

Color-Consciousness Conceptualism Defended: Pete Mandik's responses to Jacob Berger, David Pereplyotchik, Charlie Pelling, and Philippe Chuard

Before launching into the responses, I want to be clear about how totally grateful I am for the comments that these four guys provided. I thought that, across the board, they were really tough yet fair, and I had a lot of fun trying to figure out how to solve the problems raised. I also think that my thinking on these topics is now much improved due to the comments offered. Below is a snapshot of where I'm at on my current thinking. I've organized the responses as follows. First I present all of my responses to Jake and David's commentary. After that: Charlie's. And finally, Philippe's.

Responses to Berger and Pereplyotchik (2/5/2010)

1. How best to formulate the relevant version of conceptualism?

Berger and Pereplyotchik (hereafter B&P) characterize conceptualism in a way that may not be satisfactory. It seems to me that there's a danger that B&P wind up here describing a version of conceptualism that seems to be incompatible with what I would take to be a very plausible compositionality claim. The claim might be stated roughly as that, for some colors at least, a combination of concepts is brought to bear in experience. So, for example, when I see two shades of bluish-purple, one of which is both darker than the other and more blue than the other, what I bring to bear is not a single concept (TWO SHADES OF BLUISH PURPLE ONE OF WHICH IS DARKER AND BLUER THAN THE OTHER) but a combination of concepts, (say, BLUE-ISH, PURPLE, DARKER THAN, & BLUER THAN).

Anyway, as B&P point out, I'm not very explicit in the paper. However, maybe it would be better, in subsequent discussion, to work with formulations consistent with the above compositionality remarks. Perhaps something along the following lines:

- (1') For every color one can consciously experience, one possesses one or more concepts that, singly or in combination, allow one to conceive of the color.
- (2') Every time one consciously experiences a color, one deploys one or more such concepts.

Alternately, one may accommodate compositionality by allowing compositional or phrasal concepts, so, despite decomposing into component concepts, something can nonetheless count as a single concept.

Whatever way one chooses to formulate conceptualism, it should be done in such a way to be consistent with compositionality. In my subsequent remarks I will be assuming such a formulation.

2. Does the conscious experience of blue have blue1*, in addition to the concept BLUE?

My answer is mostly “yes.” The defense of this answer goes quite a bit beyond the aims of the target article, but I’m happy to here indicate, and least in a sketchy form, how I think such a defense should go. I’ll also indicate relevant publications where I spell some of this out further.

First, it will be useful here to make some clarifying remarks largely aimed at spelling out how I use certain key pieces of terminology.

At the heart of my own take on consciousness is a kind of representationalism about what we might alternately call “phenomenal character” or “the properties that differentiate conscious states with respect to what it’s like for one to have them”. What makes this a version of representationalism is the identification of phenomenal character with the representational contents of certain states. What makes this representationalism a version of conceptualism is that all of the contents are conceptual contents, the contents represented in virtue of the deployment of concepts.

It is useful to think of conscious experience as a kind of attitude. Just as the belief that the dog is blue and the desire that the dog is blue are distinct mental states despite having a common conceptual content—a content involving the common deployment of concepts adequate for the representation of the dog’s being blue—so may we regard the conscious experience that the dog is blue as a mental state that is distinct despite its common content. We distinguish these three states in virtue of what attitudes are born toward the contents. Plausibly, functional-role considerations suffice to draw the relevant attitude distinctions.

The states that B&P and I are agreeing to call “sensory impressions”, are states that do not bear conceptual content. They neither are concepts nor have concepts as constituents. Their being the causal antecedents of the deployment of concepts in certain conditions is constitutive of what makes certain attitudes count as perceptual. And, according to my Allocentric-Egocentric Interface theory of consciousness, when there are reciprocal causal interactions between the impressions and the concept-deployments, then a crucial condition is satisfied for the attitude in question to count as conscious. (see my papers “Control Consciousness” and “Phenomenal Consciousness and the Allocentric-Egocentric Interface”)

Sensory impressions are themselves the causal consequents of perceptible properties (which I take to suffice for carrying information about them) and the impressions bear resemblance relations to each other that mirror or are isomorphic to the resemblance relations between the perceptible properties. This story about information and isomorphism is part of what underlies the asterisk-based labeling system for sensory impressions. For perceptible properties blue1, blue2, etc, there are the informational, isomorphism based properties blue1*, blue2*, etc. I don’t mind calling these asterisked

properties of impressions nonconceptual content. It's a kind of informational content that is important for explaining various things such as concept learning and attitude individuation. But I'm reluctant to call it a full-blown representational content. And I certainly am not going to call it any kind of conceptual content (which is not to deny that there are sometimes concepts about it).

