


A puzzle about how things look

R. M. Sainsbury

I

Imagine you are looking at a uniformly colored object whose parts are differently illuminated. Perhaps it is someone's orange shirt: he is sitting by the window, so that the upper half is brilliantly lit by direct sunshine, whereas the lower half is not. It does not look as if he is wearing a bicolor shirt: on the contrary, it looks orange all over, the same orange top and bottom. On the other hand, the top and bottom look different: the top half looks significantly lighter than the bottom. On the face of it, we have an inconsistency in the appearances: the shirt looks uniformly orange, and yet the parts look different in color. It may be that appearances are occasionally, and under extraordinary circumstances, inconsistent: one might cite the waterfall illusion or the appearance of one of Escher's paradoxical drawings. But the present example is of the most ordinary possible kind. We hardly ever, in normal circumstances, see things which are completely evenly illuminated. It seems mad to suppose that there is hardly ever, in normal circumstances, consistency in visual appearances. This paper shows how madness can be avoided.


The relevant phenomenon is related to the fairly well understood phenomenon of "lightness constancy". Normally, as in the case of the orange shirt, the mechanism works well, ensuring that appearances are not deceived by different levels of illumination. But it can be fooled, as this example by Edward Adelson shows:



Edward H. Adelson

The original image of the illusion.

The squares marked A and B are the same shade of gray, yet they appear different.



The original image plus two stripes.

By joining the squares marked A and B with two vertical stripes of the same shade of gray, it becomes apparent that both squares are the same.

©1995, [Edward H. Adelson](http://web.mit.edu/persci/people/adelson/checkershadow_illusion.html) http://web.mit.edu/persci/people/adelson/checkershadow_illusion.html

The very mechanism which makes the orange shirt look uniform makes squares A and B look disparate, though they are the same shade of grey (check the image on the right). In this case, the visual system provides no correction, and we have an illusion: in no sense and in no way do these squares look the same grey in the figure on the left. By contrast, in the case of the orange shirt, vision appears to deliver two verdicts, both in some sense correct: that the shirt is uniform orange, and that the two regions differ in color.

The experience of the orange shirt is one we intuitively wish to count as veridical. There is no illusion, not the slightest tendency to form a false belief, and no sense of a divergence between appearance and reality. Yet on natural accounts of veridicality, we can shift the inconsistency from within the appearance into reality itself.

V1: If an object, x , looks F , this experience is veridical iff x is F .

V2: If objects x and y look different in color, this experience is veridical iff x and y are different in color.

These principles about veridicality, given our assumptions about looks stated by (1) and (2) below, transport the inconsistency from within appearance into the world of fact:

1. The shirt looks uniformly orange.
2. The upper and lower parts of the shirt look different in color.

3. The experiences which make (1) and (2) true are veridical.
4. The shirt is uniformly orange. (1, 3, V1)
5. The upper and lower parts of the shirt are different in color. (2, 3, V2)

The inconsistency between (4) and (5), a consequence of seemingly platitudinous premises, is the puzzle about how things look which this paper addresses.

It is not difficult to come up with a redescription of the situation which has no obvious gaps yet which contains no inconsistency. The shirt looks to be uniform in color, though the plainly visible different levels of illumination ensure that different shades are presented. In the end, such redescriptions hold the solution. But just as they stand, they do not tell us what is wrong with the original way of putting things: the shirt looks uniformly orange all over, yet the top part looks different in color from the bottom. In addition to providing a consistent redescription, a proper account should show how and why (if at all) the original description goes wrong.

II

Any attempt to avoid inconsistency will consider the possibility that we have ambiguity or some similar phenomenon. In particular, it might be suggested that a familiar distinction between so-called epistemic and phenomenal senses of “look” should be brought to bear on our puzzle. “Look” in its phenomenal sense is supposed to be used merely to report appearances, not things we infer from them, whereas “look” in its epistemic sense reports the result of an appearance-based inference. A standard example of the latter is

6. The neighbors look to be away

which, if true, does not involve seeing the neighbors, and so cannot be regarded as a report of their visual appearances. Rather, I see other things (the overstuffed mailbox and empty garage), and these present appearances on the basis of which I infer the absence of neighbors.

