Philosophical Studies 164, 2: 405-418 (2013)

Visual Prominence and Representationalism

Todd Ganson and Ben Bronner

Abstract A common objection to representationalism is that a representationalist view of phenomenal character cannot accommodate the effects that shifts in covert attention have on visual phenomenology: covert attention can make items more visually prominent than they would otherwise be without altering the content of visual experience. Recent empirical work on attention casts doubt on previous attempts to advance this type of objection to representationalism and it also points the way to an alternative development of the objection.

Keywords Attention, Representationalism, Determinacy of Representation, Visual Prominence

1 Introduction

Representationalism about visual experience is the view that the phenomenal character of a visual experience supervenes on its representational content.[1] Several philosophers have recently argued that representationalism about visual experience is unable to accommodate the effects that shifts in attention have on visual phenomenology (Block 2010, Chalmers 2010, Nickel 2007, Speaks 2010, and Wu 2011).[2] Of course, *overt* shifts in attention involving muscular contractions (e.g. focusing of the eye or turning of the head) do not pose any obvious problem for representationalism: the relevant muscular contractions systematically alter what inputs the visual system is responding to and thereby alter representational outputs. The issue is whether representationalism can accommodate *covert* shifts in attention, which do not alter the pattern of light falling on the retina.

Nickel's clear formulation of this type of objection to representationalism will serve as our starting point. Nickel focuses on an instance of *voluntary* covert attention. Voluntary attention is attention that is intentionally allocated with the aim of accomplishing a goal or task. It is sometimes referred to as *goal-directed* or *endogenous attention*. Attention can also be captured involuntarily by stimulus events that are unrelated to one's current goals. This involuntary form of attention is sometimes referred to as *stimulus-driven* or *exogenous attention*. Our main goal in this paper is to show that critics of representationalism do best to focus on *involuntary* covert attention.

Another distinction among types of attention that will play a role in what follows is the distinction between space-based, object-based, and feature-based attention. It has been known for some time that attention allocated to a spatial location typically bolsters performance in discrimination and reaction-time tasks with respect to the attended location. Researchers subsequently discovered two further forms of attention which enhance performance in such tasks: object-based and feature-based attention. As we shall see, all three forms raise worries about representationalism.

We begin in section two by offering empirical reasons for thinking that representationalism can accommodate shifts in voluntary covert attention. Then in section three—the heart of the paper—we develop a novel argument against representationalism which we call *the argument from prioritization*, an argument that draws on recent research on involuntary covert attention, feature-based attention, and object-based attention. In the final section we explore two promising responses to the argument from prioritization.

2 Nickel's argument

If shifts in covert attention can affect the phenomenal character of visual experience without bringing about any change in representational content, then representationalism about visual experience is false. Nickel attempts to establish the antecedent of this conditional by way of the following example. When looking at a 3 x 3 grid drawn on white paper, you can experience various subsets of the grid's nine squares as prominent. Compare two experiences which differ as a result of a difference in how you direct your attention: in the first make the four corner squares and the center square appear prominent, and in the second make the remaining four squares appear prominent. In each case the prominent squares seem to pop out at you, and yet the squares all consistently look to be equidistant from you (if you situate the drawing appropriately). Nickel claims that we have a pair of visual experiences that differ in phenomenal character without differing in content.

A familiar strategy for defending representationalism is to insist that shifts in attention affect visual phenomenology only when they also affect the content of visual experience. For example, Nanay (2010) proposes that shifts in attention always affect the determinacy of perceptual content: visual attention has the effect of *increasing* the determinacy of representational content. Returning to Nickel's two experiences, note that the prominent squares are the squares to which attention is deployed. Although Nanay is not explicit on the matter, presumably visual attention is supposed to enhance determinacy of representation with respect to the spatial properties of lines bounding attended squares vs. lines not bounding attended squares and thereby increase the prominence of the attended squares. If Nanay is on the right track, attention will alter visual phenomenology only when it affects content. Representationalism is safe.

We agree with Nanay that attention can affect the determinacy of visual representation. However, we do not accept this view for the same reasons and we reject his claim that shifts in attention *always* affect determinacy of content. These two points of contention are interconnected. We begin with our disagreement over how to defend the thesis that covert attention can affect the determinacy of representation.

