

**THE PROBLEM OF SPATIALITY FOR A RELATIONAL VIEW OF
EXPERIENCE**

BY JOHN CAMPBELL

**Forthcoming in Christopher Hill and Brian McLaughlin (eds.), Special Issue of
*Philosophical Topics on Perceptual Appearances.***

THE PROBLEM OF SPATIALITY FOR A RELATIONAL VIEW OF EXPERIENCE

On a relational view, visual experience is a relation between the perceiver and the scene observed. It's often said, quite wrongly, that relational views of experience cannot reckon with mode-of-presentation phenomena. I begin by explaining how a relational view of experience can make sense of the idea of the 'way' in which an object is given to the perceiver, in terms of the properties used to single out the object. I explain why relational views are to be preferred to approaches in terms of 'mental paint'. I then sketch the problem posed by spatial perception. Spatial perception is peculiarly difficult to characterize on a relational view, because it's difficult to characterize the way in which spatial locations and layouts are given to the subject.

1. Causal Structure in Visual Phenomenology

Suppose you're walking through grassland and as you look idly over the scene the patterning of orange screened by tall grass suddenly comes into focus and you see the tiger behind the benevolently waving tall grass. You perceive an object where you didn't see one before. And yet there's some sense in which the pictorial characteristics of your perception might be exactly the same after you seen the object as they were before. The mosaic of variously located colors and shapes in your visual field might be manifestly the

same after you see the thing as it was before. The alternating of orange and green was there in your visual field throughout.

In this transition, what we have is a causal sequence, the upshot of which was you seeing the tiger. So how should we describe the sequence? What, exactly, causes what?

Your seeing the colors – in particular, the orange against green – causally made it possible for you to see the tiger. There were those colors at those locations, that's just a fact about the external scene. You encountered those colors in your visual experience. And you seeing those colors at those locations caused your perception of the tiger. It wasn't the only cause of seeing the tiger, of course, but it was part of the causal story.

Now so far, you need not have seen the shape of the tiger at all. Your visual perception might have been resolutely indefinite so far as to the shape of the thing. But once you've reached this point, where the perception of color has caused your perception of the tiger, it's possible for you now to go on to interrogate the scene visually to find further characteristics of the thing. Is it big, or is it perhaps a rather small tiger? What is its posture, is it asleep or is it crouched to spring? The perception of colors at locations is one of the causes of your perception of the tiger, and this in turn is one of the causes of your perception of its further characteristics. Even in this simple sketch, we can see that there is some causal structure in the phenomenology of vision.

Notice that in the common-sense terms in which I've stated the points so far, the causal relations here are causal relations between states that are themselves relations between the perceiver and the surroundings. The perception of colors and locations causes the perception of the tiger. The perception of the tiger causes the perception of its various aspects, shape, size, posture and so on.

Notice also that I've described the causal story here as though it's a temporal story. It's natural to tell causal stories in that way, but in this case it doesn't seem that we need be dealing with a succession of events giving way to one another. As you see the tiger's posture, you still see the colors and locations that are causing you to see the tiger. In fact, you might argue that if you didn't perceive the colors and locations, you would no longer perceive the tiger.

A quick way to see the point here is to reflect on Lewis Carroll's story of the cat that slowly faded away, leaving nothing but its grin. Of course grins are, ordinarily, visible. But seeing a grin ordinarily causally requires seeing something of the color and shape of the thing grinning; seeing the grin comes last. That's why it's so hard to imagine seeing the grin without any of the things that make it causally possible.

Of course, it can also happen that you keep track of the tiger as it runs to become a dot in the distance. In fact, you could keep track of the object as it moves without continuing to monitor its color and shape. But the initial identification of the object, and access to its properties such as shape and size, was causally made possible by your initial perception of colors at locations. We seem able to keep track of objects in ways that don't depend on keeping track of their properties, but it doesn't follow from that, that our visual identification of objects could be initially causally generated without perception of properties at locations. Moreover, if you are to access properties such as size and shape, you seem to depend on having been able to pull out the object from its background because of your perception of its properties.

So although you can keep track of objects over time without continuously monitoring their properties, there is, as it were, something 'anaphoric' about the keeping

track here; it depends on your initial identification of the object by way of its properties. We might draw an analogy with pronouns. If I say, ‘your brother came to see me this morning. He said And then he’, the later uses of the pronoun ‘he’ have their meaning by virtue of their connection back to the initial identification of someone as ‘your brother’. Similarly in perception, as the tiger runs off, you may be perceptually identifying it ‘anaphorically’, as it becomes a dot, but your identification of it here depends on the connection back to your initial perception of colors at locations.

