red illumination *in the general vicinity*, even though this illumination is directly manifested only at *P*. As a consequence, red light emanating from adjacent matte surfaces will be discounted, and these surfaces will look less red than the light they reflect. Suppose, by contrast, that the red light coming from the very same place *P* is attributed to the reflectance of a matte surface. Then there will be a tendency for surrounding surfaces to look *more* red if they are within the same object boundary as *P*. The *Locality Thesis* is completely off-track here: one's experience of *P* does not depend solely on the kind of light that emanates from *P*. Attributing red light to the reflectance of a matte surface has quite a different phenomenological consequence than attributing it to specular reflection off a shiny surface or to a red light.

This kind of scene-parsing has an impact on visual phenomenology. The representation of illumination is not amodal, non-visual, or merely implicit. It is visually apparent to a sighted observer that she is standing in reddish illumination, or in bright sunlight, or in the dark. This impression is distinctly *voluminous*, as Hering says, quoted by William James (1890/2007):

HOW THINGS LOOK

A shady corner in an otherwise well-lighted room is full of a darkness which is not only *on* the walls and floor but *between* them in the space they include. Every sensation is there where I experience it, and if I have it at once at every point of a certain roomy space, it is then a voluminous sensation. (*ibid.*, vol 2, 137)

It is sometimes assumed that seeing colour is a matter of seeing coloured objects. Thus, it is assumed that sensing the colour of illumination is a matter of sensing the resultant colour of the objects that this illumination falls upon. This assumption is contradicted by the observation that one senses the colour of illumination and darkness as voluminous. For object colour is always surface colour.

The impression of reddish illumination that one gets at dusk is not merely a matter of seeing a lot of reddish objects. A red thing is seen at a definite distance at a definite angle relative to the viewer. The pink of dusk, by contrast, engulfs the viewer – it is not located in any surface. It has an effect on surfaces, of course, and it is visible in such stimuli as *rays*. In a ray, the motes of dust reflect coloured light – but in such a way that the illumination is sensed, less directly but nevertheless sensed, between the motes of dust. Such visual data are a result of the system extracting the illuminant and ascribing it to the observer's environs, rather than to any particular stimulus in it.

The visual impression of the dark is similarly not simply a matter of seeing nothing, or seeing things dimly: one has the impression that one would not be able to see something close by, even if there was something to see. If there is a small bright spot in the dark corner, it looks luminous – as if it were producing its own light in the midst of the gloom. The impression of *being in* darkness has its effect on how one behaves too: instinctively, one is cautious and picks one's way carefully. These are not learned responses: one does not act in these ways *because* one senses that it is dark. These action-potentials are rather part of the instinctive response to the dark; they mark the phenomenology of darkness. This phenomenology is not merely a matter seeing dark objects. Rather, it is a matter of sensing that there is no illumination.

HOW THINGS LOOK

Very much the same sort of thing holds for shape perception. To some extent, this phenomenology is multi-modal; different sensory systems are involved. Observers are proprioceptively aware of their own orientation relative to the horizon. When your head is tilted, you are aware that it is, not just by visual means but by the body's own specialized sensing system for position. Moreover, you are able to track the changes in your position as you move. As you move, the projected shapes of things around you change. Thus, you are able to track the correlation between your own movements and changes in projected shape. This gives you a great deal of information about the shapes of things around you. This information is not all muddled together. Perceptually speaking, it does not express itself as a single changing variable – the-two-dimensional-projection-of-that-thing-relative-to-my-current-perspective. Rather, it appears as two separate parameters, one of which is changing – the observer's position – and the other of which is normally invariant – the shape of the object. (Imagine something changing shape as you move, keeping its retinal projection constant: since the system is able to utilise proprioceptive information about your own movement, this thing will look as if it is changing in shape.) In addition to this, the system uses probabilistic information relating to invariants of projection in "canonical" or "generic" views. This information too enables the visual system to assign values separately to the distinct variables that contribute to patterns that objects throw on the retina.