Nanay motivates the view that attention can affect the determinacy of content by appealing to the following example:

> Attention makes the attended property more determinate. If I am attending to the color of my office telephone, I attribute very determinate (arguably super-determinate) properties to it. If, as it is more often the case, I am not attending to the color of my office telephone, I attribute only determinable properties to it (of, say, being light-colored or maybe just being colored). (Nanay 2010, p. 266)

When Nanay says that attention makes the attended property more determinate, he does not mean, of course, that the property of the phone is made more determinate by our attending to it. He means that experience attributes a more determinate property to the telephone as a result of our increased attention. The problem is that the case at hand shows no such thing. Even if we grant that covert attention makes us better able to identify the (maximally) specific shade of color before us, there are a number of competing explanations of how it is that covert attention contributes to this result. One possibility is indeed that the representational content has become more determinate, but the example at hand does nothing to rule out the alternative hypothesis that covert attention improves discriminatory capacities by reducing external noise or the further alternative hypothesis that it affects our powers of discernment by improving processing at the decisional level.

So why believe that covert attention sometimes affects the determinacy of perceptual content? In a recent review article, Carrasco and Yeshurun defend the view that covert attention improves performance in a variety of tasks by enhancing the "spatial resolution" at the attended location, allowing "us to better resolve the fine details of the visual scene at that location." (Carrasco and Yeshurun 2009, p. 65) They argue at length that the improved performance they observed could not be explained by other prominent accounts of attentional mechanisms. The improvements in question were not explicable in terms of enhancement of decisional processing, reduction of location uncertainty, or reduction of external noise. The best explanation, they submit, is signal enhancement: attention improves the quality of the stimulus representation.

Although covert attention seems to affect determinacy of representation under some conditions, we should reject Nanay's thesis that attention *always* improves the determinacy of representation. In the very same review article, Carrasco and Yeshurun note an exception to the general claim that shifts in covert attention correlate with changes in spatial resolution (and thus to changes in the determinacy of perceptual content). Their experiment employed texture patterns—arrays of line segments either (i) all oriented at the same angle or (ii) with a small patch of the lines oriented to be perpendicular to the rest of the lines. Subjects were presented with a series of cues, displays, and masks, and were asked to indicate which of the displays were of type (ii). (This testing paradigm had been used previously to measure the effects of attention on spatial resolution.) Under a variety of conditions, small but not large cues improved accuracy. Ostensibly, when attention was drawn to relatively large portions of the visual array, the deployment of attention did not improve spatial resolution. This study did not test discrimination of color, luminance, or other features, but it does provide direct evidence that, for

at least one low-level feature (namely, line orientation), the view of attention Nanay offers is too simplistic.

Yet another problem for Nanay's generalization is that it focuses entirely on determinacy of representation. A highly influential study by Carrasco et al. (2004) argues that covert attention can have a very different effect on visual content—it can boost stimulus contrast.[3] The authors write:

> [W]hen observers' transient attention was drawn to a stimulus location, observers reported that stimulus as being higher in contrast than it really was, thus indicating that attention changes appearance. (Carrasco et al. 2004, p. 309)

The subjects of the experiment were asked to make comparative judgments regarding contrast, not judgments regarding the absolute value of the stimulus, so the illusion here is akin, at least in this respect, to the Müller-Lyer illusion. Attention increased the apparent brightness of the target stimulus compared to both a neutral stimulus and an unattended stimulus. Clearly this kind of effect on sensory representation is not a shift to greater determinacy.

It is worth noting that in a subsequent study of the two other dimensions of color appearance—hue and saturation—covert attention was found to increase apparent saturation but did not affect hue (Fuller and Carrasco 2006). The authors of the study offer an attractive explanation of this difference:

> It is reasonable to speculate why attention increases the appearance of these dimensions; for instance, increased contrast and saturation facilitate the discrimination of the features of the signal, and make it easier to discriminate the signal from the background. By comparison, there is no a priori reason why attention should affect apparent hue in one direction or another. (Fuller and Carrasco 2006, p. 4043)

Increases in brightness and saturation straightforwardly facilitate discrimination in a way that modifications of hue would not.

These remarks on brightness and saturation make vivid the point that shifts in covert attention can affect how things *look*. Covert attention can affect how things appear *visually*, altering, e.g., apparent brightness. Accordingly, we should not think of attention as having "its own *sui generis* phenomenology, which is distinct from visual phenomenology" (Speaks 2010, p. 333). A difference in apparent brightness is an aspect of *visual* output and not something additional to that output.

Another lesson of the foregoing is that attention might be enhancing visual prominence in more than one way. Attention might be making certain items stand out visually by increasing determinacy of representation, making the attended items more prominent than they would otherwise have been by augmenting the amount of information received. Another way that attention might affect the prominence of things is by boosting stimulus contrast: visible features (e.g. brightness and saturation) become exaggerated enough to make the attended items more conspicuous than they otherwise would have been. It is plausible to think that the visual prominence involved in Nickel's example of the 3 x 3 grid is best explained by appeal to one or both of these potential factors.

