What is a State of Visual Perception?

Robert Hudson
Department of Philosophy
University of Saskatchewan
9 Campus Drive
Saskatoon, Saskatchewan S7N 5A5
Canada

Abstract

The goal of this paper is to re-assess the current state of debate on the issue of whether perceptual experiences states are disjunctive. According to disjunctivism, perceptual experiences contain as components the worldly objects they are about – thus, two phenomenologically identical perceptual experiences, one veridical and the other illusory, stand as different perceptual experiences since they have different objects. By comparison, the common-kind view of perception asserts that phenomenologically identical experiences, despite having different objects, are potentially the same perceptual experience. My strategy is to utilize some recent work in the empirical neuro-psychology of perception, work described in Peter Tse et al. (2005) using fMRI technology, that speaks on behalf of the common-kind view. Part of my task will be to defend a limited form of the mind/brain identity theory, talking into account discussions of the notion of a neural correlate of consciousness by Chalmers (2000) and Noë and Thompson (2004). On the basis of this limited identity theory, a philosophic basis for the common-kind view is established. In defending the philosophic significance of this basis, I examine some current criticisms of the common-kind view advanced by Campbell (2002), Martin (2002) and Millar (2007), criticisms that focus on the supposed transparency of perceptual experience. I close the paper by rejecting this supposed transparency and thus rejecting these criticisms.

1. Introduction

In November of 2005, the *Dartmouth News*, the mouthpiece of the Dartmouth College of Public Affairs, published an article with the following headline, "Finding the Mind's Eye: Dartmouth Professor Locates Areas of the Brain Involved in Visual Awareness". Peter Tse, the Dartmouth professor of concern, along with his colleagues apparently discovered using fMRI technology the place "where the [neural] correlates of visual consciousness lie" as regards a certain kind of visual illusion. Here, Tse's work is part of a larger project in current neurological science that searches for neural correlates for all sorts of psychological phenomena, such as word-finding failures, colour categories, associative memories, and many other things. It seems that such physical reductions of psychological states have now become all the rage, given the ease with which one can deploy functional Magnetic Resonance Imaging (fMRI) technology to isolate instances of neural processing in localized brain areas that respond to various environmental stimuli. Roughly, what neuroscientists are now able to do is study in relatively non-invasive fashion alert human subjects so as to locate in their brains those cortical areas that show increased (or decreased) levels of neural activity in response to various aspects of a subject's introspectively determined phenomenological state.

These increased or decreased levels of neural activity are identified by so-called BOLD (blood oxygen level dependent) signals; through magnetic resonance imaging one can discern changes in the oxygenation of blood, where (for instance) decreased levels of oxygenation indicate increased neural activity (since such activity uses up oxygen). Thus, for example, if it is the case that a subject notes the existence of some phenomenological state (such as seeing a black line, hearing a particular sound, experiencing a smell, and so on) when and only when some localized, BOLD-indicated, neural processing occurs, then the temptation is very strong to identify this phenomenological state with this localized processing. To take a particular case, where a phenomenological state is reported by a subject as "seeing a black line in a certain orientation", and where this state appears or vanishes with the presence or absence of some localized neural processing, it is tempting to say about this neural processing that it is the physical manifestation of the phenomenological state in question.

My plan in this paper is to explore some of this research into isolating the neural correlates of perceptual phenomena, specifically, the vision research performed by Tse and his collaborators alluded to above, and see where it takes us in understanding certain philosophical issues in the philosophy of perception. I will defend the claim that Tse et al.'s research casts light on the issue of the *disjunctivist's* view of visual perception, the view on which the identity of a visual state is constituted by the physical facts it is about. On this view, two observers could be in phenomenologically identical visual states but be having different visual experiences if one observer's visual state is veridical and directly reveals the facts depicted by the experience, whereas the other's visual state is not veridical and does not directly reveal the object so depicted. Following Campbell (2002), we might call the disjunctivist view the 'relational view': as Campbell describes the relational view, "the qualitative character of experience is constituted by the qualitative character of the scene perceived" (pp. 114-115). In contrast, the 'non-disjunctivist' or 'common-kind' view would assert that the two observers in the above scenario are in the same visual state since the type-identity of an experience is internal to the minds of the observers. Here, this visual state is common to the two observers, but still the observers could differ as regards the causal genesis of this state: for instance, on the causal theory of perception, the visual state of the observer with a veridical perception is caused by the object depicted by the visual state, whereas the visual state of the observer with a non-veridical perception is not so caused.

The conclusion I shall draw from Tse et al.'s work is that it provides positive support for the non-disjunctivist, common-kind view. The key to my argument will be the identification, cited above, between visual states and correlated, localized areas of neural firings. This identification is obviously highly controversial, and after presenting in section two some of the scientific details of Tse et al.'s research I spend a significant amount of time in section three defending this identification, focussing on the work of Chalmers (2000) who strives to vindicate the search for neural correlates of visual awareness, as well as the work of Noë and Thompson (2004) who forcefully criticize Chalmers' attempt. In section four I turn to examining and responding to a key criticism of the common-kind view, that it fails to account for the transparency of perceptual experience, i.e., that fact that we experience the world directly without cognisance of there being any intermediary between us and the world.

This criticism is advanced by Campbell (2002), who regards it as fatal for the common-kind view, as well as by Millar (2007), who responds on behalf of the common-kind view but whose response I find wanting. We also examine Martin (2002) who also objects to the common-kind view on the basis of the transparency problem. The approach I advocate as regards the transparency problem in section five is to dispute its cogency; pro-transparency theorists, I argue, are mistaken in how they understand visual experience. A more robust view of visual experience recognizes its fallibility and concept-dependency, and is sensitive to the fMRI evidence we have been considering in this paper, evidence that is naturally allied to a common-kind viewpoint. Having thus turned away the transparency problem, we are able to put the common-kind view of perception on firmer footing.

