Thoughts on Sensory Representation: A Commentary on Austen Clark's A Theory of Sentience

Joseph Levine, Ohio State University
levine.120@osu.edu

1. Clark's book is a detailed study of the nature of sensory representation. It is highly informed by empirical results in the psychology of perception, and philosophically rich and significant. I admire the book and learned a great deal from reading it. As it covers a wide range of topics, and as I have no overarching critique to present, in this commentary I will briefly address three issues that come up in the book: Clark's relational type-identity thesis for sensory qualities, his theory that sensory representations involve proto-singular terms referring to spatio-temporal regions in the subject's environment, and his interesting proposal concerning color to treat it as "difference coding". Some of my remarks will be critical, but others will just explore some of the implications of his view.

2. Clark distinguishes "phenomenal properties" from "qualitative properties", the former being appearance properties of things in the world (their colors, shapes, tastes, odors, etc.) and the latter being the properties of sensations by virtue of which they are sensations of their corresponding phenomenal properties. So when I see a red ball I am "directly" aware of the ball's redness and roundness - it appears red and round to me. This awareness of the ball's redness and roundness is accomplished, however, by my having a visual experience with certain qualitative properties; those that are of the sort one has when seeing something as red and round. It is these latter qualitative properties that are the subject of his relational type-identity thesis.

Before addressing that thesis, however, I want to quickly note and respond to another point Clark makes concerning the qualitative properties of sensory states. He

concludes from his characterization of phenomenal and qualitative properties that it's only the former to which we have anything like direct access, or acquaintance, but that the latter are more like inferred properties. In this connection he cites Dretske's notion of "displaced perception" (Dretske ??). But this claim about displaced perception doesn't follow merely from the characterization of the relation between phenomenal and qualitative properties. I may be directly aware of the qualitative properties of my visual experience, where what they consist in is how external objects look to me.

In general, transparency considerations are often misused in the following way. When asked to characterize what it's like to see the scene in front of me, I find I have to talk about how the various objects in the scene look to me, which means I talk about what properties and relations they appear to have. So, given this appeal to the properties and relations of external objects in the characterization of what it's like for me, it's supposed to follow that I am not directly aware of my experience, or the features of my experience, but only of the features of the external objects in the visual scene itself. However, this doesn't follow. What may follow from the need to refer in this way to the properties distal objects appear to have when describing what it's like for me, is that what it's like for me is essentially a matter of what properties distal objects appear to have. What it's like for me and how things look to me may amount to the same thing. But this alone doesn't entail that how things look to me - a state of me, not the distal object that looks that way - is not itself a matter of which I can be directly aware. (For a penetrating critique of the misuse of transparency arguments, see Siewert 2004)

The principal issue regarding sensory qualities I want to address is Clark's relational type-identity thesis. Clark argues that individual qualitative properties can only be defined in terms of their location in a qualitative structure, or a quality space. A quality space is a space of relations of n dimensions, where n is determined by the number of ways in which qualities of that kind can vary. The distance relations in the structure represent degrees of similarity/dissimilarity. What is particularly interesting about this proposal is the way it substitutes the relation of similarity for that of causation in the traditional functionalist theory of sensory qualities. Functionalism has it that we identify a sensory type (any mental type actually) by its causal role, which includes its relations to stimuli and behavior. Clark points out that what makes a visual sensation of orange, say, the kind that it is has nothing really to do with which stimuli elicit it. Furthermore, how it interacts with other mental states isn't really relevant to its identity either. What is crucial, however, is that orange is a combination of red and yellow. So Clark's idea is to ditch the causal element in functionalism, but retain the idea that a particular type of sensory quality is determined by its role in a relational structure; it's just that the relation at issue is similarity, not causation.

It's possible I don't understand exactly what he has in mind here, but to my mind his proposal doesn't really succeed in substituting similarity for causation as the fundamental relation in terms of which sensory qualities are to be defined. My worry is this. Suppose we agree that color qualities, say, are defined by the particular quality space definitive of color experience. To oversimplify, suppose it's a three-dimensional space, with the axes corresponding to hue, brightness, and saturation.

Now, considering color qualities as properties that can be instantiated in experience, I understand what it is for there to be a similarity space and individual color qualities to be identified with locations in that space. But this is all quite abstract. How do we imagine these sensory qualities being instantiated in experiences? Clearly, on Clark's account, for a sensory quality to be instantiated in an experience is for the brain to occupy a certain state. The question then becomes how the brain's states count as realizing the relevant quality space, so that any one state can be mapped onto a location in that space. Well, there are two possibilities. First, if brain states maintained relations of similarity that respected the topology of the color quality space, that would justify mapping individual states onto their respective locations in that space. But of course that isn't likely, or what I imagine Clark has in mind. There isn't going to be any well-defined notion of neural similarity that will do the job here.