The upshot is that Nickel's example of visual prominence does not pose a serious threat to representationalism. The impact attention is having on visual phenomenology plausibly coincides with a difference in representational content. In the following section we turn to evidence suggesting that there is yet another way that attention can augment visual prominence. This further case of prominence is more difficult to reconcile with a representationalist view of phenomenal character.

3 The argument from prioritization

So far we have been discussing evidence suggesting that covert attention can alter representational content in a manner that bolsters our capacities to classify stimuli. We turn next to evidence for a different kind of effect of covert attention. There is a wide variety of data showing that covert attention can improve reaction time (RT) in various tasks *without improving the accuracy of our discriminations*. The evidence in question derives from experiments investigating how voluntary and involuntary attention affect performance with respect to discrimination and RT.

In a recent review article, Prinzmetal and Landau (2010) present a powerful case in support of the claim that involuntary attention involves different mechanisms from voluntary attention. Their argument draws on experiments testing the effects of voluntary and involuntary attention on performance in various discrimination tasks, e.g. the task of discriminating line orientation. In one experiment (Prinzmetal et al. 2009) the sudden appearance of a precue attracted covert attention to some location within each visual array. In the involuntary condition subjects knew that the cue carried no information regarding the impending location of the target lines. The cue appeared at locations that were random relative to the location of the target lines that were about to appear. Although the subjects knew the cue to be irrelevant, their attention was nonetheless captured by it. In the rare cases in which the location of the cue did, by chance, coincide with the location of the target lines, reaction times improved but accuracy rates did not. In the voluntary condition subjects knew that the cue carried information about the impending location of the target lines, and under these circumstances both reaction times and accuracy improved. Prinzmetal and Landau take these experiments to provide further confirmation for the widely held view that voluntary attention "enhances the perceptual representation so that more

information is available about an attended object." (Prinzmetal and Landau 2010, p. 46) Involuntary shifts in attention, by contrast, do not seem to enhance perceptual processing. Knowing the cue to be spurious, subjects had no reason to allocate perceptual processing resources to the attended location. The cue still drew attention to its location and thereby improved reaction times when by chance the target lines were at the cued location.

It is important to note that the results just mentioned are not at all anomalous. Prinzmetal and Landau insist that there is broad experimental support for these conclusions regarding voluntary and involuntary covert attention:

> In summary, in every RT study that we have conducted, subjects were faster on valid trials than on invalid trials, regardless of whether the cue was predictive or not. In every study designed around accuracy with voluntary attention that we have conducted, subjects are more accurate on valid than on invalid trials. However, across a large number of studies with only involuntary attention, we have not observed even one study where involuntary attention improved accuracy. (Ibid., p. 50)

When involuntary attention is drawn to a spatial location that happens to coincide with the target, typically RT improves but accuracy does not.

Consider a typical subject, S, involved in the task of discriminating line orientation described above. Suppose that S, knowing the precue is uninformative, is confronted with the same display at two different times and the only difference is that at the earlier of the two times the uninformative cue happens to call attention to the location of the target to come. Suppose, further, that S's responses are typical: S's RT in response to the first exposure to the display is shorter and the accuracy of the response is the same in both cases. The question that interests us at present is the following: How do we best account for the fact that S's RT is shorter during the

first of the two exposures to the display in question? Since there is no improvement in accuracy, it is unlikely that we are dealing with a difference in content along the lines of greater determinacy or contrast. Nor is there any other obvious way that the two experiences differ in content: the display is just the same in each case and so is the focal point.

Not only does it seem unlikely that we are dealing with signal enhancement; it is also doubtful that the improvement in RT is due to a difference in decisional processing, reduction of noise, or reduction of location uncertainty. Take first the suggestion that covert attention is altering what decisional criteria are in play. Since the cues are uninformative and understood to be so, S has no reason to adjust what quantity or quality of evidence is required in making a decision. The shifts in attention here do not seem to correlate with differences in decision procedure. Second, consider the hypothesis that reduction of noise is responsible for the improved reaction time. The problem with this hypothesis is that the target was presented all by itself. No aspect of the display introduced noise external to the target. Finally, Prinzmetal et al. insist that their results are not to be explained by appeal to reduction of location uncertainty: "Participants were nearly 100% correct in localizing the targets." (Prinzmetal et al. 2009, p. 354)

So how should we account for the improved RT? We have said that there is no reason to suppose that shifts in involuntary covert attention alter what standards the subject adopts regarding quality or quantity of evidence. Further, there is no reason to think that involuntary covert attention is affecting the quality or quantity of visual evidence: there is no improvement to the evidence either in the form of an increase in information from the attended location or in the form of a decrease in noise from outside the attended location. Our alternative hypothesis is that there is a difference in the *prioritization* of the evidence issuing from the attended location. Involuntary covert attention does not change the visual evidence itself or the standards of

evidence in play; it affects the *order* in which the array of visual evidence is likely to be assessed. (For ease of exposition we assume that the processing is serial in character. If processing is instead parallel, prioritization will likely involve a difference in the rate at which the various portions of the array of evidence are assessed.[4]) Involuntary attention has the effect of prioritizing visual data for epistemic tasks.

