King's Research Portal

DOI:

[10.1080/09672559.2016.1175102](https://doi.org/10.1080/09672559.2016.1175102)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Papineau, D. (2016). Against Representationalism (About Conscious Sensory Experience). *INTERNATIONAL JOURNAL OF PHILOSOPHICAL STUDIES*, 24(3), 324-347. <https://doi.org/10.1080/09672559.2016.1175102>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Against Representationalism (about Conscious Sensory Experience)

David Papineau

1 Introduction

It is very natural to suppose that conscious sensory experience is essentially representational. However this thought gives rise to any number of philosophical problems and confusions. I shall argue that it is quite mistaken. Conscious phenomena cannot be constructed out of representational materials.

There are two rather different motivations for the thesis that the conscious features of sensory experience are essentially representational—“representationalism” henceforth.¹ One comes from cognitive science, the other from phenomenological introspection.

A number of different lines of evidence have persuaded cognitive scientists that the neural processes underlying conscious sensory experience do not simply relay the structure of sensory stimulations impacting on our bodily peripheries, but rather construct hypothetical representations of distal features of our environment. This tradition goes back to Helmholtz in the nineteenth century and has received increasing support in recent decades. Much of the focus has been on vision, but the approach has been applied to other sensory modalities too.

This tradition in cognitive science leads naturally to a representationalist view. We need only identify the conscious features of sensory experience with the representational contents of the outputs of sensory processing. According to this line of thought, we feel consciously as we do when we see a table, say, because we are in a cerebral state which represents the presence of a table.

The phenomenological motivation for representationalism is different. Here we start, not with information about brain processing, but simply with the introspectible phenomenal structure of sensory experience. When we focus introspectively on our visual experience of a table, say, is it not obvious that our conscious state presents us with a mind-independent object of a certain shape, size, colour and distance? It seems built into the introspectible nature of our experience that it lays claim to the presence of this table. And isn't this just to say, so this thought goes, that our conscious sensory experience essentially represents such a table?

The two different motivations for representationalism are often found together in the same representationalist writers. But it is worth distinguishing them, because they raise different issues. In what follows, I shall respect the first motivation, to the extent of accepting the claims about sensory representation made by cognitive science—though I shall accommodate those claims without embracing representationalism as a metaphysical thesis. By contrast, I shall argue that the ideas about representation involved in the second phenomenological motivation rest on a series of mistakes.

2 Problems of Broadness

¹ I intend this term to cover not only views that identify the phenomenal properties of experiences (their “what-it’s-likeness”) with their representational properties (their accuracy conditions), but also views that take phenomenal properties to ground representational characters, in David Kaplan’s sense, which in turn yield accuracy conditions when combined with contexts. For further discussion of the latter option, see sections 6 and 7 below.

An initial indication that something is amiss with representationalism comes from representational externalism. There is good reason to suppose that representation is broad. But it would seem odd to hold that conscious experience is broad too.

Much recent discussion assumes that broadness is an internal issue for representationalism, and that the right response is somehow to refine the way in which representationalism is formulated. But in my view the issue is a symptom of a deeper malaise. The fault lies, not in the details of different versions of representationalism, but in the whole idea that sensory experience is intrinsically representational.

Representational externalism is the view that the truth conditions of representational mental states can depend, not just on their subjects' intrinsic properties, but also on facets of their environments, histories and social milieus. Truth conditions like this are called "broad" representational contents. Broadness occurs when two intrinsically identical subjects have corresponding mental states with different representational contents.

The problem that broadness raises for representationalism about conscious sensory experience should be clear. Representationalism wants to say that the conscious properties of sensory experiences consist in those experiences representing the world to be a certain way. But if two intrinsically identical individuals can have experiences with different truth conditions, because of different environments, histories or social milieus, then it would seem to follow that those individuals must be consciously different, in virtue of representing the world differently in sensory experience. But it would seem odd, to say the least, that two individuals should be consciously different, despite their intrinsic identity, because of differences in environment, history or social milieu.

3 Examples of Broadness

The idea of broad contents was introduced to philosophers in the 1960s and 1970s with a series of examples designed to show how the truth conditions of statements or beliefs can vary across intrinsically identical subjects. So, for example, Hilary Putnam's tale of twin water aimed to show how a statement's truth condition can depend on which liquid is present in a subject's environment. Similarly, Tyler Burge's story about Alf and arthritis argued that a belief's truth condition can depend on which ailment a subject's community refers to by a certain term. And before them Saul Kripke had in effect suggested that the truth condition of a statement involving a proper name can depend on the origin of the causal chain leading up to the subject's use of the name. (Putnam 1975, Burge 1979, Kripke 1980.)

Statements and beliefs are not sensory experiences. So perhaps there is room for defenders of representationalism to allow broadness for statements and beliefs, but to deny that it ever characterizes sensory experiences. It is not hard, however, to come up with plausible examples of sensory experiences with broad representational contents, analogous to beliefs with broad contents. Here are three cases featuring pairs of subjects who are intrinsically identical, yet whose corresponding sensory states intuitively represent different things.

Particular Objects Suppose I am viewing a yellow lemon; Jane is viewing a another yellow lemon that looks just the same; and John is being manipulated by scientists to have a sensory impression as of a yellow lemon even though no lemon is present at all. Let us suppose that what is going on inside our skins is just the same in all three cases: our visual systems are engaging in just the same processes, despite our differing external circumstances. Yet on the face of things the representational contents of our states are different. I am representing that this lemon is yellow; Jane is representing that a

different particular lemon is yellow; and John's sensory experience has such singular content at all, since there is no particular object in play in his case.

Inverted Earth. On Inverted Earth the sky is yellow and daffodils are blue, and so on. You are kidnapped, drugged and taken there, but while you are drugged you have inverting lenses inserted in your eyes so you don't notice the difference when you wake up. What is going on inside your skin when you look at the sky on Inverted Earth will be just the same as what happened inside your skin when you looked skywards on Earth. But on Earth your experience represented blueness, yet (once you have been on Inverted Earth for a while) your experience there arguably represents yellowness. (Block 1990.)

Cosmic Swampbrain. Suppose that a perfect duplicate of your brain coagulates by cosmic happenstance in interstellar space together with sustaining vat, and for some while engages in just the same neural processes as your brain. Your own conscious states represent features of your Earthly environment. But the Swampbrain's conscious states arguably represent nothing at all.