2. Tse et al. 2005.

The vision research performed by Peter Tse and his collaborators (described in Tse et al. 2005 and in Macknik 2006) uses functional magnetic resonance imaging (fMRI) to examine what happens in a human subject's brain when a subject visually experiences what is termed a 'dichoptic masking illusion'. The form of dichoptic masking illusion these scientists focus on is the 'Standing Wave of Invisibility' (SWI). In this illusion, a subject observes three vertical black bars standing together in parallel. The outer two bars flicker rapidly on and off in sync, whereas the middle bar flickers in alternation with the outer bars. The basis of the illusion is that as the outer bars (the 'masks') move closer towards the inner bar (the 'target'), the subject finds that at a certain point, one where the outer bars closely flank the inner bar but do not overlap with it, the inner bar becomes invisible with only the flickering outer bars remaining. This illusion is quite striking and can be seen on-line at Macknik (2007). The illusion is said to be 'dichoptic' in that in that the masks are viewed in the opposite eye to which the target is seen. For example, one sees the target flashing in the right eye and sees the masking bars flickering in the left eye. (A 'monoptic' mask would have both the masks and the target appear in the same eye.)

Using fMRI technology, Tse and his group generated images of those parts of the brain exhibiting increased or decreased levels of neural activity in response to, respectively, the absence or presence of masks. To understand their work, it is worthwhile to keeping in mind some of the current state of knowledge of how the brain processes visual inputs. As we all know, vision starts with light hitting the retina in our eyes.

From here a neuronal signal is generated that proceeds to the lateral geniculate nucleus in the middle of the head, through to the occipital lobe at the back of the head where it activates striate cortex, conventionally labeled 'V1'. The signal then proceeds forward through prestriate cortical areas V2, V3, V3A/B and V4, which are contained within the occipital lobe. The signal then exits the occipital lobe and extends next to either the inferotemporal cortex or parietal cortex (see Van Essen and Gallant 2001 for more on this two stream model). What Tse and his group are trying to determine is where in this cortical stream one finds neural firing that is correlated with the SWI masking illusion. Here is what they found. As Macknik (2006) describes the results, we isolated the parts of the brain that showed both an increase in BOLD signal when nonillusory visible targets were displayed and a decrease in BOLD signal when the same targets were rendered less visible by visual masking. Surprisingly, only areas within the occipital lobe showed differential activation between visible and invisible targets. (pp. 193-194) In other words, there were no areas *beyond* the occipital lobe that exhibited a differential rate of neural firing correlated with the presence or absence of masking. So where in the occipital lobe does this correlated firing occur? Macknik explains:

our results show that dichoptic masking does not correlate with visual awareness in area V1, but begins only downstream of area V2, within areas V3, V3A/B, V4 and later. (2006, p. 192) From here Tse et al. concludes that "visual areas beyond V2, but within the occipital lobe, are sufficient to maintain our awareness of simple targets" (2005, p. 17178). As Macknik (2006) puts it, "we have . . . narrowed the possible areas of the brain for localizing the circuits responsible for awareness of simple unattended targets to a region between V2 and the edge of the occipital lobe" (p. 203) It is necessary to emphasize here that the conclusion being drawn by these scientists is highly circumscribed: "care should be taken", Tse et al. (2005) note, "not to generalize these results to claims about the neural correlates of awareness of objects more complex than the simple targets used here" (p. 17183; see also Macknik 2006, p. 194). We can be even more precise: what has been discovered are only the neural correlates for a particular form of dichoptic masking. In addition, the scientists are not asserting that the neural activity they have discerned through fMRI imaging is identical to the phenomenological states being considered. They only suggest that this neural activity is 'responsible for' or 'maintains' one's visual awareness of 'simple targets'. It will be my contribution to tighten up the link between the relevant neural activity and the visual states of concern, to actually assert the identity of these things. Defending this assertion is the task of the next section.[1]

3. Identifying States of Neural Activity with States of Visual Awareness

As I suggested in the Introduction, there is a strong temptation with the visual masking experiments we are considering to identify the states of neural activity indicated by the fMRI images of BOLD responses with their correlative states of visual awareness. To make the point more graphic, imagine (hypothetically) a case where a subject in the SWI masking experiment can turn a knob moving the masks closer and closer to the target bar, and in turn can watch the target bar fade in and out of awareness, while contemporaneously looking at an fMRI monitor to see how his neuronal firing increases or decreases with the level of awareness. We can easily imagine the subject exclaiming, upon watching the correlation, "Wow! There's my awareness of the target bar fading in and out!" The legitimacy of the subject's pronouncement here can be attributed in part to the highly circumscribed nature of her experience: there is nothing else going on but the movement of the masks which leads to the changing appearance of the target bar. But could the vacillating BOLD response be due solely to the movements of the masks? Tse et al. conduct experiments which involve subjects observing only the moving masks, without the presence of the target bar, and the BOLD responses found in this case occurred only in non-occipital areas, as did BOLD signals responding only to the flashing of the target bar (see Tse et al. 2005, p. 17180 and Macknik 2006, p. 203). Thus, there seems to be, at the very least, a strong correlation between the fading in and out of the target bar and the BOLD responses Tse et al. have identified.

Chalmers (2000) provides a valuable analysis of when one might be said to have found a neural correlate of consciousness, and my plan here is to utilize his insights on behalf of the stronger claim that these correlates are identical to the states with which they are correlated. Of course, none of this will be convincing if one believes mental states and physical states to be fundamentally different kinds of things, a belief one might justify on the grounds that it is possible to conceive of mental states existing in the absence of physical states. But leaving aside such dualistic considerations, there are a number of obstacles to asserting the identity of mental and physical states, obstacles raised by Chalmers (though again raised by him only with regard to the correlation thesis).

