The Puzzle of Perceptual Precision

Ned Block

This paper argues for a failure of correspondence between perceptual representation and what it is like to perceive. If what it is like to perceive is grounded in perceptual representation, then, using considerations of veridical representation, we can show that inattentive peripheral perception is less representationally precise than attentive foveal perception. However, there is empirical evidence to the contrary. The conclusion is that perceptual representation cannot ground what it is like to perceive.

Keywords
Acuity | Adaptation | Appearance | Attention | Awareness | Consciousness | Content | Contrast | Endogenous attention | Exogenous attention | Grounding | Indeterminacy | Marisa Carrasco | Perception | Peripheral perception | Precision | Reductionism | Representational content | Representationism | Salience | Tyler Burge | Unconscious perception | Vagueness | Veridicality | Visual field

Author

Ned Block
ned.block@nyu.edu
New York University
New York, NY, U.S.A.

Commentator

Sascha B. Fink
sfink@ovgu.de
Otto von Guericke Universität
Magdeburg, Germany

Editors

Thomas Metzinger
metzinger@uni-mainz.de
Johannes Gutenberg-Universität
Mainz, Germany

Jennifer M. Windt
jennifer.windt@monash.edu
Monash University
Melbourne, Australia

1 Introduction

Attention increases acuity, allowing the perceiver to see details that would otherwise be missed. In addition, for items that the perceiver does actually see, attention changes their appearance, increasing, for example, the appearances of contrast, (differences between light and dark), speed of a moving object, spatial frequency (a measure of how closely spaced light and dark areas are) and the size of a gap—as in Figure 4. But when attention makes something appear bigger or faster, does it work like a magnifying glass, trading off a gain in information at the cost of making something appear bigger or faster than it is? Or does attentive perception portray the item more as it really is? Or are both percepts veridical—or are both non-veridical? Similar issues arise with regard to inhomogeneities in the visual field. Vision in the lower visual field is about 65% more sensitive to contrast (and orientation discrimination, texture segmentation, gap size, speed, spatial frequency) than vision equidistant from fixation in the upper visual field. (See Figure 1 for examples of low and high contrast.) In addition, there is a great deal of noise in perceptual systems. Percepts involving the same area of the visual field and the same degree of attention will typically differ in visual response from occasion to occasion. So on different occasions, one can see the same object or event in the same conditions, with the same degree of attention, and from the same vantage point and it will look different in size or speed or contrast because of random factors.

What is the consequence of these facts for the veridicality of perception? One viewpoint

Block, N. (2015). The Puzzle of Perceptual Precision.
In T. Metzinger & J. M. Windt (Eds). *Open MIND:* 5(T). Frankfurt am Main: MIND Group. doi: 10.15502/9783958570726

Figure 1: Six levels of contrast. The Wikipedia caption reads "Different levels of contrast - original image top left - less contrast to the left (50%, 75%), more to the right (25%, 50%, 75%)". I take this to mean that the mid-left photo has 50% less contrast than the upper left, the lower left photo as 75% less contrast than the upper left, etc. These percentages are differences from photoshop, not absolute measures of contrast of the sort to be discussed later in the paper. Percent contrast in the sense to be discussed is the difference between the luminance of the lightest and darkest parts divided by the sum of these luminances. These images come from the Wikipedia entry on contrast. According to Wikipedia, "Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled *GNU Free Documentation License*."

says that perception is mostly slightly mistaken. We usually see length, speed and contrast non-veridically but the extent of error is small enough not to be problematic. However, this viewpoint cannot be right since it is only in virtue of a history of veridical representation both in our own lives and in the past of our species that our perceptual representations even have representational contents (Burge 2010). Without such a history of veridical representation it is not clear that perceptual representation really makes sense.

An alternative way of thinking about the issue is that perception is sufficiently imprecise in its representational content for all these varying percepts to be veridical. If a person is said to be 5 feet to 6 feet tall on one occasion and 6 feet to 7 feet tall on another, both are veridical if the person is 6 feet tall. One could put this by saying that perceptual representation is "intervalic". The intervals however would have to be pretty large given the size of these effects—notably the 65% difference between lower and upper visual field just mentioned. And it is hard to square such large differences with the phenomenology of foveal vision. Hold a piece of lined paper in front of you. You seem to see the difference between the white space and the lines fairly precisely. "Irrelevant!," you may retort, "Those differences in the visual field affect only peripheral perception; attentive foveal perception is much more precise than inattentive peripheral vision." And this resolution seems to be reflected in our phenomenological judgments: move the piece of lined paper out to 30° away from the line of sight. Doesn't your visual impression of the contrast between the lines and spaces seem, well, less precise? Surprisingly there is evidence that unattended and peripheral perception of some properties (notably contrast) are about as precisely represented in attentive foveal vision as in inattentive vision and vision in the near periphery (up to a 30° angle from the line of sight). The upshot is that the phenomenology of perception may mislead us with regard to the precision of the representational content of perception.

One might suppose that help will come from bodily action. Goodale & Murphy presented 5 rectangular blocks to subjects at various positions in the visual field ranging from 5° to 70° away from the line of sight (1997). They compared accuracy of perceptual discrimination of one block from another with accuracy of grip via a device that measured the aperture between thumb and forefinger as subjects reached out to pick up one of the blocks. Grip accuracy is roughly the same at 5° as at 70°. The fine details of action are controlled by a largely distinct system from the system that underlies conscious vision. So what this result dramatically illustrates is that the precision of bodily action is unlikely to cast any light on the precision of perceptual phenomenology.

This is the puzzle of the title. I argue that the disconnect may be real and that perceptual phenomenology may mislead about perceptual representation. Perceptual phenomenology may not be grounded in the representational content of perception. Further, there may be no "phenomenal content", that is no representational content that emerges from the phenomenology of perception.[1]

This is a very long paper so it might be useful to know what parts to focus on. You can see the basic lines of the dialectic from reading sections 1-3. Sections 4-7 concern the experimental data concerning attention and can be skimmed without losing the thread. The argument resumes with 8-10. 11 can be skipped without loss of continuity. 12 covers some of the results that the argument is based on. 13 can be skipped. 14 is the conclusion.

2 Background

This section describes some assumptions and terminology. A simple percept consists of a representation of an environmental property and a singular element that picks out an individual item (Burge 2010). The representational content is the condition of veridicality and is satisfied only if the referent of the singular element has

---

[1] Direct realists reject representational contents, holding instead that the phenomenology of perception is grounded in what properties one is directly aware of. They face a parallel set of issues with regard to the question of how precise the properties are that one is directly aware of.

the property represented by the property-representation. The precision of a representation—in my terminology— is a matter of the range of values attributed. For example, consider two visual representations of the height of a person, one representing the person as between 5’6” and 6’ tall, the other representing the person as between 5’8” and 5’10” tall. The latter has a narrower precision. Precision in the sense used here is not a matter of indeterminacy of borders but rather the size of the range.[2]

The claim that the precision of a representation is wide is a form of the claim that perception is “intervalic”. There are other measures of perception that are easily confused with precision. One of them is acuity—also known as spatial resolution. Acuity is the ability to resolve elements of stimuli. Common measures in the case of vision are the extent to which the subject can distinguish one dot from two dots, detect a gap between two figures, determine whether a rotating figure is rotating clockwise rather than counter-clockwise, ascertain whether two line segments are co-linear, distinguish a dotted from a solid line or detect which side of a Landolt Square a gap is on. (See Figure 4 for an example of a Landolt Square.)

These and other items of terminology are gathered together in a glossary at the end of the article. Of course other quite different definitions of ‘precision’ and ‘acuity’ are just as legitimate as these. Note in particular that I am not using the notion of precision as the inverse of variance or the notion of precision associated with the predictive coding literature.

Representationists (also known as representationalists and intentionalists) think that what it is like to have a perceptual experience —that is, the phenomenology of perceptual experience—is grounded in the representational content of the perception. (Not that representationists have used the notion of grounding, but I believe that it captures what they have meant.) Representationism is sometimes framed as an identity thesis (e.g., Pautz 2010; Tye 2009): what it is for an experience to have a certain phenomenal character = for it to have a certain representational content. But the identity formulation is inadequate because the phenomenology is supposed to be based in the representational content and not the other way around. Identity is symmetrical. The grounding characterization of representationism avoids this problem since grounding is asymmetrical. To say that perceptual representation grounds perceptual phenomenology is to say that it is in virtue of the representational content of a percept that it has the phenomenology it has. And *in virtue of is* asymmetrical. (See Fine 2012 on the concept of ground and my 2014a for further discussion of grounding in philosophy of mind.)

Representationism is often framed in terms of supervenience: no difference in the phenomenology of perception without a difference in its representational content. But supervenience does not capture a key motivation behind representationism: that the representational content of perception is the source of the phenomenology of perception, that it is in virtue of the representational content of the perception that it has the phenomenology it has. A supervenience formulation would entail that a difference in the precision of phenomenology requires a difference in representational content. However, on a supervenience formulation of representationism it would be a further question whether the phenomenology of perception could increase in precision without a commensurate increase—or even with a decrease—in precision of its representational content. On the grounding characterization, any change or difference in phenomenological precision is dependent on a commensurate change or difference in representational precision.

The grounding formulation of representationism rules out some but not all kinds of multiple realization. Suppose that red$_{782}$ is an example of the most fine-grained color we can experience. And suppose that the representationist theory of the experience as of red$_{782}$ is that this experience is grounded in representation of red$_{782}$. Different experiences as of red$_{782}$ can be realized by different representational states so

---

[2] As Tim Williamson noted when some of this material was presented at Oxford, the fuzziness of the borders is vagueness rather than imprecision. Ryan Perkins & Tim Bayne argue against representationism using considerations of vagueness (2013).

long as they all involve the representation of red$_{782}$.

The grounding characterization captures a representation-first view and excludes phenomenology-first doctrines that are often portrayed as representationist. Phenomenology-first views suppose that phenomenology grounds at least some kinds of representational contents (Hill 2009; Kriegel 2011, 2013; Shoemaker 2007). And it also excludes versions of representationism that treat both the phenomenology and representational content of perception as grounded in something else (Chalmers 2006; Siegel 2013). That is a plus for the grounding characterization—distinguishing between fundamentally different points of view. Although I won't talk about this much here, I think the considerations I will be raising will cast doubt on views that phenomenology grounds any kind of representational content.[3]

The reader may feel that both peripheral and unattended perception are odd and unimportant phenomena that cannot be the test of any theory of perception. However, peripheral unattended perception is ubiquitous. The fovea is the high density center of the retina. If you hold your hand at arm's length, your foveal perception encompasses about double the width of your thumb. Much of perception at any fixation occurs outside that area and a similar point applies to attention. However, even if you think that both peripheral and unattended perception are atypical, you should recognize that atypical cases often are a window into the nature of a phenomenon. The experiment in which a beam of light goes through two slits was crucial in demonstrating a wave aspect of light (Feynman 1988).

3 The inhomogeneous visual field

Although this article is mainly about differences in perception wrought by differences in attention, it will be helpful to start with a discussion of similar issues that arise independently of attention because of the massive inhomogeneities in the visual field. I will discuss the perception of contrast. The visual system is much more sensitive to differences in luminance than to luminance itself and contrast is a matter of luminance differences. (Luminance is a measure of the light reflected from a surface.) Contrast can be defined in a number of different ways, all ways of capturing the average difference in luminance between the light and dark parts of an array. The four patches in Figure 2 have roughly equal apparent contrasts if one is fixating the cross though there is substantial variation among persons in comparative sensitivities in the visual field. But the top patch has a 30% contrast and the bottom patch has a 15% contrast. (To fixate the cross is to point your eyes at it.) Vision in the lower visual field (the South) has about 65% better sensitivity than vision in the upper visual field on average along the "vertical meridian" (the vertical line through the fixation point) for points of equal eccentricity. And sensitivity is better along the horizontal meridian than the vertical meridian, that is East and West have higher sensitivity than points of equal eccentricity in the North and South. This sensitivity advantage is about 63%. Marisa Carrasco suggests that the advantage of the horizontal over vertical meridians probably has to do with the presence of more relevant information on the horizontal meridian (Carrasco et al. 2001). These differences in sensitivity manifest themselves phenomenologically in differences among patches required for equal apparent contrasts. It takes a 30% contrast patch in the North to phenomenologically match a 10% contrast patch in the East at the equal eccentricity depicted in Figure 2. Performance asymmetries along these lines have been observed for gap size, spatial frequency (roughly density of stripes), orientation discrimination, texture segmentation, letter recognition and motion perception. Performance asymmetries of this sort have been shown in comparisons between an on-screen stimulus and a stimulus from the recent

---

[3] Some of the philosophers who call themselves "representationalists", for example Michael Tye (2009), have endorsed "object-involving" representational contents. Suppose I am looking at a tomato and having an experience that represents the tomato as being red$_{782}$. You are looking at an exactly similar tomato in identical circumstances and also having an experience that represents it as having red$_{782}$. According to Tye, we are having phenomenally different experiences in virtue of looking at different tomatoes. As Burge has noted in an article on direct realism (2005), there are object-involving phenomenal types (of the sort Tye is talking about), but there are also non-object-involving phenomenal types. Representationism as discussed here is concerned with the latter types. I mentioned in footnote 1 that the same issues about precision arise for direct realism—and the same applies to Tye's view.

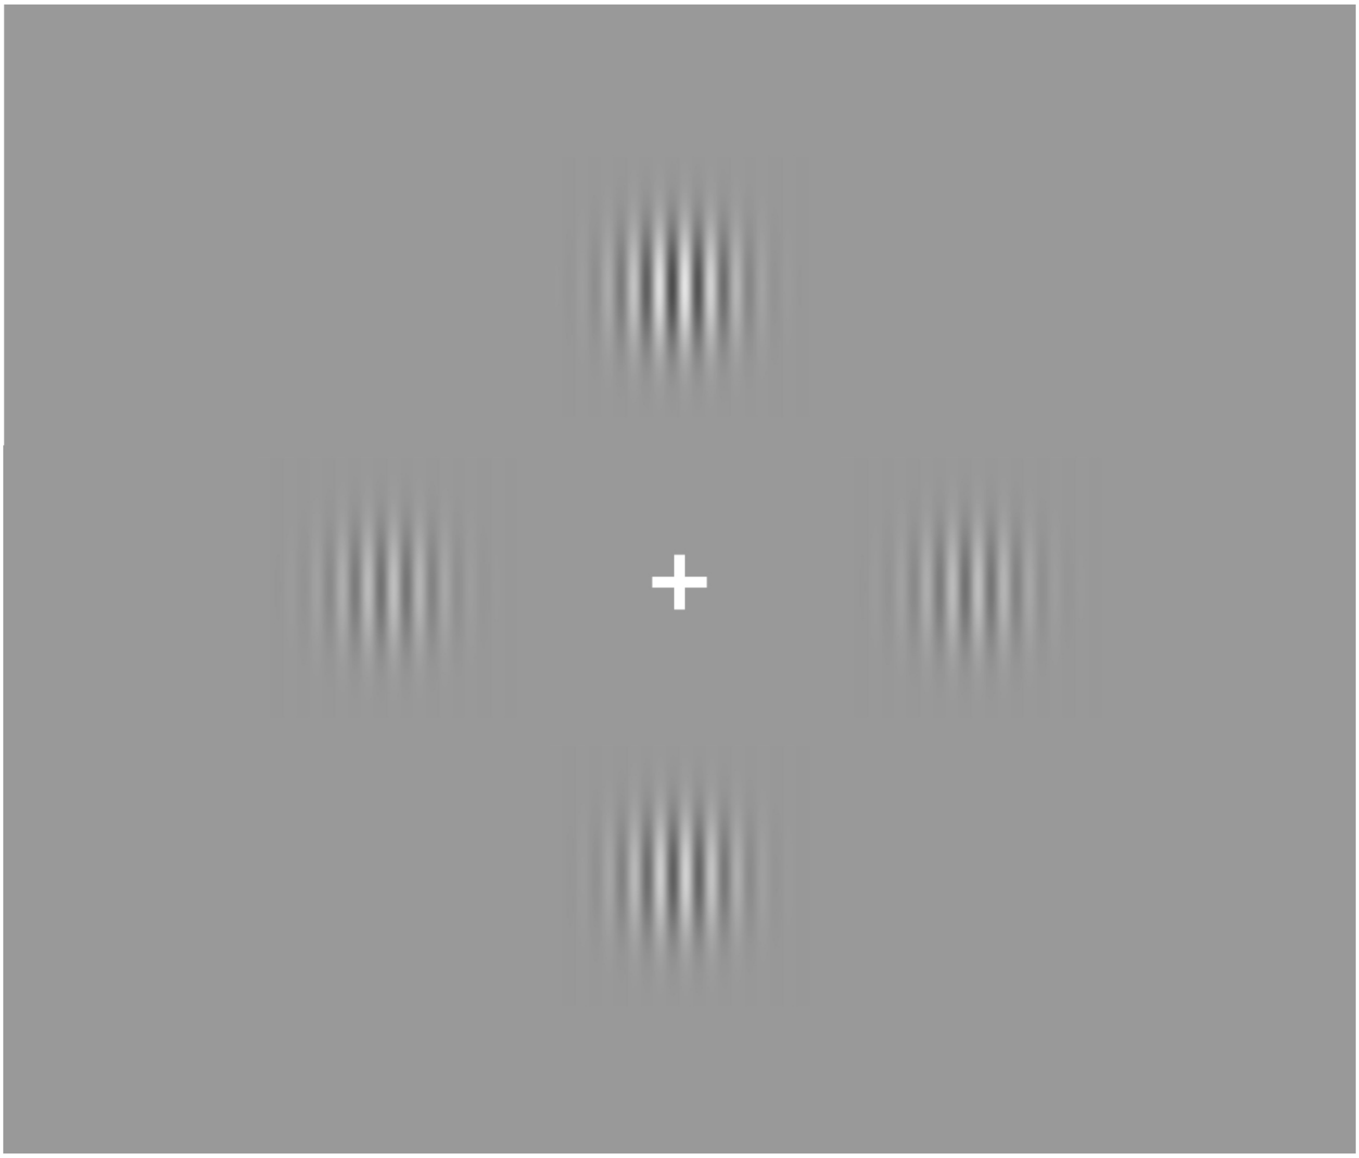

Figure 2: If you fixate (i.e., point your eyes at) the plus sign, these four different patches should look roughly equal in contrast at normal reading distance (roughly 15 inches away). The one above the horizontal meridian has twice the contrast of the one below the meridian (30% vs 15%). The two patches on the horizontal meridian have 10% contrast. It takes a 30% patch in the North to match the 10% patch equidistant from the plus sign in the East. Much of the work of investigating this phenomenon comes from Marisa Carrasco's lab. See Cameron et al. (2002) and Carrasco et al. (2001). Note that there is a large degree of variation from person to person so the patches may not look exactly the same in contrast to you. (The patches are called "Gabor patches" or sometimes just gabors.) Thanks to Jared Abrams for making this figure for me. @copyright Ned Block

past in visual short term memory for 1-3 seconds, albeit at a slightly lower level (Montaser-Kouhsari & Carrasco 2009). These differences are thought to be due to anatomical asymmetries (Abrams et al. 2012).