These examples bring out the awkward dilemma facing representationalists. Either they need to resist the natural broad interpretations which make the intrinsic identicals come out representationally different, or they have to embrace the implication that intrinsic identicals sometimes differ consciously. Neither horn seems attractive.

4 Broadness Analysed

Some philosophers are suspicious of broad contents. They are not persuaded by intuitive reactions to possible cases. In their view, there are strong theoretical reasons why truth conditions must be narrow (that is, determined by intrinsic properties of subjects). As a result, they hold that the kind of thought experiments outlined above are misleading, and the intuitive conclusions drawn from them confused.

It will be worth briefly examining the theoretical issues involved here, as it will help bring the phenomenon of representation into sharper focus.

One theoretical reason for thinking representation must be narrow relates to the phenomenological motive for representationalism aired in the Introduction above. Suppose that you think that the introspectible structure of conscious sensory experience is the fundamental source of representation. Then this itself provides reason to think that intrinsic identicals must always share representational contents. For it is natural to suppose that intrinsic identicals will always be consciously identical. And then, if representational content derives from conscious structure, it follows that intrinsic identicals will always end up representing the world the same way.

A rather different theoretical argument for narrowness relates to the explanation of action. A number of philosophers think that the essential features of mental representations are grounded in the way that they generate behaviour, from the inside, as it were. (Fodor 1908, Segal 2000.) What shows that I believe that *an apple is on the table*, say, rather than, say, *an apple is in the cupboard* is that I approach the table when I am hungry. But, if this is accepted, then broadness once more looks suspicious. Any two intrinsic identicals will surely behave the same way. So, if mental representation is constituted by its role in generating behaviour, it will make no sense to suppose that intrinsic identicals can have mental states with different representational contents.

However, it is not obvious that either of these motivations for narrowness is compelling. Note that both run counter to the natural thought that an essential feature of mental representation is the

way it relates subjects to the world around them and assists them in finding their way through it. Perspectives on representation that focus purely on the internal structure of consciousness, or on the way mental states causally prompt behaviour from within, seem in danger of leaving out this world-involving aspect of representation. After all, if our primary interest were in the internal structure of consciousness, or the internal springs of behaviour, it is not clear why we should think of mental states as ever laying claim to matters beyond the skull in the first place. Maybe broadness appears problematic if we think of mental representation as somehow limited to what goes on inside the skull. But once we think of representation in a world-involving way, then broadness can seem less puzzling.

There is a range of theories which seek to understand representation in terms of how subjects are embedded in their environments. Some such theories aim to analyse a cognitive state representing that *p* in terms of its normally being caused by p; others focus of the way such cognitive states will guide actions in a way appropriate to the presence of p; and there are also theories that invoke a mixture of these two ideas. This is not the place to assess the relative merits of these options.² For present purposes we need only observe that any such theory will render it quite unsurprising that representation should be broad. If the representational content of a cognitive state hinges on which features of the environment the subject is responding to, or orientating its behaviour to, then we should positively expect that intrinsically identical subjects embedded in different environments will be in states with different representational contents.

5 Options for Representationalists

Representationalists have two ways to go in the face of examples that purport to show that the same conscious state can represent different broad truth conditions in different intrinsically identical individuals. On the one hand, they can seek to resist the broadness, and argue that the states in question are better understood as sharing some common narrow truth condition. Alternatively, they can grasp the nettle and argue that the states in question are consciously different, in line with their differing broad truth conditions, despite the intrinsic identity of the individuals involved.

The former narrow strategy is adopted by effectively all representationalists in connection with ‘singular contents’; that is, with the putative contribution of particular objects to truth conditions, of the kind that is at issue with *Particular Objects*. Some representationalists attempt a similar narrow strategy with respect to the represented properties that are also at issue in *Inverted Earth* and *Cosmic Swampbrain*; but with such ‘general contents’ we also find representationalists who are prepared to allow that consciousness itself is sometimes broad.

This is not the place to explore all the moves that have been made in this area. From my own perspective, the whole need to make consciousness and representation line up is a problem of representationalism’s own making, and simply dissolves away once we drop the idea that conscious experience is intrinsically representational. In due course I shall give some indication of how that might work. But first it will be useful to run over a few aspects of the representationalist literature.

6 Singular Experiential Contents

There is a general reason why representationalists characteristically go narrow with respect to possible singular contents of experience. Representationalists typically adhere to the ‘common factor principle’: they hold that subjects who are perceiving veridically will share their conscious sensory properties with those who have matching illusions or hallucinations. But there will no singular contents shared across these three cases. The different experiencers in such matching cases

² See Papineau 2006 for a survey of such theories.

will be related to different particular objects, or to no particular object at all. So if representationalists want to equate the conscious property they take to be shared across these cases with some representational property, they need to find some non-singular content that the cases share.

Despite their best efforts, representationalists have not been particularly successful at locating such a shared singular content. A natural first thought is to appeal to a general existential content: that is, to take all the matching cases to be representing simply that *there is a lemon before me that is yellow*. But then there are objections involving cases where this existential claim is true by accident: imagine that there is indeed a yellow lemon in front of you, but this isn't the cause of your experience; there is in fact a screen between you and the lemon, and your experience is in fact produced by ingenious scientists stimulating your optic nerve. Intuitively, this is not a veridical sensory experience—we take the experience to be aiming to refer to some more directly related object than the lemon behind the screen, and so not to be vindicated merely by that obscured lemon being yellow.

This kind of example might suggest that we should build some causal requirement into the desired content, along the lines of *there is a lemon before me that is yellow and is causing this experience*. But this now threatens to make the content overly self-referential. It seems wrong to have experiences making meta-claims about their own aetiology. Surely it is possible to represent the world experientially without representing your own experiences.³

In response to these difficulties, many representationalists settle for “gappy contents”. The idea is that relevant sensory experiences don't in the end refer to particular objects, but merely present general properties as such. Somehow they answer to a local instantiation of yellowness, and perhaps local instantiation of lemonness, without any commitment to some specific object being supposed to possess these properties. On this analysis, the contents of sensory experience never themselves amount to conditions that can be true or false, just to something that would make up such a truth condition if combined with a particular object. (Tye 2014.)