So far, the claims I’ve made are:

- (a) perception of location and color causes perception of the tiger, and:
- (b) perception of the tiger causes perception of its shape and posture

The main point I want to consider first is the methodology of how we establish the kinds of claims I’ve been making about the causal structure of the experience of seeing. How could we establish, for example, that the perception of color and location is causally more basic than perception of the shape of the tiger? Is it really true, for example, that perception of color causes perception of the tiger and perception of shape does not? These are empirical claims, and should be established scientifically. But how do we do that?

2. What Do Vision Science Experiments Tell Us About Visual Experience?

There is a curious dissonance in what people usually say about the bearing of vision science experiments on the character of visual experience. On the one hand, this is usually thought of as an extremely vexed question. It's not uncommon for distinguished vision scientists to feel that generally they are better off without using constructs such as 'experience' or 'consciousness', and that where these ideas should ultimately figure in the analysis of vision science experiments is something best left to posterity, or retirement at any rate.

On the other hand, the fact is that the analysis of a vision science experiment usually proceeds at two levels. The authors reporting a study will typically move back and forth seamlessly between:

- (a) Describing the study as showing the relations between the subject's perceptions of various factors in the display on a computer screen, and
- (b) Providing a theoretical model to explain the results of the study in terms of the information-processing going on in the visual brain.

Now level (a) is the level at which we find talk of ordinary seeing. We say that the subject saw the coloration of some part of the display, for example, and that this is what caused the perception of an object. Talk of ordinary seeing, at this level of the relation between the person and the display on the screen, is, ordinarily, talk about the subject's experience. At level (b), however, talk about the conscious life vanishes. We are here at

best talking about neural processes and mechanisms that in some sense sustain the causal structure we find at the level of ordinary seeing.

When you think of it like this, it seems evident that vision science experiments are indeed addressing questions about the causes, effects and nature of the subject's conscious life when performing tasks relating to the display on the screen. But that is not the attitude people usually take when explicitly reflecting on what vision science experiments tell us about visual experience. It's usually supposed that to be informative about consciousness, vision science experiments would have to be telling us something about the behavior or 'qualia' or 'sensations' or 'presentations (representations richly conceived)' that lie outside the compass of ordinary vision science explanations, and yet are somehow internal to the head of the subject. Once we have stepped away from talk about ordinary seeing, on the one hand, and, on the other, from talk about information-processing being carried out in structure in the brain, then we have really very little idea of what we're talking about.

To see the difference between these two ways of thinking of the relation between vision science experiments and visual experience, consider the influential Boolean Map theory of visual attention put forward by Huang and Pashler 2007. This is an empirically very well-supported theory.

Let's start by describing things at the level of information-processing in the brain. The core idea is that we have an underlying architecture of Treismanian feature maps. At the stage of selection, a property is chosen and the totality of locations with that property are pulled out. In effect, this pulls out the object (the collection of places is 'shrink-

wrapped' to fit the object). Then the maps are consulted again to access the various properties of that complex location. This is shown in Figure 1.




Figure 1. A feature (in this case, let's suppose, orange) is used to select all and only the locations with that feature at them. A first consultation of the feature maps is made, to find which regions have red at them (in effect, this selects the tiger). This generates a Boolean map (a distinction between selected and unselected areas). The feature maps are then consulted again to determine what features are found at the region selected (is it a tiger? Is it big?), and the labeling then makes those features explicit. As you can see, Huang and Pashler locate the

level of conscious experience as somehow internal to the information-processing system. I'm suggesting that we should think of conscious experience as, rather, a set of relations to the scene observed, with its causal structure underpinned by the causal structure found at the information-processing level. (From Huang and Pashler 2007.)