Based on my identification of phenomenal character with the conceptual content of conscious states, blue1*, what B&P call a "qualitative character," will not be, on my view a phenomenal character.

This might be spelled out further in terms of a description of two distinct experiential episodes. Suppose our subject S is shown a single blue1 chip at time t1 and a single blue2 chip at time t2. Suppose further that at time t1, this gives rise to an impression with nonconceptual content blue1* which gives rise to a deployment of the concept BLUE. The resultant conscious perceptual state is going to count as a *perception* of a sample of blue1 in virtue of informational linkages mediated by the impression bearing blue1*. And it's going to count as a *conscious* perception in virtue of reciprocal recurrent linkages between the impression and the deployment of BLUE. But its phenomenal character, what it's like to be in that state, is going to be the same as the state that arises at time t2 in response to blue2. Despite the presence of an impression with blue2*, the concept deployed is just BLUE. The conceptual content is the same at t1 and t2. The phenomenal character is the same at time t1 and t2. What it's like is the same at time t1 and t2.

Despite the similarities between the conscious states at t1 and t2, there is sense to be made of a claim that at t1, there is a conscious state of which an impression bearing blue1* is a part. And this is why I am happy replying to B&P that, yes, the conscious experience of blue has blue1* in addition to the concept BLUE.

Relating these remarks to questions B&P raise about the meanings of the arrows in the flow chart (temporal successions only? or causation or property inheritance besides?) I'd say that for the purposes of my paper, I intended them as simply causal/temporal. But given my larger project, which includes an account of conscious perceptions as individuated in part by impression-involving functional roles, sense can be made of a property inheritance interpretation as well.

3. Are impressions ever conscious?

I have just sketched my account of how impressions have a role to play in the individuation of conscious states. I turn now to address B&P's question of whether impressions *themselves* are ever conscious. My answer is that they are not.

I do not suppose that there's ever a conscious state unless there's something it's like for one to be in that state. I hold that there's nothing about what it's like to be in a conscious state that outstrips one's current ability to *know* what it's like and there's nothing that one

can currently know that outstrips one's currently deployable concepts. (See my papers "Transcending Zombies" and "An Epistemological Theory of Consciousness?")

Now, this much is the sort of thing that HOT-heads (fans of the Higher-Order Thought theory of consciousness) can go along with. Further, however, HOT-heads will want to say that having an appropriate higher-order thought about an impression will make the impression *itself* a conscious state. But this is where the HOT-heads and I part company.

One concern I have is that once one has accounted for what it's like in terms of the content of the conceptual state (in their case, the HOT), it seems to me preferable to identify *that* state, the conceptually-contentful state, as the conscious one. It is, after all, the state in virtue of which there is something it's like for the subject. I see no further need to posit a reality external to the state for the state to be about (See my papers "Transcending Zombies" and "An Epistemological Theory of Consciousness?").

Another concern I have is that I don't see how *anything at all* follows about the impression from the fact that there is a HOT about it. My concern here has been spelled out at length in connection with my Unicorn Argument (see my "Beware of the Unicorn"). The gist of my concern is that since there can be representations of nonexistents, representing is not a relation and thus no relational properties can be defined in terms of being represented. But the HOT proposal under consideration is that being conscious is such a representational-relational property. Additionally, the Transitivity Principle upon which HOT-heads rely (the principle that a conscious state is a state of which one is conscious) seems to me to be methodologically dispensable (as I spell out in "Beware of the Unicorn," we can make do, for methodological purposes, with something I call *deflationary transitivity*).

4. Why not a concept for every impression?

B&P ask why I concede to Evans and McDowell that we lack as many noncomparative concepts as colors that make an impression on us. Why not have each blue1, blue2, etc which triggers a blue1*, blue2*, etc have as its downstream effects the deployment of fine-grained noncomparative concepts BLUE1, BLUE2, etc.?

One response I have to this is that it strikes me as an unmotivated profligacy of posits. I would be similarly stricken if someone posited as intermediaries between impressions and conceptual deployments what for fun we can call "con-pressions." So, external world property blue1, triggers impression blue1*, which triggers con-pression blue1\$, which triggers BLUE1. I don't have an argument against the positing of con-pressions, but until someone produces an argument *for* them, I don't feel particularly motivated to worry about them.

Since con-pressions have yet to earn their keep, let's put the cute term "con-pression" to work for a slightly different purpose, namely to name the fine-grained noncomparative concepts that B&P are wondering out loud about. Whatever con-pressions are, if there

were any there'd be as many of them as impressions, yet, unlike impressions, they'd also count as concepts.