It is fairly hard to find a clear candidate for an epistemic “look” which is followed simply by an adjective (like “orange”), or adjective phrase (like “uniformly orange”). In (6), “look” is followed by a verb in the infinitive. This might prompt the thought that in the following pair

7. The shirt looks uniformly orange
8. The shirt looks to be uniformly orange

(8) represents an epistemic rather than a phenomenal “look”. This would be offered as an interpretation of (1), (2) would be interpreted phenomenally, and our contradiction in appearance would disappear. The conflict would be of a common enough kind, between how things look and how we are inclined to judge them to be.

This easy solution does not work. First, it is clear that “looks” plus infinitive can be used to make phenomenal as opposed to epistemic reports. Seeing a Müller-Lyer puzzle, it seems quite proper to say that the lines look to be different in length, even if one has not the slightest disposition to believe that they are different in length. Second, it is clear that both (7) and (8) can properly be used to make purely phenomenal reports. The aim in both cases may properly be simply to describe the visual appearances, and not to say what one has inferred. Similar points hold for associated constructions, for example “looks like” and “looks as if”. Even when an epistemic interpretation has been set aside, we can still generate an apparent inconsistency with phenomenal readings.

The different appearance of, on the one hand, a two-colored shirt and, on the other, a uniformly colored shirt differentially illuminated is fairly hard to describe in words, and considerable painterly skill is required to represent it on canvas. (Leonardo da Vinci said: “light and shade should blend without lines or borders, in the manner of smoke” (*Notebooks* §492). Adelson (1995) says that the relatively gentle luminosity gradient generated by shadows is an important cue.) The difference, however exactly it is cued, is quite apparent: it is a difference in appearances, and is clearly phenomenally available. We can normally easily tell just by looking which case we are dealing with. In my envisaged situation, the shirt looks uniformly orange, and looks to be uniformly orange, and looks as if it is uniformly orange, and looks very different

from a two-colored shirt.¹ We cannot use the distinction between epistemic and phenomenal senses of “looks” to undermine our determination to affirm (1), understood as a phenomenal report.

We would have a flat contradiction in appearances if the following was true:

9. The shirt looks (to be) non-uniform orange.

But this seems not to be so, especially if the parentheses are deleted. (9) looks made for reporting the case in which the shirt looks multi-colored, and that is not how the shirt looks. Bearing this in mind, we need to reexamine (2) itself:

The upper and lower parts of the shirt look different in color.

We need to understand how the parts of the shirt could look that way, without the shirt looking non-uniform. The explanation seems to be this: sometimes colors appear as colors of surfaces, and sometimes they do not. When they appear as colors of surfaces, the surfaces appear to be colored. But when colors appear as, for example, colored lights, the appearance is not as of a colored surface. I suggest that the explanation of the consistency in the appearances reported by (1) and (2) is that (2) need not be understood as reporting a difference of apparent surface color, as opposed to a difference in apparent color; whereas it is very difficult to understand (1) except as reporting on an apparent surface color.

The upper and lower parts of the shirt look different in color, but do not look to be of different colors. They present different colors, but do not present surfaces as differing in color. The inconsistency in appearance is revealed as illusory. But this is not the end of the story, for the proposed interpretation of (2) does not show it to be false; on the contrary, the interpretation explains what makes (2) true. So if the veridicality principles (V1) and (V2) are correct, we still have an inconsistency, one that cannot be ignored, for it is an inconsistency in fact: the inconsistency of (4) and (5).

¹ The infinitive is needed whenever something non-adjectival follows “looks”: “looks to be a car”, “looks to be dancing”. Adjectives can feature as complements to forms of the verb “to be” (“is orange”), and can also be attached directly to nouns (“orange shirt”); this may explain why they can either directly follow “look” or else be part of the infinitive construction.

III

With the distinction between surface colors and others in hand, we are going to need more than one principle of veridicality: at a minimum, one which says that surface looks are veridical just when a surface has the apparent color in question, and another which says that the appearance of a color is veridical just when that color is present, not necessarily as a surface color.

A surface color is simply the color possessed by some surface. For color-sighted people, seeing a colored object in suitable circumstances gives rise to an appearance which is not only the appearance of a surface color, but an appearance which is distinctive of surface colors. Typically, the appearance to which a surface color gives rise is manifestly different from the appearance to which a non-surface color gives rise: when we see colored light, or see the sky, or see water underwater, or water from above at night when it is illuminated from below, or a work of James Turrell, it does not look as if we are seeing a colored surface. In seeing something whose surface is red, one typically sees the color red, and also sees the surface as red (it looks red). But one may see a color without seeing anything looking like a surface of that color: looking at the sky on a clear day, it looks blue, but it does not look as if the sky, or anything else, has a blue surface. Awareness of a color as a color of a surface guarantees awareness of that color, but the converse does not hold.

