Attention, Visual Knowledge & Psychophysics: Discriminating the Determinable

James Stazicker, *Draft* 10/2011

Attending to the things you see is a reliable way to come to know about their visible properties. Recent experiments generate a challenge to this epistemologically important idea, because they show that attention changes the way properties appear in visual experience. To meet the challenge, I argue that visual experience represents determinable properties. More generally, I argue that we need to recognize the role of determinable properties in visual experience, if we're to understand what experiments tell us about conscious vision.

1. Two roles for attention

On the face of it, visual attention has an important epistemic role. By attending to the things you see, you come to know about them. Here I focus on knowledge of visible properties. I defend the simple idea that we come to know about these properties by attending to them as they figure in visual experience.

Understood in this way, attention is a means of fixing beliefs which inherit their content from visual experience: through attention, you come to believe that things have the very properties which are represented or manifest in visual experience. You believe that the cat has a certain shape and colour, say, because that shape and colour figure in your visual experience of her. Beliefs of this kind have a distinctive epistemic status. They inherit the distinctive reliability of visual processing (Dretske 1997), or the distinctive justification or warrant provided by conscious vision (McDowell 1994; Burge 2003).

Various proposals about the nature of attention are consistent with its playing this epistemic role. Locke described attention as a '*mode of thinking*' in which sensible ideas '*are taken notice of, and, as it were, registered in the memory*' (1690/1975: II.ix.1). We might take this literally (and in a contemporary sense of 'thinking'): by turning your thoughts to what you see you register it in memory, and thereby fix beliefs about it. Alternatively, we might understand attention as a more basic form of experiential selection, a pre-requisite for both thought and belief about the things you see (Campbell 2002). Either proposal can also be expressed in terms of information-processing:

> Attention is a … selection process … which is necessary for storage in a durable working memory store or for a conscious report about stimuli.
>
> (Lamme 2004: 864)

Whichever way we think about it, attention is a significant, reliable source of visual knowledge, because it allows you to form beliefs which inherit their content from vision.[1]

Work in psychology generates a challenge to this epistemologically important idea, because it shows that attention also plays a different role. Recent experiments confirm William James's claim that attention alters visual experience, highlighting attended properties by giving them a '*more intense*' appearance (1890: 425). Yet as James remarked, '*the intensification … never seems to lead the judgement astray*' (426). The challenge is to understand how this could be. How could attention so alter visual experience without distorting the way visible properties show up in it, and making visual experience a systematically misleading source of belief? How could attention be a reliable source of visual knowledge?

[1] Although recent work focuses largely on attention to objects or locations, see e.g. Huang and Pashler 2007 for an information-processing account of attention to intrinsic properties.

The challenge is especially pressing given the contemporary theory that visual experience is a form of representation, with a psychological role that's to be understood in terms of its representational content. According to this theory, psychological evidence of changes in the conscious character of visual experience is evidence of changes in the representational content of that experience. Some philosophers think this is a necessary truth. It's also a working assumption in the empirical literature I'll discuss.[2] Here the challenge is to understand how attention could transform visual representation without making it illusory. For example, Marisa Carrasco and her colleagues found that '*attention alters appearance*'; they conclude that attention produces '*nonveridical percepts … by emphasizing relevant details at the expense of a faithful representation*' (Carrasco et al. 2008: 1162).

Against this conclusion, I argue that visual experience represents *determinable* properties – properties which admit of more specific determinations, as *red* is determined by *crimson*, and *rectangular* is determined by *square*. Because visual experience is indeterminate in this sense, attention may change the way things appear visually, without making experience illusory. One determinate property determines many determinable properties. So one determinate shape or color may take on different appearances, in veridical experiences which represent its different determinables. I show how to understand the effects of attention in these terms. I argue that this account is as empirically plausible as an account in which attention makes visual experience illusory. The empirical evidence is consistent with the idea that visual attention is a reliable source of knowledge.[3]

There is also a broader issue here, about the relationship between experiment, judgement and visual experience. Carrasco's work is distinctive partly because she connects experiments in psychophysics with claims about the content of visual experience. The psychophysical data consist in subjects' discriminations between visible stimuli, where discrimination is a form of judgement. Psychophysicists use these data to investigate processes in the visual system. But it's unobvious how we should connect their work with claims about the content of visual experience. I argue that, to do this in a principled way, we need to recognize the role of determinable properties in visual experience and visual discrimination.

In this paper I focus on spatial properties, because here the argument that visual experience represents determinable properties is comparatively straightforward. Experiences of colors, for example, present a more complex challenge, because both the philosophical theory and the information-processing theory of color experience are more controversial.

For the purposes of this discussion, I assume that visual experience represents spatial properties as actually instantiated in the subject's environment. Furthermore, I assume the following connection between the phenomenology of visual experience and its representational content: visual experience represents a property only where that property seems, subjectively, to be apparent in visual experience.[4]

---

[2] The strongest version of the theory is Tye's (1995): the conscious character of visual experience is identical with its representational content. In weaker (supervenience) versions, every variation in the conscious character of visual experience entails a variation in its representational content. The scientists' working assumption is consistent with a still weaker version: the only variations in visual experience which show up in scientific psychology are variations in representational content. Cf. Block's (1978) thought that psychology keeps track of variations in functionally specifiable representations, but not of more fine-grained variations in *qualia*.

[3] James claimed both that attention makes visible properties appear more intense, and that attention makes these properties *clearer* (1890: 403). The latter claim can be captured in terms of the idea that attentive experience represents more determinate properties than inattentive experience (Stazicker 2011). The claim about intensity requires a different treatment.

[4] This assumption is compatible with various theories about the metaphysical connection between phenomenology and representational content. See footnote 2. My positive claims could also be translated

Although I won't argue for these assumptions, it's worth forestalling a potential confusion here. If you try to understand the idea that we're visually aware of properties on the model provided by our seeing particular objects, the idea can seem absurd. For example, it can seem to require an impossible confrontation with a bare universal – the sort of perception imagined in uncharitable readings of philosophic cognition in the *Phaedo* and the *Republic*. Perhaps visual awareness of a property entails seeing a particular. Perhaps visual experience is particular in that sense. But it doesn't follow that visual awareness consists only in seeing particulars, or that visual awareness of properties should be understood on the model of seeing particulars. Rather, the idea is that we're visually aware of the general ways things are, ways which other things could be, in addition to being visually aware of the particular things which are those ways. This idea is compatible with various metaphysical theories about what properties fundamentally are. It's silent on the question whether having a property is instantiating a universal, being a member of a set, comprising a trope which is a member of a set, or something else altogether.

According to some philosophers, to experience a property visually is already to see *that* something has that property, where seeing that *p* is a species of knowing that *p* (Brewer 2006; Stroud 2009). Similarly, some cognitive scientists think '*attention is essential to conscious perception*' (Mack and Rock 1998): to see something is already to attend to it. If either of these proposals is correct, some of the explanations here are constitutive rather than causal: the seeing explains the attending or the knowing as someone's brawling explains her breaking the law. I'll assume that the explanations are causal, because that makes the issues here plainest, but if necessary my arguments could be adapted as appropriate.

Much of the empirical work on attention concerns effects fairly early in visual processing, rather than attention understood as a means of taking up what you see in thought or belief. So we have to be careful when drawing on this work to learn about the latter form of attention. But many of the experiments do shed light on this form of attention, provided we approach the data case-by-case, discriminating appropriately. Carrasco's experiments are a case in point.

In §2 I explain how Carrasco's work seems to threaten the epistemic role of attention, and I explain how to diffuse the threat in terms of the idea that visual experience represents determinable properties. In §3, §4 and §5, I argue for the claims about visual experience and visual discrimination which my interpretation of the data requires: visual experience represents determinable properties; its phenomenology and psychological role should be understood accordingly.

2. Attention alters appearance

2.1 The experiments

It's a familiar fact that movement or change in peripheral vision draws your attention to its location. A light flashes, or someone fidgets during the talk, and you find your cognitive resources directed to the place where this happened. You might turn your head or eyes to this location, or you might focus your thoughts on it. And whether or not you do these things, there's a short-term effect on your visual system. In an effect which peaks around 100 milliseconds after attention is cued and disappears after another 20ms or so, processing in the visual cortex occurs faster for stimuli at the location in question, and contrast-sensitivity is increased for stimuli there – i.e. visual-cortical responses become more sensitive to differences in those stimuli (Carrasco and McElree 2001; Treue 2000).

---

into some disjunctivist frameworks, according to which visual experience is not a form of representation (e.g. Campbell 2002).

In a series of studies, Carrasco and her colleagues argue that this effect has a phenomenological upshot: this form of attention alters visual experience, highlighting attended aspects of a scene by exaggerating them; attention increases the apparent spatial frequency of a stimulus, the apparent size of a gap in it, and various other apparent properties (Gobell and Carrasco 2005).