If intrinsic similarity among neural states doesn't provide the basis for the mapping, then what does? Presumably, it's our judgments of similarity that do the job. Given two brain states corresponding to two color experiences, it's our judgment of their relative similarity that determines their location in the abstract quality space. But now that brings the functionalist's causal relation back into play. Which color quality is being instantiated in an experience is a function of its tendency to cause certain similarity judgments. Causation isn't being replaced, it's just being constrained. That is, rather than looking at a state's total causal role, this version of functionalism restricts the class of relevant causal relations to those involving similarity judgments.

Even if I'm right about the role causation must play here, this isn't an objection to

Clark's thesis so much as a claim that he's perhaps guilty of misleading advertising. I think this question of whether similarity alone can do the job or we need causation does in fact connect with other points that I take to constitute genuine objections to type-identifying sensory qualities in this relational way. So let me move on to consider the objections directly.

Clark admits, and indeed heartily endorses, the consequence of his view that qualitative properties turn out to be relational, not intrinsic. He considers two sorts of criticism to the idea that qualia are relational properties. First, quoting a passage from a paper of mine, he notes that it at least seems conceptually possible that one be capable of having sensations of red, say, without also being capable of having sensations of green. Yet, if a sensation of red is constituted by its location in a quality space, this isn't possible. Second, he addresses the standard anti-functionalist argument from the possibility of inverted qualia.

To the first objection he just insists that the fact that qualia are constituted by locations in the relevant quality spaces is not a matter of conceptual analysis, and therefore he can happily admit the conceptual possibility at issue while nevertheless denying the corresponding metaphysical possibility. Since only the existence of the latter would threaten his thesis, he sees no problem.

His response to the inverted qualia problem is more complicated. First, with respect to vision - the inverted spectrum hypothesis - he denies that color space is actually invertible. The point is that color quality space is not symmetric, and any mapping of one set of points onto others would result in changes that would be

detectable. So the dreaded possibility isn't realizable after all.

But he admits that many functionalists, the targets of the original inverted spectrum argument, still worry about it, since, after all, it isn't hard to imagine a color system very like our own that is symmetric in the required manner. At this point Clark appeals to his distinction between the functionalist doctrine that employs causation as the basic relation and his own doctrine that employs similarity as the basic relation. He argues that while functionalists are susceptible to the inverted spectrum problem, he isn't. The idea is this. We normally express the inverted spectrum hypothesis in terms of scenarios like "red fire trucks look to you the way green grass looks to me". If indeed one were defining red and green by appeal to their functional - i.e. causal - roles, then inverting the eliciting stimuli in this way would cause a problem (if the inversion were not otherwise detectable). But Clark alleges that he can easily live with the scenario just described. Yes, you will be experiencing green while I experience red when we're both looking at a fire truck. What's the problem? To be red, after all, is to occupy the relevant location in color quality space. As he notes, what would be a problem is if you could imagine moving orange, say, from in between red and yellow over to the interval between green and blue. That would indeed constitute a refutation of Clark's doctrine. But of course that sort of inversion he feels confident in denying any sort of possibility to.

It might be thought that my remarks above concerning the role of causation immediately puts Clark's response in jeopardy, but in fact it doesn't. Clark is concerned to deny that sensory qualities are defined by their causal relations to external stimuli.

Given that the only causal relation we found necessary to appeal to above was the relation to similarity judgments, this isn't a problem. However, I do have to say I frankly don't understand his answer to the inverted spectrum problem. We are assuming for the sake of argument that a symmetric quality space for color is possible. According to Clark, to imagine an inversion I would have to imagine something like switching turquoise and orange, so that the latter is now situated in between green and blue and the former in between red and yellow. Of course that seems impossible. However, that's not the right way to pose the problem. If we have a perfectly symmetrical space, our problem is to say what it is that distinguishes orange from turquoise in the first place. If you say, well, orange is between red and yellow, then we just push the problem back to these hues. What distinguishes red from green and blue from yellow? We know we have two distinct regions in the space, but no principle for identifying one as the red-orange-yellow arc and the other as the green-turquoise-blue arc. Relationalism leaves the actual identities of the color qualia undetermined (or underdetermined).