Where we are considering token neural activity N (say, as revealed by a diminished BOLD response in a certain area of the occipital lobe) that is correlated with token visual state of awareness V (say, a subject's witnessing of the vanishing of a target bar in a dichoptic masking experiment), we ask, Is N *sufficient* for V? As Chalmers notes (2000, p. 12), in this sort of case one might take the brain as a whole as sufficient for V, which would not be too informative. Thus, he suggests we seek 'minimal sufficiency', that is, the smallest cortical area sufficient to maintain awareness (p. 13). But how small can we get here? And what about the possibility of there being more than one cortical area capable of leading to the same state of awareness? Chalmers indicates that these questions need to be answered empirically, but doesn't specify from where the empirical data will be drawn. Alas, we can now address this empirical issue: the fMRI approach we have been considering will outline the boundaries of the cortical area leading to the state of visual awareness – indeed, the boundaries can be seen in the form of fMRI images overlaying pictures of the cerebral cortex taken from different angles. In the Tse et al. article, for example, one can see pink-coloured areas overlaying the patchy green expanses of the cortex that "represent the cortical areas . . . [exhibiting] significant dichoptic masking and thus are potential candidates for maintaining awareness of simple targets" (2005, p. 17182). By this means, one could also easily see if there are multiple areas responsible for awareness by simply seeing if there are multiple, correlative 'pink' patches. If each of these areas correlates with a subject's witnessing dichoptic masking, even in the absence of the others, then each is sufficient (though not necessary) for awareness of this masking**.**

Of course, if one is claiming that neural activity is identical to a state of visual awareness, then such activity should not only be sufficient for this activity but also *necessary* for it – and as Chalmers (2000) rightly notes, "it might turn out that there is more than one neural correlate of a given conscious state [of visual awareness]" (p. 12). Where one is considering a token, conscious state present in the mind of some subject that is at one time correlated with a particular BOLD-indicated state of neural activity, it is at the limit of empirical discriminability to assess whether *this* conscious state could have occurred with some other state of neural activity. This is because we currently lack a thorough understanding of the token identity of states of visual awareness. It is seemingly an irresolvable question whether one is now in the presence of a singular, enduring state of awareness**,** or enjoying a series of type-identical, though distinct states of awareness.

The approach Tse et al. take on this issue is to consider solely the type-identity of a state of visual awareness: they examine a series of subjects with similar experiences of dichoptic masking and then assess the range of BOLD-indicated neural activity correlated with these experiences. From here the data is aggregated, removing the distinction between individuals. By this means Tse et al. hope to get a clearer picture of what type of neural activity correlates with the sort of dichoptic masking being investigated. So when we say that a certain kind of neural activity is necessary for the occurrence of dichoptic masking, we mean (taking Tse et al.'s lead) necessary for human individuals who are 'normal', which for Tse et al. means people who are "healthy" and have "normal depth perception and normal or corrected-to-normal visual acuity" (p. 17178). In these people, dichoptic masking produced by the experimental setup being considered is always correlated with neural activity in specific, localized areas in the occipital cortex. On this basis our proposal is to identify such dichoptic masking with this sort of localized neural activity.

Matters aren't quite as simple as this, though, since as Chalmers notes the correlation between states of awareness and neural states is environmentally dependent. For example, no one expects that if one were to lop out the occipital lobe of one of Tse et al.'s subjects and artificially recreate in the excised lobe the conditions it would have in a normal case of dichoptic masking, that one would have thereby a state of visual awareness of such masking (on this see Chalmers 2000, p. 14). Also, if one were to introduce some disruptive lesions, or expose the brains to some unusual stimulations, once again one might expect the correlations between the awareness and the neural activity to disappear. Chalmers (2000) spends a significant amount of time examining the issue of such non-normal states. Sometimes such interventions are designed to help identify the neural correlates of awareness, such as when awareness occurs despite the fact that a lesion has cut off the influence of a certain part of the brain, which part could then not be responsible for such awareness. The benefit of fMRI research to this extent is its non-invasive nature. It is certainly informative to investigate the properties of the brain under novel circumstances, but we might do well to understand the brain to begin with under normal operating circumstances, and fMRI is currently the closest we can get to this ideal. For this reason Tse et al. are able to say confidently say that they are investigating the responses of 'normal' subjects, as normalcy is not compromised by having their brains examined by an fMRI magnet.

But where we are considering the identity of states of neural activity with states of visual awareness, is it required that we defend this identity under all sorts of circumstances, and not just under 'normal' ones? If two things are identical, how could their identity be disrupted by what otherwise occurs in the external environment? In actual fact, I think it is fair to say that identity claims are often context dependent: for instance, in normal, modern commerce, a cheque is identical to a particular piece of embossed paper, but that is only the case in the context of the institution of cheque writing. Take away that context and the identity vanishes as well. Similarly, the Evening Star is identical to the Morning Star, but that it true only in the context of a certain linguistic usage; if the term 'Evening Star' had been used in a different sense, then the identity would likely have vanished. Thus, it is not at all unusual for identity claims to be dependent on the environmental context, which is exactly the case when we are identifying BOLD-indicated states of neural activity with states of visual awareness in experiments performed by Tse et al. These identities are asserted to hold only for normal individuals under the conditions specified by the researchers, a fact that in no way implicates the validity of these identities.