I will assume that the percepts of North and East have the same contrast phenomenologies when seen (simultaneously) in peripheral vision. Of course the fact that they don't look different does not prove that they look the same. And their looking the same does not prove that the phenomenology of each of the two patches is the same— as we know from the phenomenal Sorites problem (Morrison 2013). However, the fact that they look the same is *evidence* that they are the same phenomenologically and we would need a reason to resist that conclusion. Similar issues will be taken up later in section 8 and 10.

I take it as obvious that the North and East patches are determinately different in apparent contrast when sequentially foveated and attended. The fact that the percepts are sequential makes it unlikely that we are misled about the determinate difference by any analog of the "beats" one hears when guitar strings vibrate at slightly different pitches. (I will return to this issue in section 10.)

North and East look the same in peripheral vision and different in foveal vision. How could this be explained in terms of representational content? The only representational explanation I can think of would be based on the idea that the content of foveal representation of contrast is more precise than the content of peripheral representation of contrast. However, as I will explain below, there is evidence that the representation of contrast in the fovea is the same in precision as the representation of contrast in the periphery. So the burden is on the representationist to explain the difference between foveal and peripheral experience of contrast without appeal to a difference in representational precision. I will now turn to a much longer version of the argument which does not have the form of a burden of proof argument but which makes use of the notion of phenomenal precision.

I claim that when you fixate on and attend to the cross, both your perception of the North patch and your perception of the East patch normally veridically represent the contrasts of those patches despite the fact that one sees them only in peripheral vision. Many details cannot be seen in peripheral vision but what can be seen is seen veridically in normal circumstances. Of course the comparisons are illusory: patches that are different in contrast look the same. But the issue I am raising is whether the individual percepts of single patches are illusory. One reason to think there is no illusion is that the same kind of differences in perception caused by spatial inhomogeneities in the visual field occur in all percepts due to *temporal* inhomogeneities—that is, random noise in the visual system that differs from percept to percept. Any two percepts of the same items at the same point in the visual field with the same degree of attention are likely to differ in apparent contrast (and other properties) due to these random factors. It is hard to see a rationale for supposing that spatial inhomogeneities engender illusion while claiming the opposite for temporal inhomogeneities. And claiming that both engender illusion would make most perception illusory.

This is where my appeal to Tyler Burge's recent book comes in (2010). As Burge notes, we can explain the operation of constancy mechanisms in perception only by appeal to their function in veridically representing the distal environment. And that function precludes perception being mostly non-veridical.[4]

I will say more by way of justification of the verididality claim later but for now let us accept that claim and think about the consequences for representationism. Note that the veridicality assumption is meant to apply to non-categorical perception of properties that admit of degrees and is not meant to apply to categorical perception. Afraz et al. (2010) showed that gender neutral faces are more likely to look male in some areas of the visual field and female in others. The veridical percept in this case would represent the gender-neutral faces as androgynous so both of the percepts described are non-veridical. Many varying mag-

---

[4] The popular "predictive coding" framework (Clark 2013; Hohwy 2013) is a kind of Bayesian approach that is sometimes thought to provide a revolutionary alternative to the view of perception as constitutively involving veridically conditions. Of course all of vision science involves a background of Bayesian probabilistic processes. And prediction in the visual system is ubiquitous and important. But these approaches do not undermine the veridicality of perception. A recent review of Jakob Hohwy's 2013 book on predictive coding (Wilkinson 2014) singles out the predictive coding explanation of binocular rivalry as the parade case, claiming that the predictive coding framework "provides a very satisfying account of binocular rivalry." Clark (2013, pp. 184-185) also emphasizes the supposed explanation of binocular rivalry. Binocular rivalry is a surprising visual phenomenon in which different stimuli are presented at the same time to the two eyes, e.g., a face to one eye and a house to the other. What the subject sees however is an alternation between a face filling the whole visual field, then a house, then a face, etc. It is widely agreed in vision science that the rough outline of the binocular rivalry phenomenon is explained by a combination of reciprocal inhibition and adaptation: the competing interpretations reciprocally inhibit one another, and when one is in the ascendancy, adaptation weakens it until the other takes over. Hohwy and his colleagues more or less concede this (Hohwy et al. 2008) saying that the predictive coding framework explains why we have reciprocal inhibition and adaptation in the first place. But to the extent that this reflects what is good about the predictive coding framework, it is not a revolutionary alternative to standard vision science but rather an evolutionary gloss on it.

nitudes such as size and contrast are not perceived categorically in this way so there is no corresponding "reality check" for such magnitudes. (Some magnitudes such as orientation may mix categorical and non-categorical perception.)

A percept that attributes a property to an item is veridical only if the item has the attributed property. However, the veridical percepts of North and East (when fixating the cross) attribute the same contrast property since they look the same in contrast. Let us ask what the content of the (veridical) percepts of North and East are when one is fixating the cross and they look the same in contrast. That is, what contrast would the percepts of North and East attribute to those patches? Since East is a 10% patch and North is a 30% patch, and both are veridical, it follows that the percepts have to attribute the same contrast to them (since they look the same). What attributions would be the same and also veridical? The patches would have to be represented as having a range of contrasts between 10% and 30% at a minimum. That is, the minimal imprecision in the representation is 20%, the imprecision of a representation of 10%-30% contrast (including the endpoints).

Now let us ask what the contrast-content is when we fixate (and attend to) the East patch, the 10% patch. If the precision is the same as in peripheral perception (i.e., 20%), the percept could have a content of 10% plus or minus 10%, i.e., 0% to 20%. (A 0% contrast patch would be invisible, so presumably imprecision ranges should be weighted towards higher absolute values of the magnitude perceived. Variability in perceptual response increases with the absolute value of the magnitude perceived—one form of the Weber-Fechner Law. This is a complication that I will mainly ignore.) And for similar reasons, if the precision is the same in foveal as in peripheral perception, the contrast content of the percept of the North patch when one fixates it would be 20%-40%.

The representational precision is 20% but what about the phenomenal precision? Can we make sense of this idea? As with all that is phenomenal, no definition is possible. The best that we can do is indicate a phenomenon that the reader has to experience for him or herself. One type of example exploits the difference between an object close up and the same object at a distance. An object may look to have the same properties at both distances but with different precisions. An object may look crimson close up but merely red (and not any particular shade) at a distance.

If the phenomenology of perception is grounded in its representational content and if there is such a thing as phenomenal precision, an increase in phenomenal precision depends on a corresponding increase in representational precision. Representational precision can be indexed numerically—a representational content of the length of something as 1 inch—2 inches (i.e., between 1 and 2 inches) is more precise than a representational content of it as 1 inch—3 inches. According to representationism, phenomenal precision is just the phenomenology of the precision of representational content. We experience a percept with representational content of 1 inch-2 inches as having more (i.e., narrower, smaller range) precision—as being more phenomenally determinate— than we experience a percept with representational content 1 inch—3 inches.[5]

Note that I am not saying that we can always ask whether a certain item of phenomenology is more precise or less precise than a certain representational content (though I think there are some cases where this does make sense). What I am saying is that a representationist has to hold that a difference in phenomenal precision is grounded in a difference—of the appropriate sign and magnitude—of representational precision.

Here is the application of these ideas about precision: Foveate North and East in turn (i.e., serially). I claim that they look determinately different. According to what I mean by looking determinately different, for items to look determinately different, their phenomenolo-

---

[5] For a direct realist, phenomenal precision is just the precision of the properties we are directly aware of. We can be directly aware of properties with different precisions, for example, crimson, or alternatively red. Similarly we can be directly aware of a 10%-20% contrast property and also a 10%-30% contrast property and the difference constitutes a phenomenal precision difference.

gies cannot be almost completely overlapping. Why is lack of almost complete overlap important? The representational contents of perception can be very imprecise even though discrimination is fine grained. One might represent one patch as 10%-30% in contrast and another patch as 10.5%-30.5% and as noted by Jeremy Goodman (2013) that would in principle allow for discrimination between them. If the phenomenal precision of these percepts is also very wide, then the phenomenologies of these percepts would not be determinately different from one another—given what I mean by these terms.

Don't get me wrong: I do think that items can look different on the basis of different but overlapping contents. For example, if one is foveating a patch and simultaneously sees a patch of the same contrast in peripheral vision, the two will look different in contrast. Each of the two percepts can be veridical (even though the comparative percept is not). And being veridical and being of the same contrast, the intervallic contents have to overlap.

You may be skeptical about whether there is such a thing as phenomenal precision and whether there is such a thing as phenomenal overlap. But a representationist should not be skeptical. If one's visual experience represents one length as between 1 inch and 2 inches and a second as between 1 inch and 3 inches, then it is hard to see how a representationist could deny that the phenomenal character that is grounded in the first is more precise than the phenomenal character that is grounded in the second. And if one patch is represented as 10%-30% in contrast and another patch as 10.5%-30.5% the representationist would need a good reason to claim that the phenomenologies did not almost completely overlap. Given that representationism would seem to be committed to phenomenal precision and phenomenal overlap, it would seem legitimate to assume them in an argument against representationism.

North and East look the same when fixating the cross and determinately different when fixating (and attending) to each in turn. What does this fact tell us about representational and phenomenal precisions? The phenomenal precision of perception of contrast must be narrower (i.e., smaller range, greater precision) in foveal vision than in peripheral vision—in order to explain why North and East look the same in respect of contrast in peripheral vision but determinately different in foveal vision. Even if we cannot make sense of an absolute value of phenomenal precision at least we can make sense of differences in it. We might think of this as a phenomenal precision principle:

> If two things look the same in peripheral vision and determinately different in foveal vision, then the phenomenal precision of foveal vision is narrower (smaller range) than that of peripheral vision.

At least for one of the foveal percepts, and why would one have narrower precision but not the other? And so according to the representationist, representational precision must be narrower in foveal than in peripheral vision as well—otherwise there would be a difference in phenomenal precision that was not grounded in a difference in representational precision. (The peripheral perceptions are simultaneous and the foveal perceptions are serial. The inhomogeneities described here hold both for simulaneous and serial presentations, albeit at a slightly reduced level in serial presentations. This has been shown separately for inhomogeneities in the visual field (Montaser-Kouhsari & Carrasco 2009) and for the attentional effects to be discussed later (Rolfs et al. 2013)).

Note that as far as the doctrine of supervenience of phenomenology on representation is concerned, North and East could look the same but still be represented differently. The grounding formulation says: With qualifications to be mentioned: different representational contents require different phenomenologies; supervenience speaks to the converse only. Qualifications: there may be different representational parameters, only one of which is the ground of the relevant phenomenology. So there could be multiple representational realizations of a single type of phenomenal state where the representational differences reflect differences in the parameters that are irrelevant to grounding. And:

phenomenology might be grounded in representational content even though the grain of phenomenology is coarser than that of representational content. So there might be differences in fine-grained representational content that do not make a phenomenal difference.

No one would object to the idea that pure dispositions like fragility or solubility could be grounded in different molecular structures in the case of different substances. And physicalists about phenomenology have held that the underlying basis of a common phenomenology might be one physical state in humans and another in robots. However, I have argued that the grounding framework reveals that physicalists should not acknowledge this kind of multiple realizability (2014a). Applied to this case, the idea is that a representationist account should give us a representationist answer to the question of what it is in virtue of which the phenomenology of the peripheral percepts of North and East are the same. Phenomenal sameness requires representational sameness as a ground. And that representational sameness in this case has to be a precision range of 10%-30% or more.

Of course the notions of phenomenal precision and almost complete overlap of phenomenologies are obscure. The methodological situation we are in is that we have a well-developed science of perception but very little science of the phenomenology of perception. One response —very common until recently—is to avoid issues of phenomenology like the plague. But the time may be ripe to try to leverage the science of perception to get some insight into the phenomenology of perception. And that project cannot help but start with some vague intuitive notions.

Here is where we are: foveal percepts of the contrasts of North and East are determinately different in phenomenology but peripheral percepts of them are the same in phenomenology so the phenomenal precision of North and East, each seen foveally is narrower than the phenomenal precision of North and East seen peripherally. If representationism is to avoid a difference in phenomenal precision that is not based in a corresponding and commensurate difference (of the right direction) in representational precision, then representational precision has to be narrower in foveal than peripheral perception. That is, the representationist should hold that peripheral perception is less representationally precise than foveal perception.

Here comes the punch line: Robert Hess & David Field (1993) compared the discrimination of the locations and contrasts of patches of different contrasts. They presented triples of patches in which the middle patches could differ from the flankers in (1) locations and (2) contrasts. They asked subjects two questions concerning each triple: whether the middle patch differed from the flanker patches in location and contrast. What they found was that discrimination of locations falls off greatly in the periphery but discrimination of contrasts does not. They conclude (pp. 2664, 2666), "… we show that for normal periphery, elevated spatial uncertainty is not associated with elevated levels of contrast uncertainty at any spatial scale… A change in positional error of a factor of 14… from the fovea to the periphery has an associated contrast error that does not significantly increase over the same range of eccentricities." The graphs are striking: position error increases greatly with peripherality of the stimuli but contrast error is a flat line. See Figure 3 for one of the figures that illustrates this fact. As far as I can tell, this result is widely accepted. Even a critical reply (Levi & Klein 1996) says "Their results (discussed below) show that position discrimination is selectively degraded in the periphery, while contrast discrimination is not affected." Levi and Klein dispute the alleged explanation of the result, not the result.

Note that I am taking the fact that contrast discrimination does not diminish in peripheral vision to be evidence that representational precision does not decrease in peripheral vision. Hess & Field (1993) describe a model of the result in terms of constant "uncertainty" for contrast across the visual field but increasing uncertainty for location. Their concern is whether the subjects' visual representations produce locational errors as a result of "undersampling". They argue that undersampling should affect contrast errors too.

And because it does not, they conclude that the explanation is "uncalibrated neural disarray": "We propose that, for reasons as yet unknown, the periphery, unlike the fovea, has not undergone sufficient self-calibration to resolve all of its innate anatomical neuronal disorder…" (p. 2669). But we don't have to buy into neural disarray to accept the observation that contrast uncertainty does not decrease in the periphery.

We feel that foveal attended perception is "crisp", i.e., high in precision but for some properties—contrast and probably gap size, spatial frequency (stripe density) and speed—there is some reason to think that foveal and peripheral perception are equally precise. The resolution is that some properties—e.g., location—really are represented more imprecisely in the periphery than in the fovea (by a factor of 14). (And some properties seem to be represented *more* precisely in the periphery, e.g., flicker rate for some spatial frequencies; Strasburger et al. 2011). So we can't think of peripheral perception as imprecise in regard to all properties we can see. And for the properties that do not decline in precision in the periphery, the representational point of view doesn't seem to work very well.

Acuity is lower in the periphery than in foveal vision. Anton-Erxleben & Carrasco (2013) describe five mechanisms that jointly explain the decrease in acuity with eccentricity. Cone density and the density of the retinal ganglion cells that process cone signals decrease with eccentricity. In addition, average receptive fields are larger in the periphery. (The receptive field of a neuron is the area of space that a neuron responds to. See glossary.) So the elements of a grid will not be visible in the periphery if they are too finely spaced (i.e., if the spatial frequency is too high). To compensate for this, Hess & Field used only very coarse grids in the periphery. So what the result suggests is that contrast uncertainty does not increase in the periphery—*for grids that one can actually see in the periphery.*

But why do the behavioral results reflect on representational precision rather than phenomenal precision? The anatomical asymmetries that are the probable basis of the inhomogeneities discussed here are bound to affect unconscious perception in the same way as conscious perception.