Can we motivate the positing of con-pressions?

It's clear that there are impressions and it's reasonable that there's about as many impressions as there are discriminable properties. Further, it's easy to generate plausible explanations for why there would be far fewer individual (that is, noncompositional) color words in a language than impressions and discriminable colors. The learnability and subsequent usefulness of a word requires that it be applicable on multiple occasions and by multiple speakers. (I'm thinking here of Quine's "E Pluribus Unum" section that appears early on in his *Word and Object*.) But there will be a wide variety of impressions that a determinate shade elicits across times and across speakers. So it's clear that we can expect that there will be many more impressions than single words for colors. This difference in number is one of the crucial factors underlying our calling impressions "fine-grained" and single color words "coarse-grained."

I take it to be beyond serious dispute in the present context that there are about as many noncomparative color concepts as noncomparative color words. We posit concepts in part to explain word-involving verbal behaviors.

The question to keep in mind presently is whether there are facts about our relations to colors that require us to posit *more* concepts than just the coarse-grained ones. I'm here specifically concerned with whether the Diachronic Indistinguishability Argument forces us to posit more, and so far it seems that it does not.

5. Unease about a reidentifiability criterion

Much of my discussion proceeds without saying much about criteria on having a concept. B&P suspect that I'm assuming a reidentifiability criterion and rightly raise worries about whether reidentifiability is indeed a criterion on concept possession. I want to note that I'm reluctant to lean too heavily on reidentification. It strikes me as reasonable to connect concepts to memory. I suppose that many and perhaps most of them are learned, after all. It also seems reasonable that constraints concerning memory and learning account for the coarse-grained aspect of concepts. I'm not sure how precisely to connect concepts to memory, but the reidentifiability criterion is likely not the right way to draw the connection.

6. The explanatory insufficiency of "ability to apply comparative/noncomparative concepts"

There are several key worries that B&P raise that might collectively be summed up like this: A bare appeal to what comparative and noncomparative concepts a subject can apply seems insufficient to account for the performance discrepancies between the synchronic and diachronic contexts since it is clear that a subject could apply noncomparative concepts in both contexts and comparative concepts in both contexts. We could imagine,

suppose, that in the synchronic presentation of blue1 and blue2, the subject just applies BLUE to both. And it's clear that in the diachronic context, the conceptualization of blue2 at time t2 could just as well be BLUER THAN THE FROSTING OF ANY BIRTHDAY CAKE I'VE EVER EATEN.

No doubt the worries arise because I was insufficiently explicit that what is to do the explanatory heavy lifting is *not* a mere appeal to abilities to apply concepts. The explanations must appeal to something more complex, namely, abilities to reliably apply concepts that accurately represent the stimulus. More briefly, we might simply say that the core explanans involves abilities to apply concepts *reliably*, where it's understood here that this is more than just reliably coming up with a concept. It is reliably coming up with an *accurate* concept. (Strictly speaking, of course, it is not concepts per se that are accurate, but their deployments. However, I don't suppose that sloppiness about this is too damaging in the present discussion.)

The questions B&P raise along the lines of "why is it and what mechanism would account for the alleged fact that one can only deploy noncomparative concepts in diachronic contexts and can only deploy comparative concepts in synchronic contexts?" should be replaced with questions along the lines of "why is it and what mechanism would account for the alleged fact that one can only *reliably* deploy noncomparative concepts in diachronic contexts and can only *reliably* deploy comparative concepts in synchronic contexts?"

I here try to be a bit more explicit about what the answers to such later forms of questioning should be.

Consider the following abstract characterization of a relatively simple computing mechanism, which will have four "registers", Register1 through Register4

The input register, Register1, can have at any time only one or two of the following 200 symbols: "x=1," "x=2" ... "x=100", "y=1," y=2", ... "y=100". When Register1 has two symbols in it, there is one from the x-series and one from the y-series.

Register1 outputs to Register2. If Register1 has only one symbol in it, then register 2 responds with "x is A" for "x=1" through "x=50", "y is A" for "y=1 through "y=50", "x is B" for "x=51" through "x=100", and "y is B" for "y=51" through "y=100".

If Register1 has two symbols, then Register2 will also respond with one of the following three: "x>y", "x<y", and "x=y". It should be obvious that we can contrive to make sure that the mechanism gives an *accurate* response here. So, for example, Register1's "x=40" and "y=57" will trigger Register 2's "x<y".