Now the analysis I've been suggesting is that we should think of the causal structure of experience as constituted by the relations among the various relations studied between the subject and various aspects of the array shown on screen. We should think in terms of causal relations between the subject's perception of colors at locations on the screen, the subject's identification of an object on the screen, and the subject's perception of various aspects of objects seen on screen. This is simply a different level of analysis from the one at which we talk about neural information-processing. The important point is that the causal relations among various aspects of your experiential relation to the scene are underpinned by causal processes found in the brain. As you can see from Figure 1, though, Huang and Pashler seem implicitly to understand its implications for conscious experience in terms of the whole structure generating some further construct, 'conscious access', that lies outside their explicit study. What I'm suggesting is that we should instead see their picture as showing us the neural underpinning of a causal structure in the experiential phenomenology itself. The causal structure is one we've already noted:

- (a) perception of color and shape causes perception of the tiger, and:
- (b) perception of the tiger causes perception of its shape and posture

Now the locutions I'm using here are fully factive, they have to do with a causal structure between relations between the observer and the properties and objects in the environment. It's a causal structure in the phenomenology of experience. But it's not necessarily one that's open to introspection, you can't establish these connections merely by looking inwards at your perception. After all, when you try to look inwards at your visual experience, all you can do is focus on the scene you're observing. And there aren't causal relations between the colors and locations out there, the existence of the tiger, and it's having the shape it does. The causal relations are only relations among your perceptions. They can't be observed by introspection.

Here we're talking about general causal relations between different aspects of perception, that are established by experiment and observation. There doesn't seem to be any such thing here as tracing a causal process from one to the other, as we might trace the trajectory of a cannonball or follow someone's train of thought. Experiment establishes a causal connection between, for example, perception of colors and locations, and perception of an object. And there doesn't seem to be any such thing as 'following the process' here. At the neural level we can of course find neural processes. But at the psychological level, we just have the brute fact of causation, and its sustaining of counterfactuals, such as, 'If I weren't seeing the colors and locations here I wouldn't be seeing the object'. And the array of empirical data that Huang and Pashler (2007) give for their view of neural processing equally support the higher-level conclusions for

perceptual experience that I have been drawing, once we attend to the higher-level descriptions of what subjects are doing in their experiments ('seeing the colors in the display', and so on).

3. Mental Paint Theory

What I have been suggesting so far is that we can characterize visual phenomenology in terms of what we are seeing and the causal relations among various aspects of our perception of our environment. In a recent article, Ned Block (2010) proposes a quite different approach to how we highlight and visually determine the characteristics of objects in our surrounding. Block thinks that visual experience should be characterized in terms of a quantity that he calls 'mental paint'. Naturally, he thinks that 'mental paint' has something to do with modes of presentation. The idea is that in vision, the external scene is somehow 'painted' using this 'mental paint'. As we'll see, the big problem with this idea is to explain what 'mental paint' is.

The basic problem can be put as a dilemma. Can we explain ostensively what mental paint is? Can introspection, or an imaginative understanding of someone else's visual perception, tell you what it is? Well, a moment's trial is enough to show that this won't work. When you reflect on your visual experience to find the mental paint, you simply end up staring with a peculiar intensity at the furniture. But 'mental paint' isn't supposed to be a characteristic of the furniture. It's supposed to be a characteristic distinctively of your inner life.

Well, perhaps ‘mental paint’ is the name of a theoretical construct. After all it is perfectly legitimate to have theoretical posits, such as ‘electron’ or ‘magnetic flux’, that can’t be explained ostensively. So maybe we should regard ‘mental paint’ as being such a construct. The trouble with this way of thinking of it is that ‘mental paint’ is supposed to be an immediate characterization of visual experience. To know the facts about the mental paint in someone’s visual experience would be to know what their visual experience is like. But if ‘mental paint’ is a theoretical construct, then presumably the theory in which it is embedded could be known to a Martian psychologist, who had no variety of visual experience and had no idea what visual experience was like, though perhaps they’d often wondered about it.

So there doesn’t seem to be any way of explaining the meaning of a term like ‘mental paint’. It’s spurious, and can’t refer to anything, although the metaphor may intrigue. How does Block explain what mental paint is? Here is what seems to be his official explanation:

Are phenomenological characters of perception—e.g. what it is like to experience redness or roundness—philosophically reducible to the redness or roundness of the objects one sees or to representation of redness or roundness? If there is no such reduction, then there can be said to be mental paint.

(Block 2010, 23-24)

And that is all the positive explanation he gives of what mental paint is. Now from the standpoint of a relational picture of perceptual experience, there is no need for any such notion. We characterize the qualitative character of visual experience in terms of relations we bear to our environment, and we don't need any such inner construct. But I think we can see the intuitive motivation for some such notion as mental paint as running like this. Suppose you begin by thinking that perception should be described in representationalist terms, then you almost immediately run into two problems:

- (1) this misses out the way in which consciousness seems essential to perceptual experience, since you could have any particular type of representational state without that state being an aspect of experience,
- (2) it misses out 'perceptual presence', the sense in which redness itself, for example, or spatial distance, seem to be present in an experience of redness or distance. People sometime say, for example, that in perception, redness is 'presented', rather than 'represented'. This is intended to capture the difference between seeing redness and representing it by a verbal description. But of course, on the natural reading, this shift involves a move to a relational picture, on which the experience is a relation to the external redness that you are perceiving.