VII. Subjects and Predicates in Vision

Let us concede that when we look at a pinkly illuminated wall we are visually aware of an instance of pink in the same location as the wall. It does not follow that the *wall* appears pink. It is possible, after all, to experience two visual features in the same place without experiencing them as belonging to the thing. Suppose that you are looking through a window at a flower outside. Somebody standing behind you is reflected in the window and appears to be in exactly the same place as the flower. You would have no tendency to attribute to the flower the qualities of the reflected person: the flower may be blue and the reflected person brown, but these colours do not leak from one object to the other. Though the two objects are in the same

HOW THINGS LOOK

place, the visual system keeps their qualities separate. By the same token, one cannot assume that just because the pink we sense is co-located with the wall, it is therefore attributed to the wall. It could be the illumination that looks pink. Or *pink* could be manifested more locally – perhaps the wall looks as if it stands in pink illumination.

The question I am asking is: “Within what sort of predicative structure does vision place the instance of pink I see over there?” Or: “What is the subject of this predicate, *pink*?” The argument to this point has rested largely on how the visual system processes data; I have attempted to use this data to show that vision presents us with separated data-packages. But since the data concerns pre-conscious processes, it does not tell us about the details of visual phenomenology. The current question concerns just these details. And without some way of anchoring this issue in empirical bedrock, the question might seem excessively speculative. But what would count as evidence towards a resolution?

Let us start with this important clue. It has been known at least since the 1930s that constancy experiments depend crucially on instructions. As Jameson and Hurvich say, there is a difference between “making an adjustment to make a particular part of a display *look* identical to the same area in a differently illuminated display, as contrasted with an adjustment to make a particular surface in a display *seem identical in its surface characteristics* to the same object in a differently illuminated display” (their emphasis). Lawrence Arend *et al* (1991) similarly distinguish between matching “unasserted colour”, colour that corresponds closely to that of light coming off a surface, and “asserted colour”, the colour of the surface itself. These experimenters worked with pairs of colour-wise identical Mondrian displays (quilt-like patchworks of coloured papers) under different illumination. Experimental subjects tried to adjust the colour of a patch in one display (the *test* patch) so that it would match a patch (the *target*) in the other. To test their abilities regarding *unasserted* colour, subjects were asked to make the test patch match the hue, saturation, and brightness of the target *disregarding all other areas* in the display as much as possible. It turns out that in this task, they adjust the test display away from its true colour, so that it is sending close to the same light to the eye as the

HOW THINGS LOOK

target. With regard to *asserted* colour, they are asked to “make the test patch look as if it were cut from the same piece of paper” as the corresponding patch in the other display. In this task, they are good at matching the true colour of the target. This confirms the intuition that we have been relying on all along: on the one hand, these subjects were good at determining what colour things were, even when the colour of the illumination changed; on the other hand, there was a difference of unasserted colour in different illuminations.

The task-dependency of colour-matching is significant. Suppose that the Scene-Parsing Thesis is correct. Then one would expect that if one had to perform some task involving a seen material object, one would consult one’s store of visual data about that object – its surface reflectance, its shape, and so on. On the other hand, if one were faced with a task involving the illumination in the vicinity of the object, one would ignore object-information and consult one’s records concerning the illumination. Further, it seems that we are able to match the *total look* of different places in the visual field – we can, albeit with some difficulty, perform the visual artist’s task of muddling sources of colour together in what Arend calls “unasserted” colour. These different deployments of visual data parallel the phenomenological facts that I earlier indicated. Illumination is sensed as voluminous; the colour of objects is two-dimensionally located; hence they are sensed as independent scene-parameters. There is a parallel between these facts and the actions to which the data contribute. The phenomenology is in part an instinctive awareness of what the data allow us to do.

The Scene-Parsing Thesis implies that there is a certain separation between one’s records concerning different things. And the experimental data just cited confirms this. Now, data concerning something is just information predicated of that thing. One’s record of data concerning a material object consists of features predicated of, or attributed to, the material object; ditto data about illumination or about the light coming off the object. The work of Arend and his co-workers suggests that colour data coming from the same location is separated into separate stores or *files* (Kahneman, Treisman, and Gibbs 1992). Location does not index

HOW THINGS LOOK

these separate files; the different things *in* a location do. This relationship between data and index corresponds to predication.