The Hess & Field result shows a kind of homogeneity in the visual field in regard to contrast but as I have emphasized in regard to the phenomenon of Figure 2, the visual field is inhomogeneous in regard to contrast. How are these compatible? The inhomogeneities in Figure 1 reflect contrast sensitivity whereas the homogeneity showed by Hess & Field reflect contrast precision.

Here is the argument summarized:

1. The peripheral percepts of North and East, being the same in contrast phenomenology, are the same in contrast-representational contents—if phenomenology is grounded in representation.
2. The peripheral percepts of North and East are both veridical; that is, North and East have the properties attributed to them in peripheral perception.
3. Given veridicality and the difference between North and East in actual contrast, the representational contents of the peripheral percepts must be rather imprecise. Since North is 30% and East is 10%, and since the content characterizes both, the peripheral representational contrast-content has a precision range of at least 10%-30%.
4. Foveal percepts of North and East—one at a time— are determinately different in phenomenology
5. The phenomenal precision principle: If two things look the same in peripheral vision and determinately different in foveal vision, then the phenomenal precision of foveal vision is narrower (smaller range) than that of peripheral vision.
6. So the phenomenal precision of the foveal percepts of North and East must be narrower than that of the peripheral percepts of these patches.
7. Representationism requires that a difference in phenomenal precision be grounded in a commensurate difference in representational precision.

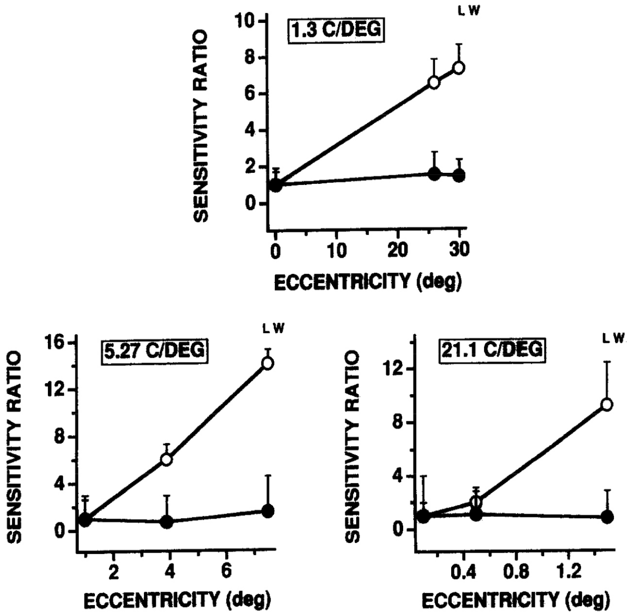

Figure 3: This is one of four graphs from (Hess & Field 1993) showing the comparison between the sensitivity to contrast as compared with the sensitivity to location. The Y-axis represents foveal sensitivity divided by peripheral sensitivity so a value of more than 1 represents greater foveal sensitivity. The solid dots represent contrast sensitivity whereas the open circles represent location sensitivity. The top graph shows sensitivity up to 30 degrees from the line of sight for a very coarse grid of 1.3 cycles per degree. The bottom two graphs show sensitivity for finer grids but at much lower eccentricities. (Coarse grids are visible in the periphery but fine grids would look like a uniform gray surface in the periphery.) Foveal discrimination thresholds are given an aribitrary value of 1. (This is referred to in the article as the values being “normalized”.) What this figure and the other 3 figures show is that contrast sensitivity for grids that are coarse enough to see is the same in the periphery as in the fovea but location sensitivity is much worse in the periphery. Reprinted with permission of *Vision Research.*

8. So representationism requires that the foveal representational precision be narrower than the peripheral representational precision. However the experimental facts suggest maybe not.[6]
9. Conclusion: there is some reason to think that the phenomenology of perception is not grounded in its representational content.
10. The same argument applies to views that hold that there is a kind of “phenomenological representational content” that emanates from the phenomenology of perception (Bayne 2014; Chalmers 2004; Horgan & Tienson 2002). If there were such a thing, it would have to be precise enough to properly reflect phenomenology but imprecise enough to handle the veridicality considerations raised here. And the argument presented here suggests that can’t happen.

The premise that I think needs the most justification is 4. Do we really have enough of a grip on what it is for percepts to be determinately different in phenomenology to justify the idea that the foveal percepts do not have almost completely overlapping phenomenal characters?

Given the problem with 4, I should remind the reader that I started with an argument that did not appeal to phenomenal precision. North and East look the same in peripheral vision and determinately different in foveal vision. How could this be explained in terms of representational content without appealing to a difference in representational precision between fovea and periphery? This argument has the usual problem of a burden of proof argument but it has the advantage of avoiding the obscurity of phenomenal precision.

Another more introspective route to the same conclusion derives from the point mentioned at the beginning that it is natural to feel that the phenomenology of seeing the contrast between lines and spaces foveally differs in precision from seeing the same lines peripherally. The foveal percept seems more “crisp” than the peripheral percept. If this intuitive judgment is correct, there is a discrepancy between the precision of phenomenology and the precision of representational content.

I think this argument gives the reader a pretty good idea of the dialectic of the paper though the paper is more concerned with the issue of change in precision due to differences in attention than with peripherality.

---

[6] This argument can be stated in direct realist terms but it would require an analog of the veridical/illusory distinction in direct realist terms. See Block (2010) and footnote 14.

Worth Boone has argued against my point of view using two-point thresholds of tactile discrimination (2013).[7] As I will explain, I think some of the issues he raises actually support my conclusion.

Boone noted that there are large differences in representational determinacy (precision in my terminology) between tactile acuity as measured by two-point thresholds at various points on the body but that—contrary to what I have said— the precision of the phenomenology matches the precision of the representational content.

First, what are two-point thresholds? "Subjective" two-point thresholds are based on one or two sharp points (e.g., pencil points) being placed at constant separations at various body parts, with subjects reporting whether it feels like there are two points or one point. Objective two-point thresholds are measured by stimulating the skin with either one or two sharp points and observing to what extent the subjects are able to discriminate between these stimuli. Objective thresholds are based on whether there actually are two rather than one point whereas subjective thresholds are based simply on the judgments themselves, independently of their accuracy. The subjective method shows extremely high variability within a single subject on the same body part for a variety of reasons. The objective method has a number of paradoxical features that I won't go into but if you are interested you can read a short article dramatically titled "The Two-Point Threshold: Not a Measure of Tactile Spatial Resolution" (Craig & Johnson 2000).

However, a recent review (Tong et al. 2013) suggests better measures of tactile acuity that confirm Boone's point that tactile acuity varies enormously from one part of the skin to another. A glance at a graph in the Tong, et. al. paper reveals that acuity on the tip of the finger is about 5 times that of the palm and about 20 times the acuity on the forearm.

We can ask: is the phenomenology of these perceptions as imprecise as the representational content? Boone says yes but he is judging the phenomenology of two-point perception. That method is doubly illicit, first because it is unclear that two point discrimination is a measure of anything tactile. Second, the two point judgments may simply reflect the representational contents rather than or in addition to the phenomenology, contaminating the verdict on the very point at issue. If you ask someone how determinate their phenomenology of a two point stimulus is, they may simply be reporting how sure they are they are perceiving two points rather than one. The latter is suggested by considerations of representational "transparency" or "diaphanousness" of experience (Stoljar 2004)[8]. As Thomas Metzinger puts it, we "look through" the experience to its object (2003, p. 173). If so, the phenomenology of judging one vs two may be contaminating the judgment of the precision of the percept.

So I will ask again: Do the differences in phenomenological precision between fingertip and palm and between fingertip and foream perception differ by factors of 5 and 20? The question is not well formed: we cannot ask about either phenomenological or representational precision without specifying what is being represented.

To get a better question, let us focus on the perception of location. Representational locational imprecision does vary with location on the body. The explanation of the variation is that the number and spatial distribution of sensory receptors that feed into a single sensory neuron (i.e., the receptive field of the sensory neuron) varies widely over the body.[9] Is there a matching change in the phenomenal precision with regard to location? If there were a massive decrease in representational precision of location from fingertip to forearm without a corresponding decrease in phenomenal precision, we would have a violation of grounding (of a different kind from those already discussed). I suggest you put a single pencil point on your finger tip

---

[7] I used Figure 2 in a talk at Pittsburgh where Boone was in the audience on November 2, 2012.

[8] G. E. Moore (1903) famously said "... the moment we try to fix our attention upon consciousness and to see what, distinctly, it is, it seems to vanish: it seems as if we had before us a mere emptiness. When we try to introspect the sensation of blue, all we can see is the blue; the other element is as if it were diaphanous ..."

[9] For an amusing account of the facts surrounding these issues see Ramachandran & Hirstein (1998).

and then on your palm and forearm. (Or if you have a helper, do them simultaneously.) If there is a five-fold or twenty-fold difference in phenomenological precision of location it should be appreciable with any stimulus. My own introspective judgment is that there is little or no difference in precision of representation despite the five-fold difference between the fingertip and palm and 20-fold difference between the fingertip and forearm. I am pretty sure that the percepts are not determinately different. No doubt people differ both in these experiences and in their introspective access to them. And with all difficult phenomenal judgments, contamination by theory is no doubt a major source of variability.

If my judgment is right, we have a case of a difference in representational precision without a corresponding difference in phenomenal precision. In the visual case just mentioned, we have evidence for a difference in phenomenal precision without a corresponding difference in representational precision. Taken together, the cases suggest a considerable disconnect between perceptual phenomenology and perceptual representation.

The conclusion of this section is that there is some reason to think that there is no representational content of perception that either grounds or is grounded by the phenomenology of perception—what it is like to perceive.

The reader may wonder how there could be such a disconnect between the phenomenology of perception and its representational content. I mentioned the fact that grip accuracy is about the same in the far periphery (70° off the line of sight) as it is close to the line of sight (5°) despite the fact that conscious vision is extremely weak in the far periphery. Conscious vision is a distinct system from the system that underlies the fine details of perceptually guided action. Though I am not alleging that the system underlying conscious perception is distinct from the system underlying perceptual representation, the upshot of this paper is that they are partially distinct.

In what follows, I will be arguing that facts about attention motivate a similar argument for a discrepancy between the phenomenology of perception and its representational content. The reason I went through the argument based on inhomogeneities first is that the issues are straightforward compared with the corresponding issues concerning attention. Attention is a complicated phenomenon about which there is a great deal of disagreement, so the rest of the paper has many twists and turns. The argument form as presented so far will not resume until section 8.

4 Attention affects appearance

William James (1890, p. 404) famously said attention "… is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others." Except for the exclusion of unconscious attention, most scientists would accept something like that characterization today. Spatial attention is attention directed to a portion of environmental space and is distinct from attention to an individual (e.g., a thing, a surface or a property instance) or to a property.

The mechanisms of attention are fairly well understood. Spatial attention boosts neural activation in circuits that process information from the spatial area that is attended, inhibiting activation in circuits that process information from adjacent areas. Feature-based attention boosts neural activation for attended features, inhibiting neural activation for other features. Object based attention does the analogous task for objects. Feature-based attention refines selectivity for the attended feature whereas spatial attention refines selectivity for the attended area of space (Carrasco 2011; Ling et al. 2014).

The main body of this paper is concerned with the effect of the modulation of spatial attention on phenomenology and representational content. Except when mentioned explicitly, I am talking about spatial attention rather than attention to a property instance or an object. My argument is based on experiments that indicate that attention affects appearance. To begin, attention affects perceptual acuity, one measure of which is whether one can detect whether there is a gap or what side it is on in a Landolt square. (For examples of Landolt sqaures, see Figure 4.)

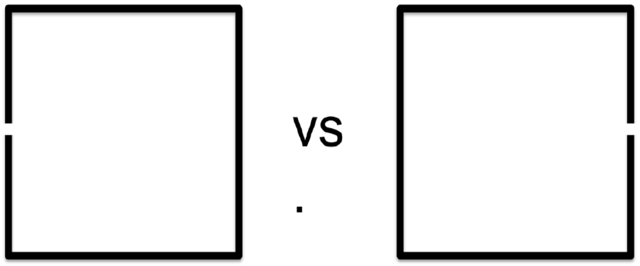

Figure 4: Landolt squares, i.e., squares with gaps. The subjects' task in the experiment diagrammed in Figure 4 was to report (via key presses) whether the gap is on the left or on the right. The squares were presented at various locations while the subject was fixating in the middle of the screen. Redrawn from Yeshurun & Carrasco (1999).

Yaffa Yeshurun & Marisa Carrasco (1999) asked subjects to press different keys depending on whether a Landolt square had a gap on the left or the right. The Landolt square could be presented at any of 16 different locations of 3 different eccentricities. In half of the trials, the square was preceded by a green bar presented briefly at the location in which the square would appear. Then after a pause, a Landolt square appeared in the same location as the line, and then a noise "mask" was presented to prevent an ongoing iconic representation of the stimulus. The subject was supposed to press a key indicating which side the gap was on. See Figure 5 for the sequence of presentations. (An icon would introduce an unwanted source of variability since "iconic memory" varies from person to person. A later experiment (Carrasco et al. 2002) obtained similar results without a mask.) The result is that subjects were more accurate and also faster when the cue indicated the location of the square than when there was no cue. Similar results were shown for other acuity tests, e.g., distinguishing a dotted line from a solid line.

This experiment involves "exogenous" attention in which the subject's attention is automatically attracted by a highly visible change, e.g., a sudden motion or disappearance of an object. A similar effect has been shown when the subject is told, for example, to attend to the right when a central bar points in that direction. This is a matter of "endogenous" attention. Exogenous spatial attention is sometimes referred to as "transient" or "bottom-up" attention, whereas endogenous spatial attention is "sustained" or "top-down". Exogenous attention is involuntary whereas endogenous attention is voluntary. Exogenous spatial attention peaks by 120 ms after the cue, whereas endogenous spatial attention requires at least 300 ms to peak and has no known upper temporal limit.

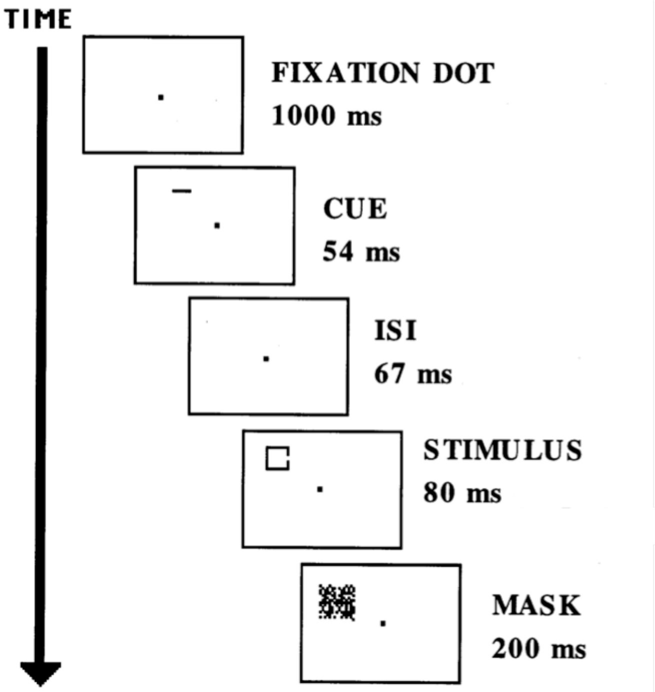

Figure 5: Yeshurun & Carrasco (1999) asked subjects to fixate (point their eyes) at a dot at the center of a screen. Then a cue appeared for 54 ms, then an "inter-stimulus interval", then a Landolt Square, then a mask. (See the text for the purpose of the mask.) Note that it takes 250 ms for eye movement to a new location, so in this and the other experiments described here the brief presentations of stimuli preclude eye movements to the cued items. I am grateful to Marisa Carrasco for giving me this figure.

Using a similar paradigm and comparing the effects of exogenous and endogenous attention (Montagna et al. 2009), researchers showed that endogenous attention decreased the minimum size of a gap that could be detected by about 35% compared to a gap on the opposite side from the cue. That is, subjects could detect much smaller gaps when they attended to the area in which they appeared.

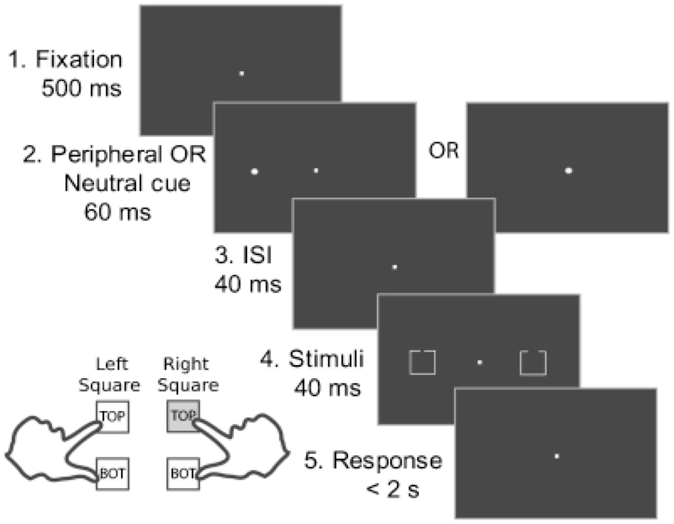

Figure 6: Experiment from Gobell & Carrasco (2005). Procedure described in text. Reproduced with permission from *Psychological Science.*

The conclusion is that attention affects acuity. This is not part of the evidential basis for the argument to come. However there is another effect that is directly relevant to my argument: attention also causes the gap to be perceived as larger. This was shown by a later type of experiment from Carrasco's lab.