At this point it is natural to supplement an appeal to 'representation' in characterizing visual experience with an appeal to 'mental paint' or 'visual sensations'. The 'paint', the 'sensations' will be what pump up the representation and add an experiential dimension

to the perception, the 'paint' will be what explains the sense in which 'red' is literally there in an experience of redness.

Now from a relationalist point of view these problems are both spurious. If experience is a relation to the environment, then there is no need to pump up something else with mental paint in order to get the experiential. And the 'perceptual presence' of redness is achieved merely by perception being a relation to the external redness of the thing.

However, if, like Block, you start out with a representationalist picture and then inquire what inner stuff has to be added in order to get visual experience, then it will seem natural to appeal to some such idea as 'mental paint'. There is, as I said, no coherent notion to be explained here: there isn't any way of explaining what this stuff is. It emerges only as the 'solution' to a badly posed philosophical problem. And 'solutions' to badly posed philosophical problems typically are incoherent. Now I just quoted Block's official explanation of what mental paint is, and it is, I think, obviously inadequate. His official explanation is what you might call the 'UFO' theory of mental paint. If there's something about visual experience that can't be explained in other terms, that's mental paint. If there are lights in the sky that can't be explained as the stars or as the lights of passenger aircraft, then there are UFOs. That 'explanation' doesn't tell you what a UFO is.

Now Block does seem to recognize that an inquiring mind might want to know a bit more about mental paint than the official explanation supplies, because a couple of pages later he says this:

I am not assuming that if there is mental paint, it is non-relational (“intrinsic”) or has no representational aspect. Since I favor physicalism, I allow that mental paint may be a relational neural property. To avoid misunderstanding: I do not claim that there is anything red or round in the head when one veridically sees a red or round thing in the world as when red pigment in a painting represents a red barn.

(2010, p. 56 footnote 2)

But, of course, this destroys the whole point of the metaphor of ‘mental paint’. The point of the metaphor of ‘mental paint’ is to allow you to explain the contrast between perception of redness and a mere verbal representation of it. When you paint you use a pigment that is literally red, so perhaps experience is like that too. If the paint isn’t literally red, then we still have no idea how a representation plus mental paint can add up to redness itself being there in the visual experience. What the various types of mental paint might be, there is no way of explaining: we can’t have a systemic vocabulary here, because of the fundamental incoherence I just remarked. You couldn’t explain the arcane characteristics of mental paint ostensibly, and you couldn’t regard them as merely theoretical constructs either. Yet still after all it is true that a visual experience doesn’t merely represent the external world, as a newspaper article might represent the external world. These qualitative characteristics of the external world, the colors and shapes and so on, are present to us in experience. Block’s need to explain that the ‘paint’ is not literally ‘red’ stems from this having been the point of the metaphor. But once you bow

to the inevitable and admit that the inner world doesn't actually have these characteristics, you give up the point of talking like that.

Still, Block argues that 'mental paint' has some essential work to do in the analysis of experiments on visual attention. Now so far I have, in effect, explained visual attention to an object as a matter of:

- (a) perception of colors and locations (e.g.) as allowing one to *select* that object in order to find out more about it visually, and
- (b) visually *accessing* (making explicit) various aspects of that object, such as its color, shape and so on.

Now Block prefers not to explain visual attention to an object in terms of perceptual relations that one bears to that particular object, but rather to proceed in terms of his notion of mental paint. If you are thinking along those lines, it's natural to emphasize the idea that attention has to do with the idea that attention somehow constitutively makes the colors brighter, the contrasts greater, acuity heightened. Let me say right away that I think there is absolutely nothing to any such idea. Suppose you're watching a beetle. It's entirely possible that elsewhere in your visual field, the colors and contrasts are being heightened, without you noticing any of that going on, your attention is fixed throughout on this unremarkable beetle. So it's hard to see how there could be any constitutive place for such phenomena in an account of visual attention. Still, Block focuses on work by Carrasco and her colleagues which claimed to show 'Transient [covert] attention increases apparent contrast for a wide range of stimulus contrasts, which in turn enabled