To my mind, all these manoeuvrings around singular contents reflect badly on the overall representationalist programme. The initial representationalist idea was to equate the conscious properties of sensory experiences with their representational ones. But as soon as we focus on the singular dimension of representation, it quickly appears that conscious properties do not cut as fine as representational ones. I take this to cast doubt on the original representationalist idea. After all, it is not as if there is any independent reason to deny singular contents to sensory experiences, apart from the need to satisfy the theoretical demands of representationalism.

7 In Favour of Singular Contents

To bring out the naturalness of singular experiential contents, and the consequent ad hocness of the lengths to which representationalism is drawn on this issue, consider a slightly different pair of examples. I see my wife Rose come through the door. I have a doppelganger in Australia whose wife Ruby looks just the same, and he sees her come through the door. It seems entirely natural to say that my experience represents *Rose* and his represents *Ruby*. The fact that I and my doppelganger share the same conscious properties (I specified that our wives look just the same) seems no reason at all to deny that our states have these different representational contents—unless, that is, we are in the theoretical grip of representationalism.

³ See Soteriou 2000 for the problems facing attempts to read sense experiences as having general existential contents.

Perhaps we should not take it for granted that all sensory experiences have the same kind of singular contents as experiences of familiar reidentifiable objects. It is one thing to say that experiences can represent well-known objects like wives, another to say that they can represent randomly encountered everyday objects, like that particular lemon. Still, even if that were right, experience of familiar objects like spouses would already be an awkward thorn in the side of representationalism. And, in any case, I see no reason not to allow the same kind of singular contents to sensory experiences in general.

There is every ecological and biological reason to suppose that a primary function of sensory perception is to enable us to track and reidentify particular objects, the better to allow us to gear our actions to their particular idiosyncracies. This aspect of perception is highlighted when the objects in question are familiar and subject is already acquainted with a rich range of idiosyncracies. But I would say that the same point applies even in the case, say, where someone sees a random lemon to be yellow, and has yet acquire any specific information about it. The truth condition of their experiential states is still that the particular lemon in question is yellow. The experience of someone who is looking at a different lemon has a correspondingly different truth condition. And, in the hallucinatory case, where no particular is in play at all, no complete truth condition has been constituted (though it is here also true that this state would have a truth condition involving yellowness if it did refer to an object).

8 General Experiential Contents

So much for singular contents. What about general contents, like the colours at issue in *Inverted Earth*, or all the properties with respect to which *Cosmic Swampbrain* is arguably representationally inferior to its earthly counterpart?

Now, one option here would be once more to seek narrow contents that are shared across the intrinsic duplicates, by analogy with the representationalist moves just explored in connection with singular contents. But of course, once we come to general properties, the strategy of “gappy contents” is no longer available, since the problem is precisely that the counterparts are now arguably referring to different properties, not particulars. So the defenders of narrow general contents are driven back to ascribing existential general contents (there is some property that . . .), perhaps augmented with a causal requirement (. . . and is causally responsible for certain effects in me). The problems that faced these moves with singular contents now arise again. Moreover, in the case of the cosmic swampbrain in particular, there is the extra problem that every property that is not intrinsic to the subject will need this treatment, arguably including the notion of *cause*, which will make the cosmic swampbrain and its earthly counterparts end up with very thin shared experiential contents indeed.⁴

In the face of these difficulties, some representationalists are prepared to resort to the other option, and maintain that the intrinsic duplicates involved would not in fact be consciously identical. According to this line, your colour experiences will be phenomenologically altered once you have been on inverted earth for a while (even though everything inside your skin is still just as it was on earth). And, in similar spirit, why suppose that the cosmic swampbrain is conscious at all, given that it has never enjoyed interaction with any real environments? (Dretske 1995 1996, Tye 1995, Lycan 1996 2001, Byrne and Tye 2006).

I do not want to dismiss these moves out of hand. Still, many philosophers will find it hard to swallow the idea that two beings can end up consciously different solely because of their

⁴ See also footnote 12 below.

environmental differences, even though everything is the same inside their skin.⁵ Perhaps once more the moral to draw from the hard choices facing representationalists is that there is something wrong with their starting point.

9 Non-Relationism

In support of this diagnosis, let me now introduce an alternative way of understanding sensory experience that avoids all the problems of broadness while preserving much of the spirit of representationalism. This alternative will respect the scientific idea that conscious sensory experiences are the outputs of processes designed to construct hypothetical representations of distal features of our environment, but will do so without embracing the metaphysical tenets of representationalism.

Consider an analogy. Written sentences are the outputs of processes designed to produce representations that will convey information to readers. It does not follow that all the properties of sentences are essentially representational. Their typographical properties are not, for instance. It is entirely contingent that this arrangement of marks on paper means what it does. In different circumstances, just that arrangement of marks could easily have meant something different, or nothing at all.

I think the same about the relation between the conscious and representational properties of sensory experiences: the former stand to the latter just as the typographical properties of sentences stand to their representational contents. It is not essential to a given conscious experience that it stand for the truth condition it does. In different circumstances, just that conscious state could have had a different truth condition, or no truth condition at all.

In effect, this is to view the consciously constituted experience as the vehicle of representation, rather than the content. With sentences, we distinguish between vehicle properties—the shape and arrangement of the letters and so on—and the representational properties—that the sentences has a certain truth condition. So with sensory experiences. The conscious features of the experience are one thing, the experience having a truth condition is another.

Note how all the problems of broadness immediately disappear once we adopt this non-relationist perspective.⁶ Just as given typographically constituted sentences can have different meanings in different languages, so can a given consciously constituted sensory state stand for different truth conditions when embedded in different environments and histories. The same narrow vehicle can have different truth conditions, or none, depending on broad circumstances.⁷

⁵ “Swampman” thought experiments are sometimes invoked to lend intuitive support to the thesis that the mental states of intrinsic identicals must have the same representational contents. In my view, they do very little to support this intuition. In the present context, however, the issue is rather whether intrinsic identicals must share phenomenal properties, and here intuition seems much more definite, for what that is worth.

⁶ I originally defended this non-relationist position in Papineau 2014.

⁷ Philosophers of physicalist inclinations are likely to start asking at this point whether the brain vehicle properties that fix phenomenal character are supposed to be strictly physical properties or (narrow) “functionalist” ones. This is a serious question, but not one that we need answer here. The central point is that either way the relevant properties will be intrinsic non-representational ones. Indeed this central point could be agreed by a dualist who takes conscious states to be metaphysically independent of physical ones. (More generally, none of the arguments in this paper depends on physicalism.)

