

Attention and intentionalism

Jeff Speaks

September 18, 2008

ABSTRACT. Many alleged counterexamples to intentionalism – the view that the phenomenology of perceptual experiences of a given sense modality supervene on the contents of experiences of that modality – can be avoided by adopting a liberal view of the sorts of properties that can be represented in perceptual experience. I argue that there is a class of counterexamples to intentionalism, based on shifts in attention from one aspect of a represented scene to another, which avoids this response. We can preserve the idea that there is a necessary connection between the contents and phenomenal characters of perceptual experiences via a distinction between perceptual phenomenology and the phenomenology of attention; but, even if this distinction is viable, these cases put surprising pressure on the thesis that phenomenal character can, in general, be explained in terms of mental representation.

The core of intentionalism in the philosophy of perception is the view that there is the following necessary connection between the content and the phenomenal character (phenomenology) of perceptual experience:

Minimal intentionalism

Necessarily, if two perceptual experiences of the same sense modality differ in phenomenology, then they differ in content.¹

Much of the interest of minimal intentionalism derives from the idea that a thesis of this sort makes room for an explanation or reduction of perceptual phenomenology in terms of perceptual representation. Of course, to establish the truth of minimal intentionalism is not to establish this further explanatory claim;² but, plausibly, a defense of the explanatory claim requires at least that minimal intentionalism be true.

¹We can distinguish intrapersonal and interpersonal versions of this thesis. Here I'll be assuming only that intentionalists endorse the intrapersonal version, and will be concerned with potential counterexamples to it.

To forestall confusion, it is worth noting that sometimes 'intentionalism' is used as a name for the view that perceptual experiences are relations to propositions. I don't think that it is straightforward that any of the intentionalist theses discussed in this paper entail much at all about the metaphysics of experience. It seems to me that just as many views of the metaphysics of belief are consistent with the idea that beliefs have contents, so many views of the metaphysics of perception are consistent with the idea that perceptions have contents. If this is so, then it is hard to see how only one view of the metaphysics of experience could be compatible with the idea that there is a necessary connection between content and phenomenal character.

²See Warfield (1999).

Most minimal intentionalists think not only that *perceptual* phenomenology can be explained in terms of *perceptual* representation, but also that phenomenology can, quite generally, be explained in representational terms.³ Such theorists would endorse intentionalism not just about perceptual states, but about mental states of any sort; accordingly, I'll call such theorists *global intentionalists*. It should be pointed out, following Byrne, that the global intentionalist might state his characteristic supervenience thesis either as an intermodal or as an intramodal claim:⁴

Intermodal global intentionalism

Necessarily, if two mental events differ in phenomenology, then they differ in content.

Intramodal global intentionalism

Necessarily, if two mental events of the same type differ in phenomenology, then they differ in content.

Intermodal global intentionalism entails the intramodal thesis, but not the converse. While the difference between these theses will be important later, for now I'll simply use 'global intentionalism' as a label for the view that the intramodal thesis is true, whether or not the intermodal thesis is. As with minimal intentionalism, these global intentionalist theses are supervenience claims rather than assertions of explanatory priority; but, if true, they might serve as a first step toward a reduction of phenomenal character to mental representation.

Many of the most prominent alleged counterexamples to intentionalism are counterexamples to global intentionalism, but not to minimal intentionalism. Examples of 'objectless' moods and bodily sensations which seem to lack representational properties are apparent counterexamples to the view that every mental event with a phenomenology also has a content and so also to the view that phenomenology can in general be explained in terms of content; but they are not counterexamples to the minimal intentionalist idea that, within a given modality of perceptual experience, there is a necessary connection between content and phenomenology.⁵ But minimal intentionalism has attracted its fair share of counterexamples as well.

Most of these are examples of pairs of perceptions which clearly differ in phenomenology, but do not differ with respect to their representation of a certain basic class of observational properties, such as the color and shape of the represented objects. Christopher Peacocke introduced the example of a pair of trees of the same height, one of which is further away from the perceiver than the other. The perceiver might represent the trees as the same height; but clearly there is a difference in the phenomenal characters of the experiences of the trees. We can gesture at this phenomenal difference by saying that one of the trees occupies more of the perceiver's visual field than the other. If this is right, then we have a putative example of a pair of visual experiences which are alike with respect to content, but differ in phenomenology: hence we have a *prima facie* counterexample to minimal intentionalism.⁶

³See, among many other places, Tye (1995b).

⁴Minimal intentionalism is as stated an intramodal claim; there is, of course, a corresponding intermodal claim, but the differences between these won't be important for what follows.

⁵For discussion of objectless moods, see Searle (1983); for a defense of the idea that some bodily sensations lack representational content see, among other places, Block (2003).


⁶See the discussion in Chapter 1 of Peacocke (1983).