observers to respond better in a discrimination task [i.e. reporting orientation].’ (Carrasco, Ling and Read 2004, 312). What I’ve been suggesting would imply that any such ‘increase in contrast’ should be regarded as a *consequence* of a shift in attention; attention should be explained as a matter of selecting an object in order to access its properties (cf. Huang and Pashler 2007). But Block ignores the problem of giving a constitutive account of visual attention in order to focus on the claim that the shifts in ‘apparent contrast’ that Carrasco finds have to be characterized in terms of ‘mental paint’. Since he does not, as we’ve seen, have anything very convincing to say about what mental paint is, it’s unsurprising that his description of the Carrasco phenomena is also unconvincing. But let’s look.

Block interprets the Carrasco findings as showing that attention affects the contrast of a seen patch. This interpretation is itself multiply problematic; but I set aside these difficulties in order to look at the theoretical work that ‘mental paint’ is being set to do. His basic framework is that perceptual experience consists of mental paint pumping up representations. Now we have to explain the Carrasco’s finding that covert attention to a patch will affect the subject’s assessment of how contrasty that patch is, when compared to an unattended patch. According to Block, in such a case, we cannot say that either the perception of the attended patch is illusory or that the perception of the unattended patch is illusory. He puts his analysis in terms of a ‘mental paint’ component, the ‘Phenomenal component P’, plus a ‘Referential component’. The idea is that we need the ‘mental paint’ component to explain the shifts in the subject’s assessment of the contrast of an external patch, depending on whether or not the subject covertly attends to it:

‘[T]here is no way to pick which distribution of attentional resources engenders veridical perception and which engenders illusion.’

(p. 45).

‘[T]he 22% patch attended (i.e. with more attention) is the same in perceived contrast as the 28% patch unattended (i.e. with less attention) ... both would be captured by

{*Phenomenal component P*; *Referential component*: <That; 22%-34%>}

(p. 53)

The ‘phenomenal component P’ part is the thing that is supposed to be a ‘mental paint’ phenomenon. What happens when we shift from the situation in which there is covert attention to one of two patches, to the neutral condition in which there is no attention to either patch? Does the Phenomenal Component P shift, or does the representational content shift?

Suppose we say that the Phenomenal Component P shifts. This seems to be the point of Block’s argument, to say that there is an impact on subjective experience of the shifts of attention that is not to be understood in terms of representational content. Then (presumably) the shift in P is what explains the shift from judging Left to be lower-contrast than Right to judging them to be similar in contrast.

Is this idea convincing? Even in its own terms, it raises some baffling questions. First, why should subjects be responding to the arcane ‘mental paint’ phenomena at all

(supposing we knew what they were)? The subjects are being asked to make judgments on the basis of the contrast of the patches. They are veridically perceiving both patches. But instead of relying on their veridical perceptions, they are causally responding to the irrelevant magnitude P (the ‘mental paint’ phenomenon). Why do that?

It’s not that what subjects do must always be normatively correct, but we do want some explanation why they should be engaging in some apparently irrelevant procedure when presumably they are trying to comply with the experimenter’s requests and respond only to the contrast of the patches they see.

Secondly, the shock value of this approach should not be understated. If this theory is right, then what the subjects are doing in Carrasco’s experiment moves them right out of the domain of normal cognitive science. Cognitive science classically tries to explain what subjects are doing in terms of (1) computational tasks that they’re performing, (2) the algorithms they use to execute those tasks, and (3) the implementation of those algorithms in the brain (Marr 1982; cf. Griffiths, Lieder and Goodman 2015). But if this theory is right, the variation from the Neutral to the Attended condition can’t be explained in information-processing terms, but must instead be explained in terms of magnitude P, the mental paint. (Of course we don’t know how this magnitude is implemented in the brain.) It’s astonishing if this is true. Carrasco’s experiment, though interesting, looks like an absolutely standard cognitive science experiment. Can it really be that the results of a standard cognitive science experiment demand explanation in terms of an exotic construct like ‘mental paint’? Particularly given that the construct seems to be incoherent, being incapable of any explanation, this conclusion is hard to accept. And of course much more straightforward explanations are

available, for example in terms of a shift in decision criterion by the subject for the contrast of an attended patch as opposed to an unattended patch (cf. Beck and Schneider in press.) Block brushes aside this kind of analysis, but he greatly underestimates the wildness of his own proposal.