When we think about the color-appearances of colored objects, we naturally think about surface color. This is what underlies (V1) (If an object, x , looks F , this experience is veridical iff x is F). Typically when we ascribe a color to an object with colored surfaces, we think of our ascription as true iff a surface has the color we ascribe. It was in the nature of our example that this did not lead to trouble with the claim that the shirt looked uniformly orange. We took the reported experience to be an experience in which surface color is manifest. Veridicality was accordingly assumed to consist in the manifest surface color being the color of the shirt's surface. For some kinds of colors with which we are presented, including those presented as non-surface colors, this account of veridicality will be inappropriate. This includes the case in which the shirt

presents different colors, though not as surface colors. These are colors which do not look to be possessed by any surface. That is how things are with the apparently lighter upper part of the shirt and the apparently darker lower part.

The point can be made in more detail by moving away from English to an extension of English which has no word “looks” but a number of word-forms on the pattern “X-looks”. Here are three of these unfamiliar words, with approximate English equivalents alongside. “F” is to be replaced by a word for a color or shade of color.

x S-looks F	x looks as if it has an F surface
x V-looks F	x looks like a volume suffused with color F
x P-looks F	x presents a color that looks F, though x neither S-looks F nor V-looks F

The table is not supposed to be complete. For example, the “V-looks” category may need subdivision, to deal with the different looks of dyed water and water illuminated by colored light; and further main categories may need to be added.

The sky V-looks blue, but does not S-look blue (no surface is visible). A ripe tomato typically S-looks red, even when nothing V-looks red. If we merely shine a magenta light on a piece of white card, the card continues to S-look white, even though it also P-looks magenta. Unless we take a lot of care, the result of our illumination is just to make the card look to be a white card with magenta light shining on it. (With care, we can make the card S-look magenta.) This example would tend to generate an apparent inconsistency in appearance when described in ordinary English, for we might incline to say that the card with the magenta light shining on it looked white (all over) and also looked magenta (all over); these colors are exclusive. The veridicality conditions for S-looking and V-looking are clearly distinct, so something (e.g. the sky) can V-look F without S-looking F (or anything else). One might at first think that P-looking cannot fail to be veridical, but there is no guarantee of this. For all I know, if one suffers from jaundice, white things present a pale yellow appearance without S-looking yellow . (This may not be how

things are in the disease, but it is not ruled out apriori.) If so, we would wish to say that the yellow P-look is non-veridical.

The envisaged X-look language enables us to describe some familiar features of appearance that are hard to describe in English. When we put on sunglasses things in a sense look darker but (as we might put it) do not look as if they are darker, a case which generates the original puzzle. The contrast is not merely that between phenomenal and epistemic looks. As I put on the sunglasses I have no tendency to judge that any object becomes darker. In addition, there is a clear sense in which the objects do not look darker. In X-language the threatened inconsistency in appearance is avoided: objects S-look the same before and after putting on the glasses, but they do not P-look the same.

P-looking is so defined as to exclude S-looking and V-looking. But is there not a more general notion, call it P*-looking, which includes the other kinds of looking, and which is the basic form of a mind's experiential confrontation with a color? We might think of it as a color being presented, but without specification of whether it is presented as a surface color, a volume color, or whatever. For every X, if something X-looks F, it P*-looks F. Then we seem to be saddled with contradiction in appearance, for it seems we must say both that the shirt P*-looks uniform in color and also P*-looks non-uniform.

I think it is wrong to say that such a general notion of looking has any part to play in describing our visual experiences, at least where colors are concerned. When the shirt looks uniform orange, it is not that a uniform orange color is presented. S-looking does not guarantee P*-looking, if "P*" connects in any natural way with the ordinary notion of presentation. What may have been overlooked is that there is a purely phenomenal notion of S-looking which does not always amount to the presentation of a color. It is not that we infer that the shirt is uniform orange: that is how it looks. Yet there is no presentation of a uniform orange color which extends across the shirt. That is just how things are with how things look; and we should welcome it, for this is what ensures that there is not really a contradiction in appearance.