One aspect of Chalmers' understanding of how neural correlates are identified should be highlighted, since it raises an issue we are going to disagree with, an issue that forms part of a criticism raised against Chalmer's approach by Noë and Thompson (2004). On Chalmers' view, in seeking neural correlates to states of visual awareness "the crucial question is whether the representational content in the neural system matches up with the representational content in visual consciousness" (2000, p. 6), for once we have this match-up we can say that the neural system of concern is a neural correlate of the state of visual awareness. But from here Chalmers further claims that this approach to identifying neural correlates "requires that we have some way of defining the representational content of a neural system *independent* of the contents of consciousness" (p. 7; my italics). To this end Chalmers gives some examples of how one might determine the representational content of a neural state directly, with a 'receptive field' definition being for him the 'simplest' which, though 'crude', is "good enough for many purposes" (p. 7). A 'receptive field', in general terms, is an "area in which stimulation leads to [a] response of a particular sensory neuron" (Levine and Shefner 1991, p. 671, cited by Krantz, 2007). In essence we are examining how single neurons respond to sensory stimuli. Those stimuli that selectively cause a neuron to fire can be said to be represented by this firing. The idea is then to match-up the content of states of visual awareness with those sets of neurons whose receptive fields mirror this content. As Chalmers describes this process:

First, we need methods for determining the contents of conscious experience in a subject, presumably by indirect behavioral criteria or by first-person phenomenology. Second, we need methods to monitor neural states in a subject, and in particular to monitor neural representational contents [in terms of receptive field contents]. Then we need to perform experiments in a variety of situations to determine which neural systems correlate with conscious states and which do not. (p. 26)

Obviously this is going to be a very complicated process, with no guarantee of success. Moreover, as Noë and Thompson (2004) argue, given the large differences between receptive field content and experiential content, the best we could hope for with this approach is an *agreement*, not a *match* in content (p. 12), for receptive field content is extremely limited, falling short of what is needed to capture experiential states such as location in egocentric space, figure-ground awareness, representing one's 'point-of-view', and so on (pp. 14-17). Noë and Thompson's general point is that perceptual experience is extraordinarily complicated, particularly if one realizes that such experience involves a "temporally extended, active, and attentional encounter with the environment" where "the content of experience is brought forth or enacted by this activity" (p. 17). As such, reduction to receptive field content is far too simplistic and bound to fail in generating meaningful neural correlates.

Noë and Thompson's point is well-taken: it is presumptuous to say that we have settled the issue of finding neural correlates in any substantive way, give especially the nascent state of research in this area (and particularly the nascent stage of fMRI research). Still, they acknowledge that the search for isomorphisms between "features of experience and features of the minimal neural substrate" can be endorsed as a "methodological constraint" (p. 26), which is presently all we need to further warrant research in this area and to set the basis for a prospective philosophic understanding of conscious visual states . Moreover, I think we can readily free of ourselves of the constraint Chalmers had imposed on us, that being to find some way of "defining the representational content of a neural system *independent* of the contents of consciousness". There is no need at all to demand such independence.

Tse et al.'s approach expressly defines the neural correlate of a perceptual state as *that* area of the cortex exhibiting BOLD-indicated neural activity; that is, the representational content of *that* neural activity is precisely the representational content of the perceptual state with which it is correlated. Thus, there is no need to construct the representational content of neural activity 'from the ground up', as it were, utilizing receptive field content, projective fields, "complex correlations with environment, patterns of behavior, activity in other cells [and so on]" (p. 7); hence, there is no need to see if the representational content thus constructed matches what we find experientially. Through the fMRI research we are considering, this match in content is built in at the beginning – though of course in only a provisional way, allowing for modification and revision as experimenters learn more about the visual system. What this means for our work in this paper is that it is possible for us to consider seriously the identification of certain states of visual awareness (concerning dichoptic masking) with correlated, BOLD-indicated states of neural activity in advance of our understanding how it possible for neural states to have such representational content to begin with.

Once we assert the identity of states of neural activity with states of visual awareness, the passage to defending the common-kind view is straightforward. For suppose two subjects exhibit the same neural activity in the same cortical areas. It will follow that they are having identical visual experiences. Moreover, they will be having identical visual experiences, even if the causes of their neural activity are different: in one case the visual experience could be caused by some external object that has the qualities depicted by the subject's visual experience, and in another case the visual experience might be caused by something which did not have these qualities. That is, one experience could be veridical and the other not, and still we would be dealing with the same experiences. There are now three points to consider. First, this approach to the common-kind view avoids the usual strategy of defending it, which is to ground the commonness of various experiences in the indistinguishability of these experiences.

It has become increasingly common in the recent literature on this topic for philosophers to be skeptical about defending the common-kind view on the basis of such indistinguishability, citing either 1) the fallibility of subjects on the question of whether their experiences *really are* indistinguishable (see Noë and Thompson 2004, p. 23), or 2) following Putnam 1999, noting that the indistinguishability of experiences doesn't translate into their identity, in light of the failure of transitivity for the property of indistinguishability as compared to the transitivity of identity (see Noë and Thompson 2004, p. 23 and Crane 2006, §3.4.1). With the approach I am suggesting, such 'indistinguishability' problems are effectively bypassed. Secondly, one might press here that, despite the fact that the same neural state is present in two subjects and that these two subjects have the same state of visual awareness (whether or not they recognize this), that still doesn't show that they are having the same experience. Yet, to respond to this, I cannot find any motivation for asserting such a claim apart from a zealous desire to advocate disjunctivism. It seems farfetched to suppose that two people could be having the same state of visual awareness corresponding to the same neural state, and have nothing else different between them as regards their states of awareness and neurophysiology, but still be said to be having different experiences.

The only possible differences would be non-intrinsic, relational differences, but the class of possible relational differences is so enormous and unregulated that it would greatly obfuscate our understanding of states of visual awareness to incorporate them without independent motivation. Finally, we can adduce some further empirical facts at this point. We have at hand the conjecture, on the common-kind view, that a veridical visual experience and a phenomenologically identical hallucination could be identical to the same state of neural activity, and so could be said to be the same experience. Is there any empirical evidence that such a situation is in fact the case? At it turns out, there is some further fMRI research that can be consulted here. Ffytche et al. (1998) describe a particular neurological disorder, termed the 'Charles Bonnet syndrome', which leads to the occurrence of spontaneous visual perceptions "identical to those associated with normal seeing" (p. 738).