With these points in mind, let us return to the pinkness that one senses upon viewing the white wall in the afternoon – the pinkness that accounts for matches of “unasserted” colour. To what is this attributed? In order to answer this question, one would need to figure out where in the perceiver’s cognitive records this sensum is, so to speak, “filed”. Does the sensed pinkness modify or augment the perceiver’s records concerning the *wall*, or is it used for some other purpose? If we can identify the stimuli that various tasks target, one will have a clue to the predicative structure of the sensory states that aid, or serve as input to, these tasks.

What kind of evidence could one hope to find about these issues? Earlier, I described how colour constancy is tested in animals. They are operantly conditioned to prefer dishes of one colour over those of another, for instance yellow over green. Then they are presented with yellow and green dishes in greenish illumination. If they explore the yellow dishes over the green ones, then they are seeing them as they really are with respect to colour. This is generally taken as evidence for the constancy of colour perception. It is important to see that it is also evidence for the attribution of constant colours *to the dishes*, i.e., to material objects. The bees can be assumed to be searching for sugar-water. And the visual states in question are the means by which these animals are able to use or to modify their factual records *about where to find sugar water* – namely, yellow dishes. (I am not aware of experiments that test animal-responses to unasserted colour.) The dishes were the target of the animals’ learned behaviour, not the illumination or anything else in the vicinity.

Now, we can return to the question we have been discussing: To what does the visual system attribute the pinkness that it detects in the vicinity of a white wall awash with the slightly pink light of the late afternoon? I am suggesting that one way into answering this question is to ask: what stimuli are implicated in the actions that the sensed pinkness triggers

HOW THINGS LOOK

(often through the intermediacy of stored knowledge)?¹⁵ To me, it seems unlikely that the sensed pinkness influences our propensities concerning *the wall* – at least if the wall is correctly perceived. Clearly, it does not influence the formation or application of expectations concerning this material object – this is what the constancy experiments on animals demonstrate. With the bees and the dishes, it is different. When it comes to applying their expectations about where to find sugar water, bees disregard the colour of light coming off a material object and concentrate on the colour of such objects. (And the same goes, *mutatis mutandis*, for shape.)

On the other hand, the sensed pinkness does make certain other tasks possible. It helps us if we are asked to adjust the illumination in such a way as to make the wall look as if it is being viewed at noon-time. It helps painters depict how the wall is illuminated. These tasks are directed toward the wall-as-it-is-illuminated, or perhaps the light that is reflected off it. This is, or is the result of, an interaction between the wall and the illumination in which it stands. My suggestion is that there is a difference, a phenomenological difference, between pinkness emerging from this interaction and its being attributed to the wall. The pinkness is experienced as ephemeral, not as a continuing property of the wall.¹⁶

Here then is my proposal. The question is: To what is a perceived feature attributed? And the answer: It is attributed to that perceived stimulus with regard to which expectations

¹⁵ Remember that we are exploring unlearned data-gathering as a source of input to learned inferential patterns – see section II above.

¹⁶ It is possible that there is a role here for Shoemaker's "appearance properties" – it could be that he meant to be pointing to the pinkness of the light reflected off the wall. (I am indebted to Todd Ganson for this clarification.) If so, I would still contest Shoemaker's attribution of the property *to the wall*.

HOW THINGS LOOK

and records are modified by the perceiving of that feature. If perceiving *F* modifies our records and expectations regarding *x* but not *y*, then *F* is attributed to *x* but not *y*.

One last observation in elaboration of this way of looking the matter. There is a large variety of visual objects in a scene (cf. Gibson 1986). Consider cast shadows – that is, shadows that are cast by an object.¹⁷ Shadows obscure the colour of the objects on which they fall. They appear as something *on* a background, a figure on a ground, such that you can see through them to the ground, but somewhat indistinctly. As far as vision is concerned, then, a shadow appears to be an object distinct from that on which it falls. Shadows do not modify the colour of objects on which they fall (as the *Variation Thesis* would hold). Nor would one naturally file away the visual information one gets from a shadow-on-a-surface as of a piebald surface – though the entire region can be seen as a patchwork of *unasserted* colour). Shadows obscure the colour of what they fall on and make the latter difficult to see.