Register 3 is a memory buffer of the prior contents of Register2

Register 4 receives inputs from Register2 and Register3 but not from Register1. In response to tokenings of "x>y", "x<y", and "x=y" in Registers 2 or 3, Register 4 just

duplicates those tokens. In response to the pairs, (“x is A”, “y is A”) and (“x is B”, “y is B”), Register 4 responds “x=y”. For other pairs, Register 4 responds with “x>y” or “x<y”, and further, it should be obvious that we can contrive for it to do so accurately (so tokening “x>y” in response to (“x is A”, “y is B”) is *accurate*).

It should be clear that Registers 1-4 duplicate key features from the boxology of my target article: Register 1 tokens are impressions, Register 2 tokens are conscious conceptualizations, Register 3 tokens are short term memories of consciousness, and Register 4 tokens are discrimination judgments.

Consider the following core cases:

Synchronic Register1 tokenings of “x=1” and “y=2”. The Register 2 response, which is accurate, will be “x<y” and this will be propagated through to register 4.

Diachronic Register1 tokenings of “x=1” and “y=2”. The Register4 response will be “x=y”, which is false.

Diachronic Register1 tokenings of “x=23” and “y=72”. The Register 4 response will be “x<y”, which is true.

We can generalize from these core cases, and make some points about reliability. The described mechanism will reliably make correct judgments for synchronic presentations of any selected pairs, regardless of whether these result in intergroup (both in A or both in B) or intragroup comparisons. However, in diachronic presentations, the described mechanism will only reliably make correct judgments for intergroup comparisons. For diachronic intragroup comparisons, the mechanism is doomed to get it wrong.

Of course, this mechanism with the four registers is a toy and comes nowhere close to being an adequate model of what happens when one consciously perceives color. This toy leaves much out, like, for instance, what it is that would make register 2 tokenings conscious. But the point of this toy is to illustrate core points about reliability and what sort of mechanism would allow for both fine-grained and coarse-grained *synchronic* discriminations but only coarse-grained *diachronic* discriminations.

So, to go back to the specific questions B&P ask, questions along the lines of “Why can’t a subject apply a comparative concept in the diachronic case?” my responses will be along the lines of “They can, but these will just be guesses or flights of fancy with no real hope of being accurate.” So, suppose a comparative concept were deployed during the diachronic case. Which one? And, when? At time t1 the subject can make some wild guess prediction that the current color is darker than the one coming up next. But there’s no reason to suppose the existence of a reliable mechanism for deploying the *right* concept. That would be clairvoyance. At time t2 the subject can make some wild guess that the previous color was darker, but I don’t suppose there to be a memory trace of what was present at time t1 and thus the subject would be no more reliable about the past than the future.

7. Successes and failures in the New Experiment

Let us use “New Experiment” for the proposed experimental design that aims to control for context effects by showing a simultaneous presentation of the color pair in both synchronic and diachronic discrimination tasks. B&P agree with the prediction of successful diachronic discrimination in the New Experiment, but express puzzlement about what the conceptualist should say in the event of diachronic discrimination *failure*.

It will help dispel the puzzlement by first saying some more about why I anticipate diachronic success. I am supposing that in the successful conditions, the subjects are told truthfully that only two colors are being utilized in the experiment. This provides a setting in which a conceptualization expressible as “the darker of the only two colors ever used in the experiment” could be *reliably* employed. That is, it can be counted on that the subject will be accurately conceptualizing which color has been presented in which spatial location across the presentations.

We might describe these situations in which the subjects are successful at diachronic discrimination as situations in which subjects have, in addition to the presented colors, crucial auxiliary information about the experimental set up. We can devise situations in which the subjects are unsuccessful by depriving them of this auxiliary information. One way to do this would be to present color pairs from a larger selection of very close shades of blue than just two. In such an experimental variation, telling the subject that there are only two colors will be telling them something false.

Responses to Pelling 2/10/10

1. On correctly characterizing the dialectic and how much experience leaves open

I’d like to raise some concerns about whether Charlie is characterizing, in a non-question-begging way, the challenges the conceptualist is obliged to meet. One might raise these concerns in terms of the question of how best to characterize the dialectic at hand. As I put it at the outset in my paper, the conceptualism I’m concerned to defend against attack is the view “that conscious experience of color is no more fine-grained than the repertoire of non-demonstrative concepts that a perceiver is able to bring to bear in perception”(p. 1). A natural way, then, to characterize the opposing side, which I’ll just call “nonconceptualism,” is as the view that conscious experience of color *is* more fine-grained than the aforementioned conceptual repertoire.