The subjects were asked to fixate on the dot that appeared for half a second (upper left in Figure 6). Then the subjects saw a dot on the left or a dot on the right or only at the fixation point in the center of the screen. Then the subjects saw two Landolt squares each of which could have a gap either on the top or the bottom (even though the figure shows the gap on the same side). The subject was then asked to report whether the bigger of the two gaps is on the top or the bottom. If the gap on the left was bigger, the subject was supposed to report the answer using the left pair of keys; mutatis mutandis if the right gap is bigger. The subject was told—correctly—that the dot did not predict anything about the size of the gaps. The subjects' instructions focus on the top/bottom difference whereas what the experiment is really about is the perceived size. The purpose of the dot was to attract the subjects' exogenous (involuntary) attention to one side or the other on some trials. What was being tested is whether attention to, e.g., the left, would cause the perceiver to treat the left gap as bigger. The result was that it does. Subjects did not discriminate between an attended .20° degree gap and an unattended .23° gap. (The gap sizes are measured in degrees of visual angle. If a distance between the eyes and the screen is specified, the degree coding can be changed into inches.)[10] Note that subjects were not asked to judge relative sizes of gaps. In particular they were not asked to judge whether an attended .20° degree gap and an unattended .23° gap "look the same". Rather, the subjects were asked to make discriminations based on apparent gap size. The result is that they are indiscriminable. And this fact about these experiments has led to disputes about what they really show, as I will explain.

The experiment diagrammed in Figure 5 shows attention increases acuity. This one shows that attention makes gaps look bigger. One of the main mechanisms by which attention improves acuity is that attention shifts and shrinks receptive fields (Anton-Erxleben & Carrasco 2013). The shifting of receptive fields is probably involved in both the increase in acuity and the larger appearance. Attention to an area of space causes neurons that were not aiming at that area of space to shift towards it. The effect is more neurons covering that area of space. More neurons covering that area increases the acuity of perception of it. And this mechanism is responsible for the increase in apparent size as explained by Katharina Anton-Erxleben et al. (2007). Their explanation depends on the "labeled-line" hypothesis that neurons in the

---

[10] This effect could be regarded as larger than the result from the Yeshurun and Carrasco paper reported in 5. In that paper, 75% accuracy was achieved by a .20° cued gap as compared with a .22° uncued gap. That difference may be because the 75% accuracy is arbitrary. Or if the difference is real, we could point to the fact that in the Gobell and Carrasco study, the comparison is between an attended square and a square from which attention has been withdrawn, whereas in the Yeshurun and Carrasco study the comparison is between a case in which something is cued and a case in which nothing is cued. Note that there is no need for a mask in this experiment since variations in iconic memory between subjects would be expected to affect equally both the square on the left and on the right.

early visual system are hard wired to code for a certain area of space. So when the receptive fields of neurons shift towards a target the brain treats the size of the target as larger.

5 Is the attentional effect perceptual?

There has been a controversy in the perception literature about whether the kind of effect I have been describing is at least in part genuinely perceptual as opposed to an effect on the decision process involved in generating a report (Schneider & Komlos 2008; Valsecchi et al. 2010).

There probably are effects of attention on aspects of decision, including on conceptualization of a stimulus (Botta et al. 2014). However, I think the case is overwhelming that the attentional effect is at least in part genuinely perceptual. One reason involves "perceptual adaptation" a phenomenon known to Aristotle in the form of the "waterfall illusion". As Aristotle noted, "...when persons turn away from looking at objects in motion, e.g., rivers, and especially those which flow very rapidly, things really at rest are seen as moving" (1955). Looking at something moving in a direction raises the threshold for seeing motion in that direction, biasing the percept towards motion in the opposite direction.

Perceptual adaptation is involved in the "tilt aftereffect". If one looks at a left-tilting patch, the neural circuits for the left direction raise their thresholds. This is sometimes described (evocatively but inaccurately—see Anton-Erxleben et al. 2013) as neural fatigue. Then when one looks at a vertical patch, it initially looks tilted to the right. (See Figure 6 of Block 2010). The reason is that the neural circuits for rightward tilt dominate the percept because of the "fatigue" of the leftward tilt neurons. Ling & Carrasco (2006) showed that attending to the adaptor increased the size and duration of a variant of the tilt-aftereffect as if the contrast of the adaptor had itself been raised. Attending to a 70% contrast grating ramped up the tilt after-effect as if the contrast had been raised from 11 to 14% (different magnitudes in different subjects). Ling and Carrasco directed subjects to attend to gratings for 16 seconds. They found a benefit of attention at first in allowing subjects to distinguish tilts, since attention increases acuity, but then as adaptation increased, discrimination of the adapted tilt was impaired. This kind of adaptation is ubiquitous in perception but does not appear to occur in cognition or decision (Block 2014b). In case anyone thought that the attentional effect was entirely an effect on decision or cognition, this experiment suggests otherwise.

But even apart from the adaptation results, there is strong evidence going back at least to the 1990s from single cell recording in monkeys and in brain imaging for the conclusion that attention increases activity in the neural circuits responsible for the perception of contrast in a manner roughly consonant with an increase in the perception of contrast. Much of this evidence is summarized in sections 4.6 and 4.7 of a review article (Carrasco 2011). My hedge "roughly" stems from debates about the exact effect of attention. There are two kinds of "multiplicative" effects. In "contrast gain" the effect is just as if the contrast of the stimulus has been multiplied by a constant factor. In "response gain" the response is multiplied by a constant factor. The balance of these effects depends on the difference between the size of the target and the size of the "attentional field" (Herrmann et al. 2010). (These ideas are very clearly explained in Chapter 2 of Wu 2014.) A further kind of amplification effect is additive rather than multiplicative: the baseline or "floor" level of activation in the circuit is increased. There is some evidence (Cutrone et al. in press) for increased input baseline as a major part of the attentional effect.

Further, there is plenty of evidence for the conclusion that attention modulates specific cortical circuits depending on what feature is attended. A recent experiment (Emmanouil & Magen 2014; Schoenfeld et al. 2014) compared brain activation when subjects attended to a surface on the basis of its motion and when subjects attended to a surface on the basis of its color. Many of the stimuli involved both color and motion but which feature was task relevant was varied. The result was that motion sensitive areas of visual cortex were activated first when motion was task relevant and color sensitive areas of visual cortex were activated when color was task relevant. In Carrasco's experiments, subjects' attention is drawn to the specific features that the experiment concerns. In the experiment diagrammed in Figure 6, subjects are directed to report the location of the bigger gap, thereby directing attention to gap size. In the analogous experiment connected with Figure 7, subjects are asked to report the tilt of the patch that is higher in contrast, thereby directing attention to contrast. In experiments concerned with color saturation, subjects are shown stimuli that vary in saturation and asked to report the tilt of the patch that is higher in saturation. Similarly for many other features—speed, spatial frequency, flicker rate, motion coherence, shape, brightness, etc. These instructions can be expected to direct attention to the indicated features with amplification in the circuits that register those features.

Schneider (2011) seems to think that when subjects are asked to report on the side of the larger gap and the gap on the attended side is . 20° while the gap on the unattended side is .23°, the subject finds that there is no difference in apparent gap size so the subject just chooses the more salient side. I will discuss salience in the next section, but there is one thing about this charge that raises a distinct issue: that subjects register the increase in apparent size only unconsciously. I now turn to that issue.

6 Is the attentional effect unconscious — like blindsight?

The experiments I have described are "forced choice" experiments in which the subjects must choose between two alternatives. In any perception experiment the issue can be raised of whether the perception is conscious or unconscious, but the issue is often especially troublesome in forced choice experiments with brief stimuli in which subjects make a conscious choice but in which the stimuli are sufficiently evanescent that subjects do not get a really good look at them (Phillips 2011). In addition, the stimuli are presented very briefly, in the experiments described above for 80 ms or less. Many subjects will say that they are never 100% sure of anything. And this can lead to the charge that what is really going on is akin to "blindsight" in which the perception, though genuine, is unconscious (Turatto et al. 2007).

Why are the presentations so brief? Brief presentations preclude eye movements, they preclude significant perceptual adaptation (the "neural fatigue" that causes afterimages), and they preclude certain kinds of strategic responding on the part of subjects. Further, it is known that the effects of exogenous attention peak at around 120 ms after the cue, so to maximize the effects of exogenous attention, brief presentations are required.

Massimo Turatto (2007) showed, using a procedure much like Carrasco's with judgments of perceived speed, that an unattended moving patch was treated as equal in speed to an attended moving patch that was slower by about 10%. However when they asked subjects for subjective judgments of moving stimuli in peripheral vision that really did differ in speed by 10% (without any attentional manipulation), subjects said they saw no difference. Turatto took this as showing that the "just noticeable difference" between the items being distinguished is above the size of the attentional effect so the effect of attention on speed is not conscious. A 10% difference in speed is well above the differences that people can see consciously when they are presented for longer periods, but Turatto argues that for these short presentations the just noticeable difference is larger—that is, it takes a larger difference to be consciously perceived.

There is a difficulty with his experimental procedure though. There are well known problems in asking subjects for same/different judgments. Whether the subjects say 'same' or 'different' depends not only on their percepts, but also on their decision processes, including how big an apparent difference has to be before they regard it as reflecting reality. These issues are nicely analyzed in (Anton-Erxleben et al. 2010, 2011). When Anton-Erxleben et al. corrected for these deficiencies in another same/different experiment, they found effect sizes that are in

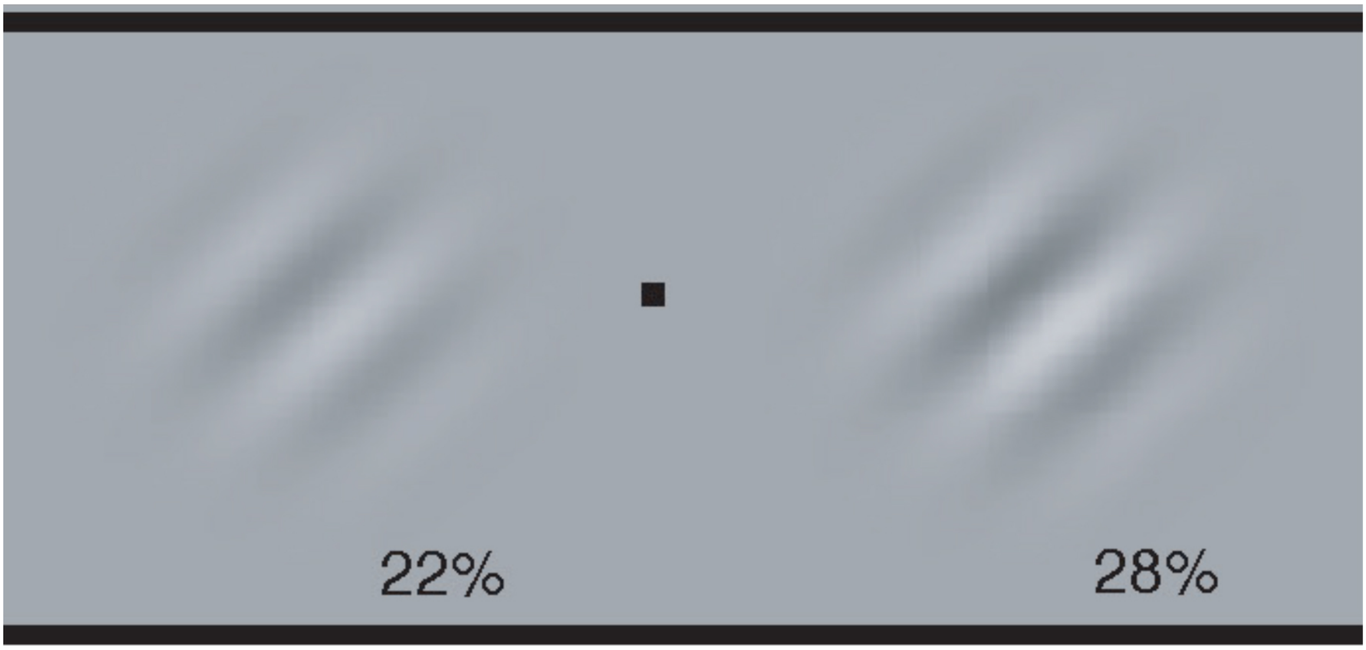

Figure 7: A version of one of the stimuli used in (Carrasco et al. (2004). Fixate at the dot in the center and move your attention to the left patch without moving your eyes. If you can manage that "covert attention", the patches should look to have about equal contrast. If you attend to where you are pointing your eyes (the center) you should be able to visually appreciate that the right patch has higher contrast. I am grateful to Marisa Carrasco for supplying this figure.

the vicinity of other paradigms from the Carrasco lab. The effect size is slightly smaller but as they note, that is probably due to inferior sensitivity of the same/different paradigm. (Similar points apply to Kerzel et al. 2010.) One of the conclusions I would draw is that the notion of a "just noticeable difference" in its usual applications is defective because *noticeability is not a perceptual property* but rather the result of an interaction between perception and cognition. I will not go into these issues further here. However, even if Turatto's methodology is flawed, the issue raised is a good one. How do we know that the effects in Carrasco's attentional experiments are in fact conscious?

The stimulus in Figure 7 was one of the stimuli used by Carrasco and her colleagues (Carrasco et al. 2004) in the first experiment that demonstrated that attention affects perception by changing the qualities of perception, in this case increasing apparent contrast. The method used was the same as described earlier in connection with Figure 6—in fact this experiment was the model for the experiment of Figure 6. Subjects were asked to report the tilt of the patch that was higher in contrast after their attention was attracted to one side or the other by a dot as in Gobell & Carrasco (2005). The result was that when the 22% patch was attended it was treated by subjects as the same in contrast as the less attended 28% patch.[11] In order to make the judgment, subjects were shown examples of higher and lower contrast. (Contrast is a measure of the difference between light and dark portions of a stimulus.)

As I mentioned, similar experiments have shown that attention increases apparent color saturation, apparent size of a moving pattern, apparent speed, apparent flicker rate, apparent spatial frequency (more about what that is below), apparent motion coherence and apparent time of occurrence—the attended event seems to appear about 40 ms before the unattended event. As I mentioned, the subjects have to take in what parameter the experimenter is talking about—saturation, spatial frequency, gap size, contrast, etc. and then decide which stimulus is greater *with respect to*

---

[11] As mentioned earlier, this methodology tells us that the two were indiscriminable and it is a further step to conclude that they actually look the same.

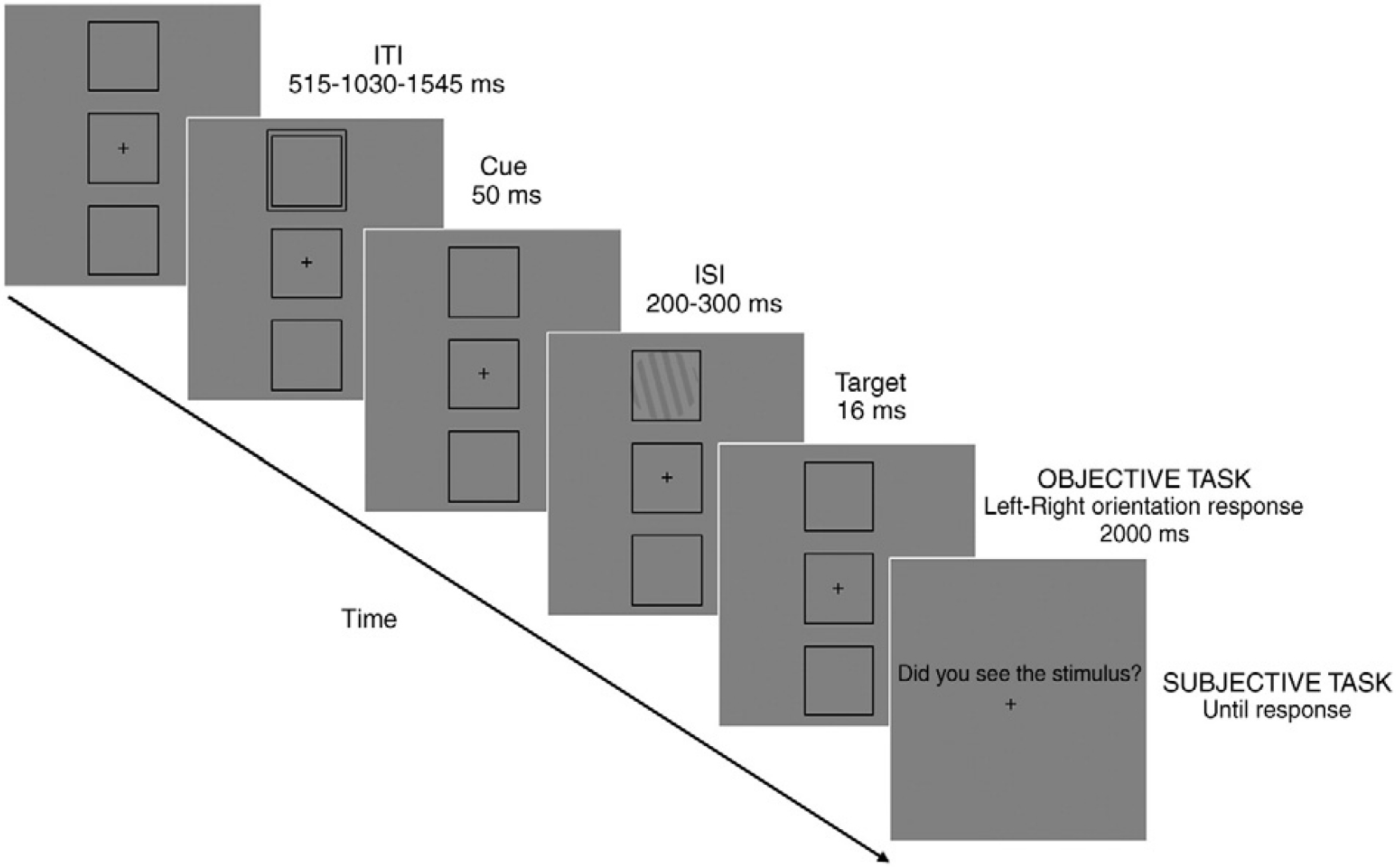

Figure 8: The sequence of events in (Chica et al. 2010) starting from the upper left. ITI = intertrial interval, ISI = interstimulus interval, in this case the period between the offset of the cue and the onset of the stimulus. Reprinted by permission of *NeuroImage*.

that parameter before they can answer the target question about orientation or side that something is on. This is more complex than any certified unconscious perception task that I know of. Further there is positive evidence, summarized in Stanislas Dehaene's recent book on consciousness (2014) that "[m]ulti-step calculations will always require a conscious effort" (p. 95).