This said, we should note that we do not keep records concerning shadows; likely, we are not conditioned to expect of them that they will behave in this way or that. (Do we track shadows, and if so how and under what circumstances? Is it the same or different than tracking objects?) Shadows provide us with information about sources of light and of objects between these and the objects on which they fall. When the colour of an object changes, our attention is drawn to the object itself. When a shadow falls on an object, our attention is drawn away from that object to the source of light, or rather to the object that obscures it. When, for example, a moving shadow suddenly falls on the ground beside a rodent or rabbit, it runs away

¹⁷ Cast shadows are phenomenologically different from implied shadows, which consist of the darker parts of objects that are further away from the source of illumination, and self-shadows, which envelop parts of objects that face away from the light. In my analysis of Shoemaker's shadowed table, in section V above, I was assuming an implied shadow. See Casati 2003 for an extraordinarily rich treatment of shadows and their phenomenology.

HOW THINGS LOOK

from an assumed predatory bird. And we too tend to look behind us in similar circumstances. Shadows are attributed, then, to the objects that throw them, not to the objects on which they fall.¹⁸

Similarly, coloured light. My desk lamp has a yellow shade and it is casting a yellow patch on the pastel-coloured wall beyond it. This makes it difficult to see the colour of the wall. But it is, at the same time, visually perfectly clear to me that the colour of the light cast by the shaded lamp is not the colour of the wall – the visual system is easily able to tell the spot of yellow for what it is because of the softness of its boundaries, and the way it interacts with shadows and the light coming from the window. Of course, the fact that the spot of yellow light obscures the colour of the wall can lead to errors of perception regarding the wall. But this does not mean that it modifies the colour of the wall, or makes it look as if it has a different colour. It's not that the light-spot makes the wall look a definite colour, different from the colour that the wall actually is (though this could happen in certain cases). I am talking rather of a case where one is aware of the colour of the spot, but at the same time one finds it difficult to tell what colour the wall is. Errors of perception are not what we are discussing.

The question, therefore, is this: To what is the pinkness cast by the afternoon sun attributed *when no error is made*? This is – or at least ought to be – the question that vexes all of the parties to the disagreement of section I. And by the test just suggested, it would seem that it is not attributed to the wall. For it is not implicated in any object-directed action.

VIII. The Isolation Argument

Back, finally, to the Isolation Thesis. The *Externality Thesis* was endorsed, it will be recalled, by proponents of both the *Variation* and the *Constancy* Theses. Those who endorse these Theses

¹⁸ This is true, of course, of *cast* shadows – not of the implied shadows, which are the dark sides of objects illuminated from the other side. Implied shadows inform us of the source of illumination.

HOW THINGS LOOK

assume, on the basis of introspection, that there is a way that the objective world looks, despite its not always being the way it looks. The *Isolation Thesis* attempts to put in doubt the coherence of this conception. What is it for something to look blue, Berkeley and Travis and Gupta ask, if the very colour-look that it presents could equally well have been presented by something white? These authors take the possibility of error about objective facts to invalidate the idea that visual experience has objective content. *Definition 1* and the *Isolation Thesis* thus relativize perceptual content to the contribution that perception makes to certainty. Since perceptual experience never gives us certainty concerning objective external facts, they conclude that it has no such content.

The approach adopted in the previous section is helpful in suggesting a different approach to how things look, one that I can only sketch briefly here. I have suggested that the visual system contributes to object-directed actions, especially (but not exclusively) epistemic actions. In short, visual appearance is a trigger for the activation of certain action-complexes, for example, operations that lead to the formation of certain expectations concerning things of a certain types. The claim is that vision groups external things together so that these expectation-forming operations enable an organism better to cope with its environment when they are used and applied consistently.