Given this understanding of nonconceptualism, it would be outright question-begging on the part of a nonconceptualist to present arguments with premises asserting a fineness of experiential grain exceeding conceptualization, since what is under dispute is the question of how comparatively fine-grained conscious experience is.

Now, I want to be clear that I am not accusing Charlie of an outright begging of the core question. I want instead to present a note of caution against doing so, for it seems that the danger may be very near. The sense of danger becomes especially acute when we look at Charlie's characterization of the conceptualist's goal. On several occasions Charlie describes the conceptualist as seeking to "account for fineness of grain". I'm uncomfortable with such a characterization of the conceptualist's goal, for it makes it sound like there's an agreement that experience is highly fine-grained, and the conceptualist needs to account for how this can be the case. I would much rather describe the dialectic as being a dispute about how fine-grained conscious experience is, and shift to the nonconceptualist the burden of showing that conscious experience of color isn't coarse-grained.

Now, *if* I may be allowed to so characterize the dialectic, *then* I think I can make a charge of begging the question stick. And it seems to me that where Charlie most clearly begs the question is in his assertion that "our experiences take a stand as to which particular shades perceived objects have." Let's suppose for conversation's sake that an object that is blue is only one of 25 determinate shades of blue (blue₁-blue₂₅). It's consistent with the conceptualism I like to defend that on a discriminating encounter with two objects that are blue₁ and blue₁₅, respectively, a subject consciously experiences them in a coarse-grained way as *one's being a darker blue than the other*. However, that's not the only way it might turn out and still be consistent with my coarse conceptualism. Other options include (1) *one's being only slightly darker blue than the other* (where the coarse-grained concept SLIGHTLY DARKER THAN is deployed), (2) *one's being some determinate degree of darkness darker blue than the other* (where the coarse-grained concepts deployed remain open on *which* determinate degree of darkness it is), and (3) *one's being some determinate shade of blue distinct from the determinate shade of the other* (where the coarse-grained concept DETERMINATE SHADE is deployed in a manner leaving open which determinate shades are present).

The nonconceptualist needs to provide some argument that our experiences do take a stand about which determinate shades are present, and thus an argument that characterizations such as (1)-(3) are inadequate for capturing the content of color consciousness. However, it's not clear that Charlie has provided such an argument.

Perhaps a charitable reading of Charlie is as presenting a phenomenological argument, an argument that has as implicit premises propositions concerning how our experiences seem upon introspection. However, such an appeal to introspection may be easily countered by the conceptualist along the lines I sketched against the Determinateness Argument. It may seem to us that our experience is of determinate shades because we deploy, in introspection, an existentially quantified conceptualization that there are some distinct determinate shades present. It may very well be the case that it seems to us in introspection that our experience takes a stand on which determinate shades are present without it being the case that there are determinate shades that experience takes a stand on.

Compare: I can believe that there is some particular man in the next room without there being a particular man that I believe to be in the next room. I hear a solitary manly voice from the next room over. I figure that it must be some particular man (what other kind of man could it be? A non-particular man?). But for each particular man I have beliefs about, I do not have a belief that commits me to that particular man being the one making the manly racket.

2. concept-possession criteria and synchronic discrimination as evidence for color concepts

There's an unanswered question hanging in the air about what concepts are. The conceptualist is well served by being committal but not *too* committal about what the answer should be. The conceptualist needs to steer between a cheap victory won by simply defining a concept as any old thing that pops into a mind and an overly stringent criterion that not all concepts actually satisfy. I've been happy to concede to the proponents of the DIA that there are important connections between concepts and memory. But I'm reluctant to say what precisely the connections need to be. And I must admit to being increasingly reluctant to subscribe to a reidentifiability criterion such as Sean Kelly argues for.

Despite my reluctance about buying whole-hog into reidentifiability, I will admit to being guided by Wittgensteinian man-with-hand-on-head worries. Not just any relation born to a height counts as knowing which height that is. Ditto for having a concept of that height as such.

Charlie suggests that criteria for color concepts can be satisfied by synchronic discrimination. More accurately, what it is that Charlie suggests is that making a synchronic discrimination of a color from another "would constitute good evidence that one possesses a concept of that color".

I must confess that I'm hesitant to follow Charlie here. Some of what I worry about is what counts as the relevant notion of *discrimination*. A magnetized needle in a compass can discriminate magnetic north from south, and a litmus strip can discriminate an acid from a base. But I don't take at all seriously the kind of panpsychism that would allow that the compass and the paper are themselves in the business of deploying concepts.