Thus with all our problem pairs. With *Particular Objects*, the same conscious vehicle refers to different particular objects in different normal cases, but to no object in the hallucinatory case. With *Inverted Earth*, the same conscious vehicle refers to blue on earth, but to yellow on inverted earth. With *Cosmic Swampbrain*, the same conscious vehicles have their normal referents in me, but no referents at all in swampbrain.

Of course, this allows that vehicle and representational properties may be tightly correlated within certain contexts. Once you fix a language, you fix a one-to-one correspondence between typographical and semantic properties (at least until ambiguity and synonymy intrude). Similarly, we are likely to find one-to-one correspondences between the conscious and representational properties of certain sensory experiences within biological species, say, or within individuals, or within individuals at given times.

10 Transparency

Given how easily non-relationism by-passes all the problems of broadness, it is surprising that it is almost entirely absent from the contemporary philosophical literature on perception. Introductions to the area will typically start with a brief mention of sense datum theory, and then quickly move on to the debate between representationalism and direct realist disjunctivism, without any suggestion that non-relationism might be a serious option.

One explicit reason sometimes offered for dismissing non-relationism is the “transparency of experience”. Imagine that you are looking at some visible scene—some fruit on a table say. Now try to turn your attention from the features of the fruit to the conscious features of your visible experience. All that will happen is that you will stare harder at the fruit and their properties, and not instead at some supposed realm of inner experience. A number of philosophers take this to argue that the properties present in your experience are ordinary properties of physical objects, like the shape and colour of the fruit, and not some special range of private non-relational conscious properties possessed by subjects rather than physical objects. (Harman 1990, Tye 2002.)

How exactly is this argument supposed to work? We can focus things by adopting a useful convention due to Christopher Peacocke. Let us refer to the conscious properties that subjects instantiate when they have sensory experience as properties*. So for example, subjects will instantiate redness*, squareness* and so on, when they see objects that are red, square and so forth. The transparency argument is then supposed to show that the properties we encounter directly in experience are properties like redness and squareness, not redness* and squareness*, as non-relationism would have it. (Peacocke 1983.)

Now, as we shall see in a minute, talk about properties being “in” experience needs to be treated with care, but let us go along with this way of talking for the moment, and moreover let us allow that non-relationism implies that the only properties we encounter directly “in” experience are properties*.

Why now is the transparency argument supposed to undermine non-relationism? At bottom the transparency argument hinges on the observation that when you try to shift your attention from the properties of physical objects to the properties of your experience, your visual phenomenology remains unchanged. But, put like this, it seems that the argument should be consistent with pretty much any account of the metaphysical nature of conscious experience. On the sense datum theory, this nature consists in my relation to some sense datum and its properties; on the representationalist theory, it consist in my relation to a representational content; on direct realist disjunctivism, it consists (at least in the veridical case) in my relation to the perceived fact itself; and

on the non-relationist view, it consists in my instantiating some intrinsic non-relational non-representational property.

On any of these accounts of conscious sensory experiences, why shouldn't my experiences remain unaffected when I "turn my attention" from their physical objects to the experiences themselves? I take it that such attentional shifts are cognitive acts, and as such there seems no obvious reason why they should have any impact at all on the mechanisms responsible for my sensory state when I am looking, say, at some fruit with my eyes open. In general we don't expect occurrent cognitive activities to alter our perceptual states, and it is not clear why we should do so here.

Perhaps the transparency argument would be a good argument against theories that take conscious sensory experience to involve "qualia" in addition to having constitutive representational properties. (For example, I take Block 2004 and Peacocke 2008 to endorse such theories.) On a portmanteau view like this, an experience of a square physical object, say, could have a squareness* property, say, due to representing the square from a certain perspective, in addition to the conscious representational property of representing the object itself to be objectively square. A view like this would indeed seem to be in the transparency argument's line of fire. Now there are two sets of properties "in" the experience, and we ought arguably to be able to shift attention from one set to the other.

But the non-relationist view I am proposing does not have this portmanteau structure. The idea isn't that somehow both the qualitative "mental paint" and the represented objective properties are "in" the experience. Rather my view is that our conscious experience is all paint, and any representational or represented features are quite external to our consciousness. So from my point of view there is no reason to expect that that we ought somehow to be able switch attention away from the other properties "in" our experience and towards the qualia. The qualia are all that were there in the first place—so the whole idea of turning away from the other features of experience and towards them doesn't get off the ground.

11 Sensory "Awareness"

Still, even if the transparency argument doesn't knock out non-relationism, doesn't it highlight its unattractiveness? Do we really want to hold that conscious experiences are constituted entirely by intrinsic non-relational qualia, and that there is no sense at all in which the properties of objects themselves are ever "in" our experience? It is certainly a natural thought, when we reflect on our conscious sensory experience, to suppose that the objects and properties that we are perceiving are somehow "in" our experience.

Well, I agree that this is a natural thought, and in the final sections of this paper will explain why. But I think it must be resisted. In the next few sections I shall explain how representationalism gets itself into a nasty tangle when it tries to accommodate this thought. But first, in this section, it will be useful briefly to make clear how the non-relationist view is at least consistent in denying that the properties of objects enter into our experience, even though this may seem initially unnatural.

You might think that non-relationism would be committed to denying that we are ever "aware of" ordinary properties of objects. And that would seem absurd. Surely I can be aware of the colour of a lemon when I look at it in good lighting?

However, I do not take non-relationism to have any such implication. I take it that we are "aware of" things when we are in mental states that represent them. In this sense, it is always physical objects

and their properties that my sensory states make me “aware of”. My sensory state represents the colour of the lemon, and thereby makes me aware of it.

My sensory state itself has a conscious property, yellowness*, which is distinct from the yellowness of the lemon. It is this property* that is conscious, not the yellowness itself. I become aware of the yellowness of the lemon by instantiating yellowness*. But I am not, in the normal course of events, aware of the yellowness*. The yellowness* is conscious, but as long as my mental states are focused on the lemon rather than my experience, I will not be aware of my conscious property.