This suggests that the significance of perception is not so much the information that is contained in it, but the instinctive epistemic actions it triggers. Both these approaches tie perception to knowledge (by contrast, for example to the sensorimotor approach of Alva Noë). But where the traditional approach takes perception to be the repository of knowledge, and therefore rejects the possibility of its possessing content that could be mistaken, the approach that I am proposing takes the significance of perception to lie in the subpersonal inferential processes it initiates. This implies that the *Epistemic Norm of Correlation* is on the wrong track. Rather, we should have:

HOW THINGS LOOK

Epistemic Norm of Use Whereas *seeing x* is the occasion for launching record keeping operations concerning *x*, so it is stipulated that:

In order for visual state *V* to count as a seeing of *x*, *V* must affect records concerning *x* thus: (. . .).

And we might replace *Definition 1* with the following:

Definition 2 *x* looks *F* to *y* if and only if how *x* looks leads the perceiver to respond instinctively to *x* in a way that is functionally appropriate for *F* things.¹⁹

Objective sense features such as *blue* or *round* can be constructed out of the similarity relations defined by *Definition 2*, provided that we discount error and bad viewing conditions, and look only to the ideal or proper uses of perception. Definitions such as the above focus on how perception is naturally used in organisms. They are still normative, because they govern the application of the concept of *seeing* and *looking similar*.

Definition 2 is not constructed, as is *Definition 1*, with an eye to solving problems that come to philosophy from scepticism. Rather, it takes up what Gibson called an “ecological approach” to perception. This approach regards perception as a facility that orients an organism to its contingent and changing surroundings, and does so in part by constructing revisable but non-momentary records concerning those surroundings. It recognizes that organisms need specific information about different kinds of object in the world around them, both in order to exploit those surroundings for their own advantage and to avoid the dangers that may be concealed within them. The Scene-Parsing Perspective assumes that this problem is not solved merely by *gathering* information. Information that is gathered must be parsed into usable data about different objects in the scene. The thesis acknowledges that the task of

¹⁹ There are instinctive responses to learned things: a conditioned response is an example of this.

HOW THINGS LOOK

data analysis cannot be performed with certainty. However, it takes it to be the function of perceptual systems to feed into usable rules of action in such a way as to put an organism into an optimally advantageous position from the Bayesian point of view of expected utilities. What I have tried to contribute to this ecological approach in this essay is to connect it with visual phenomenology – my contention has been that the ecological approach explains and illuminates visual phenomenology without thereby sacrificing a normative view of content.

HOW THINGS LOOK

REFERENCES

- Arend, Lawrence E. Jr. *et al.* 1991. Simultaneous Color Constancy: Papers with Diverse Munsell Values *Journal of the Optical Society of America A: Optics and Image Science* 8: 661-72.
- Bradley, Peter A. and Michael Tye. 2001. Of Colors, Kestrels, Caterpillars, and Leaves. *Journal of Philosophy* 98: 469-87.
- Burge, Tyler. 1986. Individualism and Psychology. *Philosophical Review* 95: 3-45.
- Casati, Roberto. 2003. *The Shadow Club: The Greatest Mystery in the Universe – Shadows – and the Thinkers who Unlocked Their Secrets* (Tr. Abigail Asher) New York: Alfred A. Knopf.
- Clark, Austen. 2000. *A Theory of Sentience* Oxford: Clarendon Press.
- 2004. Feature-Placing and Proto-Objects. *Philosophical Psychology* 17: 443-69.
- Foster, David H. 2003. Does Color Constancy Exist? *Trends in Cognitive Science* 7: 439-43.
- Frank, H. 1923 Über die Beeinflussung von Nachbildern durch die Gestalteigenschaften der Projektionsfläche. *Psychologische Forschung* 4: 33-37.
- Gendler, Tamar Szabó and John Hawthorne, eds. 2006. *Perceptual Experience* Oxford: Clarendon Press.
- Gibson, James J. 1986. *The Ecological Approach to Visual Perception* Hilldale N. J.: Lawrence Erlbaum.
- Golz, Jürgen and Donald I. A. MacLeod. 2002. Influence of Scene Statistics on Colour Constancy. *Nature* 415: 637-40.
- Gupta, Anil. 2006. Experience and Knowledge. In Gendler and Hawthorne: 181-204.
- Hamlyn, D. W. 1977. Unconscious Inference and Judgment in Perception. In *Images, Perception, and Knowledge*, ed. J. M. Nicolas, 195-211. Dordrecht: Reidel: .