The veridicality conditions for the different X-look words are obvious:

x veridically S-looks F	x has an F surface
x veridically V-looks F	x is a volume suffused with color F
x veridically P-looks F	x presents color F

It is also obvious how to extend these to comparative cases.²

The orange shirt veridically S-looks a single shade of orange all over, and veridically P-looks different shades. So the shirt has a surface which is a single shade all over, but presents different shades. What else could one expect from different levels of illumination?

IV

One might attempt to resolve the puzzle in other ways. I will mention two such alternatives.

(A) “Darker than” is unspecific or ambiguous between darker in color than, call this *darker_c*, and less brightly illuminated than, call this *darker_i*. The lower part of the shirt looks darker_i than the upper part, but does not look darker_c. So it does not look darker in color than the upper part, and so does not look different in color from the upper part; only differently illuminated. One problem with the proposal arises from the close semantic connection between the comparative and non-comparative forms: both or neither should display ambiguity or underspecification. Hence the present proposal entails that “dark” in, for example, “dark orange”, can properly interpreted as meaning “less brightly illuminated orange”. This result is incorrect: dark orange is a color, not some kind of pairing of a color and an illumination level. The proposal is also lacking in generality. Lacking the resources of the X-looks language, we can create similar puzzles that do not depend on “darker” or “dark”. The swimming pool (or the water in it) illuminated by a pink underwater light at night veridically looks pink, but it is not pink; hence we must find an alternative account of veridical looking in this case. Likewise the sky veridically looks blue, even

² There are other normative dimensions of assessment of appearances, for example, something may look just as it should (to a creature with such-and-such a perceptual system in such-and-such circumstances).

though it has no blue surface (and so counts as “not blue”, as this phrase is most often used), and an apple may veridically look red even though it has no volume suffused with red.

(B) Color words are semantically underspecific or ambiguous or polysemous: “x is orange” may be true just if x has an orange surface, or just if x is a volume suffused with orange, or just if x is an orange light. We don’t need different ways of looking, only different ways in which colors can find a place in the world. For example, instead of “x S-looks F” we should say “x looks F_S”, where “F” is to be replaced by a color adjective, and the subscript shows that it is to be taken as ascribing a surface color. Instead of “x V-looks F” we should say “x looks F_V”; and so on. There is one way of looking, but many ways color can relate to objects.

The idea has some appeal. For one thing there are certainly many ways for colors to exist, not just as surface colors. And there is independent evidence for the relevant behavior of color adjectives.³ But the proposal seems not to do justice to all the cases, including the original one: “orange” is used just once in (1) and (2), so supposing the word to be polysemous does not immediately explain the phenomenon. Another conspicuously hard case for this approach is the change in look produced by putting on sunglasses. We can describe the effect without using any specific color words (“Everything looks darker, though nothing looks to have changed color”), so polysemous color words don’t seem to be able to begin to address the case.

V

When we speak the X-looks language, there is no inconsistency, within or outside of appearances, as we saw in §III above. Nothing follows about how things are when we speak English. A simple conjecture extends the happy result for the X-look language to English: English “looks” is ambiguous between, or unspecific between, or polysemous between, the meanings of the various X-looks expressions. The upshot is (i) that we have inconsistency in appearance only when there is inconsistency in the content of “looks” sentences which have been rendered unambiguous or specific in the same way and (ii) that there is no single veridicality

³ There are well-known cases in which surface and interior colors are contrasted, as in Travis (1997) and Bezuidenhout (2002).

condition for English “looks” sentences. In the original account of how the shirt looked, the two occurrences of “looks” (in (1) and (2)) need to have different disambiguations or specifications if they are both to be true; each corresponds to a different veridicality condition. As a result, there is no inconsistency in appearance, and, afortiori, no inconsistency in fact. A further conclusion is that all arguments (and there are many) which trade on a single veridicality condition for “looks” need to be reconsidered.⁴

REFERENCES

Adelson, Edward H (1995). http://web.mit.edu/persci/people/adelson/checkershadow_illusion.html

Bezuidenhout, A. (2002). "Truth conditional pragmatics." *Philosophical Perspectives* 16: 105–34.

Travis, C. (1997). “Pragmatics.” In *A Companion to the Philosophy of Language*, B. Hale and C. Wright (eds), Oxford, Blackwell Publishers: 87–107.

da Vinci, Leonardo, *Notebooks*. Various editions and translations, for example: <http://www.gutenberg.org/etext/5000>

⁴ Many thanks to Mike Martin and Michael Tye for valuable comments.