We might modify the theory to say: ‘Phenomenal Component P shifts when we move from the Attended to the Neutral condition, but that’s not what explains the shift in people’s judgments. The shift in judgment is to be explained in standard information-processing terms.’ The trouble now is that there seems to be no point in postulating magnitude P. Carrasco’s results can be explained in terms of the processing of representations.

Or we might say: ‘Phenomenal Component P stays the same when we shift from the Attended to the Neutral case, but there’s a change in the ‘Referential Component’ of the experience, and that’s what explains the shift in judgments of relative contrast.’ However at this point there again seems to be no point in postulating the Phenomenal Component P, all the explanatory work is being done by the shifts in the Referential Component.

Let’s finally rehearse the background problem with ‘mental paint’ theory, namely, explaining what’s meant by the construct in the first place. How are we to say what the difference is in mental paint, in magnitude P, that might be required by Carrasco’s results? ‘Contrastiness’? This is the kind of idea that Block explicitly repudiates: the mental paint isn’t contrasty any more than it’s red. The fact is that we have no way of specifying ‘magnitude P’, and Block seems to acknowledge as much. He makes no attempt to say how different types of mental paint might be identified. We can’t regard

the conceptions of particular varieties of mental paint as observational, because of the transparency of perception. We can't regard the conceptions of particular varieties of mental paint as theoretical, because then you could fully grasp them without having any idea what the relevant visual experiences were like. There seems in principle to be no way of explaining the idea.

How should we proceed? We have two kinds of reasonably well-understood vocabulary for characterizing the relevant working of vision. These are:

- (1) Terms for the characteristics of the objects and properties around us: their colors and shapes, their sorts and behaviors, and so on.
- (2) Terms characterizing the information-processing taking place in the visual system in the brain – edge-detection, symmetry analysis, and so on.

Of course there are other aspects to vision – affective and aesthetic, for example – that have to come into a fuller description. But for a basic description of the role of vision in generating knowledge and action, these will do. The relational picture simply works with these plainly meaningful concepts to analyze the situation at the two levels I described earlier.

At one point Block describes 'mental paint' as a 'mode of presentation' of a representational content (p. 53). You might suppose that this gives a clue to his official explanation of the concept of mental paint. But there would then be a sharp contrast between Block's approach on the one hand, and the approach of the representational and relational views on the other. For the representationalist, there are familiar ways of trying

to explain the concept of a ‘mode of presentation’: the idea would be that an object is represented descriptively, in terms of the properties it is represented as having. For a relationalist, as I began by explaining, the idea of a mode of presentation of an object can be explained in terms of the properties whose perception allows the object to be plucked out from its background. Block’s approach would be quite different: to take the notion of a ‘mode of presentation’ as primitive, and use it to explain the concept of ‘mental paint’. This approach sheds darkness rather than light. The notion of a mode of presentation needs explanation, and mental paint theory offers none.

4. Imagination *De Re*

I want now to look at another angle on the phenomenological character of experience, that bears on both the puzzle about spatial perception that I am building up to, and on the desire people have to postulate such constructs as mental paint. This is the role of imagination in our understanding of visual experience. We characteristically understand the phenomenology of someone’s visual experience by imagining it. And we can make a distinction between different types of imagination. There is, first, and most basic, imagination *de re*. This is when you can see the scene that the other person is observing, and imagine, *of that scene*, the other person’s perception *of it*. There is another, more sophisticated kind of exercise, in which you merely imagine a scene, and then, of that imagined scene, imagine the other person perceiving it. We might call this imagining *de dicto*, since it makes no existential demands on the reality of the scene perceived.

Now as I've just explained the distinction, imagining *de re* is evidently more fundamental than imagining *de dicto*. The case in which you simply rely on your ordinary perceptions to provide you with the scene you imagine the other person's perception of seems evidently less cognitively demanding than the situation in which you have to first imaginatively construct a scene and then within that overarching enterprise, imagine experience of the scene. Developmentally, it's a natural idea that the ability to imagine other people's perception is dependent on the capacity for joint attention to a shared scene (cf. Moll and Metzoff 2011).

Of course, to say that imagining *de re* is more basic than imagining *de dicto* is not to suppose that you can only imagine someone else perceiving the same scene in the same way, or from the same perspective, as you do yourself. What develops over time is your ability to extend your imaginative understanding of perspectives other than your own. You come to get a better and better understanding of how other people's perspectives on the scene could be quite different from your own. But that does not mean that *the scene* drops out as the foundation of your understanding of the other person.