Gap size is measured in degrees of visual angle. If a gap is 5˚ wide, two imaginary straight lines reaching from its edges to the centre of the fovea form an angle of 5˚. For standard experimental stimuli, spatial frequency is measured in sinusoidal cycles of variation per degree of visual angle (cpd). In Figure 1, spatial frequency increases from left to right, in that the rate of sinusoidal variation from light to dark and back increases from left to right. Intuitively, we can think of the spatial frequency of a stimulus as its rate of variation across space. For example, where a stimulus has a higher spatial frequency, it typically has a more detailed pattern.

Figure 2 illustrates the key experimental paradigm. Subjects maintain fixation at a point in the center of the display. A cue is flashed to capture their attention, either centrally or a little to one side. Soon afterwards, two stimuli appear for 40 milliseconds, a little to the right and left of fixation. Over many trials the spatial frequency of one stimulus (the *test* stimulus) is varied, while that of the other stimulus (the *standard*) is held fixed. The task is to report the orientation of the stimulus which has the higher frequency. Carrasco uses similar experiments to investigate the effect of attention on the appearance of other properties.

The *Point of Subjective Equality* (PSE) is a value or level of a property, such as a gap size of 0.2˚ or a spatial frequency of 0.5 cpd. The PSE is defined as the value of the test stimulus such that, statistically, a subject is at chance in reporting the orientation of that stimulus: where the test stimulus has this value, the probability of a subject choosing it as higher in the relevant value than the standard is 50%. For example, when attention is cued to the point of fixation, rather than to one of the stimuli, the PSE is roughly the value of the standard.

The crucial result is that the PSE was significantly lower where attention was cued to the test stimulus. For example, with the standard gap size at 0.23˚, the PSE was 0.20˚.[5] Though the numbers are small, this is a significant, clearly visible difference in size. See Figure 3. Taking this together with the data about increased cortical sensitivity, Carrasco concludes that attention enhances or exaggerates the visual signal processed for an item, with the phenomenological effect that an attended item '*appears as if it were*' greater than it really is in gap size and so on (Carrasco et al. 2004: 311); attention produces '*non-veridical percepts*', rather than a '*faithful representation*' (Carrasco et al. 2008: 1162).

How does this threaten the epistemic role of attention? Here's an intuitive first pass at formulating the threat. In §1 I said that attention is a means of fixing beliefs which inherit their content from visual experience: through attention, you come to believe that things have the very properties which are represented in visual experience. Carrasco's experiments show that where you attend to an object, you experience it as having a property it doesn't really have – a higher spatial frequency, a larger gap etc. So beliefs formed through visual attention turn out to be systematically false. They represent a detailed, gaping world, where by comparison reality is bland and closed.

This first pass is not quite right. It conflates different phenomena all of which are reasonably called 'attention'. In Carrasco's experiments, subjects attend to *both* stimuli, not just the stimulus that's cued. They have to attend to both stimuli in order to compare them. So attention has the effect Carrasco describes only where it's cued in a certain specific way.

[5] This is the PSE defined over a cohort of subjects, but similar results were found for individual subjects.

Vision scientists distinguish between *involuntary attention*, the short-term, automatic attention captured using an exogenous cue, and *voluntary attention*, the longer-lasting attention directed by your aims and decisions. Carrasco's subjects direct voluntary attention to both the test and standard stimuli, attending to their properties as they figure in visual experience. Subjects direct their thoughts to both stimuli, and perhaps they also select both of them with a pre-cognitive experiential focus, as discussed in §1. Voluntary attention also has pre-perceptual effects, but these differ systematically from the effects of involuntary attention (Prinzmetal and Landau 2008). With respect to spatial frequency, for example, voluntary attention does not have the effect which Carrasco reports for involuntary attention (Prinzmetal et al. 1998).[6]

However, involuntary attention does affect the process of attending voluntarily which subjects exploit in their comparative judgements. Where attention is cued exogenously to one of the stimuli, this process is disrupted: visual experience is briefly distorted, and voluntary attention presents the subject with an exaggeration of the real properties before her.

So here's a more accurate way to formulate the threat. Where attention is cued exogenously, it *disrupts* the mechanism of experience and voluntary attention through which we form beliefs about visible properties. Visual experience then represents an exaggeration of the properties really present, and beliefs formed through voluntary attention do likewise. The disruption occurs when there is sudden change near a visible object. Since we inhabit a dynamic, changing environment, these circumstances often obtain. So attention routinely leads to false beliefs about visible properties. Much of what we took to be knowledge of visible properties turns out to be false belief.[7]

What's more, the threat may not be limited to cases in which involuntary attention actually distorts visual experience. In §1 I mentioned two general ways of understanding the idea that visual attention is a source of knowledge, in virtue of the fact that beliefs formed through attention inherit their content from visual experience:

(1) Beliefs formed in this way inherit the distinctive reliability of visual processing.

Here the general threat is straightforward: the mechanism of attention and visual experience turns out to be an unreliable way to form beliefs about visible properties; visual attention is not a source of knowledge about visible properties.

(2) Beliefs formed in this way inherit the distinctive justification or warrant provided by visual experience and attention.

Here the status of Carrasco's threat is more controversial. We might argue that this justification or warrant is defeated, since visual experience and attention routinely ground false beliefs. I won't develop that argument here. Even if beliefs formed through visual attention *sometimes* constitute knowledge of visible properties, Carrasco's experiments threaten to show that in vastly many cases they do not.

Note that, so construed, Carrasco's threat to visual knowledge is not an instance of the traditional Argument from Illusion. Visual knowledge of properties is not here threatened by the mere existence of cases in which vision is illusory. Rather, the threat

[6] Liu et al. (2009) found that voluntary attention has an effect similar to that of involuntary attention on experience of one visible property: brightness contrast. If further work should show that voluntary attention in fact affects experience of spatial properties in a similar way, the problem I describe would differ, but my solution would still apply.

[7] Some philosophers argue that visual beliefs always abstract from the contents of visual experience. So we might hope that, even though involuntary attention changes visual representation, the changes are too minute to show up in false beliefs. However, subjects in the experiments *do* form false beliefs, as reflected in their false comparative judgements. Their beliefs don't abstract enough from the contents of experience to avoid being mistaken.

lies in the fact that the putative illusion is pervasive: according to Carrasco, involuntary attention *routinely* leads to illusion and consequent false belief.

2.2 Cue bias and mental paint

Carrasco reports her data by saying the cued 0.20° gap '*appears as if it were*' 0.23° wide. Now there is a sense in which this is unobjectionable. The data consist in subjects' judgements, as expressed in their choices between the stimuli. When attention is cued to the 0.20° test gap, subjects judge it larger than the 0.23° standard gap in 50% of the experimental trials. This pattern of judgement mimics the pattern of judgement found when attention is not cued to either stimulus, and the test and standard are both 0.23° wide. On the basis of their visual experience, subjects judge in the same way whether they're faced with a cued 0.20° test gap, or with an uncued 0.23° test gap. In that sense, the cued 0.20° gap '*appears as if it were*' 0.23° wide.

However, it's a further question whether '*attention alters appearance*', where appearance consists in the representational content of visual experience. And it's a still further question whether attention alters the representational content of experience in a way that makes it '*nonveridical*'. We might seek alternative explanations of the pattern of judgement, explanations which don't involve illusory experience.

In this sub-part of the paper I briefly consider two such alternative explanations. I argue that neither of them successfully preserves the epistemic role of attention. Then in §2.3 I sketch my own alternative.

One natural idea here is that cuing attention might affect subjects' judgements directly, rather than via a change in visual experience (Prinzmetal et al. 2008). Carrasco claims that attention exaggerates the visual signal that's processed for the cued stimulus, with the effect that properties of the cued stimulus are exaggerated in visual experience. By contrast, on this alternative interpretation, cuing attention does not alter the visual signal or visual experience at all. Rather, just having their attention drawn to the stimulus makes subjects more likely to report on it. In the demanding experimental conditions, this is enough to explain the pattern of judgement.

Gobell and Carrasco (2005) provide convincing evidence that the pattern of judgement they found cannot be explained in this way. They conducted two control experiments. First, attention was cued to the location of the test stimulus after it had disappeared. If cuing influenced subjects' judgements directly, we could expect it to bias their judgements to the test even in this condition. But here the PSE did not shift as did in the original experiment.[8] Second, subjects had to report the stimulus which was *lower* in spatial frequency. Here the PSE shifted as it did in the original experiment: subjects were at chance in reporting on the cued test stimulus only where it had a significantly smaller gap. Cuing made subjects *less* likely to report on the test stimulus. So the result is not happily explained in terms of the idea that cuing attention directly biased subjects' judgements to the test.