Of course, I may sometimes introspect, or otherwise think about the conscious sensory properties that I currently possess. And this will make me “aware of” my sensory properties themselves, as well as of any physical properties that I am currently perceiving. But note that in this case it still won’t be my sensory experience that makes me “aware of” my conscious sensory properties, but some further cognitive state that is about those properties. The sensory state will still be about physical objects and their properties.

12 The Properties of Experience

Perhaps the non-relationist position can be cogently articulated. But many will still feel that it flies in the face of good sense.

Suppose you are looking at a bright yellow lemon. Now think of the yellowish property that you know to be present when you are introspectively aware of the nature your experience. Surely we would like to think of this property as just the same yellowness that lemons often possess, and not as some mental symbol yellowness* that bears no more relation to that property than the word “yellow” does.

Plenty of philosophers think that deliverances of introspection are unequivocal on this issue. For example, Gilbert Harman insists that, if we try to introspect a visual experience of a tree, we will find that

“... the only features there to turn your attention to will be features of the presented tree” (Harman 1990 39).

And Michael Tye, in similar vein, describing a visual experience of the Pacific Ocean, tells us that

“... what I found so pleasing in the above instance, what I was focusing on, as it were, were a certain shade and intensity of the colour blue . . .” (Tye 2002 448).

Still, while it may be initially plausible that introspection relates us directly to ordinary properties of physical objects, this intuitive idea conceals a number of hidden difficulties. It is not at all clear that representationalists have any defensible explanation of how ordinary properties of objects can be “present in” our experiences.

Note for a start that, however this is supposed to work, it is presumably not via the ordinary properties of objects being instantiated when we have experiences. As noted earlier, representationalists are common factor theorists, taking the same conscious properties to be present when I am mistakenly seeing a green lemon to be yellow as when I am veridically perceiving a yellow one. In both cases I have the property of representing the lemon to be yellow, and the conscious nature of my experience is constituted by this common fact. So now focus on the case where I have this conscious experience, yet the lemon is green. Yellowness is still supposed

somehow to be “present in” my experience. But clearly it is not there in virtue of being instantiated. Nothing in this case instantiates yellowness. The lemon is not yellow, I am not yellow, and none of my mental states is yellow.

The idea, presumably, is that the properties get into our experience, not by being instantiated, but by being represented. In experience we represent the uninstantiated property of yellowness, and this somehow constitutes the conscious state we are in when so experiencing. Some representationalists are quite explicit on this matter.

Thus Fred Dretske:

“In hallucinating pink rats we are aware of something—the properties, pink and rat-shaped that something is represented as having—but we are not aware of any object that has these properties—a pink, rat-shaped, object. We are aware of pure universals, uninstantiated properties” (Dretske 2003 73).

And Michael Tye again:

“Along with (most) other representationalists, I am happy to say that, in the hallucinatory case, the perceiver is conscious of an un-instantiated property. This seems to me to be part of naïve commonsense” (Tye 2014 304).

13 Comparison with Direct Realist Disjunctivism

I must say that I find the representationalist view hard to understand at this point. Uninstantiated properties are not located within space and time. It seems strange that a mental relation to such an abstract entity could constitute the phenomenal character of my experience. My conscious states are here-and-now, local, the kind of things that have causes and effects. How could a mental relation to an uninstantiated universal constitute this kind of state?

It is worth briefly comparing representationalism with direct realist disjunctivism on this point. Disjunctivists also hold that ordinary physical properties can be constitutive parts of our conscious experiences. But in their view this always depends on the property in question being instantiated. When we have a veridical perception, our conscious state involves the fact perceived: when we see a yellow lemon, the actual yellowness of the lemon plays a role in fixing our conscious properties. Of course, disjunctivists cannot say this about illusions or hallucinations of yellow lemons, precisely because yellowness is not instantiated in those cases. But that is all right for them, as they are not common factor theorists, and take the conscious properties in those cases to be different.

Now, you might well be uneasy about the disjunctivist suggestion that my conscious state in the veridical case depends on matters outside my skin. But the representationalist account of how ordinary properties get “into” our experiences strikes me as much more puzzling than that. It is one thing for yellowness to contribute to the conscious character of my experience in virtue of being instantiated before my eyes. It would be another for it somehow to enter into my consciousness even though nothing in my field of view or anywhere else nearby is actually yellow.

Somehow the representationalists are thinking that the yellowness is “present in” my experience, not because it is instantiated there, but in some other way. My mind reaches out and grasps the property yellowness itself, the property that is sometimes instantiated, in lemons and other things, but is not, let us take it, currently being instantiated in or around me—and this grasping is somehow

supposed to be responsible for the distinctive feel that characterizes our visual experiences as of yellow things. As I said, I find this suggestion difficult to understand.

14 Representational Properties

Perhaps I am in danger of proving too much. I am expressing scepticism about relations between thinking subjects and uninstantiated properties. But there are independent reasons for recognizing some such relations. After all, sensory experiences and other mental states do in fact represent possible states of affairs, and such representational facts do create relationships between thinking subjects and uninstantiated properties. Unless I am prepared to eliminate representational facts, this then argues that mental relations to uninstantiated properties must be legitimate after all.

This is a reasonable point. I certainly do not want to eliminate representation. Representational facts play an important role in the unfolding of the natural world. And I agree that representational facts involve relations to uninstantiated properties. However, I don't think that this is of any real help to representationalists about sensory experience. When we unravel exactly what kind of relations to uninstantiated properties representational facts commit us to, we will see that they are quite unsuitable to serve as the basis for conscious properties.

Let us assume that when someone represents that p in some mode (cognition, visual perception, audition . . .), this will involve their being in some state S that represents that p .

Further, let us assume that

S represents that p

can be equated with

S is true if and only if p , in virtue of the way that S operates as a representation⁸

and that this in turn can be equated with

S will fulfil its aim if and only if p , in virtue of the way that S operates as a representation.

I myself am inclined to understand "fulfil its aim" in this context in terms of such naturalistic categories as causation and biological design, but I intend this formulation to be neutral between different accounts of the nature of representation: after all, any such account will presumably agree that the essential feature of representational states is that they answer to some condition for their truth, and moreover that the aim of representations is, in some sense or other, to be true.

Now, if this much is agreed, then it follows from subjects representing that they will be related to properties. When a subject represents that a given lemon is yellow, for example, that subject is in a state that will fulfil its aim if and only if the lemon in question has the property of yellowness. And this in itself is a relation between the subject and yellowness.