Of course traditional functionalists have an answer, at least initially. They say that what makes one of these two points red and the other green is its causal connection to certain features of objects in the world. They use the causal connections to distal stimuli to pin down the identities. But then, as Clark notes, the anti-functionalist can just retort that for all we know what green things look like to you - i.e. what sensory quality is caused by green things - may just be what red things look like to me, so these causal connections can't do the job. But if the job of providing determinate type-identity

conditions for color qualia can't be accomplished with the aid of causal connections to distal stimuli, how is it accomplished by merely removing them from the definition?

Now, what about my thought experiment involving a creature that can only see red? Clark leans heavily on the conceivability/possibility distinction to undermine the force of the objection. Since it's not part of the concept of color experience that it's a location in a quality space, the red-only experiencer cannot be ruled out a priori. But just as it's not possible for there to be water without H_2O, even though this situation is not ruled out a priori, so too it's not possible for there to be a red-only experiencer.

I am sympathetic myself to the general methodological maxim that conceivability doesn't automatically establish possibility. However, I do think we have to take it somewhat seriously. If a situation seems possible, then we need good reasons for thinking it isn't. In the case of water we have a story to tell, and indeed our own intuitions change when we learn that the stuff in our world is really H_2O. So the question is, what discovery is it that Clark can point to that lends the kind of support to the relational type-identity thesis sufficient to overcome the apparent possibility of a red-only experiencer?

As far as I can see, there are two principal sources of support for the idea that qualia are locations in similarity spaces. As usual, let's take color as our prime example. Clark, along with many other philosophers, notes that the similarity relations in which colors stand to each other are necessary, or essential to them. As mentioned above, orange is necessarily more similar to red than it is to blue. Clearly the thesis that color qualities are locations in such a space entails that the relations are

necessary, and so explains that fact.

Second, empirical theories about how we in fact process color information cohere well with the thesis. If something like the opponent-process theory is true, then very likely a red-only experiencer is not nomologically possible at least. That what makes the red-green channel what it is is its role in the entire system makes a lot of sense when we look inside, since it doesn't seem as if anything else - except connections to distal stimuli, which Clark rejects as the basis for color quality identification - can do the job. Perhaps the point should be put this way. Only if color qualities - all sensory qualities, no doubt - are reducible to locations in a similarity space can we see how a material system like the brain could instantiate them.

With regard to the first consideration, I don't find it very compelling. Suppose we do concede that orange is necessarily related to red and yellow in the way that it is. So it's essential to orange that it be more similar to red than to green, and essential to red, for that matter, that it be more similar to orange than to green. These facts alone do not tell against the possibility of a red-only experiencer. Suppose, for instance, that to experience red is to be acquainted in some direct way with this property of objects. It is essential to being red that it maintain the particular set of relations to other color qualities that it does, but that doesn't mean that one must be acquainted with these relations in order to be acquainted with red itself. Clark's own analogy with weight, which seems like an intrinsic property but turns out to be relational, supports this point. So all one might be committed to by the fact that red is necessarily more similar to orange than it is to green is that <u>if</u> someone experiences all three, they will experience

red and orange as more similar than red and green. But this alone doesn't entail that a creature who can only experience red is impossible.

In the end, then, it seems to me that the entire weight of support must rest on the second consideration: that only if we individuate sensory qualia relationally can we make sense of how the brain realizes them. I agree, I suppose, but then to my mind that's a problem for the materialist. Sensory qualia seem to have a determinate, intrinsic character, and it seems at least on the verge of question-begging, if not outright question-begging, to assert that they don't since if they did materialism would be in trouble. Especially when you add in the problem of inverted qualia, the case for the relational theory doesn't seem that strong.

At this point Clark might insist that intuitions that a property is intrinsic shouldn't be given that much weight, again citing cases like that of weight. But there are two crucial ways in which the sensory quality case differs from the case of weight. First, there are independent grounds for accepting the theory that weight is a function of gravitational force from the earth. This explains the appearances perfectly, including what seems intrinsic about the feel of weight when we try to lift something. Secondly, and this is connected to the first point, the transformation of weight from an intrinsic to a relational property leaves us with an intrinsic property that plausibly takes its place: namely, mass. Neither of these features is present in the sensory quality case. True, the nervous system may be coding sensory qualities by reference to these dimensions of variation, but that doesn't alone, as I've argued above, show that the location in the similarity space exhausts their character. Also, nothing that is intrinsic is left after the

identification with locations in similarity space. This leaves us with precious little to appeal to in order to explain the strong sense that sensory qualities have an "absolute value" as well as a relative one; a determinate identity over and above their set of similarity relations.