It is reasonable to suppose that this prioritizing of data is *phenomenologically driven*, that attention is enhancing performance *via* its effects on phenomenology. In the aforementioned experiments testing the behavioral consequences of voluntary and involuntary attention, the subjects were responding to how things appeared visually to them. Voluntary and involuntary attention both improved RT. This improvement in RT is readily intelligible on the plausible assumption that both forms of attention are *increasing visual prominence*. More specifically, attention is systematically boosting the visual prominence of items at the attended region, making them more visually prominent than otherwise similar unattended items, and this phenomenal difference is having a causal impact on subjects' responses.

The relevance of these remarks to representationalism is straightforward. In the case of involuntary attention the boost in visual prominence does not appear to be accompanied by any change in content. That is, we seem to have a case where phenomenal character is different without any correlate difference in visual content. We will refer to this argument against representationalism as *the argument from prioritization*.

Before we offer further empirical support for the idea that attention can affect the order in which visual evidence is assessed, we want to address some potential objections to the argument from prioritization. First, one might challenge the claim that involuntary attention leaves representational content unaltered. We can grant that involuntary attention does not enhance

representational content in a manner that improves the quantity or quality of evidence regarding the character of the distal stimulus (e.g. its orientation), but this concession is compatible with the idea that involuntary attention enchances representational content in some other way that facilitates recognition of the presence of the stimulus. This line of response to the argument from prioritization would gain some teeth if we could identify a plausible candidate for a change in content that would have the relevant effect on RT—something other than change in determinacy or change in stimulus contrast, both of which improve discrimination of stimulus features like orientation. The problem is that there does not seem to be any plausible candidate on the horizon. Recall that the scene is just the same and no eye movements have occurred. It is deeply puzzling what sort of change in content could be at work here.

Let's grant, then, that involuntary attention leaves unaltered the quality and quantity of visual evidence, that it instead affects how we respond to the data. Our preferred account captures one way this might go:

> PRIORITIZATION. Involuntary attention alters the order in which evidence is likely to be assessed.

But why prefer this account to alternatives which take shifts in involuntary attention to go hand-in-hand with subtle alterations in the standards of evidence or in the weighing of evidence? Consider two such alternatives to PRIORITIZATION:

> STANDARDS. While subjects have no reason to lower standards of evidence for attended regions, they are nevertheless doing so in a systematic manner.
>
> STIMULATION. Although subjects' evidence and standards have not been affected by these shifts in involuntary attention, the precue has an adverse effect on the subjects' assessment of evidence. The precue does not alter the quality or quantity of evidence at

the moment of attending, i.e. just after the precue has flashed, but the precue does add to the total amount of stimulation deriving from the attended region over the brief period of time that includes the flashing of the precue. This greater accumulation of data from the attended region leads subjects to think that a certain threshold for evidence has been met when it has not.

Nothing we have said so far shows that PRIORITIZATION is preferable to STANDARDS or STIMULATION.[5]

Both of these alternatives are *prima facie* less attractive than PRIORITIZATION. STIMULATION is unattractive because it assigns no role to attention in accounting for the better RT. There is overwhelming evidence that attention improves performance in a wide variety of tasks. Why doubt that the improved RT in the case at hand is due to the fact that the subjects in question have their attention drawn to the location at which the target appears? And while STANDARDS allows that attention plays a role in the improved RT, it leaves us in the dark about why attention would cause us to lower our standards of evidence. STANDARDS is unsatisfying as an explanation because there is no intelligible connection between involuntary attention and lowering of standards of evidence.

The argument from prioritization relies crucially on the hypothesis that visual attention can affect the order in which visual data is assessed without enhancing or otherwise altering the data. This hypothesis is not original to us. It was introduced by Moore and Egeth (1998) to account for aspects of feature-based attention and later endorsed by Shomstein and Yantis (2002) in the context of discussing object-based attention. We begin with the former.

