Color Representations as Hash Values

JUSTIN C. FISHER

University of Arizona

The goal of this paper is to answer the following question: When we have mental states that represent certain things as being *colored*, what properties are our mental states representing these things as having?

I first state three presumptions about the notion of *representation* presupposed in this question. I then present a simple overview of potential answers to this question. In that presentation, several puzzles arise that any successful theory of color must solve. With these puzzles in mind, I present the position that I favor. I argue that color representation systems work upon the same basic principles as *hashing schemes* employed by computer scientists, and I explain how this observation enables us to answer the question with which we began and to solve the puzzles that face other approaches.

1. Three Presumptions about Content.

First, a general methodological assumption:

> (P1) A theory of representational content provides an explanatory abstraction that helps us to account for how it is that certain creatures get around in the world as well (and as poorly) as they do.

In particular, a theory of content should offer a certain sort of explanation:

> (P2) The fact that things in the world are *as represented* in normal success cases helps to explain how it is that successes came about in these cases; the fact that things are *not* as represented in many cases of failure helps to explain why those were cases of failure.

For example, the fact that only ripe bananas are yellow helps to explain the normal success of our choices of which bananas to eat. Or, if you accidentally took someone else's coat, this might be explained by the fact that fluorescent lighting caused you to mistake the coat's color.

A corollary follows from these presumptions:

> (P3) In most normal, successful uses of representations, the representations are at least very nearly accurate. Our representation-producing mechanisms needn't be perfect, but they do need to get things right often enough to be worth using.

These presumptions fit nicely with a *teleo-functionalist* theory of content (especially, Millikan 1984). They are also generally quite plausible. We commonly expect that normal successful behaviors will be explainable in part by the *accurate* representations employed by the agent(s) in question, and that *inaccurate* representations will very often lead to failure. To enable this crucial explanatory distinction, *any* theory of content must satisfy the presumptions above.

I don't wish to deny that there might be *other* notions of content that don't meet these presumptions.[1] I'm open, at least in principle, to *pluralism* about content – there may be several quite different sorts of content that might correctly and usefully be attributed to a given psychological state. Even if some possible notions of content don't meet the presumptions above, it's still an interesting and important question what sort of content would be attributed to our color representations by any theory that *does* meet the above presumptions. And that is the question I will pursue here.

2. Theories of Color.

We may divide theories of color into three categories based on their answers to two questions:

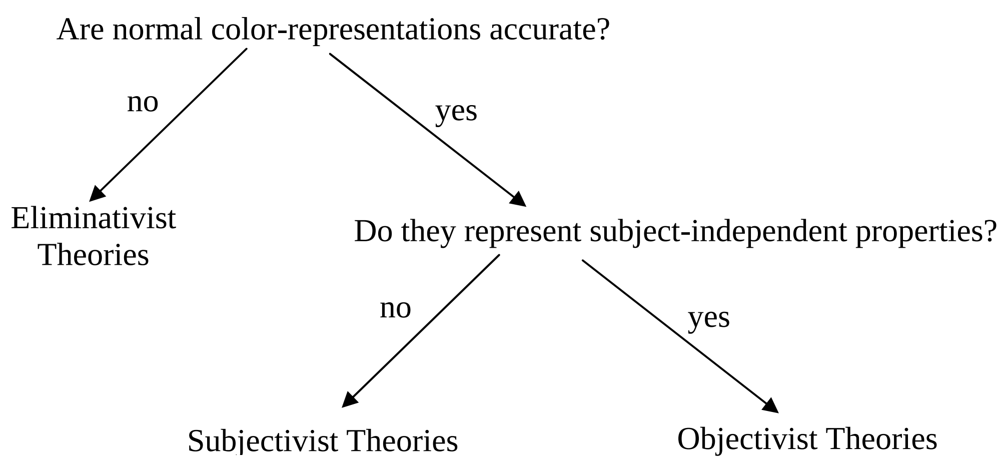

[1] Unlike many proposed understandings of content, my presumptions have not included any direct content-determining role for our common-sense conceptions of color-properties. A limited and indirect role for these conceptions is provided by the facts (1) that our common-sense beliefs about color are linked to the ways in which we actually use color-representations, and (2) that how we actually use color-representations plays an important role in the above presumptions. Insofar as common-sense belief about colors departs from what underlies *actual successful usage* of color representations, this will confer both (limited) advantages and (severe) disadvantages upon a theory of content that gives common-sense belief a direct content-determining role. Such a theory might have advantages like being well-suited to the purpose of making sense of poetical depictions of color. However, when our central goal is coming up with a scientific account of how we humans get around so well in our environment, the most centrally important facts are the ones about *how we actually do it*, and insofar as common-sense beliefs get this wrong, such beliefs won't be *directly* relevant to our project.

3. Eliminativist Theories.

Eliminativist theories hold that many of our normal color-representations are highly inaccurate.[2] Above, we presumed that a useful theory of content must explain normal successful uses of color-representations by noting their (near) accuracy, thereby enabling us to contrast these cases against the failures that come about due to inaccuracy. To get this theoretically crucial contrast, we can't employ a theory, like eliminativism, which indiscriminately holds that *all* color-representations are highly inaccurate.[3, 4]

4. Objectivist Theories.

Objectivist theories hold that our normal color-representations *accurately* depict objects as possessing certain viewer-independent properties. One challenge for objectivists is to specify precisely which properties our color-representations attribute to objects.[5]

One promising answer[6] is that the depicted properties involve surface spectral reflectances (SSR's). A surface's SSR specifies, for each wavelength λ of visible light, what fraction of light of that wavelength that hits the surface will be reflected.

> **The puzzle of metamers:** for any given color-representation, there are numerous different SSR's that will produce that color-representation in a normal human

---

[2] Eliminitavist theories may be presented as 'error' theories, 'illusion' theories, 'projectivist' theories (Boghossian & Velleman 1989, McGilvray 1994), 'Edenic' theories (Chalmers, forthcoming), or even theories that hold that our color representations do a remarkably poor job at picking out the simple physical properties (like edibility) that they are supposed to represent. Typically, eliminativists draw upon what they think is the common-sense conception of colors as simple, viewer-independent surface properties whose nature is fully revealed in visual experience. Eliminativists then typically hold that our color representations attribute such properties to objects. But, since scientific research has shown us that objects don't actually have such properties, eliminativists conclude that our color-representations are wrong – everyday objects aren't really colored after all.

[3] This is not to say that we could *never* arrive at an eliminativist conclusion about anything. E.g., it might turn out that some class of attempted representations enjoyed their limited success only by chance, and not by tracking important features in the world. (Astrology seems a likely example.) In such a case, we may safely be eliminativists about the alleged objects of these representations. In other cases, we may determine that some representations achieved their successes by (well enough) tracking important features of the world, and yet that we should abandon this representation scheme in favor of one which is much more useful. (Phlogiston theory seems a likely example.) There is a sense in which this conclusion is also eliminativist. Given our current evidence, it seems very unlikely that our color representations succeed only by chance, and it also seems unlikely that we will come up with some better candidate that could take the place of color-representations in human cognition. Hence, eliminativism about color is not an option.

[4] Of course, any true theory of color must be an 'error theory' in one way or another. Either (1) many theorists are in error about what our common sense conception of color is, or (2) this common-sense conception is in error about the nature of color-properties, or else (3) we are in error for thinking that color-properties are actually instantiated. I wish to remain skeptically agnostic regarding questions about our common-sense conception of color – I doubt that common sense really stakes clear, univocal claims regarding these matters, and, regardless, I don't think it really matters whether or not it does. Hence, I am agnostic between options (1) and (2), but I think we may rule out option (3) on the grounds cited above.

[5] Maund (2002) helpfully divides objectivist views into those which take colors to be *sui generis* and irreducible (P.M.S. Hacker, J. Campbell), those that take them to be (reducible to) microstructural properties (F. Jackson, T.Reid), and those that take colors to be dispositional light-related properties (D.M. Armstrong, J. Westphal, D.R. Hilbert).