A similar example focuses on a pair of visual experiences of a coin, one from above and one from an angle. Both visual experiences might well represent the coin as round; but there is a difference in the phenomenology of the two experiences, since there is surely some sense in which the appearance of the coin is round in one case, but elliptical in the other.⁷

The natural reply to these sorts of objections is that they rest on too narrow a view of the properties represented in perceptual experience. In Peacocke's first example, the trees may be represented as the same height; but they are not represented as in the same relative location. One is represented as farther away from the perceiver than the other. The coins may be represented as the same shape; but they are not represented as in the same relative orientation. One is represented as facing the perceiver, while the other is not. So long as there is no principled reason why relative distance and orientation can't be among the properties represented in perceptual experience, these cases pose no serious challenge to minimal intentionalism.⁸

This strategy of appealing to representation of relative distance and orientation has wide application. For instance, it can be put to work in handling some cases of 'Gestalt shifts', or cases of 'seeing as', such as the difference in phenomenology between two experiences of the Necker cube which differ with respect to the side of the cube that seems to be in front. The intuitive description of the two experiences — as exhibiting a difference in which side *seems to be in front* — makes plain that the relevant visual experiences differ with respect to the representation of the distance between the perceiver and the relevant sides of the cube.

To be sure, not all examples of 'Gestalt shifts' can be handled by invoking perception of relative distance and orientation, as is shown by the following well-known figure, which can be seen either as a duck or as a rabbit:



There is a phenomenological difference between seeing this figure as a representation of a duck and seeing it as a representation as a rabbit; but this difference does not consist in any change in the color or shape the figure is represented as having, and there is no difference in the representation of relative distance or orientation. However, despite the differences between this example and the examples of 'seeing as' discussed above, the

⁷This example has an interesting history. It appeared prominently in versions of the argument from illusion used to defend sense datum theories of perception (see, e.g., Ayer (1940)); in those cases, it was essential to the argument that we regard perception of the shape of a coin from an angle as *non-veridical*. In the present use of the example in an argument against intentionalism, it is essential to the argument that we regard this sort of perceptual experience as a *veridical* experience of the shape of the coin as circular, since if we did not regard both experiences as veridical, there would no pressure to regard them as representing the same shape property, and hence as having the same content. See, for example, Peacocke (1983) and Smith (2002).

⁸This is not to say that the representation of egocentric distance and orientation is easily handled by just any view of the nature of perceptual content; for some of the challenges, see the discussion in Chapter 3 of Peacocke (1992).

duck-rabbit example is open to substantially the same resolution. Even if representation of shapes, colors, and locations is constant between seeing the above figure as a duck and seeing it as a rabbit, there is a clear representational difference: in one case the figure is represented as duck-shaped, and in the other as rabbit-shaped.⁹ These are clearly distinct properties that an object can be represented as having, and the difference between them intuitively matches the difference in phenomenology between the two experiences.

The intentionalist might reasonably regard the discussion of the above cases as exemplifying a pattern: alleged counterexamples to minimal intentionalism, in general, rest on overly restrictive views of the properties which can be represented in experience.¹⁰ However, there is a class of counterexamples to minimal intentionalism which seem to resist this kind of reply.

These are pairs of perceptual experiences which differ only with respect to the focus of the attention of the perceiver. Consider, for example, a visual experience of the following lines on a white sheet of paper large enough to fill the perceiver's visual field:



Compare two visual experiences of these lines: in the first, the perceiver's attention is focused on the intersection of the second vertical line from the left with the horizontal line; the second differs only in that the perceiver shifts his attention to the point of intersection to the right, between the horizontal line and the third vertical line from the left. To make the case clearer, we can imagine that this shift in attention does not involve any eye movement; it is possible (even if a bit unnatural) to shift one's attention from one point in the visual field to another without foveation.¹¹ It is undeniable that one's total phenomenology differs between these two cases. No one would say that the first case is indistinguishable from, or seems the same to the perceiver as, the second; you would never say, for example, that you're not sure which of the two sorts of experiences you are having at a given moment.

So it seems that the minimal intentionalist must find some difference in content to correspond to this difference in phenomenology. But it is hard to see what this representational difference could be. Given that the background of the figure is an uninterrupted stretch of solid white, the change in focus between the two points of intersection does not bring with it a change in the representation of anything on the periphery of the perceiver's visual field.¹² Nor does the figure seem to move relative to the subject when one shifts

⁹Perhaps even more salient are local differences between the two experiences: the leftmost portion of the figure is represented first as beak-shaped, then as ear-shaped, etc. There are other, more challenging examples of 'seeing as,' including Mach's example of the square/diamond and the example of visual experiences of a 3 x 3 grid of squares discussed in Nickel (2007). I return to these below.

¹⁰It is more difficult to cast examples of spectrum inversion without misrepresentation in this mold, though perhaps one can see the use of 'appearance properties' in Shoemaker (1994) as an instance of the strategy of finding an overlooked class of properties with respect to which the troublesome pair of experiences do differ in content. Examples designed to show the possibility of spectrum inversion without misrepresentation are an important challenge to minimal intentionalism, but are beyond the scope of this essay. For a clear discussion of different versions of the 'inverted spectrum' argument against minimal intentionalism, see Marcus (2006). For discussion of the closely related challenge to intentionalism from 'shifted spectrum' examples, see Block (1999).

¹¹See among other places Schwitzgebel (2008), 255-6.

¹²And if we're really thinking of a shift in attention without foveation, there would be no change in the location of the 'periphery' anyway.

one’s attention from one’s point of intersection to another. And given the simplicity of the figure, it does not seem plausible to claim that one experience represents a given portion of the lines with more detail or determinacy.