What's more, cuing attention shifts the PSE only for some visible properties. For example, it shifts the PSE for color saturation, but not for hue (Fuller and Carrasco 2006). It's hard to see how this could be explained by the idea that cuing attention has a direct effect on visual judgement. None of this is a demonstrative proof, of course, but the evidence does support Carrasco's claim that attention alters the visual signal and visual experience. I'll work with that claim from now on.

[8] Prinzmetal et al. (2008) found the opposite result for post-cuing of a brightness-contrast stimulus: post-cuing shifted the PSE just as pre-cuing did. Carrasco et al. (2008) explain the discrepancy by showing that Prinzmetal et al.'s stimuli were near the threshold of invisibility, while their own were not. Carrasco et al. argue that cue bias *does* have a significant effect where stimuli are near-invisible, but that it does not have this effect for their stimuli.

As I said (§1), Carrasco assumes that her evidence of changes in visual experience is evidence of changes in the representational content of that experience. Ned Block (2010) disputes this. He claims that visual experience consists of a representational content together with a mode of presentation of that content. Modes of presentation play a distinctive phenomenological and psychological role in Block's account. The mode of presentation of an experience fixes its phenomenology, making the object of experience appear a certain way. For example, the mode of presentation of a subject's experience fixes a certain '*perceived gap size*' for each stimulus – say a perceived gap size of 0.23˚ (Block 2010: 31). Nevertheless, the mode of presentation of an experience is not constituted by its representational content, and the mode of presentation may vary independently of the representational content. Where the mode of presentation fixes a perceived gap size of 0.23˚, visual experience need not represent the gap as 0.23˚ wide. In that sense, visual experience is characterised by '*mental paint*', a way things appear to a subject which is not a matter of how she represents them to be. What's more, this mental paint or mode of presentation may dictate the psychological role of a visual representation. In particular, Block explains Carrasco's data in terms of changes in the mode of presentation, brought about by cuing of attention.

For the Gobell and Carrasco experiment, Block's interpretation is as follows. Where a subject's attention is cued to a 0.20˚ gap, the mode of presentation in her experience of that gap is equivalent to the mode of presentation in her experience of an uncued 0.23˚ gap. Because the phenomenology of her experience of each gap is the same, she reports the cued 0.20˚ gap larger than the uncued 0.23˚ gap, in 50% of the experimental trials. On this approach, no effect of attention on the representational content of experience is required to explain the data. We can explain the data without ascribing systematic illusion to visual experience.

However, from an epistemological point of view this is no great advance. The changing mode of presentation, brought about by shifting attention, misleads subjects into making mistaken judgements. And if the misleading phenomenology of attentive experience leads to mistaken judgements here, won't it have the same effect elsewhere? Won't the exaggerated phenomenology of attentive experience systematically mislead us into believing that things have properties they don't really have?

Block appeals to James's discussion, to argue that the change in mode of representation need not in general lead to false beliefs, outside the specific experimental conditions:

> As we rightly perceive and name the same color under various lights, the same sound at various distances; so we seem to make an analogous sort of allowance for the varying amounts of attention with which objects are viewed; and whatever changes of feeling the attention may bring we charge, as it were, to the attention's account, and still perceive and conceive the object as the same.
>
> (James 1890: 426)

James explains why altered visual experience does not lead to false belief in terms of a constancy effect: the perceiver grasps that the change is due to a shift in her perspective, a shift in her view on the world rather than in the world itself; so she does not attribute any change to the objects of experience. I don't deny that attention exhibits such a constancy effect. In fact my own interpretation is consistent with this idea. But the challenge is to explain why this effect is not at work in the experiments, if it's at work in other cases. Subjects in the experiments *are* misled.

Block comments: '*Many experiments on perception in effect encourage the subject to judge appearance*', where '*appearance*' refers to the phenomenological mode of presentation, rather than the representational content of experience (54). But if anything Carrasco's experiments *dis*courage judgements about the character of

experience, as opposed to judgements about the objects of experience. Subjects were instructed to '*report the orientation of the Gabor of higher spatial frequency*' (Gobell and Carrasco 2005: 645).[9] The instruction focuses on the actual properties of the stimulus, rather than its subjective appearance. To complete their task successfully, subjects must keep track of the objective spatial frequency of the stimulus. Furthermore, since the instruction focuses most immediately on orientation rather than spatial frequency, there's no reason for subjects to assume, confusedly, that what the experimenters are really asking about is the character of their experience of spatial frequency. So Block's proposal about mental paint does not explain how it can be that, in general, attention '*never seems to lead the judgement astray*' (James 1890: 426), while attention does lead judgement astray in the experiments.

The interpretation I propose does explain this, and it does so in terms of changes in the representational content of visual experience. On this interpretation, attention makes visual experience represent different determinable properties.[10]

2.3 Attention alters determinable appearances

Let's take stock. I propose that we accept the following three claims made by Carrasco:

(1) Cuing attention to a 0.20° gap changes the visual signal processed for it, replicating the signal for an uncued 0.23° gap.
(2) As a result, visual experience of the cued 0.20° gap replicates visual experience of the uncued 0.23° gap. In this sense, cuing attention to the 0.20° gap gives it the same appearance as an uncued 0.23° gap.
(3) This effect of attention on visual experience is an effect on its representational content: visual experience represents the cued 0.20° gap in the same way as it represents the uncued 0.23° gap.

However, it doesn't follow that:

(4) Visual experience of the cued 0.20° gap is illusory.

On some assignments of representational content to visual experience, clearly (4) would follow. Suppose that visual experience represents exact gap sizes. It represents the uncued 0.23° gap as exactly 0.23° wide. Then if cuing attention to the 0.20° gap replicates this experience, cuing attention makes visual experience illusory.

However, I will argue that visual experience does not in general represent exact or maximally determinate spatial properties. Suppose instead that subjects' experiences represent *determinable* gap sizes – gap sizes corresponding to ranges of exact gap sizes. For example, when attention is not cued to a gap of 0.23°, a subject's experience might represent it as between 0.19° and 0.27° wide. This is not to say that

[9] The type of stimulus used in known as a Gabor patch.

[10] My position here is consistent with Block's claim that there is mental paint. However, my position is inconsistent with his (2010) argument for that claim. Block makes a subtle empirical and philosophical case for denying that attention leads to illusions in Carrasco's experiments. He then argues, by eliminating alternative interpretations, that the data must therefore be explained in terms of mental paint. I argue for an alternative interpretation, on which the data are explained in terms of changing representations of determinable properties. Block allows that visual experience might represent determinables, but he gives two reasons for thinking that the data cannot be explained in these terms: (i) these representations can't capture the phenomenology of experience, the conscious awareness to which subjects in the experiments respond; (ii) these representations cannot explain the accurate fine-grained discriminations subjects make when their attention is not manipulated. I address these points in §4 and §5.

either the unit of measurement or the limits of the range are subjectively salient to the subject. Rather, we as interpreters may use this unit and these limits to define a determinable gap size which is subjectively salient to her.

Now if cuing attention to a 0.20˚ gap replicates this experience, the cued 0.20˚ gap will likewise be represented as between 0.19˚ and 0.27˚ wide. Where the two gaps are represented in this way, it's compatible with the content of experience that either gap be larger. But Carrasco's task is *forced choice*: subjects have to report one of the gaps as larger, even if they're not confident in their judgement. So where the cued test gap and the uncued standard gap are each represented as between 0.19˚ and 0.27˚, we should expect the pattern of judgement which Carrasco found: subjects report the test gap larger around 50% of the time. And on this interpretation, visual experience of the cued gap is not illusory. A gap of 0.20˚ really is between 0.19˚ and 0.27˚ wide.

We can explain the data about attention and spatial properties using an interpretation of this form. And we can also understand James's observation in this way, at least as it applies to involuntary attention. James said that attention gives visible properties a '*more intense*' appearance. On the interpretation of the data I'm proposing, we can understand this as follows. Cuing attention to a property gives it an appearance which an uncued, more intense property would have. Your visual experience represents the cued property in a way in which it would also represent an uncued, more intense property. Still, the altered appearance need not be illusory. What's altered is the determinable property which visual experience represents, and different veridical experiences of a property may represent its different determinables. For example, uncued experience of the 0.20˚ gap might represent it as between 0.16˚ and 0.24˚ wide, while cued experience represents it as between 0.19˚ and 0.27˚ wide. Both experiences are veridical.[11]

This interpretation explains why subjects are misled in Carrasco's experiments, even though, in general, attention '*never seems to lead the judgement astray*'. Subjects are forced to make a judgement which is underdetermined by their experience. The content of their experience is compatible with either gap being larger. In light of the representational content of their experiences, it's equally probable that either gap be larger. This feature of their experiences is what misleads them, but it's misleading only with respect to the specific, demanding task which Carrasco's subjects are given. By contrast, where a subject's beliefs about the gaps inherit their content from her visual experience, her beliefs are true. This interpretation is consistent with the claim that visual attention is a reliable source of knowledge.