[6] Hilbert (1987).

viewer. E.g., in normal conditions, the following two SSR's might engender the same color-representation in you. Which of these SSR's does that color-representation represent?

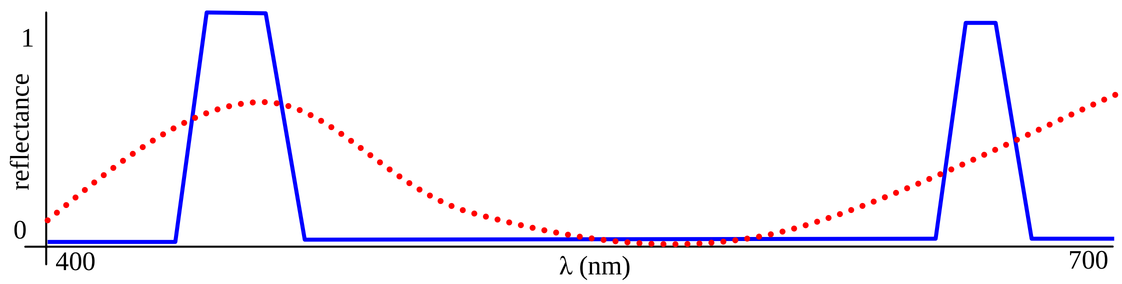

One solution would be to pick one SSR (e.g., one involving a normal distribution around a single wavelength) and hold that the color-representation represents this SSR. This solution has the problem of seeming quite arbitrary. Furthermore, it yields the result that most of our color-representations are quite inaccurate (e.g., they depict all surfaces as having a bell-shaped SSR, when in fact many surfaces have much more jagged SSR's), which violates our presumptions above.[7]

An alternative would be to hold that each color-representation represents not just one SSR, but instead a big disjunction of SSR's. (E.g., my red-representation represents the object as having the smooth SSR of a rose *or* the jagged SSR of a ruby *or* …) But which disjunctions?

> **The Puzzle of Differences.** Different color perceptual systems categorize surfaces in different ways. These differences are apparent across species, across subjects within our species,[8] and within human subjects over time.[9] Which way of categorizing surfaces – i.e., which way of forming disjunctions of SSR's – is the one that our color-representations are supposed to employ?

The most natural solutions involve the idea that the disjunctions are determined by what would be categorized together by some 'standard' visual system in normal conditions. People may disagree about whether this 'standard' system is the one possessed by an 'average' person, by an 'evolutionarily normal' person, or by the viewer herself at some appropriate stage in her life. It is puzzling which of these options we should choose.

[7] A potential response would be to say that many normal successful representations are only *somewhat* inaccurate – they attribute a curve that fits the actual SSR somewhat well, but not perfectly. This response allows that we may still explain successes in terms of the *near* accuracy of color-representations, even if these color-representations are somewhat inaccurate. Still, it is strange to say that a given color-representation of a rose (which has a smooth SSR) is more accurate than an indistinguishable color-representation of a ruby (which has a jagged SSR). It seems that both of these color-representations are doing their jobs perfectly well, and that it is completely arbitrary, unmotivated, and superfluous to say that one is representing more accurately than the other.

[8] See Block (1999).

[9] For example, with age the cornea gradually yellows, (an effect called "photoxic lens brunescence"), an effect that may be reversed by cataract surgery (Lindsey and Brown 2002; Pollock & Oved forthcoming).

Furthermore, our need to choose between these pulls us in the direction of a *subjectivist* view – one that holds that represented color-properties are somehow viewer-dependent. Let us move on to consider subjectivist views, though in the course of considering them, we will see further challenges that also face the objectivist.

5. Subjectivist Theories.

Subjectivist views hold that most of our normal color-representations accurately represent objects as bearing a particular relation to a certain sort of perceptual system.[10] The largest challenges are in specifying exactly what sort of relation is being represented and exactly what sort of perceptual system that relation is supposed to involve.