If we turn from simple gizmos like the compass and the litmus strip to a clear case of a color-discriminating concept user, problems arise about exactly what concepts need thereby to be deployed. Suppose that an English-speaking adult manages to discriminate blue₁ from blue₁₅. She may very well do so simply by conceiving the one blue as being darker (or slightly darker, or some determinate degree of darkness darker, etc.) without thereby deploying a concept of blue₁₅ *as such*. She might say "The right chip is a darker blue than the left chip" and thereby give us evidence that she has some conceptualization or other of the color of the blue₁₅ chip. But the hypothesis that she has a concept of blue₁₅ *as such* is underdetermined by this evidence. Compare: my 7 year old daughter may very

well discriminate a Boeing 747 airplane from a 1967 Chevy Impala automobile without thereby deploying a concept of a 1967 Chevy Impala *as such*. She just says “Daddy, one is a car and the other is an airplane.”

3. Phenomenal Sorites and a Synchronic Indistinguishability Argument against conceptualism

Charlie presents an intriguing argument against conceptualism, his Synchronic Indistinguishability Argument (SIA), which is based on the postulation of what I’ll here call a synchronic phenomenal sorites series. My main reaction against the SIA is to doubt whether there are such series. (Although I’ll also suggest that the conceptualist can, at very little cost, grant that there are such series.) To make my doubts look motivated instead of ad hoc, it will help to say a bit more about phenomenal sorites series, and then to say what it might mean for one to be synchronic as opposed to diachronic.

A phenomenal sorites series of colors is a set of colors ordered in such a way that each member in a pair of adjacent colors look the same, but colors at the beginning and end clearly look distinct. One example of such a series would be 34 colors, the first and last of which look unique red and unique yellow, respectively, but each of the 34 colors looks just like its immediate neighbor. The smallest phenomenal sorites series would consist of only three colors, A, B, and C.

A *diachronic* phenomenal sorites series is one in which adjacent and nonadjacent color pairs are experienced at different times. It is clear that a diachronic phenomenal sorites series presents no real problem to my kind of conceptualism. In the example of the 3-item series, A and B look the same to conscious experience by my applying the same color concept to both. At some different time, B and C look the same by my applying a different concept than before to both. There’s no obvious problem that arises in hypothesizing B being conceptualized one way at one time and a different way at another time.

Now, if there were such a thing as a *synchronic* phenomenal sorites series it would be one in which all of the colors are experienced simultaneously and would also be simultaneously experienced as bearing their various adjacency, nonadjacency, similarity, and nonsimilarity relations to each other. I think that we can motivate some serious questions about whether there can be such a series.

Consider, first, the question of whether there could be a series with very many elements, say 34 elements. See the following figure taken from Philippe Chuard’s *Phil. Studies* paper, “Non-transitive looks & fallibilism”:



Serious questions may be raised about whether foveal resolution and the capacity of attention genuinely allow for all 34 elements and their various relevant relations to enter into conscious experience all at once. It is one thing to stick all 34 colors up in front of someone's face synchronically, but the limitations imposed by overt and covert attention may force the colors and their relations to be taken in diachronically after all. The subject may be restricted to moving a limited window of attention across the spatial array and taking in various color pairs diachronically. There may thus be no color that *simultaneously* looks just like two manifestly distinct non-adjacent colors.

The natural suggestion, of course, is for the nonconceptualist to suggest the existence of a small series, just as Charlie tries to. With only three elements, it is much more plausible that all three colors may be taken in all at once. More plausible, also, is that relevant similarities and differences are also taken in at the same time. Note, however, that for a very small series, the nonadjacent colors won't be very different. They will be nowhere near as different as unique red and unique yellow, or even as different as red and orange. It would be puzzling to say of a color that it simultaneously looked just like red and just like yellow. It's puzzling because of how different red and yellow look. But if A and C look very similar to start with, it's not obvious that it's so problematic for B to be conceived of as simultaneously looking like A and like C.

Note that in the previous paragraph I said that the 3-item series is "more plausible" to regard as synchronic. But this is not to concede that it actually is plausible. With very similar color pairs, such as the ones in the figures from my target article, it takes some non-negligible amount of time and attention to see the difference between the two. Such considerations may be recruited to help raise doubts about whether even the smallest phenomenal sorites series is small enough to be synchronic.