One might object that this last point rests on the discredited view that visual experience reveals “a broad, stable field, flush with precise detail, hazy only at the borders.”¹³ But we can grant that this view is false, and that we only visually represent color and shape properties with a high level of determinacy within a surprisingly small area of the visual field — one estimate is an area of one degree of visual angle.¹⁴ We can adapt the example above so that the shape in question is smaller than that area, but still large enough that the subject can shift attention from one point of intersection to another. In such a case, the subject will, in both experiences, represent all of the lines and points of intersection in the figure, and there will still be a clear phenomenological difference between the two experiences.¹⁵

But emboldened by the success of this strategy in handling the preceding cases, the intentionalist might yet try to handle this example by finding some overlooked class of properties with respect to which the two experiences do, contrary to initial impressions, differ in their representational content. Pursuing this strategy, he might be tempted to say that the experiences differ with respect to which points of intersection are represented as prominent to the perceiver. After all, if we can make use of perceptual representation of egocentric locations and orientations, why not also representation of egocentric relations of perceptual prominence?¹⁶

However, there is an important disanalogy between the representation of egocentric relations of distance and orientation and the representation of egocentric relations of perceptual prominence. The latter, but not the former, involves representation of *properties of the relevant experiences*. After all, to say that one point of intersection on the horizontal line is more prominent than the other is just to say that one, but not the other, point of intersection on that line is attended to by the perceiver. So, to say that in one experience the perceiver represents a given point of intersection as prominent just is to say that in one experience the perceiver represents herself, or her own experience, as attending to that point of intersection.

There are at least three reasons why intentionalists should not endorse the claim that the pair of experiences described above is distinguished by their representation of the attention of the perceiver.

First, there is a worry about trivializing intentionalism. If intentionalism is to be a substantive thesis, it surely can’t be legitimate to simply slide without further argument from the premise that an experience has a certain phenomenal feature to the conclusion that the experience represents that phenomenal feature as being a property of the experience. This is the sort of built-in response to alleged counterexamples which should make us suspicious. In each case, we should require that the intentionalist should make plausible the idea that the representational property adverted to is a genuine representational property of the experience. In the present case, the idea that we not only attend to

¹³For excellent critical discussions of this view, see the discussion of the snapshot conception of vision in chapter 2 of Noë (2004), and the examples in §vi of Schwitzgebel (2008). The quote above is from Schwitzgebel (2008), 256.

¹⁴This is the finding of Eriksen and Eriksen (1974), cited in Palmer (1999), 546.

¹⁵Thanks to an anonymous referee for helping me to see the importance of the issues discussed in this paragraph, and for suggesting the clarificatory point about attention vs. foveation made above.

¹⁶The perceptual representation of perceptual prominence is also discussed in Nickel (2007); I discuss the central example of that paper below.

aspects of the represented scene but also, just in virtue of so doing, perceptually represent ourselves as so attending, does not seem to have much initial plausibility. It is hard to see what the intentionalist could say to make it more plausible.

Second, we can provide a kind of positive argument that this is not a genuine case of perceptual representation. It seems plausible that representation of any sort of property in any sense modality can be, or fail to be, veridical.¹⁷ But then perceptual representation of oneself as attending to an object would be, if genuine, the lone counterexample to this rule. What would it be like to have a perceptual experience which represented oneself as attending to an aspect of the presented scene when one was not so attending? Our inability to answer this question indicates that we don't have much of a grip on what perceptual representation of attending to a point — as something over and above simply attending to that point — could be. This seems to indicate that what we have here is just a relabeling of an aspect of phenomenal character in representational terms.¹⁸

Third, the suggestion that in the standard case we represent properties of experience undercuts one of the principal lines of argument in favor of intentionalism: the argument from the transparency of experience that phenomenal characters can't be qualities *of experiences* of which we are directly aware, since when we attend to our experiences we are not aware of anything other than the external particulars and their qualities represented as in our environment.¹⁹ But suppose that — in addition to objects, shapes, colors, etc. — our perceptual experiences represented us as attending to these objects, shapes, colors, etc. Then one would think that when we attend to those perceptual experiences, we would notice, in addition to objects, colors, and shapes, our own attention to them. But this is just what proponents of the transparency of experience deny.²⁰

It is worth adding that matters get worse yet for the intentionalist in question if he wishes — as many do — to explain phenomenal character in terms of representational content, at least if the intentionalist endorses a Russellian view of the contents of perceptual experience on which propositions have objects and properties as their constituents. For then the intentionalist would be explaining the phenomenal property corresponding to attention to a point via the perceptual representation of that phenomenal property; so, on a Russellian view of perceptual representation we would then be explaining the phenomenal property by means of a relation to a proposition of which that very property is a constituent. But that would make the proposition explanatorily prior to one of its own constituents, which is both *prima facie* odd and inconsistent with the conjunction of

¹⁷It is less that there's a straightforward argument for this sort of 'anti-infallibility' principle than that claims to the contrary seem always to be open to counterexamples. A good example here is the example of misperception of pain in a fraternity initiation which is due to Rogers Albritton and discussed in (among other places) Hill (1991), 128. For an attempt to define a quite restricted infallibility principle immune from such counterexamples, see Horgan and Kriegel (2007). Such restricted principles, however, are unlikely to come to the aid of the view that certain sorts of phenomenal features are, in *any* circumstance whatever, perceptually represented correctly.