The most common subjectivist views are *dispositionalist* views holding that color-representations represent objects as being disposed to produce a certain response in a certain sort of viewer in a certain sort of circumstances. Again, there are puzzles about who the right sort of viewer is and what the right sort of circumstances would be. There are also puzzles involving what sort of *response* the relevant dispositions are supposed to produce.

> **The Dispositionalist-Representationalist Circle.** It is attractive to think that the relevant response that colors are dispositions to produce is a response that essentially involves a representation of a color – e.g., an *idea* of that color, a *belief* about it, or a *phenomenal experience* of it. But, combining this with dispositionalism, we get the result that a color experience just is a representation of the disposition to bring about a representation of the disposition to bring about a representation of the disposition to…[11]

This circularity is at least *suspicious* if not outright *vicious*, and it would be surprising if the best way to account for how creatures get around in the world requires appeal to such a suspicious circularity.[12]

One might avoid this suspicious circularity by holding that the relevant response is the production of a particular phenomenal feel – e.g., a red *quale* – where this feel may itself be understood independently of its representational content. This avoids the circularity, but at the cost of abandoning the plausible 'representationalist' view that phenomenal feels should themselves be understood in terms of what they represent.

---

[10] Versions of subjectivism were defended by Descartes and Locke. For more recent defenses see Evans (1980), McGinn (1983), Dummett (***), McDowell (***), Johnston (1992), Thompson (1995), Lewis (1997), Tye (2000), Shoemaker (2001), McLaughlin (2001), and Cohen (2001).

[11] Shoemaker (2001) considers a very similar circle.

[12] Furthermore, if one were to embrace this suspicious circularity, one would be left with several related questions. What sets color-representations apart from other sorts of representations? E.g., if odors are also just dispositions to generate representations of themselves, why aren't odors and colors more similar? Also, how might an accurate representation of such curly-Q dispositions help to explain our successful use of color-representations? I.e., what explanatory mileage can we hope to get out of attributing these rather strange representational contents? I think the phenomenal-dispositionalist may be able to give reasonably good answers to these questions (along lines similar to the view I will propose below), but still it is good to note that these are hard questions that would need to be answered.

Alternatively, one might avoid the circularity by insisting the relevant dispositions are dispositions to produce a certain *syntactic* sort of event – e.g., an event that plays a certain role in the internal economy of a given cognitive system. This sort of view is advocated by Sydney Shoemaker[13] and, in quite a different way, in my argument below.

One further puzzle faces both subjectivists and objectivists:

> **The Puzzle of Context Dependence.** Two surfaces with the same SSR may appear to be of different colors if surrounded by different backgrounds, or if viewed in different contexts. (See Figure 1.) Which, if any, of these color-representations is accurate?

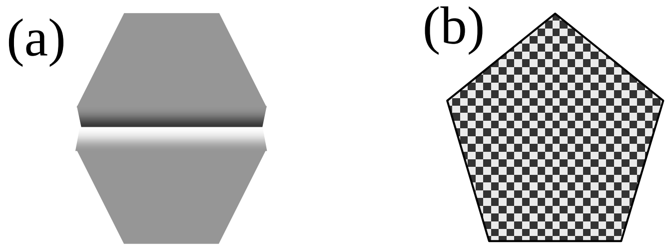

Figure 1. Your visual system probably represents the top face of the solid in (a) as darker than the bottom face of that solid, even though they have the same SSR, as will be clear if you block the central region with your finger. (This image is inspired by Beau Lotto, www.lottolab.org.) From a distance, your visual system probably represents the pentagon in (b) as gray; closer up, each region in the pentagon would be represented as either black or white; closer still, if you are viewing this image on a color monitor, each "white square" would instead be represented as a patchwork of red, green, and blue lights.

So far as I know, no one has popularized a non-arbitrary solution to this puzzle, but I will suggest one below.