So far we have been thinking of visual attention as something allocated to spatial locations, but we can also shift attention to specific features in the scene before us. Moore and Egeth distinguish two ways that feature-based attention might affect visual processing:

> Suppose that you are to meet a friend at a crowded baseball game and that you know that he will be wearing a red sweatshirt. One possibility is that looking for the red sweatshirt makes the red things in the crowd more salient (perhaps even brighter) than they would otherwise be. This would suggest that attending to a feature directly affects visual processing by directly enhancing the sensory quality of items that possess that feature. Alternatively, the sensory quality of the scene may not change as a function of feature-based attention. Looking for your friend's red sweatshirt may simply cause you to look at red items before looking at other items in the crowd. This would suggest that, rather than directly changing the sensory quality of a scene, feature-based attention results in prioritization of items that possess the attended feature. (Moore and Egeth 1988, p. 1296)

When Moore and Egeth speak of the possibility that feature-based attention might enhance the sensory quality of items possessing the feature in question, they have in mind the kinds of changes in representational content that make those items more readily discriminable. The experiments they performed spoke against this possibility and in favor of the alternative suggestion that attention affects prioritization of data for epistemic actions like shifting direction of gaze.

Moore and Egeth used visual arrays of numbers and letters to test the speed and accuracy of alphanumeric categorization and numeric identification. Sometimes the visual array remained visible until the subject responded. Under these conditions accuracy rates were uniformly high and the the response variable was speed. Other times the array was briefly flashed. Under these

conditions reaction times were uniformly short and the response variable was accuracy. Providing subjects with prior knowledge of, e.g., the color that the target will be (if it is present) improved reaction times when the array remained visible until the subject responded. However, when the array was briefly flashed, previous knowledge of target properties did not improve the accuracy of responses. Moore and Egeth believe that feature-based attention serves to "highlight the representations of stimuli that possess that feature" without otherwise changing the representation (ibid, p. 1307). On this account, "highlighting" prioritizes certain stimuli for epistemic action, and we see this highlighting at work with the improved reaction times in the time-unlimited condition. If feature-based attention *enhanced* the representation of stimuli possessing the attended feature, then accuracy should improve when the visual array is briefly flashed. But, as we said, accuracy rates did not improve.[6]

Consider next the phenomenon of object-based attention. It is widely recognized that some cognitive tasks exploit a selective mechanism that operates specifically on representations of objects as opposed to representations of spatial locations. Egly, Driver, and Rafal (1994) offer an elegant experiment that shows both spatial attention and object-based attention at work. Subjects are confronted with a display containing two adjacent rectangles oriented either horizontally or vertically. While the eyes remain fixed at the center of the display, a cue draws attention to the shorter side of one of the rectangles. In the usual case the cue coincides with the target; at other times the target appears in one of two locations equidistant from the cue—either at the other end of the cued rectangle or in the adjacent rectangle. RT was best when the cue was valid (i.e. coincident with the target)—an example of spatial attention at work. When the cue was invalid, RT was better for targets present in the same rectangle. The latter difference in RT is evidently due to something other than spatial attention, for the targets were equidistant from

the cue in invalid trials. Rather, attention seems to be operating on the representation of the highlighted object.

These and other experiments have convinced researchers that we have to let go of the image of visual attention as something like a graded spotlight illuminating a convex region of space. There is now widespread agreement that attention engages object-based representations as well as space-based representations. At the same time the underlying nature of object-based attention remains a matter of controversy. Among the most prominent models are the sensory-enhancement hypothesis and the attentional-prioritization hypothesis. According to the former, object-based attention improves performance by enhancing representational content. Shomstein and Yantis argue against this model, favoring the prioritization account according to which object-based attention "affects the order in which different regions of the scene are visually investigated when multiple attentional 'glimpses' are required." (Shomstein and Yantis 2002, p. 42) Their argument is straightforward. If attention bolsters object representation, then it should also augment the effect that distracting flankers within the same object have on performance. In fact, the presence of flankers in the same object did not have a greater impact on performance than equidistant flankers in adjacent objects.

4 Responses to the argument from prioritization

In this final section we consider two promising responses to the argument from prioritization. One response challenges the argument by articulating a rival explanation of how involuntary attention improves RT. On this alternative account, involuntary attention raises no worries about representationalism. The other response is concessive: the argument from prioritization undermines representationalism in its familiar form, but it leaves in good standing an attractive

alternative (quasi-representationalism) that is perfectly in line with the spirit of representationalism. We begin with the latter.