Their task in their research is to use fMRI technology to determine whether these hallucinations are correlated with certain neural events just as in the case of normal visual awareness, and *also* to determine whether genuine visions and corresponding, phenomenologically-identical hallucinations are correlated with *the same neural events*. As it turns out, fMRI data suggest a positive response to both research questions. Particularly, as regards the second question, ffytche et al. (1998) note that there is a striking correspondence between the hallucinatory experiences of each patient and the known functional anatomy of the occipital lobe.

In patients who hallucinated in color, activity was found in the fusiform gyrus in an area corresponding to the color center, area V4 . . . In the patient who hallucinated an unfamiliar face, additional activity was found in the left middle fusiform gyrus, an area that responds to unfamiliar face stimuli . . . In patients who hallucinated brickwork, fences and a map, activity was found around the collateral sulcus, an area that responds to visual textures. Finally, in the patient who hallucinated objects, activity was found in the middle fusiform gyrus, an area that responds to visually presented objects. These results are, to our knowledge, the first evidence of a correlation between the location of activity within specialized cortex and the contents of a hallucination. (p. 740)

What is being suggested here, on the basis of fMRI evidence, is that hallucinations that share the same phenomenology as veridical perceptions are associated with the same states of neural activity, and so count as the same experiences, in accordance with the common-kind view. Of course, the evidence at hand is not conclusive; but it is strongly suggestive, and warrants further empirical inquiry that could further enhance support for the common-kind view.

4. The Transparency Objection to the Common-Kind View

Often enough, the main sort of criticism of the common-kind view proceeds along the following lines: the common-kind view interposes an intermediary (i.e., a state of awareness) that lies between the world and a perceiver's mind, and the purported existence of this intermediary conflicts with how we introspect the nature of our perceptual experience. This introspected nature, it is said, presents us a view of the world unfettered by intermediaries, a view in which we directly apprehend objects in the world. This feature of (introspected) perceptual experience – that it lacks an intermediary – is called the 'transparency' of experience, and it is assumed by many philosophers that visual experience indeed has this feature (see, for instance, the advocacy of transparency in Martin 2002, with supportive citations from, among others, Tye 1992, p. 381, and Searle 1983, p. 388). With normal visual experience, it is claimed, we see the mundane objects of our quotidian existence perfectly directly, in all their glory, and not as a result of witnessing something else. Moreover, it is claimed that it is not a pretension on our part that we see the world this way, that it is not some sort of realist fantasy, but rather this directness is asserted to be quite genuine.

As Martin (2002) explains, the transparency of experience has been used to good effect in defeating sense-data theories of perception where sense data are asserted to be mind-dependent objects providing only indirect access to objects in the world. On this matter he quotes Putnam (1994) approvingly. Where we are considering sense-data theories, Putnam says, the traditional claim that we must conceive of our sensory experiences as *intermediaries* between us and the world has no sound arguments to support it, and, worse, makes it impossible to see how persons can be in genuine contact with a world at all. (Putnam 1994, p. 454, quoted in Martin 2002, p. 396; Putnam's italics)

Sense-data, so understood, are raw sensations and feels which the mind apprehends as such. Thus, for example, in seeing a red car one sees initially the redness of the chassis, the shape of the chassis, the blackness of the tires and their circularity, and so on, and then one's mind puts all these parts together to create the perception of a car. Here, Putnam is questioning whether, even if we assume that the mind is populated with sense-data and that in putting them together we can succeed in creating something car-like, one could really be said to be seeing the car itself. Fortunately, the form of common-kind view we are defending on the basis of fMRI data need not be committed to such a sense-data view, but can ally itself with a perspective on which the relevant intermediaries contain *intentional* content. On this approach, as Martin (2002) puts it, "the phenomenological content of [one's] experience is determined by how the experience represents the environment to be" (p. 385), as opposed to being determined by the intrinsic character of the experiences, as with the sense datum view.

As such, the experience can represent a state of affairs and lead an observer (who has the experience) to think that this state of affairs obtains, even though this state of affairs doesn't obtain. Such an approach is clearly in sync with the experimental work we just saw from ffytche et al. (1998), where evidence is presented that the same state of neural firing – i.e., the same state of visual awareness – can be associated either with a veridical experience or a hallucinatory one. Also, according to Martin (2002), the benefit of the intentional approach as compared to the sense-data view is that we restore the transparency of perceptual experience: as Martin expresses this point, with the intentionalist view when one is veridically perceiving, and there is an object for one to perceive, then that is the object of awareness, there is no other object acting as an intermediary. (2002, p. 397)

Yet not everyone agrees with Martin's assessment here. McDowell (1987) doubts that transparency (or 'openness', as it is also called) is gained via intentionalism – he believes we still have a form of intermediary (a point on which Millar 1996 offers a rebuttal; see Crane 2006, §3.3.3). Campbell (2002), too, doubts that intentionalism provides the right sort of passage to the 'real' world. As he expresses the criticism (2002, pp. 122-123), the common element view of perceptual experience (whether a sense-datum approach or an intentionalist approach) cannot explain how perceptual experience gives us thoughts about the world, whereas (he assumes) perceptual experience is assuredly the process by which we are able to generate such thoughts. On Campbell's view, intentionalism makes the regrettable move of building real world access into experience conceptually right from the first, thus demoting perceptual experience from its rightful place as the means by which one explains this access.