⁸ In Papineau 2014 I advocated this neo-Davidsonian way of formulating representational facts as a means of avoiding existential commitment to propositions as abstract set-theoretical objects. But since then I have been persuaded that propositions need not be thought of in this way, and in truth are no more ontologically objectionable than properties (indeed we can think of them as 0-adic properties of the world). (See Rumfitt 2014.) In the present paper the point of portraying representational facts in neo-Davidsonian terms is merely to bring out their complex conditional nature.

But note how indirect and conditional a relation this is. In particular, note that a subject can bear this relation to yellowness even if yellowness is not instantiated anywhere in the subject's vicinity, as when the lemon being represented is not in fact yellow. The way that the subject's state is hooked up to yellowness, so to speak, does not demand that yellowness be presently instantiated.⁹ It only imposes the conditional requirement that the state will do its job if and only if the lemon instantiates yellowness—that is, either it does its job and the lemon is yellow or it doesn't do its job and the lemon is not yellow. And in some cases—where the representation is false—this disjunction will be made true by the latter disjunct, and no actual instance of yellowness will currently be in play at all.¹⁰

Given this, it seems very strange to hold that the representational relation to yellowness can account for the “presence” of yellowness, that very property that some surfaces possess, in our conscious experiences. Defenders of standard representationalism are committed to this (“... only features *there* to turn your attention to...” Harman *ibid*, “It seems to me that what I found so pleasing *in* the above instance...” Tye *ibid*, my italics in both cases), but the nature of the representational relation, once clearly spelt out, does nothing to substantiate the thought that the properties of objects are to be found “in” our conscious experience.

Consider an analogy. I harbour certain antibodies X whose job is to protect me against some antigen Y. They will fulfil their aim if and only if they repel an infection by Y. As it happens, I am not currently infected by Y, and so antibody X isn't fulfilling its aim. It take it that nobody would want to say on this account that nevertheless the antigen Y is currently “present” in me, in virtue of my harbouring X, whose aim involves Y. Yet this is effectively what representationalists say about represented properties. The represented property Y is “present” in my consciousness, in virtue of my harbouring S, whose aim involves Y, even when that aim isn't being fulfilled.

We have been considering the suggestion that representationalism about sensory experience is preferable to non-relationism because it respects the intuitive thought that in conscious experience we are acquainted with ordinary properties of physical objects, like colours and shapes, and not just with properties*. But this suggestion has not stood up to examination. There is no good way to make sense of the idea that ordinary properties of objects are somehow present in conscious sensory experiences.

So on this score, representationalism turns out to fare no better than non-relationism. It offers no real alternative to the view that the only properties of conscious experience with which we can make introspective contact are properties*, intrinsic properties of subjects that have no essential connection with the objectual properties that they contingently represent.¹¹ Moreover, given that non-relationism also avoids all the difficulties that broadness poses for representationalism, we would seem to have ample reason to prefer it.

⁹ Indeed it is doubtful that representing a property requires that the property ever be instantiated. Perhaps I am mistakenly representing the lemon to have a particular shade of yellow that no object has or will ever possessed. Of course, there are specific questions about how mental states can get to refer to such never-instantiated properties, but I take them to be answerable.

¹⁰ I haven't forgotten that representationalists take conscious sensory experience to represent “gappily”. However this only makes the relationship between a subject and any sensorily represented property Q even more indirect: the subject houses a mental predicate which, if combined with a mental name of a particular, would yield a mental state which would be true iff p... where p involves Q.

¹¹ It is of course consistent with non-relationism that we normally refer to experiential properties indirectly, by invoking their contingent properties of representing certain objectual properties (as in “an experience of yellowness”). Reference via contingent description is a common enough linguistic phenomenon.

15 Phenomenal Intentionality

Let me now return to the phenomenological motive for representationalism mentioned at the beginning of this paper. This appealed to the idea that introspection can show us directly that conscious sensory experience is representational. Many philosophers who are moved by this thought side with non-relationism in holding that intrinsic properties of subjects suffice to fix their conscious properties. Where they differ is in holding that these intrinsic properties are by their nature representational: conscious sensory states are not like the typographical words that we use to write English sentences, which could well have meant other things, or nothing at all; rather, introspection shows that our conscious sensory states necessarily have a definite representational content, necessarily represent the world to be thus-and-so.

Let us call this position “phenomenal intentionalism”, following Uriah Kriegel (2013). Phenomenal intentionalists typically point to certain introspectible feature of sensory consciousness to support their contention that it is intrinsically representational. I am happy to agree that sensory consciousness displays the features in question. But I do not accept that these features suffice to establish representationalism.

In the first instance, the relevant features consist of certain constancies that are displayed by interlinked sensory experiences as we move through time and space. As I move my head, or walk around, or stand up and sit down, my successive sensory experiences will have a number of salient common elements, corresponding to the ordinary physical objects in my environment, such as chairs, table, trees, people and so on. What is more, the relationships between my successive experiences will mean that these common elements maintain a constant position in my visual space (or a continuous trajectory in those cases where the corresponding objects are moving). In addition, my sensory experience will contain constant elements corresponding to various properties of the relevant objects, including their shapes, colours, facial characteristics and so on. (See Farkas 2013, Masrour 2013.)

It is no doubt these structural feature in sensory experience that makes it so natural to suppose that properties of ordinary physical objects can be found “within” experience. But, as we saw earlier, there are fundamental difficulties facing any representationalist who wants to understand things in this way. The alternative is to take the constancies found within sensory experience to be intrinsic features of experience itself. There may be a genuine chair-ish entity in my experience all right, in the sense of a sensory item that maintain its visual position, shape and colour, even as I move around, shift perspective and undergo changes in illumination. And, given such structural feature of experience, we might usefully talk of “phenomenal objects” and their properties, and even acknowledge that they display a kind of “mind-independence”, in that they maintain certain constancies even as we walk around and bob up and down. But nothing in this requires us to think of these objects and properties as anything more than modulations of the intrinsic structure of experience.

Sometimes philosophers speak of experience being “intentional” rather than “representational”. It is not always clear what this commits them to. If all they mean is that our sensory conscious experience contains “phenomenal objects”, in the sense just outlined, then I am quite happy to agree that sensory experience is “intentional”. There is no doubt that sensory experience has the internal structural features in question. But if it is supposed to be part of sensory “intentionality” that sensory experiences have essential correctness conditions, and thereby lay claim to the world being a certain way, then I deny that sensory experiences are intentional.