Representationalism about visual experience, as we have defined it, has no obvious place for what is sometimes referred to as the *assertoric character* of visual experience.[7] Visual experience has the role of *informing/reporting*, and it is plausible to suppose this role helps to shape the phenomenal character of visual experience. On the assumption that this functional role of informing makes a distinctive contribution to visual phenomenology,[8] we can interpret the experimental data reviewed above as suggesting that differences in *amplitude of assertion* can affect visual phenomenology: two visual experiences can report one and the same content with a difference in emphasis, a difference that affects how prominent things appear. After all, once we allow that some experiences have an assertoric character, it is a small step to acknowledge differences in *assertoric strength*, i.e. differences in amplitude of assertion. The cases of involuntary attention discussed above seem to have the effect of amplifying visual representation, and they do so without affecting the determinacy of perceptual content or any other aspect of perceptual content. What we have here is a difference, not in what is reported, but in what is stressed. Not a difference in veridicality conditions, but a difference in emphasis. Unlike the representationalist, the quasi-representationalist[9] can acknowledge factors additional to veridicality conditions that determine phenomenal character. The quasi-representationalist can allow that part of what shapes the phenomenal character of visual experience is the degree of emphasis with which truth-evaluable content is reported.

Quasi-representationalism may well be worth pursuing further, but as it stands the view is too underdeveloped to serve as a clear, well-motivated response to the argument from prioritization. While the idea that sensory states have the function of reporting/informing is

reasonably clear, the further suggestion that these sensory reports can differ in emphasis is more difficult to make sense of. If James (1890) were right that "Accentuation and Emphasis are present in every perception," then we would expect differences in emphasis to be perfectly familiar. In fact, the idea of emphasis in vision is fairly obscure. What, exactly, would it be like for two visual experiences to differ from one another merely with respect to emphasis?[10]

A more straightforward strategy for the representationalist is to try to locate a weakness in the argument from prioritization. This argument against representationalism begins with the claim, supported by considerable experimental data, that involuntary covert attention can improve RT in a variety of tasks without improving accuracy. For reasons discussed at length in section two, the differences in RT do not seem to correlate with differences in representational content. Now, representationalism is in trouble only if involuntary attention is like voluntary attention, only if it improves performance thanks to its effect on how things look—making items at the attended location phenomenologically more salient than they would otherwise be and thereby improving RT. The worry about representationalism would diminish greatly if we could discover a satisfying, alternative explanation of the difference in RT that does not posit any phenomenal difference.

Consider once again our subject S involved in the task of discriminating line orientation described above. Assume for the moment that the contents of visual experience are *rich* in the sense that they often contain more than the subject notices at a given time. Now, even if S is enjoying an experience rich in information about the target, it is doubtful that S is in a position to *utilize* this information about the target in forming a person-level judgment about it until S *notices* the seen target.[11] That is, S is in a position to issue a response only once S has made the target an object of attention. Presumably when the precue is predictive/valid, S has no need to

shift attention to a different spatial location in the display and his RT is typically faster as a result. The precue affects the order in which S assesses the visual evidence for the simple reason that it affects how attention is directed. While a valid precue serves to direct attention to the target, an invalid cue amounts to a kind of distraction. On this way of explaining the difference in RT, there is no obvious need to posit changes in phenomenology corresponding to shifts in involuntary attention.

This explanation relies on a familiar, folk-psychological strategy for addressing failures to act on a target stimulus that is in plain sight. Even though the subject is otherwise in a position to discriminate the target, we say she is not yet acting because she has yet to notice the target. Precisely what principle is at work in such explanations is a matter that we need not settle here. It is sufficient for present purposes to note that we commonly find an omission to act on a target intelligible as long as the subject has yet to notice the target.

Invoking this strategy in the present context is not unproblematic. In everyday contexts this way of explaining failures to act on a target in plain sight makes sense to us, in part, because we think that the target is not "popping out" in the right way for the subject. That is, the subject's failure to respond appropriately seems to be partly grounded in phenomenology. We can reasonably wonder whether we still have a fully satisfying explanation once this phenomenological aspect is removed from the explanation. (Note that the argument from prioritization relies on an explanation of the difference in RT that is substantially the same as the current proposal with one important difference. The difference is that the former explanation retains the idea that visual prominence is playing a causal role in determining the order in which evidence is assessed. This *assessing of visual evidence* is just a way of attending to the evidence:

a subject is not in a position to respond to the discrimination task until she has assessed the evidence relating to the target.)

One move available to the representationalist at this point is to deny that the contents of visual experience are rich. Suppose that, in order for S to see the target, S must attend to/notice the target. In that case the difference in RT is, once again, due to the extra time required, in the case of the invalid cue, to shift attention to the location of the target. Prior to this shift in spatial attention there will presumably be a difference in phenomenology between the two cases (valid vs. invalid cue), but this difference will be accompanied by a difference in content. In the valid case one is seeing the target; in the invalid case one is not yet seeing the target.[12] If this account of the difference in RT is correct, there is no reason to worry that involuntary attention is having an impact on visual phenomenology without altering the content of visual experience.