Millar (2007), for his part, expresses sympathy with the sort of concern raised by Campbell: he remarks that one motivation for approving of the disjunctive view is the idea that the intrinsic character of the experiences implicated in perception should make it evident why having such experiences amounts to cognitive contact with objects or facts. (p. 194)

Unfortunately, experiences on the common-kind view lack this 'intrinsic character' as they can occur in a situation where they misrepresent the world. Thus, Millar, as someone who wishes to support the common-kind view, feels the need to supplement this view to accommodate the insight that it is through experience that one gains 'cognitive contact' with the world. His strategy in achieving this accommodation is to associate with perceptual experiences 'recognitional capacities', where such capacities denote the ability of an observer to reliably pick out worldly objects (see Millar 2007, pp. 190-195). For example, suppose an observer has an experience of a sparrow. We grant on the common-kind view that the observer could have that very same experience, even if there were no sparrow there. We now add to this person's experience the fact that she would not form a belief that there is a sparrow there unless there really was a sparrow there – i.e., she has the appropriate recognitional capacity. Millar's assertion is that it is now the case that her experience really does provide her with cognitive contact with the world.

One might question Millar's strategy on the grounds that, unbeknownst to a perceiver, a recognitional capacity may or may not apply, dependent on the environmental circumstances at hand. Thus, a perceiver of a sparrow may or may not be in cognitive contact with a real sparrow, despite the appearance of such (which appearance is the same in both the veridical and illusory cases), dependent on the environmental circumstances and whether the appropriate recognitional capacity is in place. Either the perceiver is reliably seeing a real sparrow, or she isn't; either the perceiver has the relevant recognitional capacity, or she doesn't. There is no middle ground here on which one has the same recognitional capacity in both alternatives. Accordingly I submit that on Millar's view we have disjunctivism after all, one level up at the level of recognitional capacities. This is how he gains the cognitive contact for which he strives in sympathy with disjunctivism. But now we lose the sense in which we have a defense of the common kind view. For consider the purported common kind of thing that is shared in both veridical and illusory perceptual experiences.

The point of supposing that there is such a common kind is that one could be fooled about the veridicality of one's experience – that one could have the very same experience one has in the veridical case, but be in a state of illusion. But if one has the relevant recognitional capacity, then, if one is in a veridical state of perception, one *couldn't* be in a state of illusion; that's the point of saying that one has a recognitional capacity. Conversely, where one doesn't have the relevant recognitional capacity, then in the same veridical state one is in a perceptual state that *could* be illusory. But one and the same perceptual state can't both be 'possibly illusory' and 'not possibly illusory'. So the two perceptual states, one in which the recognitional capacities hold and one in which they don't, turn out to be distinct states, and so we have disjunctivism all over again, first at the level or recognitional capacities and then, as a result, at the level of perceptual experiences themselves.

Martin's view, again, is that intentionalism is a success as regards the issue of the transparency of perceptual experiences. However, he argues in his (2002) that our assessment is less positive for intentionalism when we consider sensory *imaging*. Sensory imaging involves imagining objects and their sensory properties (for example, with vision, sensory imaging becomes 'visualizing'). As regards sensory imaging, Martin advocates the 'Dependency Thesis' according to which to sensorily imagine X is to imagine sensorily *experiencing* X (2002, p. 404).

So, for instance, if I sensorily imagine in visual terms (i.e., visualize) a tree, then according to the Dependency Thesis this means I am imagining myself seeing (experiencing visually) this tree. Martin submits that with disjunctivism, as regards an imagined scene, one has an *actual* attitude toward that scene, that it contains an X which is a constituent of one's imagined experience of X (2002, p. 414). On the other hand, with intentionalism, when imagining a scene one's *imagined* attitude toward the scene is that it contains an X, but one's *actual* attitude to the scene is one of neutrality – either the X could be there, or not. One's actual attitude is one of neutrality since, in imagining an experience, one can imagine this experience occurring with the presence of the intentional object (i.e., it seems that the object is there) but without the presence of the actual object. Thus, whereas with perceptual experience intentionalism satisfies the transparency requirement (since the representational qualities of experience form no barrier to the world), things are different with (for example) visualizing, for in imagining a visual experience of X the representational features of the imagined experience become apparent in the context of one's *actual* attitude towards X and we thus lose the transparency of perception in the imagined situation. Arguably this describes the sort of situation neuroscientists such as Tse et al. find themselves in when imagining the visual experience of a research subject.

These scientists presumably bring to mind a situation where inside this subject's brain there is a localized set of neural firings (traced by a BOLD response), such firings being correlated with experiences possessed by the subject. As such, this set of neural firings could be said to represent whatever it is that the subject's experiences represent, and do so intentionally in accordance with the intentional view. Moreover, this set of neural firings along with the relevant associated experience could be described as an intermediary between the subject and the world (even if the experience has intentional qualities that make it appear to the subject as if she is viewing the world transparently), since such firings could obtain while fallibly representing the world. So we are in effect agreeing with Martin that intentionalism has the result he says it has: that in a case of sensory imagining one's *imagined* attitude toward the imagined state is transparent (perceivers 'pretend' to transparency, we might say), whereas one's *actual* attitude to the imagined state, informed as it is by fMRI evidence, fails to preserve transparency.

Accordingly, with the work of Campbell (2002), Martin (2002) and Millar (2007), the common-kind theorist faces something of a dilemma: either she asserts that perceptual experiences lack transparency (in alliance with what seems to be the current trend in empirical, neuro-psychological, fMRI research, but contra what seems to be the prevailing view of a majority of philosophers), or she affirms transparency contra this neuro-scientific trend, whilst being forced to entertain the risk (considering Millar 2007) of re-engaging the prospect of disjunctivism yet again, at the level of recognitional capacities (and so once more at the level of perceptual awareness). In the next and final section, I argue that the common kind theorist's best option is to take the former option, for there are a number of reasons to resist the transparency of perceptual experience.

5. Rejecting Transparency.

It is to be granted that perceptual experience has the appearance of being transparent. But that is not to say that perceptual experience *really does* have the property of being transparent. Martin, who suggests that the appearance of transparency is arrived at through introspection (2002, pp. 418-419), acknowledges the fallibility of introspection in this regard, so that for him it is at least possible that perceptual experience is not transparent, after all.