This point brings us back to Clark's insistence that causal connections to distal stimuli are not part of the identity conditions for sensory qualities. For Representationalists about qualia, what it is for me to experience a sensory quality of a particular sort is to occupy a state that represents a particular feature of distal stimuli. On this view, sensory qualities are not defined by their similarity relations to other experiences, but by the features they represent. It seems to me that this saves more of the intuitive data than Clark's view, though of course it runs into trouble of its own. At this point, in fact, it makes sense to look at what Clark has to say about the controversy over color itself as a feature of distal objects.

3. As everyone knows, there is a raging controversy about the ontological status of color - what Clark would call the phenomenal property, not the sensory quality. Realists maintain that colors are identifiable either with objective features of surfaces, say their spectral reflectances, or with dispositions to cause certain sensations in observers. For reasons that should be clear from the discussion above, Representationalists about the sensory qualities instantiated in experience tend to be realist-objectivists about the corresponding properties of distal stimuli. Irrealists maintain that physical objects are not actually colored, though we have chromatic experiences.

The difficult problem for realists is dealing with the rampant relativity of color

experience. How something looks is a function not only of its surface spectral reflectance, but also of the particular mix of light reflecting off of it, the reflectances of the objects around it, the state of adaptation of the observer's visual system, certain hard-wired features of each person's visual system, and so on. What makes the problem especially difficult is the fact that there doesn't seem to be any natural way to pick any values of these various dimensions of relativity and set them as the "normal" ones in terms of which an objective color could be defined. So when different people in different situations see objects as different colors, there doesn't seem to be a principled basis on which one can judge that one is right and one is wrong.

On the other hand, colors certainly seem to present themselves in experience as objective features of external objects. Irrealists usually describe what's going on in color perception as a kind of projection from certain features of their experience onto the perceived objects in the world. But if this is the case, then we seem to be accusing perception of making a systematic error. While it's not obviously absurd to adopt an error theory for color perception, it does seem to be a position to avoid if possible.
So we have come to this. Color perception seems to present us with information about the external world - that objects have certain properties - but we can't find a place for these properties in our theory of the world. Okay, so maybe colors are illusions, like witches and phlogiston. Our perceptual system is guilty of describing the world in terms of a false theory - albeit a useful fiction. But of course we are loathe to convict perception of such a systematic mistake, so we try to find an alternative way of looking at it, one that grounds a claim of perceptual veridicality. What might that be?

Well, one interesting idea is the one presented by Clark. He presents his own theory of color in the last chapter of the book. He argues against what he calls "chromatic atomism", the view that each type of color experience represents a particular objective physical property of external objects. Against this he favors the idea of "difference coding": that what chromatic experience does is detect differences in reflectance properties, but doesn't represent any particular reflectances. So, for instance, consider a red patch against a white background. On the atomist view, the red patch is represented as possessing some particular spectral reflectance - or at least one among a list of alternative reflectances. However, on the difference view, the visual experience is telling you nothing about the nature of the reflectances within the patch except this - that, whatever they are, they are redder than those in the white background.

He draws an analogy with so-called "false color" photography in astronomical research. Data concerning various properties of stellar emissions is represented with colors, so that one can look at such a photograph and immediately read off of it where the radiation gradients are. A red patch will signify high intensity, say, with a blue patch representing a cooler region. The colors in themselves are not meant to represent anything about the intrinsic quality of the properties being tracked, but the color differences do represent genuine differences in the levels of intensity.

Notice that when it comes to differences, relativity isn't a problem. Two surfaces with very different reflectance properties will match under some combination of illuminant, contrast, and state of observer, so that they match is no information that their

reflectances have anything in common except that we can't tell them apart in that particular circumstance. However, that two samples do not match, that they look different in the same circumstance, is definitive evidence that their reflectances differ. What's more, Clark argues, this is really all we need color perception for, since it is a way of guiding us concerning the identification of the borders of objects, the identification of objects under conditions of occlusion, etc..

But does this mean we don't have to appeal to any systematic illusion in color perception? Unfortunately, I think it doesn't. Two points are relevant here. First, notice that it isn't pure difference alone that is represented by border regions. We don't just say, this is different from that, but we describe the difference in terms of a quality - our visual system says, "redder here than there". But what's being redder? Again, in perception it's presented as an objective relation among the patches - not merely that they are different, mind you, which <u>is</u> an objective relation - but that the difference is a matter of being redder than. The problem is that it still seems as if the only interpretation we can give to the relation of being redder than is that it has more of what it takes to excite the red channel in the visual system than the other patch. Again, if we were content to let a state of our brain be what it is ultimately about, then we could have stopped earlier. To the extent we were unsatisfied with that for being red, I don't see why we should be satisfied with that for being redder.