Representation, as I am understanding it, requires a mental state to lay claim to something other than itself. Something beyond the state is required for the state to be true. Representation requires the representer to reach out beyond itself, so to speak, in an attempt to hook up with some putative fact. It is not immediately obvious how the mere presence of phenomenal objects in our conscious sensory states could bring this about. Those objects are intrinsic features of conscious experiences, features that the experiences have in themselves, independently of anything else. It is difficult to see how such features on their own could ensure any representational powers.

16 Checking for Accuracy

Even so, some philosophers are explicit in maintaining that the intrinsic features of sensory consciousness suffice to determine everything needed for representation. Terence Horgan and John Tienson, in their influential paper “The Intentionality of Phenomenology and the Phenomenology of Intentionality” 2002, argue that the experiences of any two “phenomenal duplicates” will have the same truth conditions, independently of their environments, histories or anything else. (“Phenomenal duplicates” are beings who are consciously the same; Horgan and Tienson agree that my cosmic swampbrain counterpart, for instance, will be a phenomenal duplicate of me.)

Thus Horgan and Tienson:

“Consider any creature who is a complete phenomenal duplicate of yourself—its mental life is phenomenally exactly like yours. Assume nothing *else* about this creature. . . . suppose that you have the experience of seeing a picture hanging crooked. Each of your phenomenal duplicates has a phenomenally identical experience. Some of these experiences will be accurate and some will be inaccurate. . . . Thus, the sensory-phenomenal experience, by itself, determines conditions of accuracy: i.e., a class of ways the environment must be in order for the experience to be accurate. In order for such an experience to be accurate, there must be a picture before oneself, and it must be crooked” (2002 225).

I see no reason to accept this. From my perspective, conscious sensory experiences only represent contingently. Whether a given conscious experience represents a picture, or something else, or nothing at all, depends on factors beyond itself, such as historical correlations to feature of an environment, and is not fixed by its phenomenal nature.

Horgan and Tienson offer an immediate argument for their view.

“That these phenomenally identical experiences all have the same truth conditions is reflected in the fact that each of the experiences is subject in the same way to investigation as to whether it is accurate. For example, you and your phenomenal duplicate each might have the experience of seeming to oneself to be testing one’s perceptual experience for accuracy by making measurements or using a level. You and your phenomenal duplicate each might have the subsequent experience of seeming to oneself to discover that the picture merely appears to be crooked because of irregularities of the wall, or tricks of light” (226).

However, this argument does not serve. I agree that my phenomenal duplicates will go through the same motions, so to speak, in checking their experiences for veridicality (though of course my cosmic swampbrain counterpart won’t literally go through any motions, as opposed to initiating motor signals that terminate at its cerebral boundaries). But this is no reason to suppose that my duplicates’ states represent, as opposed to accepting that my duplicates think that their states represent.

In my view, my sensory states do not represent essentially, but they certainly represent contingently, and it does not require too much sophistication on my part to figure this out. All I need to do is to reflect on such facts as that: any given type of conscious sensory state will normally be caused by a given type of fact in my immediate environment, and will incline me to behave in ways appropriate to that fact; while at the same time that type of state will occasionally be produced in the absence of the relevant fact, but even then will still cause me to behave in the same way. A few simple considerations like these seem quite enough to lead me to regard my experiences as representing the possible facts that they stand proxy for—and no doubt on occasion to wonder whether they are representing accurately, and to take steps to check this.

And, if I can come to think that my states represent, and as a result be moved on occasion to check them for accuracy, then so can my phenomenal duplicates, including my cosmic swampbrain counterpart. After all, that counterpart is intrinsically identical to me, and so will have states corresponding to my belief that my sensory experiences represent, and to my intention to check whether the picture is indeed crooked, and so on. Of course, by my lights it isn't true that the sensory states of my swampbrain counterpart represent, and so there isn't any point in its trying to check them for accuracy. But that does not alter the fact that it will go through the same mental motions as I do, even though its sensory states do not represent anything.

To digress for a moment, there is of course a question of whether the cosmic swampbrain's cognitive states really represent anything, analogous to the issue of whether its sensory states represent anything. Let us assume that cognitive states like occurrent thoughts, beliefs, and so on, have a phenomenology—that is, that there are conscious properties that we instantiate when are in such states. This is of course contentious, but it is something that will be agreed by most phenomenal intentionalists. Now, are these conscious cognitive properties essentially representational? This is just the same question that we have been asking about conscious sensory properties, and I want to give just the same answer. In the actual world, these conscious cognitive states do indeed represent, but only contingently, in virtue of their environmental and historical embedding.

So from my point of view the cosmic swampbrain's cognitive states won't actually represent. They will feel just like my cognitive states, but lack representational content. In particular, the "thought" prompted by "reflection" on its "crooked picture" sensory state won't actually have the truth condition that this sensory state represents a crooked picture. It will only feel the same as the cognitive state which has this content in me.

Still, this by itself is enough to answer Horgan and Tienson's argument. What must be conceded to their argument is that my phenomenal duplicates are in a position to form cognitive states which correspond to my (true) beliefs about the representational contents of my sensory states. But it does not follow that the sensory states of my duplicates have representational contents, nor even that their cognitive states have representational contents.¹²

¹² It follows from the commitments of phenomenal intentionalism that the correctness conditions of sensory (and cognitive) states must be narrow. Phenomenal intentionalists thus face all the difficulties about specifying narrow contents raised in sections 6 and 8 above. It is not clear to me that these difficulties are always fully appreciated. Thus Horgan and Tienson, discussing states with the phenomenology of ordinary cat thoughts, suggest that "You, your Twin Earth doppelganger, and your Cartesian duplicate all have phenomenally identical thoughts with the same narrow truth conditions. For all three of you, these thoughts are intentionally directed toward certain small, common furry critters that meow, rub legs, drink milk, etc." (2002 229). But of course the referential value of concepts of furry, meowing, legs and milk cannot themselves be assumed in this context. The familiar Newman-style objections to Ramsifications of scientific theories are relevant here.