Block comments that '*the change invoked by changing attention does not look like a change in the world*' (2010: 53). This is surely right. When your attention is cued to a gap, the gap does not seem to grow. My interpretation predicts this. If you represent something as having two different but compatible determinable properties, over the course of a visual episode, this is not ordinarily grounds for supposing that the object has changed.

I don't suppose that the particular assignment of representational content I used is the correct one. But the interpretation works for various assignments of content. The basic constraint is that experience must be sufficiently indeterminate: the determinables represented must correspond to a sufficiently broad range of

[11] So the solution is *not* that attention makes visual experience represent more determinate properties (although that sometimes occurs – Stazicker 2011). That explains neither James's observation nor Carrasco's data. We can make the distinction precise (for all visible properties) as follows. Coming to represent more determinate properties can be modeled as a *dilation* in property space: the region of property space corresponding to the determinable property represented becomes smaller. By contrast, if attention replicates experience of more intense properties (e.g. greater gap sizes), this can be modeled as a *translation* in property space: the region corresponding to the determinable represented shifts to a new set of locations in property space, independently of any variation in its dimensions. Both transformations may occur simultaneously, but I appeal only to the latter here.

determinate gap sizes. For example, if the uncued 0.20° gap were represented as between 0.19° and 0.21°, cuing attention to it could not replicate the appearance of the 0.23° gap in the way I've described, and remain veridical.[12]

My interpretation of the data turns on three general claims about visual experience and visual discrimination, together with a specific application of each claim to Carrasco's experiments:

(i) The spatial properties which visual experience represents are determinable, rather than maximally determinate. In particular, it's empirically plausible that subjects' experiences in Carrasco's experiments represent determinable properties corresponding to fairly broad ranges of determinates.

(ii) Where visual experience represents determinable properties, these properties figure in the phenomenology of visual experience. In particular, it's plausible that the visual appearances to which subjects in the experiments respond consist in the relevant determinable properties being subjectively apparent, or figuring in the phenomenology of experience.[13]

(iii) Comparative judgements about spatial properties are based on experiences of determinable properties. In particular, Carrasco's subjects' powers of discrimination can be explained in terms of their experiencing determinable properties which correspond to fairly broad ranges of determinates.

In the remainder of the paper I defend (i), (ii) and (iii). I don't claim to prove that my interpretation of the experiments is correct. I just argue that it's as empirically plausible as an account in which attention makes visual experience illusory. The empirical evidence is consistent with the idea that visual attention is a reliable source of knowledge. Because visual experience represents determinable properties, attention may change the way things appear in visual experience, without generating illusions.

3. *Indeterminacy*

Indeterminacy: The spatial properties which visual experience represents are determinable.

First I'll make *Indeterminacy* precise. It's often said that to have a property which determines a property *P* is to have *P* in a specific way. But to be red and square, say, is in a perfectly good sense to be red in a specific way. Being red and square does not determine being red. So this is not a sufficient condition on determination.[14]

---

[12] There is also a further constraint: the determinate size of the gap must not be near the bottom of the range corresponding to the determinable represented in uncued experience. For example, if the uncued 0.20° gap were represented as between 0.19° and 0.27° wide, intensifying visual experience in the way I've described would result in an illusion. See §5.

[13] (ii) might be thought dispensable. But I assume that visual experience represents a property only where that property seems, subjectively, to be apparent in visual experience (§1). Moreover, as James noted, attention seems to have an effect on the phenomenal character of visual experience. We should prefer an account which explains the data in terms of such an effect.

[14] Yablo (1992) works with a purely modal condition: where '*P*' and '*Q*' express a determinable and its determination respectively, $\Box((\forall x)(Qx \rightarrow Px)) \ \& \ \Diamond((\exists x)(Px \ \& \ \sim Qx))$. Necessarily, everything crimson is red, but possibly – indeed actually – there are red things which aren't crimson. As Yablo notes, this condition is not sufficient for determination as it's traditionally understood (cf. Johnson 1921; Prior 1949). It doesn't distinguish determinations from conjunctive properties. Necessarily, what has a conjunctive property has each of the properties conjoined, and it's possible for something to have one of the properties conjoined without having the conjunctive property. Perhaps there are no conjunctive properties. But if so, we need to

I'll work with a summary version of Eric Funkhouser's (2006) analysis of determination. We can think of a determinable property as having *determination dimensions*. A determination dimension is a variable along which a property is determined. For example, being colored has the determination dimensions hue, saturation and brightness. Being colored is a structurally simple case: to have any determination of being colored is just to have values within some range for hue, saturation and brightness. Other properties are more complex, in that they're determined along a varying number of determination dimensions. Being (2-D) shaped has determination dimensions for number of sides, length of each side, and size of each interior angle; the number of dimensions for side-length and angle-size is a function of the number of sides. Still, being colored and being shaped each fix at least the *schematic* determination dimensions I've described (Funkhouser 2006: 556).

Property P determines property Q if and only if P differs in nature from Q only along the schematic determination dimensions of Q, such that the values along these variables consistent with instantiating P are a proper subset of the values consistent with instantiating Q.[15] For example, being scarlet differs in nature from being red only in that the range of hue, saturation and brightness consistent with being scarlet is a proper subset of the range consistent with being red. Being triangular differs in nature from being shaped only in that the number of sides and range of side-lengths and angle-sizes consistent with being triangular is a proper subset of those consistent with being shaped. By contrast, being red and square differs from being red in squareness, which is not a value for any schematic determination dimension of being red.

On this analysis, a property P may determine many determinables – all those determinables instantiation of which is consistent with a range of determination-dimension values that's a super-set of the values consistent with instantiating P. So where an object has P, two veridical representations of the object may represent it as having two different determinable properties.

Indeterminacy says that visual experience represents properties corresponding to ranges, rather than absolute values, along the determination dimensions of spatial properties. Empirical justification for *Indeterminacy* lies in the limited and varying *spatial resolution* of vision. The spatial resolution of a representation is given by the maximum spatial frequency to which it's sensitive. Spatial frequency consists, roughly speaking, in the rate of variation of a phenomenon across space (§2.1).

We represent distant objects in lower spatial resolution relative to nearby objects: more rapid variations across space are visible the nearer the object is.[16] For example, if you walk a little distance away, you won't see the variations on the right of Figure 1. Spatial resolution also decreases away from the fovea at the center of the eye. Early in post-retinal visual processing, information about the scene you see passes through varying spatial filters, attuned to various spatial frequencies. As the locations of these filters get further from the fovea, they're attuned to progressively lower spatial frequencies: higher-frequency detail goes unprocessed. So peripheral vision has a progressively lower spatial resolution. Spatial resolution is limited even for foveated objects: even when you look right at something, there are spatial frequencies across it to which early visual processing is insensitive (De Valois and De Valois 1988).

Because visual processes are insensitive in these ways to high rates of variation across space, they're insensitive to absolute values along the determination dimensions

---

understand what distinguishes determinations from these spurious entities – or alternatively what distinguishes determinables from (spurious) disjunctive properties.

[15] This is only the gist of Funkhouser's detailed analysis and mathematical model. As stated here the analysis is obviously circular, albeit informative. It's not clear to me whether Funkhouser's analysis avoids this circularity. That depends on the precise role of his mathematical model.

[16] Strictly speaking, this is true only where spatial frequency is given in terms of a standard measure of distance, rather than degrees of visual angle.

of spatial properties. Consider an (artificially simple) example. You're presented with a line that's clearly visible, but a fair distance away. Call your visual response to the line's length *R*. Conditional on *R*, there is a high probability that the line has any of a range of lengths, say between about 9.995cm and 10.005cm. Conditional on *R*, the probability that the line has a length around the edge of this range is lower, and the probability that it has a length much outside the range is negligible.[17] This is a corollary of the fact that, at the distance at which the line is presented, visual responses are not sensitive to variations which occur across less than about 0.01cm.

As a result, if the line is in fact exactly 10cm long, *R* is not reliably correlated with this exact length. Rather, *R* is reliably correlated with a range of lengths between about 9.995cm and 10.005cm. This is consistent with the relationship between stimulus lengths and visual responses being roughly linear. Overall, as the length goes up, the level of visual response does likewise. But for a specific level of response, any length within a range is probable. Note also that the ranges of properties to which different responses are sensitive in this way may overlap. For example, conditional on a response slightly greater than *R*, any of a range of lengths between 9.996 and 10.006cm might be probable (Figure 4).

Now it's a basic principle of interpretation that, when we assign semantic content to a representational system, we should do so in such a way that the contents come out by and large veridical (Dennett 1975). This is consistent with the possibility that many representations in the system are not veridical. For example, interpreting vision in this way is consistent with the possibility that attention causes widespread visual illusions. What's required is just a background of veridicality.