What further can be said about whether the effect is conscious? I would be remiss if I did not mention that when you look at a good reproduction of the Carrasco stimuli (Figure 7) you can just see the effect for yourself. (Don't stare for more than a second or two though since adaptation will set in.) It can take a bit of practice to learn to do "covert attention", i.e., to move your attention without moving your eyes though. (In my 2010 paper I included a figure, Figure 2 on p. 32, one of whose purposes was to give the reader practice in covert attention.) Of course you have as much time as you like to see the effect, whereas in the experiments described you have very little time. Still, what counts for the argument I am making is the effect itself, not its timing. A further difference between just seeing the effect for yourself and the experiments described is that they utilize different types of attention, endogenous for your personal demonstration and exogenous in the experiments described. Endogenous and exogenous attention have been shown to produce roughly comparable effects in Carrasco's experiments, though in some paradigms some exogenous attention is required for endogenous attention to be efficacious (Botta et al. 2014).

I think many people are convinced of the effect because they can just experience for themselves. Not everyone can though as with almost any visual phenomenon. Of course we all know the dangers of relying too heavily on introspective judgments since they are easily manipulated. There is a line of experimentation that addresses part of the issue.

Ana Chica and her colleagues (Chica et al. 2011; Chica et al. 2010) have done a series of experiments that directly address visibility.

The Chica et al. experiment (the 2010 version) presents subjects with tilted patches that are designed to be on the threshold of conscious perception and subjects were explicitly asked whether they saw the target (after making an orientation judgment). Subjects were strongly encouraged to be conservative in saying they saw the target. They were supposed to avoid “false alarms”, i.e., saying there was a target when there was no target, and they saw periodic messages indicating how well they had been doing in avoiding false alarms. In 25% of the trials there was no target.

First subjects saw a fixation point inside the middle of 3 boxes (pictured on the upper left side of Figure 8), then there was a brief cue consisting of a square around one of the boxes. Then the target—a patch oriented either to the right or the left—could appear for 16 ms (even briefer than in Carrasco’s experiments). Next, subjects had to indicate by pressing keys—within 2 seconds—which way the patch was oriented. They had to choose one of the keys whether they saw something or not, i.e., this was a “forced choice” experiment. Then they indicated whether they saw the target or not. The experimenters adopted a procedure—tailored to each subject’s perceptual abilities—to make sure the target was at the threshold of visibility—for that subject. They started each subject with a patch of sufficiently high contrast to see the stimulus. Every 16 trials they lowered the contrast until the subject was not detecting at least 25% of the patches (by the “Did you see the stimulus” test). If the percentage of avowedly seen patches went below 60%, they increased the contrast.

The main result was that the proportion of avowedly seen patches was much higher for “validly cued targets,” i.e., when the cue was on the box that had the patch than when the cue was invalid, i.e., on the box on the opposite side or neutral (when the cue was on the middle box where no target ever appeared). In addition, the reaction time for the cued patches was much shorter than for uncued patches. Chica et al. also collected brain imaging data that suggested unsurprisingly that the valid cues attracted attention to the cued side, and more interestingly, that when the subjects saw the patch despite invalid cuing (i.e., the cue was on the opposite side), the cue had often failed to attract attention.

This experiment suggests that attention can affect whether a target is consciously visible or not. The subjects were not probed, however, on the issue of whether they actually made their judgments on the basis of the consciously visible tilt. However, when *subjects reported not seeing the target, they were at chance on reporting the tilt. And when subjects reported seeing the tilt, they were substantially above chance.* This is not the profile one sees in blindsight or in unconscious priming where subjects report not seeing the stimulus at all; but more significantly the tight relationship between consciously seeing the stimulus and being able to judge the tilt does suggest that they were reporting the tilt on the basis of the conscious perception.

Chica’s experiments are relevant to the consciousness of Carrasco’s stimuli in another way. Chica’s stimuli were presented very briefly: 16 ms in the experiment just described. Carrasco’s stimuli were presented for longer, up to 100 ms in some experiments. In addition, Chica’s contrasts were very low, as befits stimuli that were supposed to be at the threshold of visibility. The experiment described above does not report contrasts but in other papers with somewhat more complex experiments along the same lines (Botta et al. 2014; Chica et al. 2013), the contrasts required for 50% detection were about 3%; high detection seems to require up to 10% contrast. In Carrasco’s experiments, much higher contrasts are almost always used. I conclude that the reasons against the “blindsight” analogy in Chica’s experiments apply even more strongly to Carrasco’s methodology.

Given the high rates of conscious vision of 16 ms stimulus presentation even at lower contrasts than most of those used in Carrasco’s experiments, I will ignore the issue of brevity of stimulus presentations in the discussion to follow.

Keith Schneider (2011, 2006; Schneider & Komlos 2008) has argued that Carrasco's results are based on salience rather than perceptual variables such as perceived contrast, gap size, flicker rate, spatial frequency, etc. (Recently, Schneider and Jake Beck have written a draft of a paper on this topic. Rather than ascribe any specific view to a paper in draft, I will discuss the issue of salience—stimulated by their remarks—but from my own point of view.) I believe that the Carrasco lab is correct in their experimental and methodological disagreement with Schneider (Anton-Erxleben et al. 2010, 2011), however it would be digressive for me to discuss the issues involved in any detail here. I believe though that it is possible to get some insight without going into those issues.

A crude version of a salience objection treats salience as a "response bias" in the sense of a behavioral disposition to respond (in the basic Carrasco paradigm illustrated in Figure 7) by choosing the attended item. The idea is that when faced with a choice between gaps, the subject is disposed to choose not the gap that looks larger but rather the attended gap. This account is ruled out by a control in many of the Carrasco experiments of asking the subject to report the properties of the smaller gap or the patch that is lower in contrast. The attended side is still boosted in apparent contrast or gap size though the effect can be slightly smaller in magnitude so there is a small effect of "response bias" together with a main effect on perception. Carrasco also showed that choosing the lower contrast patch or smaller gap did not take any extra time, ruling out a version of the behavioral disposition objection that adds on an "inversion of response".

A more sophisticated salience objection alleges a "decision bias" in the sense of a post-perceptual feature of the cognitive process involved in making a decision of how to respond. All such accounts that I know of are ruled out by the fact, mentioned above, that the effect is substantially perceptual in nature. In addition, Carrasco showed that the effect works for some properties but not others (Fuller & Carrasco 2006). As I have mentioned a number of times, it works for saturation but not hue. Both experiments involved a procedure like that in Figure 6. In the saturation version, subjects were asked to report the tilt of the more "colorful" stimulus, where the stimuli differed in color saturation. In the hue version, subjects were asked to report the tilt of the "more bluish" stimulus, where stimuli differed along a blue/purple continuum. The result: there is an attentional effect on saturation but not hue. A cognitive decision bias should equally affect both saturation and hue. If the subjects are not aware of any difference in hue between the attended and unattended sides, it would seem that the "salience" perspective would say they would choose the attended side. But they don't. Another possibility is that the bias is perceptual in some way, say a matter of perceptual prediction (Hohwy 2013). In either case, the conclusion is that the effect is substantially perceptual and cannot be due simply to any kind of a cognitive decision bias toward choosing attended stimuli.

Whatever understanding of salience the salience objection appeals to, salience must be or be associated with a perceptual property, i.e., a property that is genuinely represented in vision. Some say (Prinz 2012) that the perceptual properties that are involved in vision are limited to a small set whose basic low level representations are products of sensory transduction: shape, spatial relations (including position and size), geometrical motion, texture, brightness, contrast and color. In other words, according to this "lean" theory of perceptual properties, though we speak loosely of seeing something as a face or as a case of causation, in reality seeing-as is limited to a small list of properties that are the output of peripheral sense organs. Others (Block 2014b; Siegel 2010) argue for a more expansive list of genuinely perceptual properties.

How do we know which properties are perceptual? We know that contrast, size, speed, spatial frequency (roughly stripe density), etc. are perceptual properties because they participate in perceptual phenomena, for example in perceptual adaptation and perceptual popout. I mentioned the waterfall illusion in which staring at a moving stimulus makes a stationary item seem to move in the opposite direction. And I'm sure every reader is familiar with color afterimages. Adaptation is a ubiquitous perceptual phenomenon that can be used to show that size, speed, stripiness, etc. are perceptual properties. Note that I am not trying to *define* the notion of a perceptual property in terms of …perception. The point rather is that perception is a natural kind and the perceptual nature of a representation is revealed in participating in phenomena in that kind (Block 2014b). By these tests, for example, there is evidence that certain face and facial emotion-related properties are perceptual. Viewers seeing an array of objects including one face can pick out the face very quickly on the basis of "parallel search", just as they can pick out a red object in a sea of green objects. Similarly there are many adaptation effects for faces and facial expressions.

Is salience a perceptual property in this sense? Attention is important to both cognition and perception, but attention can be perceptual. In order to explain the effect of attention on increasing the duration and magnitude of the tilt after-effect (and the improvement followed by impairment in discrimination) as described earlier, the visual system would have to track or register attention or where or what one is attending to in addition to being affected by attention. As I will explain in the next section, there is an open question of whether the visual system does much by way of tracking attention.

In discussions of salience there is often a conflation between salience as a perceptual property and the genuine perceptual properties that are involved in attracting attention, like high contrast or speed or sudden changes in position. We commonly speak of a saliency map in the sense of the map of locations in the visible environmental layout in terms of whether they are likely to attract attention. The perceptual properties here do not involve salience itself but rather differences with nearby locations in visible feature dimensions (Itti & Koch 2000), for example in visible motion or appearance or disappearance. People also speak of a saliency map in the brain, meaning the increased activations that correspond to attended areas. If salience is supposed to be something other than attention itself, that is If Beck and Schneider are giving an explanation that is a genuine alternative to Carrasco's, they have to show that salience is the kind of perceptual property that is registered in the visual system and that can combine in an additive fashion with contrast to affect adaptation as in the tilt after-effect. I know of absolutely no evidence for such a thing.

Many sources of evidence contribute to our knowledge of the fact that attention increases apparent contrast. John Reynolds et al. have shown that attention boosts responses in individual neurons in monkeys (2000). They developed a model of the mechanisms of this and more complex effects involving multiple stimuli (Reynolds & Chelazzi 2004). Many brain-imaging studies have shown similar effects. See sections 4.6 and 4.7 of Carrasco (2011) where much of this work is summarized. At the behavioral level, attention increases sensitivity roughly as if contrast were increased and similarly, attention can mimic the effects of increased contrast on making a stimulus visible (as in the Chica experiment mentioned earlier). And as mentioned earlier, attention increases adaptational effects as if contrast were increased. Every result involving "salience" that I have seen is just a redescription of effects of the sort mentioned.

Here is a way of seeing the emptiness of appeals to salience: As mentioned earlier, at the neural level, there are two main types of response functions by which attention increases the firing rate of neurons, multiplicative and additive. I mentioned a recent paper that compares simulations of neural responses of these sorts to behavioral data in order to ascertain which of the types of amplication are mainly being used by the visual system. Though the multiplicative models work pretty well, one additive model works very well. Thus we have strong evidence for the functional relation between attention and the increase in contrast responses in in the visual system. For simplicity, let us focus on the multiplicative mechanisms: In multiplicative gain, the response of the neuron is multiplied by a constant factor. For example, a neuron that responds to orientation will give a large response to its preferred orientation and a smaller response to other orientations to the extent that they are distant from the preferred orientation. (For example, a neuron that likes vertical lines will give a large response to vertical lines.) Since multiplying a larger number by a constant produces a larger effect, multiplicative gain is most effective for the preferred orientation. A second multiplicative response function is response gain in which the sensitivity of the neuron is multiplied by a constant factor. The effect is one of ratcheting up the response to stimuli of every orientation. The widely accepted normalization model of attention (Reynolds & Heeger 2009) explains the balance of these two mechanisms (and of additive gain) in terms of factors such as the relative size of the target and the attentional field. The attentional function of a given neuron can show more multiplicative gain or more response gain or more additive gain depending on these factors. Here is my point. We can answer the question of what the difference between these response functions is with respect to increasing apparent contrast. For example, multiplicative gain increases the apparent contrast more for the preferred orientation and response gain increases apparent contrast more for unpreferred orientations. What is the answer to the corresponding question for salience? Does multiplicative gain increase salience more for preferred or unpreferred orientations? Is it the same as for contrast? If so, then maybe "salience" is being used as a synonym for contrast. A different answer would be: multiplicative gain and response gain are equally mechanisms of salience. In this latter case it looks as if "salience" is just being used to mean attention. To repeat the general point: Those who advocate a "salience" explanation of the phenomena have to show that there is a property that is (1) perceptual, (2) not contrast and (3) acts in the ways indicated above.

Sometimes the issue is put in terms of "phenomenal salience" (Wu 2014). I think this way of talking just muddies the waters. Perceptual properties can operate in both conscious and unconscious perception. (At least: it would be an amazing discovery that there is a perceptual property that only appears at the conscious level.) Attention—at least exogenous attention —operates in unconscious perception in a similar manner to conscious perception (Chica et al. 2011; Kentridge et al. 2008; Norman et al. 2013). Further, it has recently been discovered using optogenetic methods that top-down activation of visual area V1 is about the same in awake and anesthetized mice (Zhang et al. 2014). This top-down activation involved feedback from a brain area in the mouse that corresponds to a locus of voluntary attention in humans. If salience is a perceptual property, it should be operative in unconscious perception. So the salience issue is an issue about perception, not about just conscious perception.

The upshot is that it is not at all clear how a salience objection would work, so the burden is on those who advocate it to explain it. I raised the issue of whether we are aware of where we are attending in connection with whether we are aware of salience, so I now turn briefly to that question.

7 Are we aware of where we are and are not attending?

There are a number of ways of approaching this issue, none of them very satisfying. We are certainly aware of some aspects of voluntary attention—when we "pay attention" to one thing rather than another ("endogenous attention"). But much of attention is involuntary ("exogenous"). Any perceptibly sudden movement, appearance or disappearance or sound will be likely to attract exogenous attention. Subjects in perceptual experiments can try to ignore sudden movements and sounds but they attract exogenous attention nonetheless. Exogenous attention ramps up more quickly (120 ms vs 300 ms) and dies off more quickly. Eye movements can also be voluntary or involuntary. Awareness of where the eyes are pointing is a rough index of awareness of attention. There is some evidence that people are not very aware of the time and direction of their "saccades", the quick ballistic eye movements that occur when we are visually exploring our environment. Heiner Deubel et al. showed that subjects seem to "have no explicit knowledge about their…eye position" and often don't "notice the occurrence of even large saccadic eye movements" (1999, p. 68).

However, this is not conclusive evidence that they don't know where they are attending since they may confuse movement of attention with movements of the eyes. And the visual system could track attention even if subjects are not aware of where they are attending.

Perhaps more illumination can be achieved from work on the "landscape of attention" (Datta & DeYoe 2009). Brain imaging shows a complex rapidly shifting map of spatial attention in the visual system. Spatial attention can be "focused" at one location even though there is almost as much attention at a number of other locations and some attention throughout half the visual field. The attentional field often has a "Mexican hat" shape with amplification at the center surrounded by a ring of inhibition and then an increase outside that ring. Certainly no one is aware of all that dynamic detail though I have been unable to find any specific study addressing the issue. I think it is safe to say that in normal perception there is no phenomenology that specifies much of the detail of where one is and is not attending—nor how much one is attending. So any attempt to explain Carrasco's results that appeals to our awareness of where we are attending takes on the burden of showing that we do have sufficiently fine-grained awareness of where we are attending.

8 Veridicality and representationism

In Carrasco's experiments, an attended .20° gap is not discriminated from an "unattended" .23° gap. I think the best conclusion is that attention changes perceived size and contrast. (Recall that I am talking about spatial attention rather than attention to a property instance or an object.) Do the gaps just fail to look different or do they look the same?