6. Hashing Schemes Introduced.

I will now introduce a useful idea from computer science. Below, I will explain how attention to this idea might help us understand colors and other secondary qualities.

A *hashing scheme* is a way of organizing information so that it will be quickly retrievable. Suppose we have ten drawers in a big filing cabinet in which we wish to distribute a separate file for each of a hundred different items; and suppose that each time we'll want to store or retrieve the file for a given item, we will have available some canonical information about that item. One solution would be to employ some quick-

[13] At great risk of confusion, Shoemaker (2001) terms the physical/syntactical responses "qualia" and distinguishes these from the "phenomenal character" of our color experiences. He holds that the "phenomenal character" of a red-experience is its property of representing (normally correctly) an object as having manifested a disposition to produce certain sorts of physical changes (certain changes in "qualia") in the subject's brain.

and-dirty 'hash-function' which associates each possible set of canonical information with one of ten 'hash-values'.[14] We can then label each drawer with a hash-value, and store the file for each item in whichever drawer matches that item's hash-value. This way, we'll always have a quick way to determine which drawer to look in to find the file for a given item.[15] (See figure 2.)

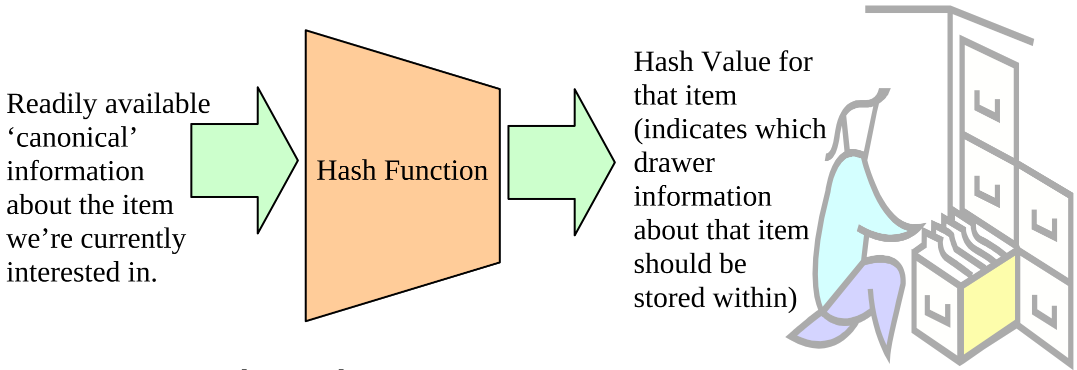

Figure 2: A Hashing Scheme.

Such a hashing scheme can work very well even if the hash-function is quite arbitrary. So long as the hash-function consistently assigns the same hash-value to a given item, it doesn't really matter *how* it goes about assigning that value. In particular, it doesn't matter whether the categorization of items by hash-values matches up with any natural, interesting way that some scientifically-minded observer would choose for categorizing the items in question. It wouldn't matter, for example, if a hash-function stored information about such disparate objects as ripe bananas, canaries, school buses, and the sun under the same hash-value – so long as these items are always given that hash-value (and most other items are given other hash-values), the hashing scheme will successfully reduce the number of files that need to be searched in order to retrieve information about these items. Part of the beauty of a hashing scheme is that it allows a designer to use pretty much *any* fast hashing function,[16] rather than requiring her to have an independently interesting categorization of the items that her system will encounter.

Of course, an ambitious designer might hope to build into her system the capacity eventually to *learn* predictively useful ways of categorizing the items it has encountered. For this purpose, it would be nice to have a hash-function assign hash-values that aren't *completely* arbitrary, but instead are at least somewhat predictively useful in the environment in question. E.g., it would be nice to have a hash-function assign the same hash-value to *all* canaries and to *all and only* the ripe bananas. But, even in these cases, the general success of this scheme doesn't require that ripe bananas, canaries, etc, receive