This principle is an important heuristic device, irrespective of the approach we take to *constitutive* conditions on content in the system. For example, we should employ the principle even if the representational content of visual experience supervenes on local states of the brain. Further considerations might ultimately trump this principle. But absent specific reasons for taking a different approach, we should say that a visual representation *R* represents a property *P* only if the occurrence of *R* is reliably correlated with the instantiation of *P*. The reliable correlation should hold across the range of circumstances such that *R* occurs in a normal way. If *R* is reliably correlated with *P* only under exceptionally ideal or peculiar circumstances, then *R* will not be by and large a veridical representation of *P*. Equally, if *R* sometimes occurs in abnormal circumstances where *P* is not instantiated, this shouldn't count against the claim that *R* represents *P*.

So the notion of visual sensitivity that's relevant for our purposes is:

Sensitivity: In every background condition such that *R* occurs in the normal way, if *P* had not been instantiated, *R* would not have occurred.

On a standard semantics for the counterfactual (Lewis 1973), and where the instantiation of *P* is a contingent matter, *Sensitivity* is true iff:

> For every possible world *w* at which background conditions are such that *R* occurs in the normal way, there is at least one world *x* at which *P* is not instantiated and *R* does not occur, such that *x* is closer to *w* than every world at which *P* is not instantiated and *R* does occur.[18]

---

[17] Typically, the probability distribution will be roughly Gaussian or bell-shaped. At any rate it's continuous, so there is no sharp cut-off between probable lengths and improbable lengths.

[18] The condition of normalcy excludes cases in which there is a very low probability that *P* is instantiated, conditional on *R* (and *vice versa*). Normalcy is notoriously hard to specify. There may be no way of specifying it which marks a determinate threshold in the probability distribution, below which *R* is not sensitive to *P*. In that event, there will be vagueness in what counts as being sensitive to a property. I

Which spatial properties and visual representations meet this condition depends on the spatial resolution of vision. For example, take an object with maximally determinate length *L*, seen by someone whose (normally occurring) visual representation of the object's length is *R*. Whether or not the object has *L* does not make a difference to whether *R* occurs in the observer. There are possible worlds arbitrarily close to actuality, such that the object has a length which differs from *L* by an arbitrarily small magnitude. At at least one of these close worlds, *R* still occurs.[19]

This gives us good reason to deny that the visual representation *R* represents *L*. On that assignment of representational content, *R* would be veridical only by astonishing accident. Since *L* does not make a difference to whether *R* occurs, there is no reason to expect *R* to be reliably correlated with *L*. By the same reasoning, there are various less determinate lengths which *R* does not represent.

By contrast, there is a determinable length to which *R* is sensitive. In the example above, this might be the length: *between 9.995 and 10.005cm*. There are worlds at which *R* does not occur and the line is longer than 10.005cm by an arbitrarily small magnitude. These worlds are closer to actuality than every world at which *R* does occur and the line is not between 9.995 and 10.005cm long. This gives us a reason to say that, if *R* represents a length, it's the length *between 9.995 and 10.005cm*: *R* is reliably correlated with this length.[20]

To be clear, I am not claiming that visual experience represents the probability distribution in virtue of which it is sensitive to a property. Psychophysicists interpret visual processes as encoding such probability distributions, and I don't dispute that claim about the processing. But I assume that conscious visual experience represents spatial properties as *actually* instantiated (§1). If visual experience is sensitive to the spatial properties it represents in this way, then these properties are determinable.

In Carrasco's experiments, it's plausible that the determinable properties represented may be fairly broad, because it's plausible that visual resolution may be quite limited. Because the experiments are designed to probe the short-term effects of involuntary attention, stimuli are presented extremely briefly – for only 40ms. Neurons in the visual cortex become sensitive to increasingly higher spatial frequencies over the course of a normal fixation of around 200ms (Frazor et al. 2004). 42ms after onset of a stimulus, the maximum spatial frequency to which cells in V1 are sensitive is 2½ times lower than that to which they're sensitive 74ms after onset (Bredfeldt and Ringach 2002). And vision scientists connect these facts about early vision with conscious vision, arguing for a distinction between vision 'at a glance' and 'vision with scrutiny' (Hochstein and Ahissar 2002). On this approach, conscious vision is at first characterized by the schematic, categorizing information encoded in higher visual areas; only later, once feedback from higher areas tunes processing in the lower areas, do details of the scene reach consciousness. None of this *demonstrates* that subjects in the experiments experience determinables as broad as my interpretation requires. But it does make that an empirically relevant possibility, provided that we can explain

---

propose *Sensitivity* as a necessary, not sufficient condition on visual representation of a property. So vagueness in a representation's sensitivity might or might not entail semantic vagueness in that representation. Note that taking *Sensitivity* to be a condition on visual representation of a property does not commit us to thinking that there is an analogous condition on propositional knowledge, such that knowledge that *p* must be sensitive to the proposition that *p*. See Williamson 2002, Ch.7 for criticism of that idea.

[19] At at least one, because *L* might in principle be at the upper or lower bound of lengths beyond which *R* would not have occurred, in background conditions such that *R* occurs in the usual way.

[20] We could make a different argument that vision represents determinable properties, if we accepted a causal condition on representation, together with Yablo's (1992) argument that determinable properties are causes, in virtue of the fact that they make a difference to effects. Nothing here turns on that argument.

subjects' powers of discrimination in terms of such experiences. In §5 I show how we might meet that demand.

I said that *Sensitivity* should fix our interpretation of a representational system, absent specific reasons for taking a different approach. For some representational systems there might be such reasons. Take the case of natural language. In some contexts though not in others, you can truly assert that France is hexagonal (Austin 1962). On one (controversial) way of understanding this, the semantic content asserted remains the same in these different contexts; what changes is the degree of precision required for truth. On this approach, the assertion that France is hexagonal represents a property to which it's not sensitive: the property of being strictly-speaking hexagonal. The assertion may nonetheless be veridical, if veridicality need not be *precise* veridicality. You might motivate this approach by noting that the term 'hexagonal' is shared between this assertion and others – for example the assertion that a shape is hexagonal if and only if it's a closed plane figure with six sides. If the latter assertion fixes an interpretation for the term 'hexagonal', and we assume that semantic content supervenes on linguistic expression, then perhaps we have a principled reason for taking the assertion that France is hexagonal to represent a property to which it's not sensitive.

There is no parallel reason for taking visual experience to represent spatial properties to which it's not sensitive. There is no identifiable vehicle of visual representation that's shared between different cases in a parallel way. However, you might propose a different, phenomenological reason for taking visual experience to represent spatial properties to which it's not sensitive. It's tempting to think that visual phenomenology consists in being presented with maximally determinate spatial properties. I address this issue next.[21]

4. *Indeterminacy* and phenomenology

The spatial properties which visual experience represents are determinable, and these properties figure in the phenomenology of visual experience: they seem, subjectively, to be apparent in visual experience. Why should we think this implausible, as compared with the claim that maximally determinate spatial properties figure in the phenomenology of visual experience? To represent the maximally determinate spatial properties of an object is to specify those properties in terms of spatial points.[22] Why would one think, upon phenomenological reflection, that visual experience represents these properties, or that these properties figure in the phenomenology of visual experience? You might just think that's introspectively obvious. But it's a general, theoretical claim, and we should be suspicious of appeals to introspection to justify such claims in an immediate way. To motivate the claim, we would need careful phenomenological descriptions of particular cases. And we can't expect people to describe the shapes they see in terms of spatial points. In fact, we usually think descriptions of visual experience underspecify its content.

I suspect that the lingering role of *impressions* in our thinking is an important source of the idea that visual phenomenology presents us with maximally determinate properties:

[21] Another reason for taking representational content to outstrip *Sensitivity* is the view that a representation's content is fixed by its evolutionary function (Millikan 1984). On this approach, a representation may be systematically falsidical, where the evolutionary advantage of its being so outweighs the cost. I take it there would be no such advantage in representing spatial determinates falsidically.

[22] A spatial point is in principle indivisible. In some metaphysical accounts space is gunky: there are no spatial points. Nothing here rests on that view.

> 'tis confest, that no object can appear to the senses; or in other words, that no impression can become present to the mind, without being determin'd in its degrees both of quantity and quality. The confusion, in which impressions are sometimes involv'd, proceeds only from their faintness and unsteadiness, not from any capacity in the mind to receive any impression, which in its real existence has no particular degree nor proportion.
>
> (Hume, 1740/1978, I.1.vii)

According to Hume, visual phenomenology consists in impressions being present to the mind. Like ordinary objects ordinarily conceived, impressions instantiate determinable properties only where they instantiate determinations of those determinables. Your impression is shaped only if it's shaped in a determinate way. And for an impression to have a property, that property must be present to the mind. So if a determinable property is present to you, a determination of that property is also present to you.[23]

Few today would endorse this line of reasoning, but talk of impressions persists (notably in Block 2007 and Carrasco 2004). And once we reify visual experience in this way, it's tempting to confuse determinacy in experience with the determinacy of its objects. It's tempting to think that, if the things we see have maximally determinate spatial properties, visual experience must represent such properties. In principle, you could reify visual experience without drawing this conclusion, since determinate properties of an image may represent less determinate properties of reality. (Think of shading in a sketch.) However, in practice talk of impressions tends to obscure the possibility that visual experience represents determinable properties.