In Carrasco's main paradigm, subjects are forced to choose which stimulus is bigger (or faster or higher in contrast). In the case of an attended .20° gap as compared with an "unattended" .23° gap, subjects are as likely to choose one option as the other. In this sense these options are not discriminable. However, I mentioned that when subjects are asked instead whether the items are the same or different, the effect of attention is slightly smaller. And that may suggest that there is substantive daylight between not looking different and looking the same. (Of course this difference matters very much in some contexts, for example, as mentioned earlier, the context of the phenomenal Sorites issue; Morrison 2013.) As I mentioned earlier; Anton-Erxleben et al. and her colleagues argue persuasively that the smaller effect is due to the same/different paradigm being a less sensitive measure (2011). In what follows I will assume that the attended .20° gap looks the same in respect of size as the "unattended" .23° gap.

I put the "unattended" in quotes because I mean no commitment to the improbable claim that there is no attention on the .23° gap. There is no agreement on whether there can be conscious perception or even unconscious perception with zero spatial attention or whether zero spatial attention is even possible.[12] Indeterminacy in our concept of attention may even make this an unanswerable question. Still, I will adopt the abbreviatory convention of referring to stimuli that are not focally attended as "unattended".

Since the apparent size of the gap differs depending on where one is attending, the question arises as to which of these various percepts of the gap gets its size right (or most nearly right) and which gets its size wrong (or most nearly wrong). Veridicality is a matter of getting things right and veridicality in perception is a matter of the world being as it appears to be. There would be no good reason to decide that the veridical percept of the gap is one in which one is attending to a spot one inch away from it; why pick one inch rather than one centimeter? (Recall that the attentional landscape of amplification and inhibition varies from place to place and from moment to moment.) The most obvious candidate for a non-arbitrary answer to the question is: the veridical percept of the gap (if there is any veridical percept) is the one in which one is attending to the gap itself.

---

[12] Spatial attention does not require feature-based attention or attention to an object. See Wayne Wu's book on some of these issues (2014).

I think of veridicality as all or none, but for the sake of accommodating different opinions I can countenance degrees of veridicality. One innocuous use of such a phrase is that if one represents a .19° gap as .20°, the percept is more veridical than if one represents it as .21°. Also we could say that other things equal, a percept that attributes a higher probability of being .21° to a .21° gap is more veridical.

But once it is stated that the most veridical percept of the gap is one in which one is attending to the gap, one wonders why one should believe this hypothesis rather than the *opposite* hypothesis that attention distorts by magnifying, illusorily, for the purpose of getting information and that the attended item is seen illusorily. Is the perception of the gap with less attention really illusory in the sense of a discrepancy between stimulus and perception?

In an article on Carrasco's discovery, Stefan Treue (2004, p. 436) says this:

> In summary, this study provides convincing support for an attentional enhancement of stimulus appearance. It completes a triangle of converging evidence from electrophysiology, functional brain imaging and now psychophysical findings, which argues that attention not only enhances the processing of attended sensory information but manipulates its very appearance. …attention turns out to be another tool at the visual system's disposal to provide an organism with an optimized representation of the sensory input that emphasizes relevant details, *even at the expense of a faithful representation of the sensory input.* (italics added)[13]

I quote Treue not because I agree with him but in order to get a statement of that view on the table. There is no sufficient reason to accept the view that an attended perception of a gap allows us to see it as it really is rather than the view that attention in perception is like a magnifying glass, distorting for informational purposes at the cost of illusion. I can imagine considerations that might incline one towards adopting one or the other of these positions—that attention falsifies or that attention "veridicalizes"—but the adoption would be for purposes of one or another kind of utility, *not as a principled reason to think that the highest degree of veridicality is really to be found in that case.*

The challenge is to find a principled reason for regarding seeing a thing or place with a certain degree of attention to be more veridical than seeing it with a different degree of attention. Sufficiently decreasing attention to something can move the perception below the threshold of visibility. But not seeing something that is too small to see or to faint to see need not be a matter of illusion.

Chris Hill (in conversation) and Sebastian Watzl (forthcoming) have argued that there is an optimal level of attention and perception with all other values engender illusion. Watzl's version of this view appeals to the idea that the function of attention is to make perceptual representations usable—as opposed to the function of perception of veridically representing the world. These functions will conflict in normal circumstances. The optimal level of attention for fulfilling the function of perception—veridicality—will be achieved in an idealized scenario of no attention, or one of equal attention to everything. This is an interesting point of view, but is contradicted by the point made earlier that veridicality conditions require a history of veridical representation.

Epistemicists about vagueness think that there can be an unknowable fact of the matter as to a sharp border between bald and not bald, a number of hairs that a bald man can have even though adding a single hair will make the man not bald. But if there can be a fact about a border even though there could be no principled reason to regard any particular border as the real one, why can't there be a fact about what degree of attention engenders veridicality that no one could have a principled reason to accept? Epistemicists should not regard the cases as analogous since they think there is a principled reason to hold there is a fact about a border and a principled

---

[13] Carrasco (Carrasco et al. 2008, p. 1162) has been interpreted as agreeing with Treue by Stazicker (2011a) and Watzl (forthcoming). Carrasco tells me she did not mean to endorse the Treue view.

explanation for our ignorance (Sorenson 2013).

It may be objected that there is good reason to accept Treue's point of view since after all, attention to the .20° gap makes it look, illusorily, to be the same size as the unattended .23° gap. But why not blame the illusion on the percept of the unattended gap rather than the attended gap? One can blame the mismatch, but that does not help in deciding whether attention to an individual item engenders veridicality or illusion. I think the issues are clearer when one avoids the comparative perception and just asks, say of the situation in Figure 5, whether perception of the gap can be veridical when it is cued and one is attending to it or when it is not cued and one is attending elsewhere. There is no adequate justification for one answer over the other. Some may wish to abandon the notion of veridicality as applied to perception but that would be to abandon the notion of representational content as applied to perception and so to abandon representationism. The representational content of a perception is—constitutively— the veridicality conditions. There is a strong a priori case for perceptual representation (Siegel 2010). And in any case the science of perception makes essential use of veridicality (Burge 2010).

In the discussion of the analogous issue with regard to inhomogeneities in the visual field, I noted that the sort of differences in perception caused by spatial inhomogeneities are paralleled by differences due to temporal inhomogeneities—that is variation from percept to percept due to random factors. Any two percepts of the same items at the same point in the visual field with the same degree of attention are likely to differ in apparent contrast (and other properties) due to these random factors. It is hard to see a rationale for treating spatial inhomogeneities differently from temporal inhomogeneities and it is hard to see a rationale for treating either of them differently from the inhomogeneities due to distribution of attention. Claiming that all engender illusion would make most perception illusory.

We are considering the question of whether the veridical percept of the gap is the attended one or the unattended one. But is there a well formed question here? Is it endogenous attention that counts? Or exogenous attention? We are talking about spatial attention but what if feature-based or object based attention goes counter to spatial attention? That is, one can be attending to a place but also to a property that is instantiated in another place. And is it the absolute or relative value of spatial attention that matters? That is, is it some absolute attentional value or is it the most attended place that is seen veridically? Talking on a fake cell phone drains away spatial attention, causing the subjects to miss seeing objects in the centers of their visual fields (Scholl et al. 2003). (Scholl et al. used a fake cell phone to avoid the unnecessary source of variability of features of the responses from the other end of the line.) If it is absolute value that counts then when talking on a fake cell phone (and presumably a real one too), all vision would be illusory. That is a conclusion that we would have to have some very good reason to accept.[14]

One caution: I am speaking oversimply in a number of respects in asking whether attention engenders veridicality or illusion. I mentioned the issue of whether veridicality is graded. And there is an independent issue of relativity to what property one is talking about. In the experiment pictured in Figure 5,

---

[14] It may be thought that the issue of which percept is veridical is avoided by forms of direct realism that hold that there are no perceptual illusions. For example, Bill Brewer holds that in the Müller-Lyer illusion (so called) in which lines of the same length look to be of different lengths, what one is seeing is a resemblance between the situation in front of one's eyes and what he calls a paradigm of different lengths. The idea is that that is what equal lines look like when surrounded by opposite–facing arrowheads. And the way equal lines look in that circumstance is like pairs of unequal lines one has seen. On this form of direct realism, the "illusion" to the extent that one can speak of such a thing is in the mistaken inference that the lines in front of one's eyes are of different lengths. They resemble pairs of lines of unequal lines but one should not conclude that they are unequal.
However, Brewer requires differentiating between cases in which one sees a property instantiated before one's eyes that is not a resemblance to something unseen and the cases in which one sees a resemblance. In effect, the cases in which one sees a resemblance to something unseen is a pseudo-illusion category that he has to recognize. So the question arises of whether this pseudo-illusion arises in the case of attention or in the case of the lack of it. That is, is one seeing a resemblance to a non-existent thing when one attends or when one does not attend? And this is an unanswerable question for the reasons explored in this section.

the attended percept is certainly more likely to be veridical in respect of which side of the square the gap is on. And in the experiment pictured in Figure 6, the comparative percept —that is, the percept of the comparative size between the right and the left is distorted by attention to one side and improved by attention to the fixation point. So veridicality is certainly affected by attention—though in different ways for different properties. The question I am asking about gap size is whether a single gap is perceived more—or on the contrary, less — veridically if it is attended. More generally, there are certain properties—I have mentioned size, contrast, color saturation and others—for which attention to an individual item changes appearance of that property. Which is more veridical, the pre-change or post-change appearance?

One might think that there is a simple way to get at the issue of whether attention magnifies, illusorily, for the purposes of getting information, or whether attention makes things look more as they really are. You could just ask people how contrasty a patch is or how big a gap is and then consider whether those answers correspond better to reality when perception is attentive or inattentive. But the human ability to make such absolute judgments for at least some relevant dimensions is remarkably poor, certainly orders of magnitude worse than our ability to discriminate stimuli (Chirimuuta & Tolhurst 2005a). In particular, the uncertainty of absolute identification (absolute in the sense of the ability to say what the contrast is in percentage terms) is far larger than the effects of attention. Even if there were some sort of statistical advantage or disadvantage to conditions of attention in estimating contrast or gap size one would have to ask whether the advantage could be ascribed to better perception or to better inference from a percept that did not differ in veridicality.

I will assume in what follows that attended and unattended perception can both be veridical. Considerations of the same sort mentioned here also apply to the veridicality of perception in both the upper and lower visual field.

9 Indeterminate contents and the phenomenal precision principle

As I mentioned, an attended .20° gap looks the same size as an unattended .23° gap. Of course the comparative percept—the gaps looking the same—is illusory. But what about the percepts of each gap, considered separately? Is the percept of the attended .20° gap illusory? Is the percept of the unattended .23° gap illusory? I argued that we would need a better reason than we have to suppose that one but not the other is illusory. And I claimed that we should not suppose that both are illusory. The option I have argued for is that both are (or rather can be in normal circumstances) veridical. As I mentioned, the simplest perceptual representations contain two elements, a singular element that represents an individual item and a perceptual "attributive" in Burge's terminology that attributes a property to that individual item (2010). In the gap-size case, the perceptual attributive attributes sizes to gaps. A veridical percept attributes a size to a gap only if the gap has that size. In order for the attributed property to apply to both gaps, that property will have to be "intervalic", i.e., a somewhat imprecise property—for example, the property of being within the range of .20° to .23° (inclusive of endpoints). Since both gaps are in that range, both percepts are veridical (in respect of gap size).

Perhaps a probabilistic treatment of these ranges of values is in order? But how can one justify one probability distribution rather than another without making assumptions about whether the attended gap is seen more veridically than the unattended gap? For example, to say that the unattended percept of the .nn° gap represents the gap as most likely to be .nn°, whereas the attended percept represents the same gap as most likely to have some other value is to regard the unattended percept as more veridical than the attended percept. A probabilistic treatment would perhaps pass the sufficient reason test though if the same probability were attributed to both ends of the range.

It will be useful to move back to the example of contrast. The data portrayed in Figure

7 comes from the bottom right of Figure 9. Figure 9 contains four comparisons, each of which is keyed to one of the four little squares between the patches. If one fixates on one of the squares, the patch to the left of the square attended is the same in apparent contrast as the patch on the right unattended. The 22% patch can be unattended—in which case it has the same apparent contrast as the 16% patch when it is attended, or the 22% patch can be attended in which case it has the same apparent contrast as the 28% patch when it is unattended. So different veridical percepts of the 22% patch could represent it as the same as patches that are 6% more or 6% less in contrast.

Consider the contrast phenomenology of an attended percept of the 22% patch. That phenomenology is the same as the phenomenology of a 28% patch unattended. Assuming that there is not normally a phenomenology that specifies what one is and is not attending to, a matter discussed above in section 7—the phenomenology of a 22% patch attended does not carry the information of whether it is the phenomenology of a percept of a 22% patch or of a 28% patch. So in order for both percepts with that phenomenology to be veridical, the representational content would have to be at a minimum 22%-28% (inclusive of 22% and 28%).

However, there is a determinate difference in phenomenology between percepts of the 22% patch and the 28% patch when serially fixated and attended as you can verify by looking at Figure 9. (There are larger differences of this sort to be described later and as I mentioned in section 3, inhomogeneities in the visual field produce larger differences of this sort.) I believe that this determinate difference is appreciable if one moves one's attention while fixating the little square but the difference is even more obvious if one moves fixation as well as attention.

The 22% and 28% patches look determinately different if one is attending to and foveating (looking right at) each in turn. So if representationism is true, there can be veridical representational contents of 22%-28% only if the phenomenal precision of the percepts of the patches seen attended and foveated is narrower than the phenomenal precision of at least one of the percepts seen in the periphery with only one attended. This is a version of what I called the phenomenal precision principle in section 2. If two things look the same (veridically) when seen in peripheral vision with at least one unattended, but the same two things look determinately different—also veridically— when seen foveally and attentively, then the phenomenal precision of the attended and foveal percepts must be narrower than at least one of the prior percepts. And we can guess that it is the unattended prior percept that has to be less precise.

Recall that perceptual representations that are imprecise in that they attribute ranges can still be fine-grained. Suppose for example that a percept attributes a broad range of sizes to a gap of .10°-.50°. That is a different representational content from .11°-.51°, and that is different from .12°-.52°. So our ability to see small differences can be based on absolute representation even if perception is imprecise. But if the representational contents of the foveal percepts almost totally overlap, as with .11°-.51° .12°-.52°, how could those representational contents ground the determinately different phenomenologies?

Consider an analog for inattentive peripheral perception of color in which there is a red patch on one side and a blue patch on the other and the subject fixates in the middle. Suppose—and as far as I know this is science fiction—that there is some distribution of attention such that the two patches seen briefly and inattentively in the periphery can look the same and look red-blue and have the representational content red-blue. I don't mean reddish blue. I mean indeterminate as between central red and central blue or in between. They could be red, they could be blue, or they could be in between. Since attentive foveated percepts of red and blue in normal conditions are determinately different from one another (and from other colors) in phenomenology, the representational contents of red and blue seen foveated, attentively (and leisurely), would have to be more precise than the supposed contents seen peripherally, inattentively (and briefly). Otherwise there would be increasing precision in phenomenology without increasing precision in representational content and representationism cannot allow that.

In short, representationism requires that inattentive peripheral perception be less precise representationally than attentive and foveal perception.

Now here is the striking fact: there is evidence that peripheral inattentive perception of many properties is not less representationally precise than foveal attentive perception. This conclusion conflicts with the application of the phenomenal precision principle to the cases at hand. I have already discussed the peripheral vs foveal aspect of this point and I will go over some of the evidence for the attentional component in section 11 below.

I will explain the argument just sketched in more detail. But first I must discuss a piece of the puzzle, the notion of a just noticeable difference.

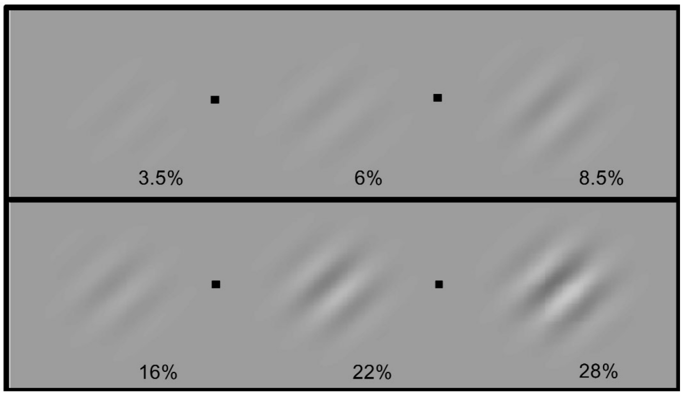

Figure 9: If one maintains fixation on one of the 4 little squares while varying attention to the patches on either side of the square, the patch to the left of the square seen with attention has the same appearance as the patch to the right without attention. I am grateful to Marisa Carrasco for this figure.

10 Just noticeable differences

A ubiquitous feature of perception is that perceptual discrimination is more fine-grained than perceptual identification. Even those with absolute pitch can identify perhaps 100 pitches (depending on exactly how absolute pitch is defined), but can discriminate many thousands of pitches from one another (Raffman 1995). Those who have absolute pitch are no better than other musically literate people in pitch discrimination (Levitin 2005, 2008). Given the disparity between identification and discrimination one might wonder whether our ability to make fine grained perceptual discriminations misleads us as to the precision of our perceptual representations. It certainly seems to us that each of those thousands of pitches has a distinct phenomenology but maybe that judgment feeds more off of the phenomenology of discrimination of differences than off of the phenomenology of individual pitches.