The elegance of this response to the argument from prioritization cannot be gainsaid. There are, however, two points to bear in mind. First, this response is an attempt to *accommodate* the phenomenon of involuntary attention, but the denial of richness may well turn out to be *incompatible* with the phenomenon. It could turn out that in standard cases of involuntary attention a stimulus captures attention *because it stands out visually*, that things can grab our attention *because of the way they look to us*. This sort of causal explanation evidently presupposes that *our seeing things* is one thing, *our consciously noticing them* another. Second, this kind of response is not available to Dretske, Tye, and other representationalists who defend the view that the content of visual experience is rich. It remains unclear whether these philosophers can tell a fully satisfying story about the effects of involuntary covert attention—not to mention feature-based and object-based attention.

Acknowledgements Special thanks to Neil Mehta and an anonymous reviewer for insightful comments. Thanks also to Benj Hellie for useful comments on an antecedent to this paper.

References

Block, N. (2010). Attention and mental paint. *Philosophical Issues* 20, 23-63.

Carrasco, M. (2011). Visual Attention: The past 25 years. *Vision Research* 51, 1484-1525.

Carrasco, M., S. Ling, and S. Read. (2004). Attention alters appearance. *Nature Neuroscience* 7, 308-313.

Carrasco, M. and Y. Yeshurun. (2009). Covert attention effects on spatial resolution. *Progress in Brain Research* 176, 65-86.

Dretske, F. (2007). What change blindness teaches about consciousness. *Philosophical Perspectives* 21: 216-230.

Chalmers, D. (2010). *The character of consciousness*. Oxford: Oxford University Press.

Egly, R., J. Driver, and R. D. Rafal. (1994). Shifting visual attention between objects and locations: Evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General* 123, 161-177.

Fuller, S. and M. Carrasco. (2006). Exogenous attention and color perception: Performance and appearance of saturation and hue. *Vision Research* 46, 4032-4047.

Hall, R. (2008). If it itches, scratch! *Australasian Journal of Philosophy* 86, 525-535.

James, W. (1890). *Principles of Psychology*, v. 1. New York: Henry Holt and Co.

Klein, C. (2007). An imperative theory of pain. *The Journal of Philosophy* 104, 517-532.

Lycan, W. (1996). *Consciousness and experience*. Cambridge: MIT Press.

Matthen, M. (2005). *Seeing, doing, and knowing: A philosophical theory of sense perception*. Oxford: Oxford University Press.

Moore, C. M. and H. Egeth. (1998). How does feature-based attention affect visual processing? *Journal of Experimental Psychology: Human Perception and Performance* 24, 1296-1310.

Nanay, B. (2010). Attention and perceptual content. *Analysis* 70, 263-270.

Nickel, B. (2007). Against intentionalism. *Philosophical Studies* 136, 279-304.

O'Reagan, J. K, and A. Noë. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* 24, 939-1011.

Prinzmetal, W., V. Long, and J. Leonhardt. (2008). Involuntary attention and brightness contrast. *Perception & Psychophysics* 70, 1139-1150.

Prinzmetal, W., A. Zvinyatskovskiy, P. Gutierrez, and L. Dilem. (2009). Voluntary and involuntary attention have different consequences: The effect of perceptual difficulty. *The Quarterly Journal of Experimental Psychology* 62, 352-369.

Prinzmetal, W. and A. N. Landau. (2010). Dissecting spatial visual attention. In V. Coltheart (Ed.), *Tutorials in Visual Cognition* (pp. 43-66). Hove: Psychology Press.

Pryor, J. (2000). The skeptic and the dogmatist. *Nous* 34, 517-549.

Schneider, K. A. (2006). Does attention alter appearance? *Perception & Psychophysics* 68, 800-814.

Shomstein, S., and S. Yantis. (2002). Object-based attention: Sensory modulation or priority setting? *Perception & Psychophysics* 64, 41-51.

Smithies, D. (2011). What is the role of consciousness in demonstrative thought? *The Journal of Philosophy* 108, 5-34.

Speaks, J. (2010). Attention and intentionalism. *Philosophical Quarterly* 60 (239), 325-342.

Tumulty, M. (2009). Pains, imperatives, and intentionalism. *The Journal of Philosophy* 106, 161-166.

Tye, M. (2006). Nonconceptual content, richness, and fineness of grain. In T. S. Gendler and J. Hawthorne (Eds.), *Perceptual Experience* (pp. 504-530). New York: Oxford University Press.