The key idea behind transparency is that an observer directly apprehends the outer world. But there is a certain degree of naivety about such a suggestion. Michael Tye asserts (in Tye 1992, as recounted in Martin 2002, p. 381) that when he experiences a blue ocean, he experiences the blue ocean itself and not a representation of such. Similarly, John Searle asserts (in Searle 1983, as recounted in Martin 2002, p. 388) that when he experiences a yellow station wagon, he experiences the yellow station wagon itself and not a representation of such. But neither philosopher knows whether the blue or yellow they see is the same as the blue or yellow you or I see, or whether the blue or yellow they see really are properties of outer things, or even whether outer objects have colour properties at all. Consider, too, that anyone with even a passing familiarity with the history of science knows how the realistic assurances of one generation of thinkers about the nature of the observable, physical world are often overturned with the next generation of thinkers. For this reason, too, there is prudence in asserting that, whenever we are working with convictions about the nature of the observable world, we are working in every case with fallible *representations* of the world, and not with the worldly objects themselves.

In this respect, a form of anti-realism about the observable world can come to seem quite enlightened. It is almost a mark of intellectual maturity to realize that in our conceptualizing of the world we are only working with representations of the world. For example, as parents and elders we often advise younger people that how they think the world to be is not necessarily how the world is. Similarly, as scientists in the wake of a scientific advance we often advise earlier scientists in the grip of the former view that how they think the world to be is not necessarily how the world is, and that they too are working with representations.

To be sure, I am here considering only conceptualizations of the world – not perceptions of the world. So perhaps perception provides us with somewhat more direct access to the world, and this is indeed the focus of the pronouncements of disjunctivists and direct realists. But it is a further aspect of our current philosophic enlightenment that we learn that perception is theory- or concept- laden, and that what we see is very much informed by what we believe about what we see. This enlightenment came as a result of reflection on the positivist's image of observation which viewed observation as involving the direct apprehension of properties of outer objects, apprehension unmediated by conceptualization. This concept-free view of perception was firmly and convincingly rejected by N. R. Hanson, Mary Hesse and others decades ago, and not much has changed since then to return to the earlier positivist conception of observation.

Further counsel in this respect is disseminated in numerous psychology and sociology classes around the world where students, almost by rote, learn that their beliefs and perceptions are often an artefact of various psychological and sociological determinants beyond their control. That is, not only are our perceptions concept-laden, but these concepts themselves are 'laden' on our minds by unconscious forces. This is certainly what we find with Tse et al.'s experiments described above. The subjects witnessed the fading in and out of a target bar, as synchronized with the correlative in and out movements of the masks – but of course there was really no fading in and out of the target bar. These subjects are suffering an illusion, one that persists even in the case where they know it is an illusion. Of course, the ways in which minds can create their own phenomena is well-catalogued. We add to this the fact that, where we have neural firing in a case of veridical perception, this neural firing could presumably be instigated artificially, and the subject would not know the difference in terms of their perceptual phenomenology (allowing, in deference to the views of Noë and Thompson, that such an achievement may take a formidable degree of technological advance). It is worthwhile, too, taking into account the work of ffytche et al. (1998) that supports the claim that the same neural firing occurs independently of whether what is seen is actually before the mind or is a form of hallucination. These forms of empirical considerations speak on behalf of being cautious about asserting a realism about what is seen in the world; there is good empirical support for asserting that we only see representations that may or may not be accurate, and that could be the *same* representations, whether or not they were accurate.

At this stage, Campbell's approach to defending disjunctivism in his (2002) arguably takes the right strategy. Campbell's idea, once more, is that for us to explain how we can conceptualize worldly objects, we need to presuppose some form of non-conceptual direct access to them. As a result, the intentionalist strategy is a non-starter since it builds conceptual details into perception from the beginning. Put another way, we need to have some sort of 'bare bones' access to the world to get the project of accessing the world off the ground. If perceptions are conceptually endowed all the way through, then our perceptions of the world are no better than elaborate conceptions and we thus lose the distinction between perception and conception. At this stage Campbell makes clear that, even if we have conceptions all the way down, the hazard here isn't skepticism: these conceptions might be perfectly accurate and the subject of knowledge.

The problem, rather, is that presumably it is our direct, unmediated, perceptual access to the world that ultimately allows us to think about the world, a fact lost on the common-kind view. For instance, how do we find out that physical objects are subject to the force of gravity? Roughly, the process is that we note that physical objects tend to fall to the ground when released and, after noting this phenomenon a number of times, we arrive at the idea that the object has a property called 'weight'. In effect, our experiences form the ground on which we can accurately conceptualize objects. Campbell's suggestion seems to be that, on such a basis, our concepts about the world are built up. And let us admit that there is a strong epistemological attractiveness to such a view: its builds into our experience of the world the fact that objects don't necessarily conform to our conceptualizations, which in turn explains how it is possible for us to genuinely learn things about the world. Put another way, our conceptualizations of objects can turn out to be mistaken because these objects aren't necessarily imbued with these conceptualizations.

But similar to my above critique of these sorts of direct access views, I submit that Campbell's approach ignores the fact that the world is very likely a more complicated place that our naive perception and direct realism makes it out to be. For instance, for Campbell's approach to work we need to track the same object over time and gain knowledge about it inductively. But when can we be said to be tracking the same object over time? Objects, over time, grow or decompose or alter chemically or have different uses or are added to other objects (to name just a few changes they can undergo); and it is a topic of pressing metaphysical concern to arrive at an accurate idea of what the identity conditions of objects are over time. This is a topic that preoccupies the investigative energies of countless philosophers, and before our (inductive) investigations into the nature of objects can begin it is surely the case that we need to advance some preliminary answers to these metaphysical questions.