¹⁸The opposite direction holds as well. By hypothesis, the intentionalist cannot allow the possibility of shifts in attention which are not represented as such (for this would entail a change in phenomenology without any corresponding change in representational content). So perceptual representation of attention is not just infallible, but, so to speak, omniscient — no shift in attention can fail to be represented as such by the perceptual experience. But, again, it seems implausible that any sort of perceptual experience could be omniscient, in this sense; it seems plausible that, for any property sometimes represented by experience, that property could obtain either perceived or unperceived.

¹⁹An argument of this sort is laid out in much more detail in Tye (2002).

²⁰Of course, the intentionalist might reply simply by denying that experience is, in this sense, transparent. The resulting position is not inconsistent; the present point is just that making this move undermines what many take to be a central motivation for intentionalism.

(i) the idea that the essence of a proposition is given by a specification of its constituents along with their relations to each other and (ii) the plausible principle that nothing can be a part of the explanation of its own essence.²¹

If we agree that we cannot appeal to perceptual representation of our own attention to objects and properties, then it seems fairly clear that there is no representational difference between the pair of experiences described above. So it seems that we have here a genuine example of a pair of visual experiences with the same content but different phenomenology.

One can see in hindsight why defenders of intentionalism have overlooked the possibility of this kind of counterexample. As just noted, intentionalists are often motivated by the transparency of experience, which is often glossed as the view that the only things we can notice when we attend to our experiences are the objects and properties that experience represents as in the environment of the perceiver. From here it seems but a short step to the conclusion that any noticeable *difference* between experiences — i.e., any difference in phenomenology, or phenomenal character — must correspond to some difference in the objects and properties presented as in the perceiver’s environment, and so also to a difference in content.²² What this step misses is that a difference in phenomenal character can be generated not just by a change in *what* is noticed when one attends to one’s experience but also by a change in *where* one’s attention is focused.

The most straightforward response to this sort of example is that minimal and, a fortiori, global intentionalism are false, and so that there is no internal connection between the content and phenomenology of perceptual experiences. On this sort of anti-intentionalist view, content and phenomenology are two distinct sorts of properties of perceptual experiences, which may be systematically, but only contingently, correlated.

However, a second kind of response is possible. There is a clear sense in which these kinds of shifts in attention are not part of specifically visual phenomenology at all: similar cases can easily be generated for any of the other sense modalities. (Imagine listening to a duet in an otherwise silent environment while shifting your focus of attention from one voice to the other.) So perhaps the minimal intentionalist should respond to these cases not by trying to find some representational difference between the two experiences, but by thinking of attention as having its own, sui generis phenomenology, which is distinct from visual phenomenology, auditory phenomenology, and the phenomenologies specific to the other sense modalities.²³ Attention is not itself a modality of sense experience, so differences in the phenomenology of attention which are unaccompanied by differences in content would, on this view, be no more a threat to minimal intentionalism than are alleged examples of pains or moods which differ in their phenomenal characters without differing at the level of content.

On this view, we could compare our two experiences of the grid to a pair of visual experiences identical but that one is accompanied by a toothache. While there will be a clear difference in the total phenomenology of the subject during the two visual experiences without a difference in the content of the two visual experiences, this is no

²¹This parallels the argument against dispositionalism about value in Johnston (2001). For an extension of this argument to cover a neo-Fregean view of the contents of experience, see Johnston (2001), pp. 196–199.

²²The classic discussion of the transparency of experience in this context is Harman (1990). For arguments for intentionalism on the basis of the transparency of experience, see Byrne (2001); Tye (2000). I discuss the above line of argument in Speaks (forthcoming).


²³One challenge to this separation of the various modalities is to emphasize cases in which experiences in one sense modality affect the phenomenology and content of experiences in other modalities. See, for example, the discussion of cross-modal illusions in Chapter 11 of O’Callaghan (2007).

counterexample to minimal intentionalism, since there is also no difference in specifically visual phenomenology. Of course, visual attention is constrained by visual perceptual experience in a way that toothaches are not, but from the point of view of minimal intentionalism — which is a thesis only about the supervenience of perceptual phenomenology on perceptual representation — the two examples are equally irrelevant.

This distinction between perceptual phenomenology and the phenomenology of attention also gives the minimal intentionalist the resources to handle two of the most challenging recent apparent counterexamples to the view.

The first of these is Mach’s example of seeing a box as a square, and then as a diamond. There is a clear difference in phenomenology between the two experiences, but no obvious difference in representational content. The natural line of response for the intentionalist is to say that some properties of the shape are represented by the experience in which the subject sees it as a square, but not in the experience in which the subject sees it as a diamond, and vice versa. For example, Peacocke suggests that when the figure is seen as a square one’s visual experience represents a symmetry about the bisectors of the sides of the shape, whereas when the figure is seen as a diamond one’s visual experience represents a symmetry about the bisectors of the angles of the shape.²⁴ However, as Fiona Macpherson has pointed out, it is possible to see a box as a square while visually representing the symmetry in the bisectors of the angles of the shape.²⁵ I suggest that the intentionalist should respond to this case in the same way as she should respond to the example of the intersecting lines: she should say that the difference between seeing the box as a square and seeing it as a diamond is not a difference in visual phenomenology, but rather is a difference in the phenomenology of attention. When the box is seen as a square, one is attending to symmetries involving the sides, whereas when it is seen as a diamond one is attending to symmetries involving the angles.