Wu, W. (2011). What is conscious attention? *Philosophy and Phenomenological Research* 82, 93-120.

Yeshurun, Y. and M. Carrasco. (2008). The effects of transient attention on spatial resolution and the size of the attentional cue. *Perception & Psychophysics* 70, 104-113.

---

[1]We are assuming that human visual experiences have accuracy conditions (representational contents). For a recent defense of this way of thinking about the output of the human visual system, see Burge 2010.

[2]Wu is targeting the stronger, reductive claim that experiences are phenomenally similar or different *in virtue of* their similarity or difference in representational content—not the weaker supervenience claim.

[3]Some have challenged the results of Carrasco et al. 2004. See Schneider 2006 and Prinzmetal et al. 2008.

[4]We are assuming with Tye (2006) and Dretske (2007) that the content of visual experience is *rich* in the sense that it often contains more than the subject notices. For a very different view, see O'Reagan and Noë 2001.

[5]Prinzmetal et al. (2009) distinguish PRIORITIZATION and STIMULATION, but do not discuss STANDARDS. Prinzmetal and Landau (2010) argue that there is experimental data which favors STIMULATION over the serial version of PRIORITIZATION. They do not

discuss the parallel version of PRIORITIZATION, and they fail to realize that the data discussed actually count against STIMULATION as much as the serial version of PRIORITIZATION. For their data and argument, see Prinzmetal and Landau 2010, pp. 52 – 55. According to their accumulator model (STIMULATION), the cue is taken as evidence for a target at the cued location. Otherwise, evidence accumulates at the same rate for the attended and non-attended locations, and the threshold is the same for each location. It follows that, for the 2-target trials, the non-cued location should *never* win the race to threshold, yet the target at the non-cued location is reported 20% of the time. In order for their accumulator model to work, we need an additional stipulation: when the non-cued location wins, it does not simply "win the race to threshold." Instead, the cued location reaches threshold first but (somehow) this event fails to produce a response; meanwhile, the non-cued location continues to accumulate evidence and reaches threshold at a later time, producing a response. With this additional stipulation, however, the experimental evidence no longer counts against the serial version of PRIORITIZATION. We could say that subjects tend to begin by assessing the evidence for the cued location. On occasion the evidence passes threshold without triggering a response. In these cases the subjects assess the evidence at the non-cued location and respond to the target there. The serial version of PRIORITIZATION is now consistent with the timing of subjects' responses. It may seem implausible that evidence would pass threshold and yet not trigger a response. Note, though, that, if we disallow this possibility, the experimental evidence is inconsistent with STIMULATION as well as with the serial version of PRIORITIZATION. On the other hand, if we embrace this possibility, the experimental evidence is consistent with both of these models.

[6] The idea that feature-based attention does not affect the accuracy of perceptual judgments has been challenged by more recent research, as reviewed in Carrasco 2011. Our challenge to intentionalism rests, however, on the results of Prinzmetal and colleagues rather than those of Moore and Egeth. Moore and Egeth 1998 is still noteworthy for two reasons. First, it provides a model for the prioritization-based account we offer for the results of Prinzmetal et al. Second, it suggests the possibility that, under certain circumstances, feature-based attention affects the phenomenology but not the content of visual experience, a possibility that warrants further investigation.

[7]For the idea that experience can have an assertoric character, see Pryor 2000 and Matthen 2005. Of course, the senses do not literally assert anything: the output of vision is *like assertion* insofar as it plays the role of informing/reporting.

[8]It can be difficult to identify the assertoric character of experience. A brief comment on the nature of itches may help. It is somewhat doubtful that itches merely report or inform. After all, we have a hard time understanding how itches might be accurate or inaccurate. It seems more plausible to suppose that itches *command* scratching here now (Hall 2008, Klein 2007, Tumulty 2009). On the face of it, what is commanded by the itch, the content of the itch, is something that can be common to a tactile perception. We can also perceive that there is scratching here now, as when the command is satisfied. Attending to the felt difference between an itch and a correlate tactual perception of scratching is useful for picking out the assertoric character of the latter: we are evidently attending to something other than truth-evaluable content, which is roughly the same for both states.

[9]This term is due to Bill Lycan, who writes: "Now, I am myself a quasi-representationalist at best…because I have no official reason to deny that *part* of phenomenal character, even within a single sense modality, is constituted by functional rather than representational properties." (Lycan 1996, pp. 134-5)

[10]We thank an anonymous reviewer for raising this worry about quasi-representationalism.

[11]Cf. Smithies 2011.

[12]We thank an anonymous reviewer for pressing the importance of pursuing this kind of response to the argument from prioritization.