As we suggested earlier, some of the answers will be psychologically hardwired into our minds, or could be due to sociological conventions. Alternatively, some of these answers might be the result of the imaginative inventions of scientists and philosophers. But however this goes, some answers will always be needed ahead of the inductive game, and these answers will speak directly to how we go about conceptualizing – that is, to how we go about representing – objects in the world. Thus, even if we grant the a priori attractiveness of Campbell's argumentative strategy, it is ultimately impracticable as regards how a mind can work, and particularly impracticable as regards the fallible nature of the human mind.

Let me be clear that the sorts of considerations I have been raising here are not at all novel. In essence, I am trotting out familiar constructivist and anti-realist conceptions of knowledge that are in no way indisputable. My point, rather, is that this anti-realism needs to be squarely faced if one is a disjunctivist, and is a natural ally to the common-kind view. This is because disjunctivism in embracing transparency is in effect embracing a direct realist view, and a direct realist view is, to be candid, an extreme position in the philosophy of science and knowledge. It is, in effect, ruled out by the fact that we do not have non-conceptual access to the world by means of perception. Thus, however attractive and reassuring the transparency thesis is, it surely must seem too easy for us to have such direct access to the physical world. Such direct access is betrayed by the continuing fallibility of human knowledge, and over-simplifies the mind/world relation as it is currently being shaped by empirical neuroscientists. It follows that, having blunted the force of the transparency objection, our defense of the common-kind view on the basis of the neuro-scientific evidence we have adduced is made stronger.

Acknowledgement

The author gratefully acknowledges support from the Social Sciences and Research Council of Canada, grant number 410-2006-214.

Endnote

1. The question may arise why visual masking experiments are particularly valuable for the purposes of identifying a neural correlate to visual awareness. Thought this is not in fact a question Tse et al. address, one could argue that in these masking experiments the phenomenon being investigated – the observed fading in and out of the target bar – is less influenced by confounding background conditions than would be the case if a subject were observing an objective black bar fading in and out, as the masking illusion occurs solely, as it were, in the mind of the subject (i.e., the target bar does not *really* fade in and out). In this respect, consider the following comments from Wiens (2006) relating to the use of masks in experiments investigating the perception of the emotional states of other people: because perception varies with small changes in stimulus parameters, it is important to keep testing conditions as consistent as possible over time and participants. However, variables such as background illumination, viewing angle, and dark adaption might be difficult to equate completely, particularly so in brain imaging research. Because these variables introduce undesired "noise" in the data and may confound effects from other variables of interest (e.g., state anxiety), these considerations suggest that perceptual awareness ought to be assessed rather than assumed for each participant. (p. 675) My conjecture, which I cannot defend here, is that dichoptic masking experiment described by Tse et al. is one way to avert the effects of such 'confounding' variables.

References

Campbell, J. (2002), *Reference and Consciousness*. Oxford: Clarendon Press.

Chalmers, D. (2000), "What is a Neural Correlate of Consciousness?", in *Neural Correlates of Consciousness: Conceptual and Empirical Questions*, edited by Thomas Metzinger. Cambridge Mass.: MIT Press. Page references are to the online version at http://consc.net/papers/ncc2.pdf, accessed November 16, 2007.

Crane, T. (2006), "The Problem of Perception", *The Stanford Encyclopedia of Philosophy (Winter 2006 Edition)*, Edward N. Zalta (ed.), http://plato.stanford.edu/archives/ win2006/entries/perception-problem/, accessed November 16, 2007.

Krantz, J. (2007), "Receptive Fields Tutorial", http://psych.hanover.edu/Krantz/receptive/index.html, accessed November 9, 2007.

Levine M.W. and Shefner, J.M. (1991). *Fundamentals of sensation and perception*, 2nd ed. Pacific Grove, CA: Brooks/Cole.

Macknik, S. (2006), "Visual masking approaches to visual awareness", in S. Martinez- Conde *et* al. (eds.), *Progress in Brain Research*, 155, pp. 177-215.

Macknik, S. (2007), Macknik Lab Page, http://macknik.neuralcorrelate.com/node/6, accessed November 9, 2007.

Martin, M. G. F. (2002), "The Transparency of Experience", *Mind and Language*, 17, pp.376-425.

McDowell, John, 1987, "Singular Thought and the Extent of Inner Space", in J. McDowell and P. Pettit (eds.), *Subject, Thought and Context*, Oxford: Oxford University Press.

Millar, Alan, 1996, "The Idea of Experience" *Proceedings of the Aristotelian Society* 97:75–90.

Millar, A. (2007), "What the Disjunctivist is Right About", *Philosophy and Phenomenological Research*, 74, pp. 176-198.

Noë, A. and E. Thompson, (2004), "Are There Neural Correlates of Consciousness", *Journal of Consciousness Studies*, 11, pp. 3–28.

Putnam, H. (1994), "Sense, Nonsense and the Senses", *Journal of Philosophy*, 91, pp. 445-517.

Putnam, H. (1999), *The Threefold Chord: Mind, Body, and World.* New York: ColumbiaUniversity Press.

Searle, J. (1983), *Intentionality*. Cambridge: Cambridge University Press.

Tse, P., S. Martinez-Conde, A. Schlegel, and S. Macknik (2005), "Visibility, visual awareness, and visual masking of simple unattended targets are confined to areas in the occipital cortex beyond human V1/V2", *Proceedings of the National Academy of Science*, vol. 102, pp. 17178-17183.

Tye, M. (1992), "Visual Qualia and Visual Content", in Crane (ed.), *The Contents of Experience*. Cambridge: Cambridge University Press.

Van Essen, D. C. and J. Gallant (2001), "Neural Mechanisms of Form and Motion Processing in the Primate Visual System", in W. Bechtel et al. (eds.), *Philosophy and the Neurosciences*, Blackwell: Malden, Mass., pp. 209-224.

Wiens, S. (2006), "Current concerns in visual masking", *Emotion*, 6, pp. 675-680.