Against this suggestion, Macpherson might argue that one can see the box as a square even while *attending* to the symmetry in the bisectors of the angles.²⁶ But this seems not to be the case. Consider the following figure, discussed by Macpherson:



It is plausible that, as Macpherson says, visual experiences of this figure represent the angle bisector symmetry and that it is possible to see this figure as a square; and, as Macpherson says, this seems to count against Peacocke’s idea that one sees the box as a diamond when one visually represents the symmetry about the bisectors of the figure’s angles. But now attend to the angle bisector symmetry by attending to the intersecting dotted lines. This shift in attention generates the Gestalt shift to seeing the figure as a diamond. This is strong evidence that the difference in phenomenology between seeing the box as a square


²⁴See Peacocke (1992).

²⁵See §7 of Macpherson (2006).

²⁶This is suggested by her claim that “it seems perfectly possible to see a square as a square while focusing intently on its angle bisector symmetry.” (103)

and seeing it as a diamond is, like the difference between the two experiences of the intersecting horizontal and vertical lines above, due to a shift in attention. If this is correct, and if (as suggested above) it is legitimate to regard attention as a type of non-perceptual state with its own associated phenomenology, the square/diamond example is no threat to minimal intentionalism.²⁷

The second proposed counterexample to minimal intentionalism is due to Bernard Nickel, who asks us to consider two perceptual experiences of a 3 x 3 grid of squares like the following



which differ only in which groups of squares appear as prominent.²⁸ In one experience, the corner and center squares appear prominent, and in the other the remaining four ‘side’ squares appear as prominent. There is, as Nickel says, a clear difference in phenomenology here, and Nickel argues convincingly against a number of different attempts to find a difference in content between the two experiences. Intuitively, though, it seems that we have the same phenomenon here as in the case of the intersecting lines and the case of the square/diamond: the relevant phenomenological difference is generated by a shift in attention from one group of boxes to the other. So, as above, if we can think of attentional states as non-perceptual states with their own phenomenology, this sort of example need not worry the minimal intentionalist.²⁹

²⁷Macpherson also argues convincingly that similar phenomena arise with figures which lack the symmetries which seem central to the square/diamond example. To adapt the above view to these cases, what is needed is to find properties of the relevant figures which can serve as the focus of the relevant attentional shifts; there’s no reason why these properties should have to involve symmetries.

This discussion of the square/diamond case leaves open the question of whether it is a threat to global intentionalism; this depends on how we understand the relevant attentional states. I return to this issue below.

²⁸Nickel (2007), 284.

²⁹Nickel himself notes that it “seems apt” (289) to describe the difference between the two experiences in terms of a shift in attention, but doesn’t consider the present way of using this point in defense of minimal intentionalism.

Though I think that the distinction between the phenomenology of perceptual experience and of attention is ultimately the right way to handle this example, I’m less convinced that Nickel’s example forces this distinction on the intentionalist. It seems intuitively plausible to say that in the experiences of the 3 x 3 grid, the squares which seem prominent also seem — as in the case of the Necker cube — slightly closer to the perceiver. This would, as in the case of the cube, provide the needed difference in content between the two experiences. Nickel objects (288) that this response neglects the fact that if one object appears closer than another, and each subtends the same angle of the visual field of the perceiver, the closer object will appear smaller. However, I’m inclined to think that, contra Nickel, the prominent squares do appear slightly smaller. This, to be sure, leads to a further oddity, since it makes the content of the relevant experience necessarily false, given that one cannot have a 3 x 3 grid of squares which share an outer boundary but which is such that (for example) the middle and corner squares are all smaller than the others. This consequence does not strike me as implausible, but insofar as one thinks that perceptual experience of the sort exemplified by Nickel’s discussion can’t have necessarily false propositions as their contents, this might provide a reason for thinking that the present response to Nickel’s case is unsatis-

One might worry that this way of handling the cases runs the risk of trivializing minimal intentionalism; we don't want the appeal to the phenomenology of attention to be an unprincipled way for the minimal intentionalist to simply relabel any proposed counterexample as a mere shift in the phenomenology of attention without any change in the phenomenology of the relevant perceptual experiences.

I think that there are two responses to this worry, one intuitive and one a bit more theoretically loaded. The intuitive point is just that each of the three cases really do seem to essentially involve attention. It is difficult even to describe the initial example of the intersecting lines without describing it as a shift in attention. The same is true of the two experiences of the box pictured above with dotted lines marking bisectors of the shape's angles; to generate the two different experiences, one directs a subject 'to change his focus of attention' from the dotted lines to the symmetry of the sides, and back. Analogous remarks apply to Nickel's 3 x 3 grid. It is also telling that in each case it is natural to describe the two experiences by talking about what aspects of the relevant figure are 'prominent.' By contrast, it would *not* be natural to extend this model to other putative counterexamples to minimal intentionalism: the differences between visual experiences of Peacocke's trees, the two perspectives on the round coin, the views of a brightly lit scene at the end of a dark tunnel with one and with two eyes open, and blurred vision and seeing an object as blurry, have, intuitively, nothing to do with shifts in attention or focus.