I have argued that the phenomenology of perception does not allow for a large degree of imprecision.[15] I appealed to the "just noticeable difference" of contrast of 2%. I said:

> The representationist may retort that the point is not that the contents are fuzzy or represented indeterminately but that they are abstract relative to other contents, as determinables are to determinates, for example as red is to scarlet. But this line of thought runs into the following difficulty: the variation of 6% due to attention is way above the 'just noticeable difference' threshold, which for stimuli at these levels is approximately 2%. (Or so I am told. In any case, just looking at the stimuli in Figure 4 [Figure 9 here] shows that the difference is easily detectable. And you may recall that in the discussion of the tilt aftereffect, there was evidence that at higher levels of contrast, the increase due to attention was as much as 14%.) The point is that there is no single 'look' that something has if it is 22% plus or minus 6% in contrast. By analogy, consider the supposition that something looks as follows: rectangular or triangular or circular. That disjunctive predicate does not describe *one* way that something can look—at least not in normal perceptual circumstances (Block 2010, p. 52).

Jeremy Goodman (2013) has criticized my reasoning. He says:

---

[15] In Block (2010). Actually, I spoke of "indeterminacy" and—mistakenly—of "vagueness" of perception.

> Ned Block, when considering the hypothesis that perceptual appearances are 'abstract relative to other contents, as determinables are to determinates, for example as red is to scarlet', objects that 'the variation of 6% due to attention is way above the "just noticeable difference" threshold, which for stimuli at these levels is approximately 2%' (p. 35).

Goodman goes on to speak of "Block's objection that our discrimination thresholds place an upper bound on the unspecificity of perceptual appearances…" (2013, p. 35). Although it may have sounded that way, I did not intend to claim that discrimination places an upper bound on either representational or phenomenal imprecision. But I think that a certain kind of discrimination is relevant to both imprecisions as I will explain.

An ability to discriminate between two observable magnitudes does not prove that one's percepts of the magnitudes actually differ (in either representational content or phenomenology). One example that I have used to illustrate this point (Block 2007, p. 540) is the phenomenon of "beats" (alternating soft and loud sounds) caused by interference between guitar strings of very slightly different pitches even when the two pitches are phenomenally the same on their own. (The frequency of beats in response to two pure pitches is the difference in frequencies.) Another is the color border effects that allow one to see that two colors are different even when the colors themselves would not be distinguishable if separated slightly. Even for achromatic objects, slight differences are amplified by a well known border phenomenon that is responsible for the famous "Mach Bands" illusion. Goodman uses the example of two trees that look to have slightly different heights because of how far they stick up above the tree canopy. His point is that vision might represent overlapping but slightly different intervalic values, but one could also use the example to illustrate heights that don't look at all different when seen separately while nonetheless allowing one to see a difference when seen next to one another.

It is intuitive to think that the way the visual system detects differences between one thing and another is by registering the properties of each thing separately and comparing those registrations. But this is not always the case: Differences are often detected via different processes than the processes that register the entities or properties that differ. Beats are produced by interference between two sound waves, allowing one to detect differences between sounds that would otherwise be inaudible.

As the examples just given illustrate, discrimination may be possible without any difference in the phenomenology of the individual percepts. Two pitches can be indiscriminable even if one knows they differ because of beats. However, there is no reason to think that specialized discrimination mechanisms are at work in the experiments described. Specialized discrimination mechanisms can be expected to depend on the specific features of the perceptual situations and so not robust to changes in the situation of the perception. For example, if you change your angle of view you might see the full vertical length of the trees but not their differential protrusion above the canopy. Border contrast effects are fragile—move the color samples just a bit apart and the effect vanishes. (This is nicely illustrated in the Wikipedia entry for "Mach Bands".) However, the attentional effects I have been talking about apply to color, speed, size of a moving object, spatial frequency (stripe density), time of occurrence, flicker rate, motion coherence (the extent to which many moving items are going in the same direction), as well as to contrast and gap size. What is the chance that there is some specialized discrimination method at work for all these magnitudes? Most impressively, these effects can be exhibited in visual short term *memory*—that is, they don't even require simultaneous perception. This was shown by Martin Rolfs & Marisa Carrasco using a different experimental paradigm than the ones so far discussed (2012; Rolfs et al. 2013). I won't describe it except to say that the patches are compared in respect of contrast by comparing a patch seen earlier with a currently seen patch, and with similar results to those already described. As I mentioned in section 3, a similar experiment shows that a perceived patch at one location in the visual field can be compared with a remembered patch at a different location with results that show the inhomogeneities in the visual field (Montaser-Kouhsari & Carrasco 2009). The likelihood that there is some method of comparison that does not depend on the individual percepts themselves but survives all these variations does seem slight.

So the kind of discrimination that is not based on specialized mechanisms of detecting differences independently of registering absolute value can be used to make a better case.

But even if we can make very fine grained discriminations and even if the percepts involved in the discriminations are distinct, it does not follow that the percepts are not highly imprecise—as mentioned earlier. Suppose for example, that perceptions of contrast are so imprecise as to cover nearly all the range of contrasts. Consider a representation of contrast of 4%-98%. Still, 4.1%-98.1% would be another equally imprecise content that is nonetheless distinct from the first one. And so more generally discriminability has little in the way of immediate consequences for imprecision.

The notion of a just noticeable difference is not very useful for my purposes. Discrimination can be finer than absolute registration as in the case of beats. And strong ability to discriminate is compatible with a high degree of imprecision. Further, the notion of a just noticeable difference combining as it does, perception with cognition, allows the possibility of a difference in conscious percepts that is not cognitively accessible.

11 Absolute representation

The phenomenal precision principle tells us that if the phenomenology of perception is grounded in its representational content, then peripheral unattended perception must be more imprecise than foveal attended perception. This result applies to contrast, size, spatial frequency and some other properties but not location. However, experimental results to be described in the next section suggest that contrast perception is as precise in foveal attended perception as in peripheral unattended perception. But what this evidence does not tell us is how precise they both are, i.e., whether both are relatively precise or relatively imprecise.

I mentioned a study by Mazviita Chirimuuta & David Tolhurst (2005a) that is relevant to the issue of how precise absolute representations of contrast are in foveal attended perception. Chirimuuta and Tolhurst have a behavioral result that shows that performance in classifying contrasts falls off sharply after 4 contrasts. They have a neural model of contrast identification that suggests that the brain is capable of representing only 4-5 contrasts and that this limit is compatible with very fine-grained discriminations. Chirimuuta's view is that the response probabilities in the visual system for contrasts are very broad, with the tails of every distribution covering much of the span of possible contrasts. (That is, there is a non-zero probability across almost the whole range of contrasts.) Contrasts can only be identified when the response is near the peak of the probability distribution but two responses can be compared when responses are in the tails so long as the tails do not overlap much.

I'll start with the behavioral result. She presented subjects with a number of patches of up to 8 grades of contrast that were labeled “1” through “8” in each sequence of trials. Subjects looked at the contrasts and labels for as long as they liked and could have a refresher any time in the midst of the experiment if they liked. They had to hold the pairs of digits and contrasts in working memory and assign numbers to contrast stimuli. Then, patches were presented for half a second and subjects had to try to give the digit label. Performance was good up to 4 items and fell off drastically for larger sets.

Performance on 4 contrasts was near perfect. Then when new contrasts outside the original range were added, performance fell off, even for the original 4 contrasts. This is a pattern often seen in working memory experiments. For example, wild monkeys participated in an experiment in which an experimenter sets up two buckets and ostentatiously places, one at a time, a number of pieces of apple in each bucket. For example, there might be 4 in one bucket and 3 in the other. The result is that for numbers of slices of 4 or less, monkeys reliably go to the bucket with more but with more than 4 items, performance falls off to chance (Barner et al. 2008; Hauser et al. 2000). Human infants show similar results with a limit closer to 3 (Feigenson et al. 2002).

The number 4 figures in working memory experiments in which subjects are asked to remember digits but are given another simultaneous distraction task to prevent overt strategies of "chunking" digits into units. Subjects can typically remember about 4 digits. In a completely different paradigm, George Sperling showed subjects a grid of letters briefly (1960). Subjects often said they could continue to see all or almost all the items faintly after the patch disappeared. (This kind of image has been called a visual "icon".) When the grid had 3 rows of 4 items, and subjects were asked to recite as many letters as they could, they could name 3-4 letters. However Sperling gave subjects a cuing system: a high tone for the top row, a medium tone for the middle row and a low tone for the bottom row. When cued, subjects could report 3-4 from any given row.

In a different paradigm, honeybees were trained on a maze in which they had to choose to go either left or right at a T-junction to get a reward. At the entrance of the maze there were dots on each side and the bees had to choose the side with more dots to get the reward. The bees could learn to choose 4 rather than 3 but not 5 rather than 4 (Gross et al. 2009).

The working memory significance of roughly 4 items is so ubiquitous that it stimulated an article called "The magical number 4 in short-term memory: A reconsideration of mental storage capacity" (Cowan 2001). Up until 5-10 years ago, "slot" models of working memory were popular. I think it would now be agreed that roughly slot-like behavior emerges from an underlying working memory system in which there is a pool of resources that is distributed over items differently depending on number and complexity (Ma 2014). George Alvarez & Patrick Cavanagh (2004) suggested that there might be a limit of around 5 items of ideally simple structure but Alvarez's recent work suggests a more complex picture in which there are a variety of components of working memory that may independently fit a more slot-like or a more pool-like structure (Suchow et al. 2014). Slot-like working memory depends on simple stimuli that are hard to confuse with one another. Stimuli that have shown slot-like behavior include alphanumeric characters, horizontal/vertical rectangles and colors that differ substantially from one another. (I am indebted to conversations with Weiji Ma on this topic.)

So I would suggest that Chirimuuta's behavioral result probably depends on the fact that subjects had to hold a number of pairs of digits and contrasts in mind in order to categorize the next contrast. (You could try it yourself for say 5 lengths.) They did well up to 4 such pairs and then performance declined radically. The article contains an anecdote that further supports this idea:

> DJT [one of the subjects and experimenters] performed an experiment in which 4 contrasts of grating were chosen that were close together whilst still allowing near-perfect identification performance over 50 trials of each: 1, 8, 18 and 27 dB. [Note from NB: this is a different way of quantifying contrast than the percentages used here.] In the 50 trials of each contrast, 1 error of identification was made for each of the 8 and 18 dB gratings. Then, two more contrasts were added to the stimulus set at the lower end (40 and 50 dB); contrast 40 dB should have been easily discriminable from 27 dB. In fact, addition of contrasts 40 and 50 dB resulted in an increase in the errors of identification of the original set of four contrasts over 50 trials of each (8dB – 2 errors; 18dB – 9 errors; 27dB – 6 errors). (Chirimuuta & Tolhurst 2005a, p. 2965)

There are two notable aspects of this anecdote: first, performance over 50 trials of each of 4 contrasts were near perfect despite the fact that the gratings covered only part of the spectrum of contrasts. This suggests that the limit of 4 does not have to do with representations of contrast per se. The second aspect is that in this case as in so much of the work on working memory, adding more possibilities to a set of 4 decreases performance in the original set of 4. I conclude that the behavioral result probably has more to do with working memory than with any limit on perception.

Chirimuuta's second result, the one that motivates the idea that visual representations of contrast are so indeterminate that only 4-5 levels of identification are possible, is the modeling result based partly on data from monkey V1 neurons. (V1 is the first cortical area that processes vision, the lowest level of the visual system.) The striking fact about this result is that it does not concern working memory at all or indeed any kind of memory. It is only concerned with perceptual representation in V1. The model of V1 neurons comes from another paper that is concerned with the "dipper function", a notable curve shape in which one contrast stimulus is "masked"—diminished by the processing of another stimulus that follows right after it (Chirimuuta & Tolhurst 2005b). The model predicts that V1 can represent 4 contrasts perfectly with a sharp fall-off at 4, with a capacity to represent slightly more than 5 items.

However, the model based on V1 neurons gets some important facts wrong, for example it predicts poorer performance at high and low contrasts, whereas people actually do better at high and low contrasts. A version of the model with some postulated features that are not based on anything neural can get that right. However, this "curve fitting" approach deprives the model of the neurophysiological support that motivated the original model. Another problem with the model is that what is predicted is "mutual information" shared between contrast stimuli and V1 responses of 2.35 bits. Mutual information is a measure of shared information—in this case between stimuli and V1 neurons. A mutual information value of 2 bits would allow 2^2 (=4) contrast identifications; a mutual information value of 3 bits would allow 2^3 (=8) identifications. This shared information, as Chirimuuta notes (Chirimuuta & Tolhurst 2005a, p. 2968), is "essentially looking at perfect, 100% performance." For this reason, mutual information is not very useful as a psychophysical measure. And as Chirimuuta notes, its utility is limited for another reason: it is a compressive measure and so large increases in neural activity can be expected to make small differences in information. The issue of 100% performance is especially troublesome since in perceptual systems no performance can be perfect. In particular the convention for a "just noticeable difference" is distinguishability 75% of the time. So it is difficult to know how to compare the absolute identification level of 2.35 bits with a more visually sensible visual identification level.

Further, our experience seems to conflict with the idea that we have distinct visual representations of only 4-5 contrasts. A good reproduction of Figure 8 seems to reveal 6 phenomenologically different contrasts even though the figure covers only a third of the range of contrasts. And the Carrasco results apply to many different parameters, gap size, spatial frequency, etc. You might test it out if you happen to be near a brick wall. Look at the height of one brick, two bricks, three bricks and four bricks. If you are close enough so that those sizes look different from one another, ask yourself whether there are other sizes that look different from all four of those sizes. If Chirimuuta's result applies more widely, the answer is no. It has to be said though that that sense of distinctness could be due to discriminatory abilities.

Whatever the facts are about how precise foveal attentive perception is, the next section presents evidence that it is not more precise than inattentive peripheral perception.

12 Attention may not increase representational precision

I said that if 2 things look the same when seen in peripheral vision with at least one unattended, but the same two things look determinately different when seen foveally and attentively, then the phenomenal precision of the attended and foveal percepts must be greater than at least one of the prior percepts. (As I mentioned, the assumption of veridicality is required to justify the imprecise representational contents of the peripheral percepts.)

It is common for philosophers to claim that attention increases "determinacy" of perception (Boone 2013; Nanay 2010; Stazicker 2011a, 2011b, 2013; but not Speaks 2010). The relevant kind of determinacy as I have been saying is precision. But it will be useful to distinguish precision from other forms of determinacy. Responses to attended stimuli are certainly less variable than responses to unattended stimuli. And attention increases acuity in the sense of spatial resolution, e.g., the ability to distinguish one dot from two dots. I will argue that spatial attention may not increase precision even if it reduces variability and acuity, and that further, in a rationally designed system spatial attention would not be expected to increase precision.

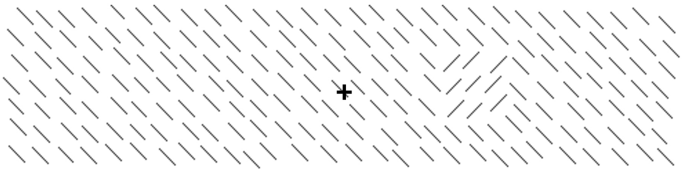

Figure 10: A textured figure used by Yeshurun & Carrasco (1998). Using stimuli like this one, stimuli were presented in which the square immediately to the right of the plus sign could appear at different eccentricities. When the resolution was low in peripheral areas, attention increased the subjects' ability to detect the square. But when the resolution was high—nearer to the fixation point--attention *decreased* the subjects' ability to detect the square because the increased resolution obscured the forest in favor of the trees.

Yeshurun & Carrasco (1998) showed that attention can increase resolution, making subjects (paradoxically) less likely to see the attended stimulus. For textured figures like the square to the right of the fixation plus sign in Figure 10, there is an optimal degree of resolution. If resolution is too high, the subjects miss the forest for the trees, failing to see the larger scale textured figures. Too low a resolution can cause subjects to miss the trees as well. Yeshurun and Carrasco presented textured figures at varying degrees of eccentricity. Since resolution is better for stimuli that are closer to the fovea, this had the effect of presenting the figures at varying degrees of resolution. They also varied resolution by manipulating where subjects were attending, using cues of the sort described earlier. Putting together the contributions to resolution from eccentricity and attention, they were able to show that there were different optimal degrees of resolution for different figures.

One neural mechanism by which attention increases resolution is shrinking of the "receptive fields" of neurons in the visual system. Recall that a receptive field is the area of space that a neuron responds to. Resolution increases when neurons respond to smaller areas. Another mechanism is the shifting of receptive fields from adjacent areas that was mentioned earlier.

As I mentioned, the sensitivity of high "spatial frequency" channels increases—probably as a result of these mechanisms. Recall that spatial frequency in the case of a stripy stimulus like the "Gabor patches" used in many of the figures In this article (e.g., Figure 15) is a measure of how dense the stripes are. Boosting the sensitivity to high spatial frequencies makes resolution higher, thereby improving perception of textured figures when the resolution is too low and impairing perception when resolution is too high. The Yeshurun and Carrasco paper concerns exogenous attention. Later work (Barbot et al. 2012) shows that endogenous attention is more flexible, raising or lowering the sensitivities of high spatial frequency channels so as to improve perception.

I mention the increase in resolution and the sensitivity to high spatial frequencies in order to be sure that the reader is distinguishing these matters from an increase in precision.

Representational precision is a matter of how wide a range of values is allowed by the representational content, what values are compatible with the veridicality of the percept. (Phenomenal precision is a matter of "crispness" of the appearance.) One dot and two dots may look the same in peripheral vision even though we can clearly see the difference in foveal vision. That is a difference in acuity rather than a difference in precision. Increasing precision for representation is sharpening the representational content.