If we lack *positive* phenomenological evidence that visual experience represents maximally determinate spatial properties, perhaps a more promising strategy is to argue negatively, to object that *Indeterminacy* somehow gets the phenomenology wrong. One such line of objection is as follows: 'Representing a determinable property is equivalent to representing the disjunction of that property's determinations. It's phenomenologically implausible that visual experience represents disjunctions, so it's implausible that visual experience represents determinable properties.'

Now if representational content is given extensionally, in terms of the possibilities logically compatible with it, it's true that representing determinables is equivalent to representing disjunctions of their determinations. But this approach to representational content gives up the aim of capturing phenomenology in the structure of the contents we ascribe. This shows up in the fact that, on such an approach, representing any content *C* is equivalent to representing the disjunction of *C* and a necessary falsehood. So this line of argument can't be used to motivate the view that visual experience represents maximally determinate properties. Such representations are disjunctive by the same criterion.

Furthermore, on the analysis I'm working with, there is a distinctive metaphysical relation connecting determinables with their determinates, which does not connect mere disjunctions with their disjuncts (§3). As a result, where several possibilities all involve instantiation of the same determinable property, these possibilities are unified by more than mere disjunction. So even if we use these possibilities to capture visual contents, there is a unity in representations of determinables which is lacking in merely disjunctive representations.

We can understand this unity as having a phenomenological dimension, in the following way. Suppose we're using phenomenology as a guide to the representational contents we assign to someone's visual experience. What must her experience be like,

[23] For Hume the relevant maximal determinations are *minima sensibilia*, the most determinate properties discernible by our senses (*Treatise* I.1.ii). So Hume could allow that spatial points are not visibly present, where point-hood is measured according to a physical standard.

subjectively, for us to say that its content is a set of possibilities each of which involves instantiation of the same determinable shape? Well one idea is that her experience has an explicitly disjunctive import: just through the experience, she understands that the object she sees has one of a set of different determinate shapes, and she understands what those possible determinate shapes are. But there is an alternative to this unrealistic idea. Given that determinables are not merely disjunctive properties, we need not find any explicitly disjunctive import in the subjective character of experience, in order to assign the content in question. We can assign this content on the basis that experience seems to make apparent a single, unified shape, without making apparent which maximally determinate shape it is. Experience may then be silent, in its subjective import, on the more specific disjoint possibilities consistent with the object's having this single, unified shape.

Compare a different sort of case. If Liz thinks that Pierre is in France, we can model the content of Liz's thought as a set of several possible worlds – worlds at which Pierre is in Paris, worlds at which he's in Lyon, worlds at which he's in the Camargue, and so on. In doing so, we don't require that these more specific possibilities be subjectively salient to Liz. She might not know of the Camargue at all. What justifies the model is that France is subjectively salient to Liz, and that France in fact includes these locations. Similarly, visual representation of determinable properties need not introduce disjoint possibilities into visual phenomenology. Visual experience of a determinable property may have the phenomenology characteristic of experience of a single, unified shape.

It's also important to distinguish representation of determinable properties from some other forms of representational indeterminacy. For example, John Searle (1987) finds it phenomenologically implausible that the contents of conscious episodes should be indeterminate. But what's objected to here is the idea that conscious contents might be subject to *Quinean* indeterminacy about reference. This is quite different from the indeterminacy for which I've argued. Quine's thought was, very roughly, that there is no principled reason for saying that a cognitive response represents one aspect of what stimulates it rather than another; the response does not determinately represent any aspect of the stimulus. By contrast, the indeterminacy to which I'm appealing falls within the scope of the representation: a state that's determinately a representation of blue represents a determinable property; in that sense, the representation is indeterminate between royal blue and navy blue. My proposal is consistent with the idea that there's a determinate fact about the content of every conscious state.

One might similarly object to the idea that conscious visual representations are *vague*. A representation is vague where its content is not fixed precisely, so that there may be no determinate fact about whether it represents things as they are. For example, you might claim that it's indeterminate which determinable property visual experience represents. You might claim that it's indeterminate whether visual experience represents a line as between 9.995cm and 10.005cm long, or as between 9.997cm long and 10.007cm long. If the line is in fact 9.996cm long, it will then be indeterminate whether visual experience represents things as they are. My approach is compatible with, but does not require, this sort of vagueness in visual representation.[24]

A different proposal about vagueness in visual representation takes it to be indeterminate which maximally determinate spatial property visual experience represents. On this approach, it's determinate that a visual experience represents *some* maximally determinate property within the range to which the experience is sensitive, but it's indeterminate which of these maximally determinate properties the experience represents. Where the object of experience has one of these properties, it's then indeterminate whether the experience represents things as they are. Call this proposal

[24] See footnote 17.

Vagueness. *Vagueness* is an alternative to *Indeterminacy*, since according to *Vagueness* visual experience does represent some property to which it's not sensitive. By contrast, according to *Indeterminacy* visual experience represents a determinable property to which it's sensitive; there is no indeterminacy about which more determinate property is represented, since more determinate properties are not represented at all.

To the extent to which Searle's intuition is compelling, we have a reason to reject *Vagueness* which is not a reason to reject *Indeterminacy*: according to *Vagueness*, it's indeterminate what the content of conscious experience is. However, I adopt a more conservative strategy against *Vagueness*. To motivate *Vagueness*, we would need a good reason for saying that visual representation outstrips *Sensitivity*. I'm arguing that visual phenomenology, at least, does not provide such a reason.

Finally, we should distinguish between the visual representation of determinable properties and blurry vision. The visual representation of determinable properties is a form of inexact vision. Michael Tye identifies blurry vision with inexact spatial vision:

> In these cases, one simply loses information. … In seeing blurrily, one undergoes sensory representations which fail to specify just where the boundaries and contours lie. Some information that was present with eyes focused in now missing. In particular, there is less definite information about surface depth, orientation, contours, etc.
>
> (Tye 2002: 147-8)

Now it's true that blurry vision '*loses information*' in this way. But blurry vision does not consist only in this loss of information. Blurry vision is not *identical* with inexact spatial vision. Consider the case of seeing a distant object. Here vision loses '*definite information about surface depth, orientation, contours, etc*', as compared to when you see the object close up. But distance vision need not be blurry. I take it that this is grounds for rejecting Tye's account of blurry vision, rather than grounds for thinking that visual experience represents equally exact spatial information about an object, however far away the object is.[25] Visual experience of determinable properties need not be blurry.

In §3 I cited empirical work which makes it plausible that subjects in Carrasco's experiments experience determinable properties corresponding to broad ranges of determinates, because the stimuli are presented so briefly. This is also plausible from a phenomenological point of view. For example, Merleau-Ponty remarked that normal-length visual episodes are '*a passage from the indeterminate to the determinate*' (1945/1962: 36). And we need not take Merleau-Ponty's word for it. Try Carrasco's task for yourself: http://philosophy.berkeley.edu/people/page/99. Even at 100ms, 2½ times longer than Carrasco's subjects get, it's hard to make out determinate detail.

In principle, one might insist that the short temporal window limits only the determinacy of judgement, rather than the determinacy of visual experience. But we've seen that neither the empirical evidence nor phenomenological reflection supports that view. It's consistent with both that subjects in the experiments experience determinable properties corresponding to broad ranges of determinates.

5. Discriminating the determinable

The task facing Carrasco's subjects is to discriminate the spatial properties of the stimuli – to judge how they differ. Given *Indeterminacy*, how should we understand spatial discrimination?

[25] Citing Tye's discussion, Block (2010: 51-2) argues that determinable properties cannot figure in the phenomenology Carrasco's subjects enjoy, because subjects don't see blurrily. Block doesn't consider the case of distance vision.

Suppose you're presented with two clearly visible lines. Your visual experience represents line *A* as between 9.995 and 10.005cm long, and it represents line *B* as between 10.008 and 10.018cm long. In this situation, you can come to know through visual experience that line *B* is longer. If your beliefs inherit their content from visual experience, it's deducible from these beliefs that *B* is longer. Every determination of the determinable length you represent *B* as having is longer than every determination of the determinable length you represent *A* as having. I'll call this the *canonical method* of visual spatial discrimination.[26]

On the other hand, where your visual experience represents two determinable properties which share some determinations, these properties are indiscriminable to you by the canonical method. This will often be the case even where your experience represents two *different* determinable properties.[27]

On the interpretation I propose, Carrasco's subjects are likely to be in the latter position, in trials where the stimuli are similar. For example, if a subject represents the 0.20˚ gap as between 0.16˚ and 0.24˚ wide, and the 0.23˚ gap as between 0.19˚ and 0.27˚ wide, she cannot discriminate them by the canonical method. Yet subjects discriminate fairly accurately when their attention is not cued to either stimulus. Overall, they discriminate accurately on 9/10 trials. How do they do it?