The second line of response involves the voluntariness of at least some attentional shifts. Typically, changes in, for example, visual phenomenology cannot be brought about at will; one cannot go from a view of the coin in which it has a round appearance to one in which it has an elliptical appearance without changing the relative position of the coin. But in each of the cases above, we can effect the relevant phenomenal change without bringing about any such 'external' change in the scene perceived; this makes it more plausible that these changes are due to attentional shifts since, in at least many cases, we can shift our attention from one element of a represented scene to another at will.³⁰

If the foregoing is correct, then minimal intentionalism is well-positioned to handle all of the examples discussed thus far. But this sort of separation of the phenomenology of attention from the phenomenologies of the senses is, obviously, of no help to the global intentionalist's project of reducing phenomenology to representation. Whether or not shifts in attention are thought of as changes in *perceptual* phenomenology, they are still changes in phenomenology for which we have so far found no corresponding changes in content.

If we stop here, and endorse minimal intentionalism without global intentionalism, the project of a comprehensive reduction of phenomenology to content is dead. If this were the end of the story, then for many minimal intentionalism would lose much of its interest. I would suggest that, if we are interested in understanding the nature of perceptual experience, this is a mistake. If even minimal intentionalism is true, we have a necessary connection between the representational and phenomenological properties of the various modalities of perceptual experience. This puts a substantial constraint on theorizing about perception: for surely one of the things that we should want any theory

factory. But, in any case, once the intentionalist sees that the example of the shift in attention from one point of intersection in the grid to another requires the distinction between the phenomenologies of visual experience and attention, there seems no reason not to deploy that distinction here as well.

³⁰This is not to say that there is no involuntary aspect to these phenomenal changes; as Macpherson (2006) (§4) argues, there is. The point is just that the relevant phenomenal shifts can, in the standard case, be initiated at will without any change in stimulus from the environment.

of perception which hopes to make sense of the representational and phenomenological aspects of perceptual experience to explain is why this necessary connection obtains.

But perhaps we needn't stop here; perhaps the distinction between visual phenomenology and the phenomenology of attention can also come to the rescue of the global intentionalist. Since attention is a state with its own associated phenomenology, the global intentionalist ought to say that it is a state with representational content. If this is correct, then perhaps, even if there is no difference in the content of visual experience between the first visual experience of the intersecting lines and the second, there is a difference in the content of the two attentional states.

This leads to the question of what sorts of things are represented by such states; and there seem to be two sorts of answers to this question: either they represent properties of the relevant perceptual experience, or they represent aspects of the scene represented by that perceptual experience. As we'll see, either route forces some difficult choices on the global intentionalist.

Let's consider the first option, according to which the attentional states with which we are concerned are higher-order representational states, which represent properties of the subject's perceptual experiences.³¹ On one way of developing this view, the phenomenal difference between the two experiences of the intersecting lines is to be explained by the fact that the attentional state which accompanies the first visual experience represents the subject's experience as representing the left point of intersection, whereas the attentional state which accompanies the second represents the subject's experience as representing the right point of intersection. (Of course, in both cases the subject's visual experience is representing both points of intersection; it's just that in each case the subject's attention only represents part of what is represented by the corresponding visual experience.)

This seems to be a coherent view, and one which suits the global intentionalist's purposes. But it is in some ways an odd view. When I shift my attention from one point of intersection to the other, it seems to me that if I am representing anything, I am representing aspects of the figure represented by my perceptual experience, and not that experience itself. Indeed, this sort of intuition is one of the 'transparency' intuitions often used to generate intuitive support for intentionalism. The following quote from Michael Tye is representative:

"Intuitively, you are directly aware of blueness and squareness as ... features of an external surface. Now shift your gaze inward and try to become aware of your experience itself, inside you, apart from its objects. Try to focus your attention on some intrinsic feature of the experience that distinguishes it from other experiences, something other than what it is an experience of. The task seems impossible: one's awareness seems always to slip through the experience to blueness and squareness, as instantiated together in an external object. In turning ones mind inward to attend to the experience, one seems to end up concentrating on what is outside again, on external features or properties."³²

This is explicitly a claim about attention, and the claim is that even if when try to attend

³¹This view corresponds roughly to 'inner eye' model of attention, discussed in, e.g., Armstrong (1968), Lycan (1987), Hill (1991). It is a bit hard to connect the distinction between these two views of the content of attention with much current discussion of attention, since much of that discussion does not explicitly attribute content to attentional states at all — an example here is Campbell (2002). Of course, if attentional states don't have contents, then, if the foregoing argument is correct, it follows immediately that global intentionalism is false.

³²Tye (1995b), 30.

to properties of our perceptual experiences, we end up attending to properties of the scene represented by that experience. But this is at least in tension with the claim that attentional states represent properties of perceptual experiences, rather than properties of the things represented by perceptual experiences.³³

This view also runs into a sort of problem which we mentioned above in connection with the suggestion that perceptual experiences represent the subject as attending to this or that. Ordinarily, we expect types of representational states — even states which represent other mental states — to be such that some of their tokens represent their target correctly, and some incorrectly. But the present view makes attention a counterinstance to this rule. After all, how could attentional states misrepresent, on the present view? Presumably by representing the relevant perceptual experience as representing something which it does not, in fact, represent. But what would such a state be like? Suppose (for *reductio*) that an attentional state represented an experience as representing an object as purple but that the experience did not, in fact, represent the object as purple. Would the subject’s total phenomenology include the phenomenal character typical of visual experiences which represent things as purple, or not? If so, then it looks like, *contra* our supposition, the subject’s visual experience *would* be representing the relevant object as purple. But if not, then it seems that (again *contra* our supposition) the attentional state is *not* representing the experience as representing the object as purple. (In virtue of what would it be representing the experience as representing the object as purple, rather than as some other color?)