---

[14] Where does the name "hash function" come from? The verb "hash" means to chop. Pictorially, a hash function quite arbitrarily 'chops up' the space of possible items into categories.
[15] Of course, we will still have to work out some way of searching among the various files *within* that cabinet, but since there will be relatively few files in any given cabinet, this task will be much easier than the hard task we were originally faced with.
[16] There are some limitations. E.g., it will generally be more efficient to use a function that will assign a similar number of items to each cabinet. Also, as mentioned immediately below, it may be preferable to use a hashing function that reflects some statistically relevant features in the environment.

their hash-value *for the same reason*, nor that they have any independently interesting feature in common. Such a scheme may be predictively useful just by reflecting a number of quite unrelated statistical regularities in the world, even while, overall, the scheme is quite arbitrary.[17]

7. Representation in Hashing Schemes.

One final question, before returning to the case of color: *When a hash-function assigns a hash-value to an item, what properties does it thereby represent that item as having*?

One (bad) way of answering this question would be to look at what sorts of 'canonical' features would lead the hash-function to yield that output. (This would be the simplest information- or indicator-based approach to naturalizing representation.) This approach would lead us to conclude that the hash-function's output represents an item as having some messy disjunctive property, namely the property of having *some cluster or other* of features which would be sufficient to get the hash-function to yield that output.

The most significant problems with this simple approach are that it attributes representational contents which (in schemes involving quite arbitrary hash-functions) are quite disjunctive and messy, and which (regardless of the hash-functions' arbitrariness) are not *directly* explanatorily relevant to the success of the hashing scheme. As noted above, what is most directly relevant to explaining the success of a hashing scheme is just the fact that the hash-function consistently gives the same hash-value to any given item, and not the sundry details involving how exactly it goes about doing this.

Hence, when a hash-function assigns a hash-value to an item, it is plausible to construe this as representing *just* the fact that that item would normally be assigned that hash-value and *not* the messy disjunction of all the clusters of sundry details that might normally engender that hash-value. The hash-value is best understood as the hash-function's way of announcing "Hey, that item is such that it would normally receive this hash-value", and not (directly) as a way of announcing that that item possesses some messy disjunctive property.[18]

---

[17] From an evolutionary design perspective, it is not surprising that mother nature would choose a hashing scheme that does reflect useful statistical regularities. For building a hashing scheme involves the heavy cost of building and maintaining the sensory apparatus to detect the 'canonical' features that will be used for hashing. This cost will be much easier to bear if that sensory apparatus is immediately useful for other purposes as well – i.e., if it picks up on statistically relevant cues in the species' environment, e.g., cues that distinguish between nutritious berries and poisonous berries (Mollon 1989).

[18] What about cases where hash-values are used not just for filing data, but also as ways of predicting features of the world? E.g., what about the case in which a system learns that all the banana-shaped things that receive the same hash-value as canaries are good to eat, and therefore predicts that this ripe banana will be good to eat? Even in this case, it isn't (directly) explanatorily relevant that ripe bananas typically have some messy disjunctive property. Instead, what is explanatorily relevant is that they normally receive a certain hash-value, one which the system has learned to be correlated (within the domain of banana-shaped things) with being good to eat. Hence, even in this case, the most appropriate explanation (at the level of description appropriate for a theory of content) of the system's success attributes to the hash-function's output *just* the content that this item would normally receive this hash-value, and *not* that the item has some messy disjunctive property.

8. Color-Representations as Hash-Values.

My positive proposal is that our theories should account for color-representations in much the same way that they should account for hash-values.

Like hash-values, color-representations are used, in the first instance, to facilitate recognition of various items (and kinds) and to enable storage and recall of information about them. The success of this scheme depends, largely, upon the fact that any relevant[19] item will typically engender the same (or close enough to same) color-representations whenever it is encountered in normal circumstances. The success of this scheme does not require that a categorization of items by the color-representations they engender should match up to any independently interesting categorization of items in the world.

Like hash-values, color representations may also be used in a second, predictive, way. One may learn that all canaries and ripe bananas engender *yellow*-representations (in good viewing circumstances), and use these facts in successful bird-watching or fruit-eating. Such successes depend upon the facts, respectively, that all canaries engender (close enough to) the same yellow-representation and that, among bananas, all and only the ripe ones engender (close enough to) the same yellow-representation.