HOW THINGS LOOK

- Hardin, C. L. 1988. *Color for Philosophers: Unweaving the Rainbow* Indianapolis: Hackett.
- 2003. A Spectral Reflectance Doth Not a Color Make. *Journal of Philosophy* 100: 191-200.
- Hatfield, Gary. 1990. *The Natural and the Normative: Theories of Spatial Perception from Kant to Helmholtz*. Cambridge MA: MIT Press, Bradford Books.
- Helmholtz, Hermann von. [1921] 1977. *Epistemological Writings*, trans. Malcolm Lowe Reprint, Dordrecht, Holland: Reidel, *Boston Studies in the Philosophy of Science* vol 37. (From the Helmholtz Centenary Edition edited with comments and notes by Paul Hertz and Moritz Schlick.)
- Hoffman, Donald D. 1998. *Visual Intelligence: How We Create What We See*. New York: W. W. Norton.
- Kahneman, Daniel, Anne Treisman, and Brian J. Gibbs. 1992. The Reviewing of Object Files: Object-Specific Integration of Information. *Cognitive Psychology* 24: 175-219.
- James, William. [1890] 2007 *The Principles of Psychology*. 2 vols. Reprint, New York: Cosimo Classics.
- Jameson, Dorothea and Leo M. Hurwich. 1989. Essay Concerning Color Constancy. *Annual Review of Psychology* 40: 1-22.
- Jolicoeur, Pierre and G. Keith Humphrey. 1998. Perception of Rotated Two-dimensional and Three-dimensional Objects and Shapes. In Walsh and Kulikowski: 69-123.
- Koffka, Kurt. 1935. *The Principles of Gestalt Psychology*. London: K. Paul, Trench, Tubner.
- MacLeod, Donald I. A. 2003. The Verriest Lecture: Colour Discrimination, Colour Constancy and Natural Scene Statistics. In *Normal and Defective Colour Vision*, ed J. D. Mollon, J. Pokorny, and K. Knoblauch, 188-217. Oxford: Oxford University Press.

HOW THINGS LOOK

- Martin, M. G. F. 2005. Perception. In *The Oxford Handbook of Contemporary Philosophy*, ed Frank Jackson and Michael Smith, 701-738. Oxford: Oxford University Press.
- Mollon, J. D. 2003. The Origins of Modern Color Science. In *The Science of Color*, ed. S. K. Shevell, 1-40. Amsterdam: Optical Society of America, Elsevier.
- Neumeyer, Christa. 1998. Comparative Aspects of Colour Constancy. In Walsh and Kulikowski: 323-51.
- Noë, Alva. 2004. *Action in Perception* Cambridge MA: MIT Press.
- Schellenberg, Susanna. 2008. The Situation-Dependency of Perception. *Journal of Philosophy* 105: 55-84.
- Shoemaker, Sydney 1994. Phenomenal Character. *Noûs* 28: 21-39.
- (2006) On the Ways Things Appear. In Gendler and Hawthorne: 461-80.
- Travis, Charles. 2004. The Silence of the Senses. *Mind* 113: 57-94.
- Turner, R. Steven. 1994 *In the Eye's Mind: Vision an the Helmholtz-Hering Controversy* Princeton: Princeton University Press.
- Walsh, Vincent and Janusz Kulikowski, eds. 1998. *Perceptual Constancy: Why Things Look as They Do*. Cambridge: Cambridge University Press.
- Wertheimer, Max. [1924] 1938. Über Gestalttheorie. An address given before the Kant Society. Trans. Willis D. Ellis in *Source Book of Gestalt Psychology* ed. W. Ellis. New York: Harcourt, Brace and Co. Available at <http://gestalttheory.net/archive/wert1.html>.
- Woodworth, Robert S. and Harold Schlosberg. 1954. *Experimental Psychology* 5th edition, New York: Holt, Rinehart, and Winston.