With that in mind, consider another line of thought that can seem to support the 'mental paint' idea, which isn't quite explicit in Block's discussion. I think this is again a case of a badly put philosophical problem leading to an incoherent construct. But I also think it's a fairly popular line of thought, so I want to make it explicit. The idea runs something like this. The phenomenological character of an experience has to do with what that experience is like. The 'what-it-is-likeness' of an experience is what we have to characterize. We can't characterize this what-it-is-likeness in terms of the external

environment, and we can't characterize it in terms of representations the subject forms. Therefore, we have to appeal to mental paint.

This kind of reasoning is usually taken to be perfectly straightforward, and something that the writer need only gesture at before proceeding to formulate some construct. But I think it needs some discussion. Let's begin with 'what it is like'. The phrase 'what it is like' was introduced to the philosophical literature by Thomas Nagel in the paper that's the foundation of today's consciousness studies, 'What is It Like to be a Bat?'. Now Nagel was quite explicit about how he was using the phrase. It was correlative with the possibility of *imaginative understanding* of another subject. As he put it:

... the fact that an organism has conscious experience at all means, basically, that there is something it is like to *be* that organism.

At present we are completely unequipped to think about the subjective character of experience without relying on imagination - without taking up the point of view of the experiential subject.

The 'phenomenological character of experience' is what we know about by means of imaginative understanding, and not (currently, at any rate) in any other way. Now the key move that theorists make is to suppose that 'imaginative understanding' here has to be understood as relating to imagination *de dicto*. That is, people take it that imaginative understanding of another person's perception, in particular visual experience,

encompasses only 'internal' aspects of that person, and does not encompass aspect of the subject's environment. That's why we need 'mental paint', to characterize the inner 'what it is likeness' of the experience. Maybe add some representation too. But, of course, if what I have been saying about imagining *de re* being the most fundamental type of imagining here, then this line of thought is simply a mistake. The fundamental way you understand the qualitative character of someone else's visual experience is by seeing the scene perceived.

This point is, incidentally, also implicit in the 'two level' analysis I gave of vision science experiments. I said that at one level, we have the relation between the experimental subject and the tasks being performed on the screen, and I said that this is the level at which we find visual experience. That is the level at which we use imagination *de re* to understand the subject's experience, to imagine, of the screen, what it is like to experience that. The level of neural processing is a different level. One of the unfortunate side-effects of supposing that imagining *de dicto* is basic is getting the idea that the subject's experience must somehow be 'inside the head' too, and consequently must be discernible somewhere as a stage in neural processing at which qualia, or mental paint, are squeezed out like toothpaste.

Of course, you have to accommodate the other person's perspective on the scene. Now one way that can happen is the way we began. It can be that you and the other person see just the same objects, but that the properties of the objects that are visible are different between the two of you. So it can happen that the properties of the object that you and the other person use to pull out these objects from their backgrounds are different. Maybe you focus on the difference in color between the object and its

background, whereas the other person uses slight movements of the object to discern it. You can imagine that happening; you can imaginatively understand how the other person might, in that way, be experiencing the object differently than you do. But there is a more basic difference that there can be between you and the other person. Typically you will, of course, be at different places. So perspectival locations you experience the object and its properties as having will be different between the two of you. So this is another, and more basic, 'mode of presentation' problem. We have to understand how one and the same place can be presented in different 'ways' to different people, or to the same person at different times. And here the strategy I began with, explaining 'mode of presentation' in terms of the distinctive properties used to pull out an object from its background, will not help. We don't, in general, visually identify places by using their distinctive properties to pull them out from their backgrounds. Rather, our basic way of identifying places seems to be presupposed by any such procedure. So how can we make sense of the idea of a 'way' of identifying places, on a relational picture of experience? I want to end by amplifying on this problem.

5. The Spatiality of Experience

Visual experience seems to be deeply spatial. We see the spatial relations between the objects around us. The objects themselves fill space. And it's not just as if we see only spatial relations, or a space defined by space-filling objects. There seems to be a sense in which the space of visual experience is absolute. To use an example of Wittgenstein's

(1975), suppose you imagine a pitch-dark night, and the only things visible are two stars circling one another. The spatial relations between things in the scene are unchanging. But there's a manifest change of position here, in a space that isn't defined by the visible objects in it. As Wittgenstein put it, it's as if a set of axes flared up in the visual space and the stars were located with respect to these axes. Only, if there were such a set of axes, it would itself be oriented, as if in relation to a further, unseen set of axes. How should we characterize this spatiality of experience, the sense in which locations in your visual space seem to be identified as places in an absolute space? At this point we have two familiar positions, with two familiar sets of problems:

- (1) In representational terms: we characterize a frame of reference that is used to identify particular places and spatial relations between particulars.