17 Mind and World

We need to be careful that we are not seduced by the following line of thought:

Sensory experience, whether veridical, illusory or hallucinatory, presents us with properties that ordinary physical objects can possess, such as colours and shapes and so on. But the presence of these properties in experience does not guarantee that they really are possessed by any physical object, or even that such an object exists. So sensory experience by its very nature poses a further question, of whether there really is an independent physical object with the properties we are experientially presented with. That is, experience by its nature is representational.

This line of thought would indeed be compelling, if only the initial idea that ordinary objectual properties are present in experience were granted. However, as we saw earlier, this idea does not stand up to examination. The property of yellowness is not “in” our experience when we have an experience as of a yellow lemon. Rather the property that we know introspectively to be instantiated in such cases, whether veridical, illusory or hallucinatory, is yellowness*, a conscious property of mental subjects, not a surface property of physical objects. These properties* might represent object properties, but in themselves they are like typographical words, items that have no constitutive tie to what they contingently represent.

This point isn’t altered by the sense in which sensory experiences do contain “phenomenal objects” with constant features. These “objects” and their features are still on the side of properties*, aspects of experiences that contingently represent, not the kinds of things that are so represented. These aspects may display structural features that invite us to characterize them as displaying a kind of “mind-independence”, but this doesn’t mean that they are the kinds of things that can exist outside experience.

It is of course very tempting to think of the properties that are present in experience as the same properties that physical objects might or might not have.¹³ And it would indeed follow from this that experiences are intrinsically representational. (For they would intrinsically pose the question: are those properties also present in reality?) But the temptation must be resisted. The properties in experience are properties*, which have no constitutive connection with the ordinary objectual properties they contingently represent. So something beyond experience itself is needed to establish representation relations between experience and the rest of the world.

One final thought. I have been assuming throughout this paper that most of the world that we represent in experience is mind-independent in a strong metaphysical sense (it would still have existed even if humans with perceiving minds had never evolved). An alternative would be to view the world itself as made of idealist materials, as some kind of construction with sensory constituents. In the context of this idealist alternative, the sharp distinction that I have drawn between sensory properties* and ordinary objectual properties would need re-examination. If lemons are made of the same fundamental material as minds, then perhaps they can possess just the same properties as experiences after all. And then perhaps sense experience could be shown to be constitutively representational, via the line of thought that I have been considering in this

¹³ Thus consider Horgan and Tienson: “. . . sensory-phenomenal states . . . present an *apparent* world full of *apparent* objects that *apparently* instantiate a wide range of properties and relations . . .” (2002 225, my italics). Where do the “apparents” come from here? There is nothing apparent about the property my experience possesses when I have an experience as of a yellow lemon. It is what it is. To see this property as “apparent” is already to assume that it is the same property that a real lemon might or might not have.

section. But that is all a topic for another paper. For now it will be enough if I have shown that, on any non-idealist metaphysics, sensory experience is not essentially representational.

References

Block, N. 1990 "Inverted Earth" Philosophical Perspectives 4 53-79

Block, N. 2004 "Mental Paint" in Hahn, M. and Ramberg, B. eds Reflections and Replies: Essays on the Philosophy of Tyler Burge Cambridge MA: MIT Press 165-200

Burge, T. 1979 "Individualism and the Mental" Midwest Studies in Philosophy 4 73–121

Byrne, A. and Tye, M. 2006 "Qualia Ain't in the Head" Noûs 40 241–55.

Dretske, F. 1995 Naturalizing the Mind Cambridge MA: MIT Press

Dretske, F. 1996 "Phenomenal Externalism, or If Meanings Ain't in the Head, Where Are Qualia?" in Villanueva, E. ed. Philosophical Issues 7 Perception Atascadero CA: Ridgeview Publishing 143-58

Dretske 2003 "Experience as Representation" Philosophical Issues 13 67-82

Farkas, K. 2013 "Constructing a World for the Senses" in Kriegel, U. ed. Phenomenal Intentionality Oxford: Oxford University Press 99–115

Fodor, J. 1980 "Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology" Behavioral and Brain Sciences 3 63-109

Harman, G. 1990 "The Intrinsic Quality of Experience" in Tomberlin, E. ed Philosophical Perspectives 4 Action Theory and Philosophy of Mind Atascadero CA: Ridgeview Publishing 31–52

Horgan, T. and Tienson, J. 2002 "The Intentionality of Phenomenology and the Phenomenology of Intentionality" in Chalmers, D. ed. Philosophy of Mind: Classical and Contemporary Readings Oxford University Press 520-33

Kriegel, U. "The Phenomenal Intentionality Research Program" in Kriegel, U. ed. Phenomenal Intentionality Oxford: Oxford University Press 1-26

Kripke, S. 1980 Naming and Necessity Oxford: Basil Blackwell

Lycan, W. 1996 Consciousness and Experience Cambridge MA: MIT Press

Lycan, W. 2001 "The Case for Phenomenal Externalism" in Tomberlin, J. ed. Philosophical Perspectives 15 Metaphysics Atascadero CA: Ridgeview Publishing 17-35

Masrour, F. "Phenomenal Objectivity and Phenomenal Intentionality" in Kriegel, U. ed. Phenomenal Intentionality Oxford: Oxford University Press 99–115

Papineau, D. 2006 "Naturalist Theories of Meaning" in Lepore, E. and Smith, B. eds Oxford Handbook of the Philosophy of Language Oxford: Oxford University Press 175-88

Peacocke, C. 2003 Sense and Content: Experience, Thought, and Their Relations Oxford: Oxford University Press

Peacocke, C. 2008 "Sensational Properties: Theses to Accept and Theses to Reject" Revue Internationale de Philosophie 62 7–24

Putnam, H. 1975: "The Meaning of 'Meaning'". In Keith Gunderson, K. ed. Minnesota Studies in the Philosophy of Science 7 Language, Mind and Knowledge Minneapolis: University of Minnesota Press 131–93

Rumfitt, I. 2014 "Truth and Meaning" Aristotelian Society Supplementary Volume 88

Segal, G. 2002 A Slim Book About Narrow Content Cambridge MA: MIT Press

Soteriou, M. 2000: "The Particularity of Visual Perception" European Journal of Philosophy 8 173–89

Tye, M. 1995 Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind Cambridge MA: MIT Press

Tye, M. 2002 "Representationalism and the Transparency of Experience" Nous 36 137–151

Tye, M. 2014 "What is the Content of a Hallucinatory Experience?" in Brogaard, B. ed Does Perception have Content? Oxford: Oxford University Press