I end here with one more quick description of a way out for the conceptualist. Even if it is granted that such series can be plausibly argued to be synchronic, there's no obvious harm done to the conceptualist in allowing that the conceptualizations deployed in conscious experience are contradictory. Of course, I don't mind supposing that *reality* has no room for contradictions. Something cannot at one at the same time *be* just like A and not just like A. But it's much less problematic allowing that there are contradictory *representations*. There is, for example, the following sentence: "B is a color that is simultaneously just like A and not just like A." That sentence gets on just fine being contradictory. Perhaps analogous *mental* representations exist while being analogously contradictory. Of course, when the representations in question are beliefs, and the believers are rational, and the contradictions are very simple and obvious, many philosophers will want to say that there's some sort of problem here. But the conceptualism on offer is not committed to conscious experiences being beliefs.

Conscious experiences need only be similar to beliefs in the following manner: they are attitudes toward contents exhausted by deployed concepts.


Responses to Chuard 2/19/10

1. What *sort* of response do I offer against the Argument from Fineness of Grain (AFG)?

Philippe’s characterization of AFG and my response to it is fair, but I would like to offer a slight modification of how best to describe my response. My response is not *just* to deny (1) **fineness of grain**, but also to give a species of resistance against (2) **conceptual limitation**. Re (2), while I think we lack finegrained *noncompositional* color concepts for all the colors we perceptually discriminate, there need be no limitation on the compositional conceptualizations we may bring to bear on experience, and thus, sometimes, represent maximally determinate shades. But I should emphasize that I don’t think Philippe is unfairly representing the view as I present it in the target article: it’s fair to characterize the response as leaning on (1).

2. What’s the proper role of context effects in my argument?

The general point that I wanted to make is that different contexts give rise to differences in color consciousness. In the specific illustration I use, the figure with the masked and unmasked “Rubik’s cubes”, there is the same color (gray) represented differently at different times (gray with mask, yellow and blue without). Philippe correctly points out that it would be more directly relevant to discussing DIA to present a case of different colors looking the same at different times. Fair enough. I think such evidence can be provided. I don’t intend to just offer it as a natural assumption or a derivation of the previously illustrated context effects. So, consider this illustration involving colored regions arranged in three rows:



If one were to foveate at one time at the leftmost third-row color, and at another time at the rightmost third-row color, they are each consciously perceived at those times in the same way. Similar remarks apply to the leftmost and rightmost second-row colors. However, there is a clear difference between the two first-row blues. By the way, each row contains the same pair of differently shaded blues. We have here the requested illustration: different colors looking the same at different times. In each row, leftmost blue and rightmost blue are different shades that look the same at different times.

3. Further worries about context.

Philippe raises a concern about distinct colors that look the same in the same contexts. Consider, for example, a blue square on a white background at one time, and a slightly darker blue square on the same white background at another time. The problem posed is that some explanation is needed for why the different colors are represented in the same way in consciousness. But it seems to me that this is not a particularly worrying problem. There are many plausible mechanisms that may be appealed to by the conceptualist whereby the difference between blue1 and its white context is insufficiently different from the difference between blue2 and its white context to trigger different concepts.

4. Philippe's case for diachronic fineness

Philippe tries to raise a problem concerning 2 sets of 5 shades of green each where the 2 sets are diachronically distinguishable from each other and the synchronically presented shades are synchronically distinguishable. I don't see that any special challenge is posed here. The conceptualist might say that the contents are something like:

at t1, 5 shades of green each darker than the next and 5 other shades each darker than the next AND darker than the first 5.

Philippe suggests that the sorts of conceptualist maneuvers in response to such cases come at a cost of needing to embrace coarse-grained imagery and memory. However, I don't see that this is so costly. That imagery and memory are coarse grained strikes me as independently plausible. If I close my eyes and attempt to imagine the determinate shade of yellow that the grey regions of the rubik's cube illustration appeared to be, it's not obvious that there is some determinate shade of yellow that I am imagining to be on that square. It strikes me as plausible that I am doing something different, which is imagining that there is some determinate shade of yellow on the top of the square. The difference at stake here is a difference in scope. It's a difference between (a) there being an x such that x is a determinate yellow shade and I'm imagining x, and (b) my imagining that there is an x such that x is a determinate yellow shade. It's not obvious to me that (a) better captures the phenomenology than (b).

5. A couple of things Philippe has convinced me of

Philippe's (v) and (vi) under his Worry #1 strike me as largely correct. Re (v), I agree that I overstated what assumptions underlie DIA. Re (vi), everything Philippe says here seems right to me. For more of my thoughts on this, see my responses to Jake and David re the "New Experiment".