So the idea of attentional states misrepresenting the associated perceptual experiences of the subject seems incoherent; it appears to be a necessary truth that, if attentional states represent perceptual experiences as having certain contents, they must always represent those experiences correctly. But this is mysterious. Why should attention, lone among representational states, be incapable of false representation?

I think that it would be preferable for the global intentionalist not to be tied to a view which is in tension with transparency intuitions, and which involves commitment to a class of infallible representational states. So let’s consider the second option for the global intentionalist, on which attentional states are representational states, but ones which represent aspects of the scene represented by the experience rather than aspects of that experience. On this view, attention — like the various perceptual modalities — represents objects and properties in the environment of the subject. As applied to the above case, a proponent of this view might say that in the first experience, the subject’s attention represents the left point of intersection, and in the second experience, the subject’s attention represents the right point of intersection.

The odd thing about this sort of response on behalf of the global intentionalist is that there’s a clear sense in which attention does not add anything to the total representational state of the subject. The subject already visually represents both points of intersection; so why should adding an attentional representation of one of those points make any difference? It is tempting to reply to this question by finding some property of the relevant point of intersection that is represented in attention, but not in visual experience. But this reply faces a dilemma: either the properties represented by the attentional state will be properties of the perceptual experience of the subject, or they will be properties of the figure perceived. The first option involves a return to the view of attention as a

³³One could claim that attentional states have ‘two layers’ of content, one which concerns aspects of the represented scene, and one which concerns aspects of the perceptual experience which represents that scene. This seems *ad hoc* to me, but perhaps it could be developed in a way which would make the view seem plausible.

higher-order representational state, which is what the present view aims to avoid. But the second option seems like a nonstarter, since there simply are no plausible candidates for properties of the figure which are represented by the subject's attentional state but not represented by the subject's visual experience of the figure. Indeed, it seems that it is in principle impossible to attend to any aspect of the scene before one which is not represented by one's perceptual experience.³⁴

So it seems that the global intentionalist who wants to avoid a higher-order view of attention has some difficulty explaining, in terms of representational content, why an attentional state whose content simply duplicates an aspect of the content of the subject's visual experience can affect the subject's total phenomenology. What she seems forced to say is that it is simply that the relevant content *is the content of an attentional state* which explains the difference in total phenomenology. But saying this involves abandoning the strongest and most popular version of intentionalism.

Recall that we can distinguish between *intermodal* and *intramodal* global intentionalist theses:

Intermodal global intentionalism

Necessarily, if two mental events differ in phenomenology, then they differ in content.

Intramodal global intentionalism

Necessarily, if two mental events of the same type differ in phenomenology, then they differ in content.

Most philosophers who would endorse intramodal global intentionalism would also endorse the intermodal thesis.³⁵ The reasons for this are, I think, less straightforward arguments than considerations of theoretical economy. If our aim is to naturalize phenomenal character via intentionalism plus a materialist reduction of mental content, then intermodal intentionalism is what we should want — otherwise we still need to supply an account of what it is for a content to be, for example, visually or auditorily represented. And the simplest reduction of all — one which simply identifies phenomenal properties with certain representational properties — clearly entails intermodal intentionalism. But it seems fairly clear that if attentional states represent aspects of the scene around the subject which are already represented by the subject's perceptual experiences, but nonetheless bring about a change in the subject's total phenomenology at the time of the experiences, intermodal intentionalism must be false, since on this view differences in modality would alone have phenomenal effects.

It's important to see that the present problem for intermodal global intentionalism is, in some respects, more serious than a more familiar problem for intermodal intentionalist theses. This problem is based on the fact that some properties, like shapes, can be represented in more than one sense modality. This fact seems to pose a problem for the

³⁴If this were possible, then it seems as though attention would be a modality of perceptual experience. An anonymous reviewer pointed out that this principle — that attention can represent only what is represented by the sense modality in question — can also be used to pose a problem for the view that attentional states are higher-order representational states. After all, the higher-order view of attention, along with this principle, entails that perceptual experiences represent themselves as having certain properties. But many people will want to avoid the view that perceptual experiences can represent themselves as having certain properties. Important exceptions here include Searle (1983) and Siegel (2006).

³⁵See, for example, Tye (1995b), Dretske (1995), Tye (2000), Byrne (2001).

intermodal intentionalist: after all, visual and tactile representations of a curved surface clearly differ in phenomenology, but seem to be representations of just the same property. So doesn't this show that the sense modality of an experience, as well as the content of the experience, makes a contribution to phenomenology?