But the second point is that we're not really done with being red either. True, we've identified a function for the visual system's "color coding" that serves to track genuine objective properties, or relations, out there in the world, namely, spectral

difference. When a patch all looks the same, according to Clark, all this means is that our visual system is saying, in effect, “no difference detected here”, even though in many cases different parts of the apparently same-colored patch might indeed have significantly different reflectance properties. It’s just that our visual system is not sensitive to this difference. But again, it doesn’t just say to us - “same old same old” - it says how it is the same, it’s red, or red31. It looks a certain particular way in its sameness, and different from the way that other patch looks in its sameness.

This brings us back to Clark’s analogy with the false color photography. It’s hard to tell just how seriously Clark wants to push the analogy. However, it seems from the very logic of his position that he should take the analogy quite seriously indeed. What I mean is this. For him the color appearances in experience should be treated as no different from the colors in the “false color” photographs. Redness, as it appears, stands for nothing more than “there’s more in this area of what stimulates the red channel than there is in the surround”. If this is right, then of course the next question is what to make of the property we in fact experience - the redness in the experience itself. Put it this way. It’s fine to talk about color coding of other properties as a form of representation, since after all the vehicles of that representation - the regions in the false color photograph - really have the colors being used to do the color coding of levels of radiation. But in the end, what does it mean to say that what the visual system is doing is color coding spectral reflectance differences in the world? What has the colors in this case? As in the old Pepsodent commercial, you wonder “where the yellow went?”

It seems to me that Clark's solution to the problem of color and his relational type-identity thesis concerning sensory qualities are of a piece. The general idea is to rid the world of intrinsic qualities, whether of objects or states of mind, of the sort that do not fit easily into a natural physicalist framework. I myself am inclined toward a subjectivist account of color, but cannot buy the relational account of the sensory qualities that correspond to color. But now, let me turn to what is arguably the most significant issue raised in the book, Clark's account of the nature of sensory representation. At the end of the section I'll connect the discussion back to the issue we've been talking about.

4. On Clark's view, sensations have representational content; in particular, something like a singular term-general term form. One argument he gives for why there has to be a singular term component to sensory representations is that otherwise we couldn't distinguish a sensation as of a red square and a green circle from a sensation as of a red circle and a green square. We need "collecting principles" to determine that redness and squareness go together. What serves that function is representational structure that allows one's sensory system to represent that it is the very same object that is red that is square. For this one needs not only predicates, but singular terms to which the predicates attach.

The sensory predicates represent the phenomenal properties of objects: redness, saltiness, the odor of a rose, and the like. We've already discussed the nature of these properties and relations above. But what about the sensory singular terms, or "proto singular terms"? According to Clark, the objects in question are space-time

regions of finite extent. It is not a solid, enduring object of which one predicates “red” in the first instance - i.e. in sensation - but rather a region of space-time. He calls this the “feature placing” hypothesis. The idea is that the messages that our sensory systems send to higher cognitive centers are more like “reddish and squarish at 2 o’clock” than like “some red square object is at 2 o’clock”, where the latter is to be interpreted as saying that it just happens to be located there, but it, the very same thing, could move.

That sensory representation must involve some singular term role strikes me as evidently plausible. Without some such representational device it really is hard to see how the so-called “binding problem” could be solved. However, it certainly isn’t obvious that the choice of individuals for sensory singular terms should be spatio-temporal regions. As Jonathan Cohen notes in his contribution to this symposium, Zenon Pylyshyn (2003) has argued persuasively for taking the primitive visual objects to be individuals that occupy locations, not locations themselves. Pylyshyn points to a number of experiments involving multiple object tracking and the like to support his hypothesis. I needn’t rehearse the arguments in favor of the view, as Cohen has done that admirably. However, I did want to put in my two cents here, given that Cohen was nice enough to share his paper with me in advance and I also found Pylyshyn’s account extremely compelling. What I want to do is explore a bit the implications of the two views - Pylyshyn’s and Clark’s - for the question of demonstrative reference.