Of course, our use of color-representations differs in many ways from the simple hashing schemes employed by computer programmers. Human memories are stored, linked, and recalled on the basis of many more factors than just color-coding. I think there is an important (but difficult to spell out) sense in which human memory cross-references information on the combined basis of many more hash-values than just colors.[20] Despite these complexities, it's clear that the general principles that explain how hash-values work in simple hashing schemes also apply to the case of color-representations in human cognition.

---

[19] This proviso is very important. Color hash-values are very useful in storing memories of items or kinds that tend to have the same color whenever they are encountered. They are less useful in storing memories about items (like Cleo the chameleon) or about kinds (like *cat*) which display different colors on different encounters. In these cases, it is important that we also use other hash-values and other ways of organizing our memories. Still, the important point is that color-representations make their living, in large part, by helping us to recognize items and kinds that *do* tend to display the same colors in different encounters.

As a useful thought experiment, imagine a world in which all things continuously change their SSR's. In this world, our manner of color vision would not evolve. But imagine this world in such a way that at one moment, everything is colored as it is in our world, but everything's SSR's continually rotate at the same rate around the visible spectrum. I.e., if we visited this world, a rose would turn from red to orange, to green, to blue, and back to red. In this world, there would be evolutionary rewards for a perceptual system that effectively subtracts off the rotational component, and assigns hash-values to items that track their SSR's at some canonical point in the rotation. To such a perceptual system, the rose would retain a constant 'color' throughout the rotation. What we perceive as color change would seem to it to be a paradigm case of 'color constancy'.

[20] And perhaps partly upon some further bases not understandable on the model of hashing.

This enables us to solve the questions and puzzles raised above.

What properties do color-representations attribute to objects? A color-representation attributes to an object the disposition to engender that (syntactic) type of representation in a normal sort of way in normal encounters with that item.

In a case of metamers which representation is accurate? Both representations accurately represent their respective surfaces as being such as to normally produce a particular (syntactic) type of color-representation.

Which sort of perceptual system is the 'standard' one that determines which surfaces with different SSR's can be accurately depicted by the same color-representation? The most centrally important answer is that what is relevant is that *the cognitive system in question* would assign the same hash-value to the represented item on other occasions. In certain circumstances, this answer may be plied apart into strands that come into tension with one another. E.g., when a yellowed cornea is removed in cataract surgery, the subsequently-formed color-representations correctly match how things will look again in the near future, but they fail to match how things looked in the recent past.[21] Hence, there is a sense in which these post-operative color representations are accurate, and another sense in which they are inaccurate.[22]

How can we evade the suspicious circularity without giving up representationalism? The represented dispositions are dispositions to engender certain syntactic changes in a cognitive system, and a specification of these syntactic changes does not require any reference back to the dispositions.

What explanatory mileage do we get out of attributing dispositional properties as the content of color-representations? The fact that an item is disposed consistently to produce a certain hash-value is highly explanatorily relevant to the success of a hashing scheme, and color perception operates on the same basic principles as hashing schemes.

In which context can I view an item to accurately see its color? So long as your visual system is working normally, you'll see a thing's true colors if you view it in the same context as your normal encounters with that thing.[23] This is, after all, what is demanded by the general principles that explain the normal successes of your color-representational

[21] Lindsey & Brown (2002), Pollock & Oved (forthcoming).

[22] Given our many *social* uses of color-talk, the success of some color-representations requires also that one assign hash-values (color-labels) in a way that corresponds appropriately to the way in which such labels have been and will continue to be assigned by other members of one's speech community. In strange cases, these requirements might be plied apart from one another and/or from the factors required for the hashing scheme's *intra*-personal success.