Psychophysicists don't usually interpret their data in terms of visual experience or its representational content. Typically they use *signal detection theory*. Signal detection theory models visual judgement as the output of a decision process, the input to which is a response in the visual system. Visual processes encode the (subjective) probability that the stimulus with which you're presented has a certain character, conditional on a certain visual response. The process by which you reach a judgement or decision about the character of the stimulus is modeled as a probabilistic inference, based on the visual response. In Carrasco's experiment, the gap in each stimulus produces a visual response. Depending on the relative levels of these responses, the visual system encodes a probability that the test gap is larger, and a probability that the standard gap is larger. On this basis, the system computes a decision as to which gap is larger (see MacMillan and Creelman 2005: 113-4).

Signal detection theory makes no mention of visual experience, but it's open to us to model visual experience as a visual response. From this perspective, we can see that subjects need not be able to discriminate stimuli by the canonical method, in order

---

[26] This is not to say that objects are discriminable *wherever* visual experience veridically represents them as having incompatible spatial properties. Consider this case: *a* has the length *10.0049cm*; *b* has the length *10.0051cm*; visual experience *E* is sensitive to the length *between 9.995 and 10.005cm*, and *E* represents *a* as having that length; visual experience *F* is sensitive to the length *between 10.005 and 10.015cm*, and *F* represents *b* as having that length. Circumstances must be odd, since *E* is produced by a length at the upper bound of its sensitivity, while *F* is produced by a length at the lower bound of its sensitivity. So things could easily have gone wrong: had circumstances differed just slightly, *E* and/or *F* would have been illusory. If to discriminate is to know, then someone judging *a* and *b* distinct in length here does not discriminate them. She fails standard conditions on knowledge. If she had judged in this way in slightly different circumstances, she would have judged *a* and *b* distinct in length even though they were not – so she fails the safety condition of Williamson 2002. And if *a* and *b* had not differed in length, she would have judged them distinct all the same – so she fails the tracking condition of Nozick 1981.

[27] Cf. Hellie 2005 on color-discrimination. If visual experience is transparent, in the strong sense that you know nothing about its content beyond what it seems to reveal about its objects, then this proposal entails the falsity of a strong Fregean constraint on visual experience: your visual experiences differ without your knowing that they do, when visual experience represents different but compatible properties. I don't think this counts against the proposal. First, experience may not be transparent in this strong sense. Let's grant that spatial experience is introspectible only to the extent to which the objects of experience seem to have certain spatial properties. Since different determinables may be compatible, it doesn't follow that *differences* among experiences of spatial properties are introspectible only where the objects of experience seem to differ. Second, the strong Fregean constraint does not obviously follow from the demand that differences in visual representation be cognitively significant. Perhaps different visual representations lead to different patterns of judgement, even where you don't know that those representations differ.

to discriminate them accurately. Take a subject who experiences the 0.20˚ test gap as between 0.16˚ and 0.24˚ wide, and the 0.23˚ standard gap as between 0.19˚ and 0.27˚ wide. If she's to discriminate accurately, it must be probable that the standard gap is larger, conditional on these responses, and she must be sensitive to this fact.

Conditional on these responses, it *is* probable that the standard gap is larger. Recall that, where a visual representation *R* represents a property *P*, *R* is reliably correlated with *P* in circumstances such that *R* occurs in the normal way. For example, where visual experience represents the test gap as between 0.16˚ and 0.24˚, there are different possible circumstances in which this experience occurs in the normal way, such that the stimulus has any gap size within this range. And where experience represents the standard gap as between 0.19˚ and 0.27˚, there are different possible circumstances in which this experience occurs in the normal way, such that the stimulus has any gap size within this range. Taking all these possible circumstances together, there are more possibilities such that the standard gap is wider than there are possibilities such that the test gap is wider.[28]

Carrasco's subjects are trained in the task. Before results are recorded, they go through large blocks of trials, getting feedback on each trial as to whether they judged correctly. And their performance improves radically after training. Training gives subjects the opportunity to become sensitive to the probability that a particular gap is wider, conditional on their visual experiences of the gaps. For example, where you represent one gap as between 0.16˚ and 0.24˚ and the other as between 0.19˚ and 0.27˚, you'll judge correctly most of the time if you say the latter gap is wider; feedback about your judgements informs you of this fact. We should not think of subjects as explicitly or consciously learning about the relevant probability-distributions. In general, psychophysicists don't take us to be aware of the subjective probabilities on which our perceptual judgements depend. Rather, these probabilities are encoded implicitly in sub-personal systems. From the subject's point of view, each gap has a certain visual appearance. On the basis of these appearances, she estimates which gap is larger, even where she's not at all confident in her judgement. Signal detection theory gives us a way of understanding how this could be a reliable method of discrimination.[29]

Consider again the case in which subjects' judgements are systematically *in*accurate: where their attention is cued to the test stimulus. Carrasco claims that subjects experience the cued 0.20˚ gap in the same way as they experience the uncued 0.23˚ gap. For example, a subject's visual experience might represent both as between 0.19˚ and 0.27˚ wide. Conditional on these responses, it's equally probable that either stimulus is wider. If she judges in the way learned during training, the subject will judge the 0.20˚ gap wider in roughly 50% of trials. That explains the shifting PSEs brought about by cuing attention. Yet the subject's experience is veridical.

This is of course just a rough sketch of how we might explain the accuracies and inaccuracies of subjects' judgements. Certainly I haven't offered a statistical argument for any specific reading of the data. But the interpretation I've proposed works for various assignments of representational content. For simplicity's sake, I used an example in which the determinate gap size of stimulus (0.23˚) is at the centre of the range of gap sizes consistent with the visual representation (0.19˚-0.27˚). The interpretation does not turn on this. Realistically, subjects will represent different determinables in different cases, with the actual property of the stimulus at various

[28] Strictly speaking, the probability distributions are probably Gaussian, and so weighted to widths near the centre of the range. But it's fair to assume that this weighting is similar for the test and the standard. So it's still more probable that the standard gap is wider, than it is that the test gap is wider.

[29] Just anecdotally, subjects in similar experiments at Michael Silver's Berkeley lab said they were often guessing, even when they judged correctly. That was my own experience too. One researcher told me that, when he first programmed the experiment, he thought he'd done it wrong because the task seemed impossible. He hadn't, and trained subjects 'guessed' very accurately.

points within the range. In cases in which this point is near the edge of the range consistent with their visual representation, subjects are liable to judge inaccurately, even when attention is not cued. For example, if you represent the 0.20° test gap as between 0.19° and 0.27° wide, and the 0.23° standard likewise as between 0.19° and 0.27°, you're liable to get it wrong about half the time. But statistically these cases are unlikely to be the norm.

Similarly, where the determinate gap size of the stimulus is near the bottom of the range consistent with uncued visual representation, cuing attention is liable to produce illusions. For example, if uncued experience represents the 0.20° test gap as between 0.19° and 0.27° wide, intensifying visual experience in the way I've described will produce an illusion. But again, statistically these cases are unlikely to be the norm. Even if attention occasionally causes visual illusions, we need not suppose that it does so routinely.

Signal detection theory is relevant to these questions about veridical experience, only because visual experience represents determinable properties. For example, suppose that visual experience represented maximally determinate gap sizes. If that were right, then only a misrepresentation could make it subjectively probable that the cued 0.20° gap was wider than an uncued 0.23° gap.

In general, to understand what psychophysical experiments tell us about the content of visual experience, we need to recognise the role of determinable properties in visual experience and visual discrimination. Psychophysics reveals that various contextual factors other than the character of a stimulus affect the course of visual processing. In demanding tasks, these effects sometimes lead to systematically inaccurate judgements (see Carrasco et al. 2008). If we assume that visual experience represents maximally determinate properties, the experiments seem to reveal systematic illusion. The epistemologist's idea of visual experience looks like a naïve mistake. In fact, however, different veridical experiences may represent one unchanging property differently, where they represent different determinables of that property. As a result, contextual effects may alter the content of visual experience, leading to inaccurate judgements in demanding tasks, without experience being illusory. Where beliefs simply inherit their content from visual experience, experience may be a reliable source of knowledge.