6. A bunch of greens

Philippe asks for detail concerning how one might conceptualize a presentation of 6 or 12 greens. I'm inclined to regard as satisfactory relatively coarse grained representations like:

A row of greens, each darker than the next

or

Between 4 and 16 greens in a row, each darker than its neighbor

I'd like to hear considerations in favor of regarding the experience as being more finegrained than that. See also my responses to Charlie about phenomenology and burden of proof.

7. Does conceptualism make experience *too fine grained*?

Most of the discussion so far has been whether conceptualism can be finegrained *enough*. But Philippe also raises some really interesting problems concerning whether conceptualism might introduce *too much* fineness of grain.

One way this might arise is when one color is subjected to two different comparisons to two other colors. B might be conceived as lighter than A on one occasion and darker than C on another. I take it that Philippe thinks that the conceptualist is committed to saying that B is experienced differently on the two occasions and that this is implausible.

Another way this excessive fineness might arise is if “x is darker than y” and “y is lighter than x” express distinct concepts. If they do, then there can be two different conceptualizations of synchronic presentations of A and B: *A is darker than B* on one occasion and *B is darker than A* on another. I take it that Philippe regards as implausible the suggestion that experience admits of this much fineness of grain.

In response to the ABC case, I'm not sure I see much of a problem here. First, consider an analogy concerning thought. Suppose on one occasion I think of Mark as hairy but not as hairy as Luke. Suppose on another occasion I think of Mark as hairy and much hairier than John. Am I thinking of Mark in different ways on the two different occasions? Well, I am thinking of him as hairy on both occasions, so they aren't completely different ways of thinking of Mark. But I am thinking of Mark in different ways, ways that leave open which determinate degree of hairiness Mark is. It doesn't strike me as super-implausible for analogous remarks to apply to colors A, B, and C.

In response to the AB case, my first inclination is to question whether there genuinely is a conceptual difference between x's being darker than y and y's being lighter than x. But I don't want to get hung up on the particular example. Surely there's some plausible example or other of concepts that differ despite being coextensive. Philippe's concern may be raised in a general way: is it really plausible that concepts that share extensions can nonetheless introduce different contents into conscious experience? I don't mind saying that there can be such differences. For an extended treatment of this sort of thing, see my (2006) **The Introspectability of Brain States as Such**. In: Keeley, Brian (ed.) *Paul M. Churchland: Contemporary Philosophy in Focus* Cambridge: Cambridge University Press. pp. 66-87.

8. The No Introspection Problem

I'm not entirely sure I feel the force of the problem. Perhaps what I'm hung up on centers on (3). How exactly should this be interpreted. Is the claim that one can perceive without introspecting? That seems harmless and easy to accommodate: perceiving and introspecting are two different attitudes. Perhaps instead the claim is that one *can't* introspect that the percept has conceptual content. That strikes me as overly strong. Why should I grant that? Is it supposed to be phenomenologically obvious? It isn't to me. I just glanced over at my coffee cup, and it's clear to me that I'm perceiving/conceptualizing it as being blue, shiny, reflecting two coins and a wedding ring, and having a smudge of chapstick on it. Ugh: it's dirty and needs to be washed. And this is all apparent to me upon introspecting my visual experience. Now, not everyone who has conscious percepts is necessarily able to introspect that concepts are present. That takes a bit of theory and a bit of practice.

I'm aware that my remarks here are considered controversial in many circles. I don't offer them as obviously correct. But I take it that the point here under discussion is whether they are obviously incorrect.

9. The Extra Layer Problem

I think premise (4) in Philippe's Extra Layer Problem can be shown to be false. It doesn't follow from there being introspectible differences between thoughts and experiences with common conceptual contents that the difference is due to *nonconceptual content*. However, I will grant, as we'll see in a moment, that there is a difference which is nonconceptual. Here's how I think of what's going on. Let the content in question be singular: George is blue and involve the deployment of two concepts, one for George and one for blueness. A thought and an experience, I'll suppose for conversation's sake, can have exactly the same content: I can think that George is blue and I can consciously perceptually experience that George is blue. Thought and experience, however are different attitudes, I assume. There are different attitudes born to one and the same content. Enter introspection. Introspection is a third attitude. But if I introspect a difference between thought and experience even though they the thought and the experience have common contents, there's an easy way to explain this. The introspective attitude involves the deployment of different concepts, for example, the concept of a thought and the concept of an experience. So the conceptual content of my introspecting a difference between thinking George is Blue and experiencing George as blue is due to the deployment of more than just concepts for George and blueness. So, there's a difference between thought and experience which is a difference that is nonconceptual: it's attitudinal. However, this is consistent with the content of introspection being exhaustively conceptual. The attitudinal differences are conceptualized via concepts of those attitudes.