[23] As in note 22, accounting for the social uses of color-talk introduces a sense in which one also represents things as being such as to produce corresponding hash-values *in the members of one's speech community*. If you normally see your mate's hair in the morning sun, while his co-workers normally see it in fluorescent light, then there is one sense in which all of you see it correctly (i.e., as you each are likely to see it on other encounters with him), and another sense in which the content of your representations is indeterminate (because it is slightly indeterminate which conditions are 'normal' for members of your speech community to see his hair).

system. Hence, if you'll normally see the two faces of the solid depicted in Figure (1a) simultaneously, then the top face really is darker than the bottom face, and if you'll normally see the pentagon in (1b) from a normal reading distance, then that pentagon is really black and white, rather than gray or multi-colored. Of course, for many items, there is no viewing context in which you would normally see them. There is no determinate fact about exactly which 'true color' these items have; though it probably is the case that the range of indeterminacy is quite small. This does not alter the fact that your color-representations of these items attribute to them the property of normally bringing about certain responses in you, nor does it rule out the possibility that your color-representations of these items are (determinately) very-nearly-accurate, even if it is somewhat indeterminate precisely how inaccurate they are.[24]

References.

Block, N. 1999. "Sexism, Racism, Ageism and the Nature of Consciousness" in Richard Moran, Jennifer Whiting, and Alan Sidelle (eds.) *The Philosophy of Sydney Shoemaker*, *Philosophical Topics*, 26, 1 and 2.

Boghossian, P. & J.D.Velleman. 1989. "Colour as a Secondary Quality," *Mind*, Vol. 98, No. 389. pp. 81-103.

Chalmers, D. forthcoming. "Perception and the Fall from Eden." in (T. Gendler & J. Hawthorne, eds) *Perceptual Experience*. Oxford UP.

Cohen, J. 2001. "Subjectivism, Physicalism, or None of the Above . . .", *Consciousness and Cognition*, 10, pp. 94-104.

Evans, G. 1980. "Things Without the Mind", in Z. v Straaten, *Philosophical Subjects*, 10, pp. 76-116.

Hilbert, D. 1987. *Color and Color Perception*. Stanford: C.S.L.I.

Johnston, M. 1992. "How to Speak of the Colors," *Philosophical Studies*, Vol. 68, No. 3. pp. 221-63.

Lewis, D. 1997. "Naming the Colors." *The Australasian Journal of Philosophy*, Vol. 75, No. 3. pp. 325-42.

Lindsey, D. T., & Brown, A. M. 2002 "Color naming and the phototoxic effects of sunlight on the eye". *Psychological Science*, Vol. 13, No. 5. pp. 506-512.

McGilvray, J. 1994. "Constant Colors in the Head," *Synthese*, Vol. 100, No. 2. pp. 197-239.

[24] I am thankful for helpful comments from David Chalmers and an audience at the Society for Philosophy and Psychology.

McGinn, C. 1983. The Subjective View: Secondary Qualities and Indexical Thoughts. Oxford: Clarendon Press.

McLaughlin, B. 2001. “The Place of Color in Nature” in R.Mausfield and D.Heyer (eds.), *Color Percepton: From Light to Object*, New York: Oxford UP.

Millikan, R. 1984. *Language, Thought, and Other Biological Categories.* Cambridge, MA: MIT Press.

Mollon, J. D. 1989. "'Tho' She Kneel'd in That Place Where They Grew . . .' The Uses and Origins of Primate Colour Vision," *Journal of Experimental Biology*, Vol. 146. pp. 21-38.

Pollock, J. & Oved, I. (forthcoming) “Vision, Knowledge, and the Mystery Link” in *Epistemology: New Essays*, Quentin Smith (ed), Oxford UP.

Shoemaker, S. 2001. “Introspection and Phenomenomal Character.” *Philosophical Topics*. Reprinted in D. Chalmers, ed. *Philosophy of Mind: Classical and Contemporary Readings*. Oxford UP.

Thompson, E. 1995. *Color Vision*, London: Routledge.

Tye, M. 2000. *Consciousness, Color, and Content*, Cambridge, Mass.: MIT/Bradford.