The intermodal intentionalist has a few replies available. She can point out that the two experiences *do* differ in content, even if there is some overlap; for example, in at least the simplest versions of the case above, the visual experience but not the tactile experience of the curved surface will represent the color of that surface. She might also claim that the two experiences represent the relevant property at different levels of detail, and so differ representationally even with respect to the property represented by both experiences.³⁶ But neither of these ways of finding a representational difference to explain the difference in phenomenology seem as plausible in the case of changes in the phenomenology of attention. After all, if attentional states have contents, it seems to be in the nature of such states that their contents are parasitic on the contents of the relevant experiences. So attentional states don't represent properties which aren't represented by the relevant visual experience, and don't represent the properties which are represented by the visual experience with more or less detail than the visual experience itself. The contents of the attentional states really just are parts of the contents of the relevant visual experiences. So it does seem that in this case, it is hard to avoid the conclusion that the difference in phenomenology generated by attentional shifts must be explained by the fact that a given content *is the content of an attentional state*.

To sum up: shifts in the focus of a subject's attention present problems for the intentionalist which cannot be handled simply by appeal to a relaxed view about the properties which can be represented in perceptual experience. To explain the differences in phenomenology which accompany attentional shifts, the intentionalist has to appeal to a distinction between perceptual phenomenology and the phenomenology of attention. But even if this move saves minimal intentionalism, it does so only by introducing another state type — attentional states — with an associated phenomenology, which in turn raises the question of whether the phenomenal characters of states of this type are explained by the contents of those states.

If we answer this question in the negative, then we must concede that global intentionalism is false, and hence that phenomenal character cannot, in general, be explained in terms of representational content. The global intentionalist does have the resources to resist this challenge, and explain how the contents of attentional states might determine the phenomenal characters of those states; but in order to do so she must either say some *prima facie* odd things about attention, or abandon intermodal global intentionalism.³⁷

References

- David Armstrong, 1968. *A Materialist Theory of the Mind*. London: Routledge & Kegan Paul.
- A. J. Ayer, 1940. *The Foundations of Empirical Knowledge*. London: MacMillan & co.
- Ned Block, 1999. Sexism, Racism, Ageism and the Nature of Consciousness. *Philosophical Topics* 26(1-2).

³⁶For discussion, see e.g. Tye (1995a) and Block (2003).

³⁷Thanks for helpful discussion of these issues to Casey O'Callaghan and the participants in my graduate seminar on mental content at Notre Dame in the spring of 2007.

- Ned Block, 2003. Mental Paint. In *Reflections and Replies: Essays on the Reflections and Replies: Essays on the Philosophy of Tyler Burge*, edited by Martin Martin Hahn and Bjorn Ramberg. MIT Press.
- Alex Byrne, 2001. Intentionalism Defended. *Philosophical Review* 110:2:199–240.
- John Campbell, 2002. *Reference and Consciousness*. New York: Oxford University Press.
- Fred Dretske, 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- B. A. Eriksen and C. W. Eriksen, 1974. Effects of Noise Letters Upon the Identification of a Target Letter in a Nonsearch Task. *Perception & Psychophysics* 16(1):143–149.
- Gilbert Harman, 1990. The Intrinsic Quality of Experience. *Philosophical Perspectives* 4:31–52.
- Christopher Hill, 1991. *Sensations*. Cambridge University Press.
- Terry Horgan and Uriah Kriegel, 2007. Phenomenal Epistemology: What is Consciousness That We May Know It So Well? *Philosophical Issues* 17:123–144.
- Mark Johnston, 2001. The Authority of Affect. *Philosophy and Phenomenological Research* 63:181–214.
- William Lycan, 1987. *Consciousness*. MIT Press.
- Fiona Macpherson, 2006. Ambiguous Figures and the Content of Experience. *Noûs* 40(1):82–117.
- Eric Marcus, 2006. Intentionalism and the Imaginability of the Inverted Spectrum. *Philosophical Quarterly* 55(224):321–339.
- Bernard Nickel, 2007. Against Intentionalism. *Philosophical Studies* 136:279–304.
- Alva Noë, 2004. *Action in Perception*. MIT Press.
- Casey O’Callaghan, 2007. *Sounds: A Philosophical Theory*. New York: Oxford University Press.
- Stephen Palmer, 1999. *Vision Science*. Cambridge, MA: MIT Press.
- Christopher Peacocke, 1983. *Sense and Content: Experience, Thought, and Their Relations*. Oxford: Oxford University Press.
- Christopher Peacocke, 1992. *A Study of Concepts*. Cambridge, MA: MIT Press.
- Eric Schwitzgebel, 2008. The Unreliability of Naive Introspection. *Philosophical Review* 117(2):245–273.
- John Searle, 1983. *Intentionality*. New York: Cambridge U.P.
- Sydney Shoemaker, 1994. Phenomenal Character. *Noûs* 28:21–38.
- Susanna Siegel, 2006. Subject and Object in the Contents of Visual Experience. *Philosophical Review* 115(3):355–388.

- A. D. Smith, 2002. *The Problem of Perception*. Cambridge, MA: Harvard University Press.
- Jeff Speaks, forthcoming. Transparency, intentionalism, and the nature of perceptual content. *Philosophy and Phenomenological Research* .
- Michael Tye, 1995a. Blindsight, Orgasm, and Representational Overlap. *Behavioral and Brain Sciences* 18:268–9.
- Michael Tye, 1995b. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, MA: MIT Press.
- Michael Tye, 2000. *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
- Michael Tye, 2002. Representationalism and the Transparency of Experience. *Noûs* 36:137–151.
- Ted Warfield, 1999. Against Representational Theories of Consciousness. *Journal of Consciousness Studies* 6(1):66–69.