In particular, because visual experience represents determinable properties, attention may change the way properties appear in visual experience, without generating illusions. James claimed that attention gives visible properties a '*more intense*' appearance. Where visual experience represents determinable properties, we can understand this consistently with the view that attention is a means of fixing beliefs which inherit a distinctive reliability or justification from visual experience. Attention to a property gives it an appearance which an unattended, more intense property would also have. Nonetheless, the attentive appearance need not be illusory. For all the data show, attention may well be a reliable source of visual knowledge. Here I have made this case only for spatial properties. Applying the argument to other properties is a more complex challenge.[30]

[30] Work on this material was presented at John Campbell's Fall 2009 seminar at Berkeley, the 2010 Berkeley-Stanford-Davis Graduate Student Philosophy Conference, the 2010 London-Berkeley Graduate Student Philosophy Conference, and the NYU Consciousness Project in September 2011. Thanks to participants for discussion, and to Bert Baumgartner and Craig French for useful comments. For other helpful discussions, thanks to Ned Block, Michael Caie, John Campbell, Fabian Dorsch, Mark Eli Kalderon, Tamar Lando, Geoffrey Lee, MGF Martin, Carlos Montemayor, Bence Nanay, Alva Noë, Ariel Rokem, John Schwenkler, Michael Silver, and Barry Stroud.

REFERENCES

Austin, J. 1962. *How To Do Things With Words*. Oxford: Clarendon Press.

Block, N. 1978. Troubles with functionalism. *Minnesota Studies in the Philosophy of Science* 9: 261-325.

Block, N. 2007. Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences* 30: 481–548.

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

Bredfeldt, C. & Ringach, D. 2002: Dynamics of spatial frequency tuning in macaque V1. *Journal of Neuroscience* 22(5): 1976-1984.

Brewer, W. 2006. Perception and content. *European Journal of Philosophy* 14(2): 165-181

Burge, T. 2003. Perceptual entitlement. *Philosophy and Phenomenological Research* 57(3): 503-548.

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

Carrasco, M. & McElree, B. 2001. Covert attention accelerates the rate of visual information-processing. *Proceedings of the National Academy of Sciences* 98: 5363-5367.

Carrasco, M., Fuller, S., & Ling, S. 2008. Transient attention does increase perceived contrast of suprathreshold stimuli: A reply to Prinzmetal, Long and Leonhardt (2008). *Perception & Psychophysics* 70(7): 1151–1164.

Carrasco, M., Ling, S., & Read, S. 2004. Attention alters appearance. *Nature Neuroscience* 7: 308–313.

Dennett, D. 1978. Intentional systems. In *Brainstorms*. Montgomery VT: Bradford Books.

DeValois, R. & DeValois, K. 1988. *Spatial Vision.* Oxford: OUP.

Frazor, R., Albrecht, D., Geisler, W., & Crane, A. 2004. Visual cortex neurons of monkeys and cats: temporal dynamics of the spatial frequency response function. *Journal of Neurophysiology* 91: 2607–2627.

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

Funkhouser, E. 2006. The determinable-determinate relation. *Nous* 40(3): 548-569.

Gobell, J., & Carrasco, M. 2005. Attention alters the appearance of spatial frequency and gap size. *Psychological Science* 16: 644–651.

Hellie, B. 2005. Noise and perceptual indiscriminability. *Mind* 114(455): 481-508.

Hochstein, S. & Ahissar, M. 2002. View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron* 36: 791-804.

Huang, L. & Pashler, H. 2007. A Boolean map theory of visual attention. *Psychological Review* 114(3): 599–631.

Hume, D. 1740/1978. *A Treatise of Human Nature,* ed. R. Selby-Bigge / P. Nidditch. Oxford: OUP.

James, W. 1890. *Principles of Psychology*. New York: Henry Holt.

Johnson, W. 1921. *Logic* (Vol. 1). Cambridge: Cambridge University Press.

Lamme, V. 2004. Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness. *Neural Networks* 17: 861–872.

Lewis, D. 1973. *Counterfactuals*. Oxford: Basil Blackwell.

Liu, T., Abrams, J., & Carrasco, M. 2009. Voluntary attention enhances contrast appearance. *Psychological Science*, 20(3): 354–362.

Locke, J. 1690/1975. *An Essay Concerning Human Understanding*. Ed. P. Nidditch. Oxford: OUP.

Mack, A. & Rock, I. 1998. Inattentional blindness: perception without attention. In R. Wright (ed.), *Visual Attention.* Oxford: OUP.

Macmillan, N., & Creelman, C. 2005. *Detection Theory: A User's Guide* (2nd ed.). Mahwah, N.J.: Lawrence Erlbaum Associates.

McDowell, J. 1994. *Mind and World.* Cambridge MA.: Harvard University Press

Merleau-Ponty, M. 1945/1962. *Phenomenology of Perception*, trans. C. Smith. London: Routledge & Kegan Paul.

Millikan, R.G. 1984. *Language, Thought and Other Biological Categories.* Cambridge MA.: MIT Press

Nozick, R. 1981. *Philosophical Explanations*. Cambridge MA: Belknap Press.

Prinzmetal, W. & Landau, A. 2008. Dissecting spatial visual attention. In V. Coltheart (ed.), *Tutorials in Visual Cognition.* Hove: Psychology Press.

Prior, A. 1949. Determinables, determinates and determinants (Part I). *Mind* 58(229): 1-20.

Prinzmetal, W., Long, V. & Leonhardt, J. 2008. Involuntary attention and brightness contrast. *Perception & Psychophysics* 70(7): 1139-1150.

Searle, J. 1987. Indeterminacy, empiricism, and the first person. *Journal of Philosophy* 84(3): 123-146.

Stazicker, J. 2011. Attention, visual consciousness, and indeterminacy. *Mind & Language* 26(2): 156–184.

Stroud, B. 2009. Scepticism and the senses. *European Journal of Philosophy*, 17(4): 559–570.

Treue, S. 2000. Neural correlates of attention in primate visual cortex. *Trends in Neuroscience* 24: 295-300.

Tye, M. 1995. *Ten Problems of Consciousness*. Cambridge MA: MIT Press.

Tye, M. 2002. Representationalism and the transparency of experience. *Nous*, 36(1): 137–151.

Williamson, T. 2002. *Knowledge and its Limits*. Oxford: OUP.

Yablo, S. 1992. Mental causation. *Philosophical Review* 101(2): 245-280.

Figure 1

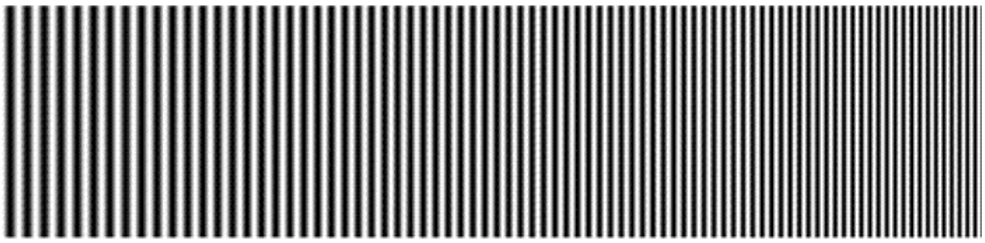

Figure 2

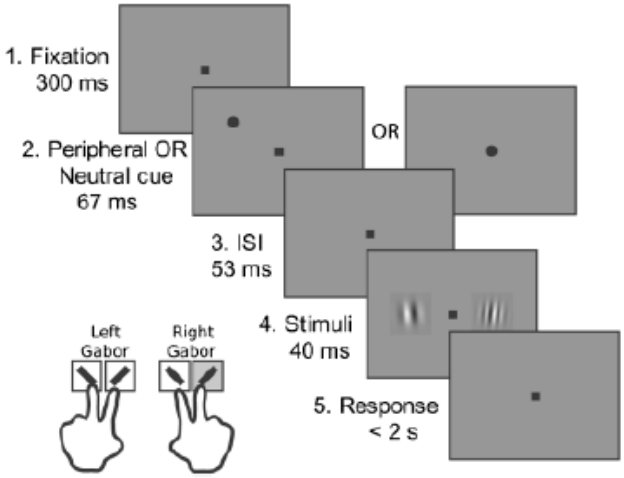

Figure 3 PSE Standard stimulus

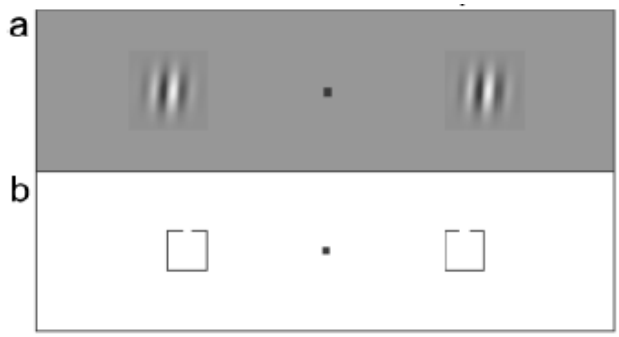

Figure 4