There is clearly an important connection between demonstratives and perception. Though it is possible to use demonstratives to pick out objects other than by way of perception, there is a certain primacy to the perceptual case. I see a coffee

mug on my desk and say “isn’t that the coffee mug I was looking for?”, or hear someone coming up the stairs and say (possibly to myself) “who is that, I wonder?” My use of “that” in each of these cases is clearly dependent on my perceptual experience in a crucial way. I see or hear what I’m referring to by “that”. The interesting question is how to understand the connection between the percept and the use of the demonstrative, itself a constituent of a non-perceptual representation.

What might seem like the most straightforward connection between thought and perception here is to identify the referent of “that” in “is that the coffee mug I was looking for?” with the “visual object”, as Pylyshyn calls it, that is the relevant individual in the perceptual representation constituting my perceptual experience. But if the demonstrative in my thought picks up its reference directly from the demonstrative in my percept, how is the reference of the latter accomplished?

Pylyshyn insists in his account of the connection between perceptual “indexes” and visual objects that the indexes are not to be understood as definite descriptions. He makes the same kind of points against the definite description account here that Kripke does for names; that the reference to the visual objects survives changes in any of its descriptive properties. We can index an object through changes of location, color, and a host of other visual properties. So the index can’t be equivalent to something like “the shape that has such-and-such color and such-and-such location”. Rather, as with Kripke’s account of names, some causal connection does the work.

But if you take the analogy with Kripke on names seriously, it’s clear we need more. After all, on Kripke’s account the causal chain by which standard uses of a name

refer to the relevant person (or object) begins with a reference-fixing event: either a dubbing or a definite description. Dubbings usually involve perceptual demonstratives. The reference fixer for the name, or the verbal demonstrative, is then the percept. So what fixes the reference of the percept? What connects the index to the object?

One possibility is attention. There is normally a lot going on in a perceived scene, so if we expect some mental index to be assigned to an object, it makes sense to employ the mechanism of attention. In fact, in Pylyshyn's theory of perception attention does play a prominent role. But again, this can't do the job, as Pylyshyn himself admits. Attention is a mechanism that involves allocating some mental resource (maybe just pointing the eyes) to an object, which clearly can't work unless the object in question is already identified. So the indexing to visual objects must be a preattentive process.

At this point it strikes me that Clark's view can make an important contribution. A plausible hypothesis about how reference-fixing for visual indexes takes place is that an object is posited, and labeled with an index, where a spatio-temporal region instantiates a certain constellation of properties. Just which constellation of properties triggers the application of a visual index is an empirical issue, of course, but it's plausible that various considerations having to do with contiguity and saliency of borders play a dominant role. The crucial point, however, is that the visual system first represents what's happening in various regions of spacetime, and only then gets around to assigning labels to objects. If you like, one might say that both Clark and Pylyshyn are right. At the most primitive level of representation, the perceptual system assigns

features to spatio-temporal regions, and then at a somewhat higher level - though still within the system of early vision that Pylyshyn is at such pains to distinguish from higher cognition, reference to location-independent objects emerges. True, Pylyshyn himself dismissed the idea that indexes could be assigned on the basis of predications to spatio-temporal regions. But he was thinking of these along the lines of definitions, not reference-fixers. So long as they only function as reference-fixers, his concern about the logical independence of visual objects from locations and features - the fact that they can survive changes in these features - is fully addressed.

Two questions still remain. First, how does reference to a spatio-temporal region get established? Is it by some isomorphism between a system of spatio-temporal representation and spacetime itself? Is it by some brute causal connection? Someone might wonder, if this question of reference-fixing is still on the table - and clearly it has to bottom out somewhere - why it couldn't bottom out with visual objects themselves. Perhaps. But to me it seems more natural and plausible that if there is a kind of primitive reference grounded in perception, upon which all other reference-fixing depends, that it should be based in space and time itself. That perception is at its core in the business of telling us what is happening where and when seems quite plausible to me, and it seems like the basic idea behind Clark's thesis.

Second, what is the connection between the various levels of perceptual representation and what's present in conscious experience? When I consciously attend to the looks and feels of things, is this what theories like Pylyshyn's and Clark's are tellimg me about? It's not clear from Clark's discussion just what his view about this is,

though what comes across is that his thesis concerning the structure of sensory predications is supposed to capture how things seem to me when I perceive the world. No doubt this is not a question that has an easy answer.

REFERENCES

Pylyshyn, Z., 2003. Seeing and Visualizing: It's Not What You Think. MIT Press: Cambridge, MA.

Siewert, C., 2004. "Is Experience Transparent?", Philosophical Studies 117, 15-41.