- (2) In terms of 'mental paint', or, as Christopher Peacocke put it, 'sensations'. He said: 'we should not be afraid of acknowledging that there are sensational spaces, and their spatiality cannot be explained away as derivative or artificial. The sensational spaces are not, of course, populated by the sensations themselves. Rather, for each intrinsically spatial sense-modality, there is a space such that all experiences in that modality have intrinsic sensational properties characterized by reference to that sensational space. The sensational spaces of different modalities do not bear spatial relations to each other: it makes no sense to ask of a location in visual *sensational* space what spatial relation it bears to a location in tactual sensational space (e.g., the tactual location when a certain place on the neck is

touched). In this, the sensational spaces stand in sharp contrast with the space – everyday, public, physical space – which the representational content of experiences in different modalities concerns. It is the very same relation ‘closer to’, concerning the very same space, which enters the representational content as of someone walking closer to one and an auditory experience as of someone walking closer to one.’ (Peacocke 1983, 51-52).

At this point, there are by now familiar problems with both of these ideas:

- (1) If the notion of representational content is understood in causal-teleological terms, then the fact that it’s *experience* we’re dealing with seems not to have been recognised.
- (2) We seem to have no way of explaining the notion of a sensational space: we can’t do it ostensively, and we can’t regard it as merely a theoretical construct (though we can oscillate between the two).

So it seems worth exploring a relational picture of spatial experience: spatial experience is constituted by the space the experiencer is located in. Spatial experience is a relation to that space. But the external space does not in itself have any particular perspective or ‘frame-of-reference’ structure, so how can we regard the perspectival space of experience as constituted by it?

The problem is that the natural way to reinstate the idea of ‘axes’ in the perceptual field is to appeal to something representational, or at worst to an internal sensational space. But either of these appeals threatens to subvert the whole idea of a relational characterization of visual experience. The properties and objects seen now have to be located in this representational or sensational space, and that wrecks the idea of regarding them as located merely in the ordinary external space.

We can put the problem like this. When we think of visual experience as a relation that you bear to an external scene, we have to acknowledge that this is at least a three-place relation, between you, the external scene, and the point of view on that scene which you are occupying. But then, in characterizing the ‘point of view’, it’s not enough merely to specify a particular location that you are occupying. This has to be thought of as ‘the visual perspective from which you are experiencing the space’, bringing with it the effect of a set of axes that are being used to identify locations. But how we are to explain the sense in which locations are being identified in one ‘way’ rather than another is not obvious. I think there are a number of options you might explore here, but at the moment, I do not see that any of them are compelling.

REFERENCES

Beck, Jacob, and Keith Schneider. In press. 'Attention and Mental Primer'. *Mind and Language*.

Block, Ned. 2010. 'Attention and Mental Paint'. *Philosophical Issues* 20, 23–63.

Carrasco, M., S. Ling and S. Read. 2004. 'Attention Alters Appearance'. *Nature Neuroscience* 7, 308-313.

Griffiths, Thomas L., Falk Lieder and Noah D. Goodman. 2015. 'Rational Use of Cognitive Resources: Levels of Analysis Between the Computational and the Algorithmic'. *Topics in Cognitive Science* 7, 217-229.

Huang, Liqiang and Harold Pashler. 2007. 'A Boolean Map Theory of Visual Attention'. *Psychological Review* 114, 599-631.

Marr, David. 1982. *Vision*. San Francisco: W.H. Freeman.

Moll, Henrike and Andrew Meltzoff. 2011. 'Perspective-Taking and its Foundation in Joint Attention'. In Naomi Eilan, Hemdat Lerman, & Johannes Roessler (eds.),

Perception, Causation, and Objectivity. Issues in Philosophy and Psychology. Oxford: Oxford University Press.

Nagel, Thomas. 2002. 'What is it Like to be a Bat?'. In David Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings.* New York: Oxford University Press.

Peacocke, Christopher. 1983. *Sense and Content.* Oxford: Oxford University Press.

Wittgenstein, Ludwig. 1975. *Philosophical Remarks.* Edited by Rush Rhees; translated by Raymond Hargreaves and Roger White.