The relation between variability and precision is more complex. Imprecision is often cashed out in terms of reliable correlation between a representation and the world (Stazicker 2013). In that sense, since attention decreases variability it must increase precision. However, there are different sorts of noise. As we will see in the first experiment to be described below, attention may decrease noise across the whole spectrum without affecting what might be thought of as intrinsic variation in the signal and thus not increasing a kind of systematic precision. As I will explain, the experiment to be described helps us to precisify what precision comes to.

I will describe two experiments that will help to make the notion of precision more precise or at least concrete and will suggest that spatial attention does not narrow representational precision. Before I do that, let me say briefly why one should expect that spatial attention will not make the attended properties any more precise. Increasing precision normally involves suppression of responses outside the expected range. It would not make sense for a system to be designed to suppress some values without some indication of the irrelevancy of those values. Spatial attention tunes for spatial area, suppressing responses to other spatial areas (Montagna et al. 2009). So spatial attention can be expected to increase precision for spacial location but not for contrast, size, spatial frequency or speed.[16] For feature-based attention, the opposite is true. If one is looking for the red thing, it makes sense to suppress sensitivity to other colors. Spatial attention should tune for space only and feature-based attention should tune for the property attended to.

The first experiment uses the "attentional blink", a phenomenon in which there is a series of stimuli and two targets amid distractors. In part of the experiment, the targets were squares and the distractors circles. The general finding is that if the subject consciously sees the first target square, and if the second target square is presented 200-400 ms after the first square, the subject will be much less likely to consciously see the second square. The mechanism has been shown to depend on the first target absorbing the subject's attention so that there is insufficient attention to consciously see the second square. The second square is described as "blinked", where the blinking deprives the square of attention. Asplund et al. (2014) used this technique with a paradigm in which the target squares were colored and in which subjects had to report the color of the second square by moving a mouse to click on a color wheel that had 180 colors on it. The idea is that the effect of attention on how intervalic the perceptual representation is could be assessed by examining the effect of the presence or absence of attention on the precision of subjects' identifications of the color using the color wheel.

The experimental procedure is diagrammed in Figure 11. The subject saw a fixation point (lowest square on the left). Then there were 7-13 colored disks, then a target, T1, a square that was either black or white (RSVP = rapid serial visual presentation), then some number of colored disks, then another square, then 3 more disks. Then subjects reported the color of T2 using the color wheel. They got immediate feedback in how far off they were on identifying the color (in degrees on the color wheel) for 500 ms, then they indicated whether T1 was white or black. If the subject got T1 wrong, that trial's report of T2 was disregarded. This design allowed for comparison of precision of reporting the color of T2 between trials in which attention was maximally reduced (T2 presented 2 items after T1, described as "lag 2") with trials in which the lag was so long or so short that there was no attentional blink at all. The key result is that although the lag time was strongly correlated with the average correctness of the response (as always in the attentional blink), *the precision of the responses that were not random was not affected significantly*. The same experiment was done with faces using a slightly different form of the attentional blink. T1 was one of two faces that subjects had to recognize and the response wheel for T2 had a

---

[16] This is oversimple since attention increases sensitivity to high spatial frequencies (Barbot et al. 2012).

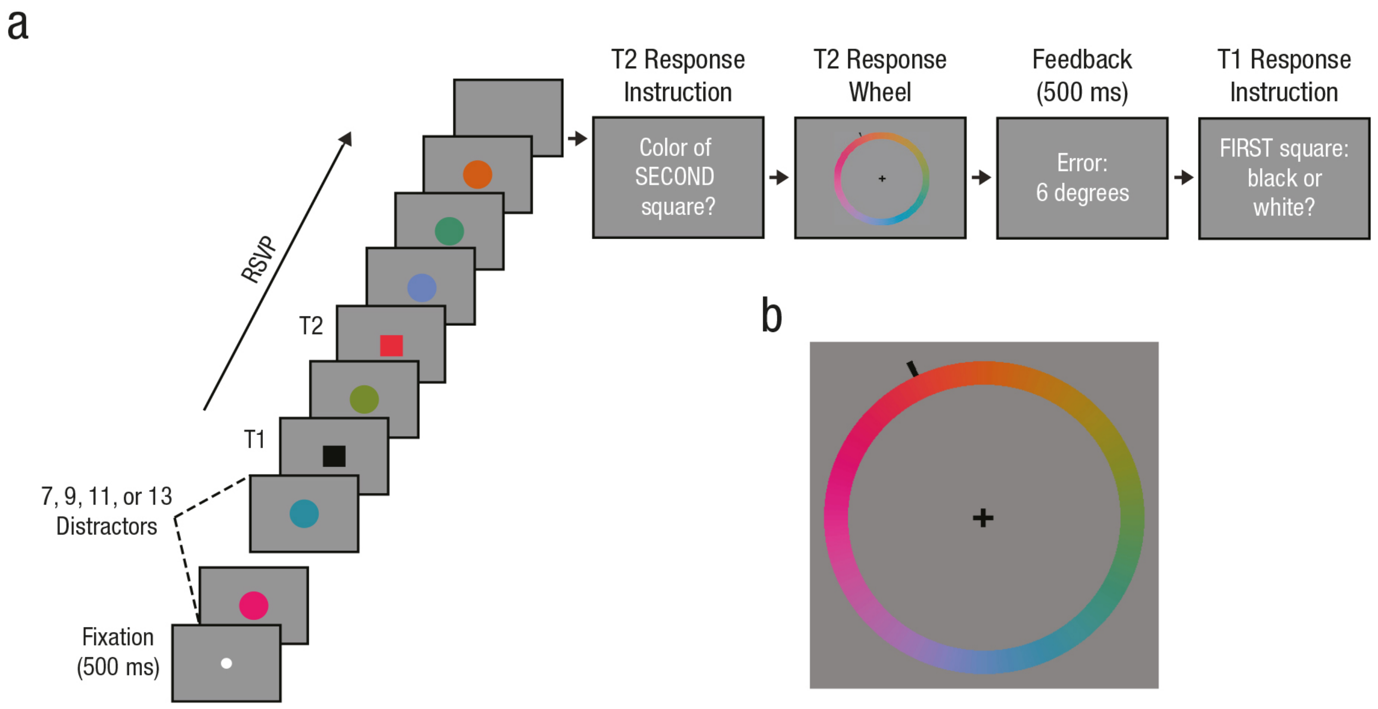

Figure 11: Procedure from Asplund et al. (2014). Understanding of this diagram is aided by color reproduction. Thanks to Chris Asplund for supplying this figure.

series of 150 face morphs based on 3 faces, with 49 intermediate faces interposed between them. The results were the same with faces as with colors. The key result for both studies is that the identification of T2 was either random (much more likely at the critical "lag 2" for an attentional blindness effect of 200-400 ms) or just as precise at lag 2 as at lag 8. Note that the experiment does not directly test the precision of any single percept. The assumption is that the precision of representation of a blinked color will be reflected in how tightly clustered the different responses are. Asplund et al. (2014) conclude (p. 6): "Across both stimulus classes (colors and faces) and experimental designs ..., we found that the reported precision of a target item is not affected in the AB [attentional blink], even though our paradigms had the sensitivity to detect such effects."

But wait, you may ask: "Didn't I say that attention decreases variability? And why is there supposed to be a difference between (the inverse of) variability and precision?" (Indeed, the inverse of "variance", one measure of variability, is a common notion of precision.) The answer is that if you look at the raw data in this experiment, the blinked color identifications are much more variable than the ones that are not blinked. However, the authors were able to show via modeling that the response distribution was a superimposition of two very different distributions. One distribution was uniform over the whole color wheel with no clustering around one color, whereas the second distribution was tightly clustered around the correct color, *just as tightly as when the color stimulus was not blinked*. They reasoned that the first (random) distribution represented cases in which the subject simply did not see the stimulus. However, when the subject did see the stimulus, the precision of the response was just as if it had not been blinked. (They considered and rejected a "variable precision" model that predicted the data less well; van den Berg et al. 2012). So overall variability of response is not a good guide to the precision of the representation. And this shows an important flaw in crude correlational approaches to precision. The precision of a perceptual representation should not be taken to be a matter of how well perceptual representation correlates with stimuli since what is really

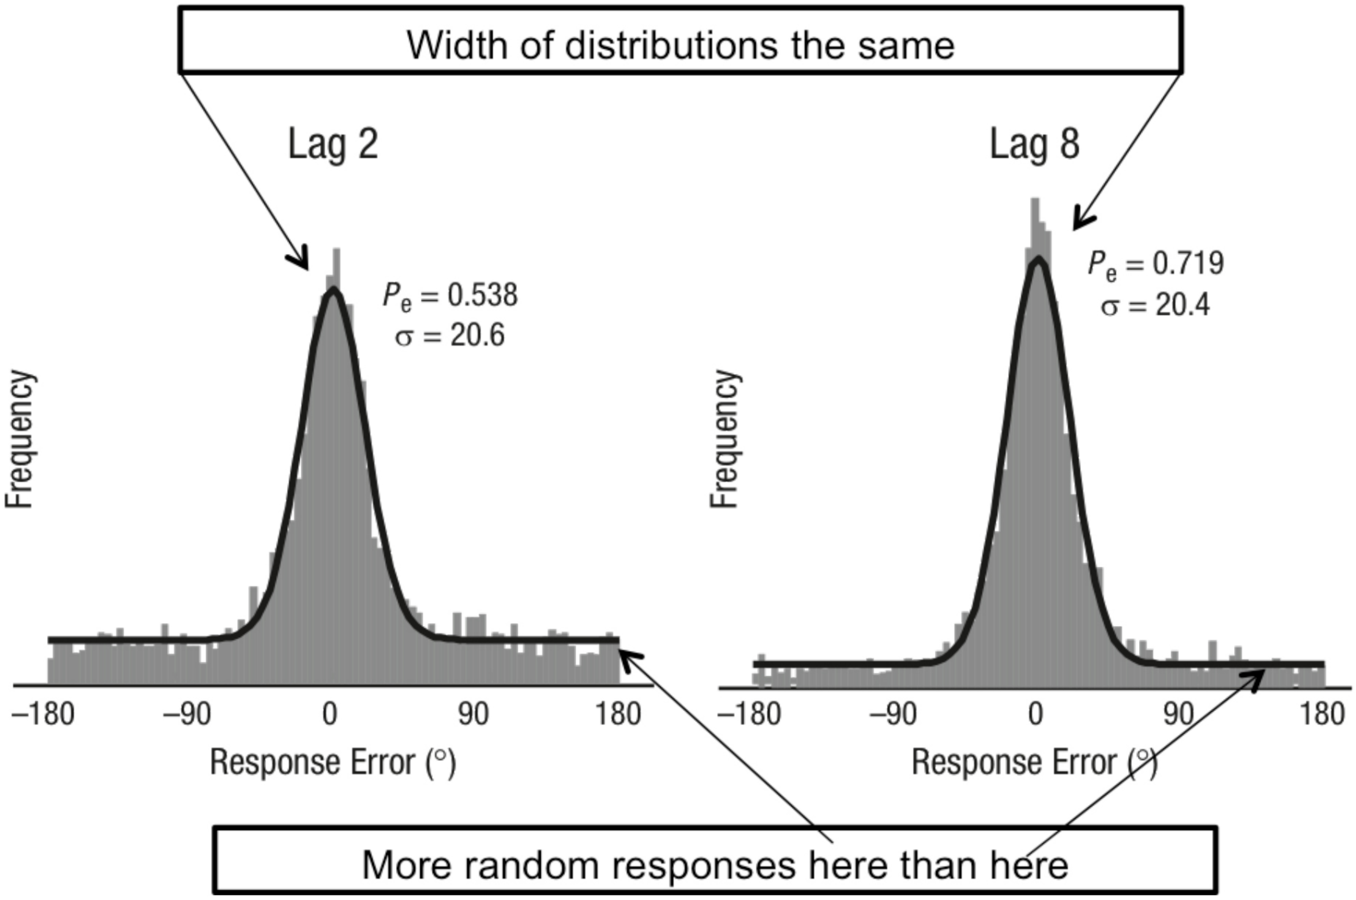

Figure 12: This is a modified form of a figure from Asplund et al. (2014). The figure compares response errors for lag 2—the value with the maximum effect of the attentional blink with lag 8—the value with the minimum attentional blink. What the figure shows is that the precision of the responses in which the subject actually saw the stimulus was the same. And the figure shows an increase in random responses for the blinked stimulus. Thanks to Chris Asplund for supplying the figure which has been modified here.

relevant is the cases in which the subject actually sees the stimulus.

This point is illustrated in Figure 12 in which the response error profile for lag 2 in which the attentional blink is most powerful is compared with the response error profile for lag 8 in which the attentional blink is least powerful. The widths of the distributions are the same. What differs is the number of random responses as indicated by the higher “tails” of the distributions.

Note the difference between precision and veridicality in this experiment. Precision is a matter of how tightly the responses cluster and veridicality is a matter of whether the responses cluster around the value of the item that was seen regardless of how tightly they cluster. If the color seen was focal red, responses could pick out focal green in a very precise manner, but be non-veridical nonetheless. Conversely, the average of the responses might be the color seen (focal red) and thus the responses are on the average veridical even though the intervalic content is very wide.

An objector might say that the cases in which the blinked stimulus is reported in a non-random manner might be cases in which it was not in fact deprived of attention by the first percept. Imaging studies of the attentional blink do suggest a general deprivation of attention of the blinked stimulus (Sergent et al. 2005) but I don’t know of one that looks specifically at this issue. There are always potential confounds and the general remedy is to approach the same issue in more than one way. In the case of this result, the same conclusion has been reached by approaches that do not share vulnerabilities.

Another approach recorded from single neurons in a monkey visual area (V4) that is known to be sensitive to shape and form (David et al. 2008). Orientation tuning was not narrowed by spatial attention, but it was narrowed by attention to a specific orientation—feature-based attention. A recent review (Ling et al. 2014) summarizes this approach as follows:

> Although initial physiological reports suggested that directing spatial attention to an item sharpens the band-width of orientation-selective cells in macaque visual area V4…, this was later shown not to necessarily be the case. Follow-up studies using a more sensitive measure for tuning bandwidth found no effect of spatial attention on the width of the orientation tuning function… Rather, these studies instead only found changes in the responsivity and baseline firing rate of neurons coding for the spatially attended location. Thus, the neurophysiological evidence appears to indicate that spatially attending to a location leaves a neuron's feature tuning unaffected.

A psychophysical study came to the same conclusion, that spatial attention boosts activation but not precision.

Ling et al. (2009) contrasted spatial and feature-based attention. The stimuli were random-dot cinematograms in which dots move in one direction or another for short distances. In the low noise condition shown on the left side of Figure 13, the dots show a high degree of coherence in that most of them move in the same direction. As noise increases, motion coherence decreases. Subjects had to make a series of judgments of the orientation of overall motion of these cinematograms. In the spatial attention version, they could be cued as to the place the stimulus would appear. In the feature-based attention version, they were cued to one of four directions of motion and had to report the observed motion as a clockwise or counterclockwise deviation from the cued motion. Ling et al. were especially interested in comparing two different models for how attention boosts performance in detecting the direction of motion, using stimuli that could move in different directions. See Figure 14.

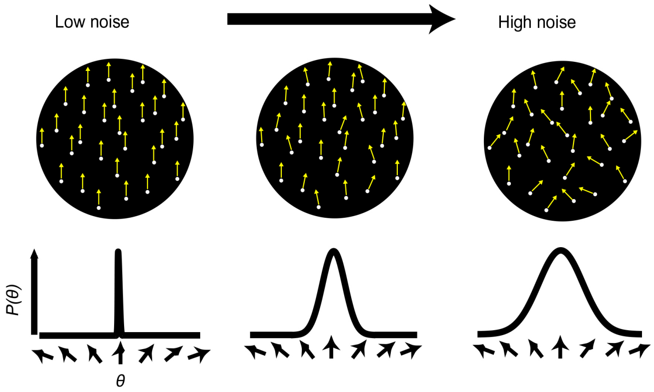

Figure 13: Random-dot cinematograms in which dots exhibit local motion. In the low noise condition, most of the dots are moving in the same direction. As noise increases, the spread of directions increases and motion coherence decreases. From an experiment comparing spatial attention with feature-based attention. With permission of *Vision Research*.

According to the gain model of (a), the response to the stimulus is increased as if the volume—i.e., the signal strength—were simply turned up equally for all movement-direction detectors. (Orientations of motion are indicated by arrows along the x-axis. The signal strength was turned up in the sense that the signal strength prior to attention is multiplied by a constant factor. For the values that are already high—i.e., at the peak—the multiplying a large value by a constant factor has a bigger effect than at the tails of the distribution where multiplying the constant factor times a zero yields zero.) According to the sharpening model of (b), the effect of attention is not to turn up the response but rather to suppress the irrelevant noise in the stimulus, narrowing the intervalic range of the response profile. These two models make different predictions for "threshold vs. noise" curves pictured in the bottom of Figure 14. The gain model predicts an increase in discriminability only when the external noise is low compared with internal noise. When external noise is low, there is a benefit to turning up the volume— even though the volume increases both signal and external noise—since the effect of turning up the volume is to "swamp" the internal noise. (Internal noise is a blanket term for variation in the visual system, whatever makes visual responses vary even when the external signal remains the same.) As external signal and noise dominates the percept, internal noise decreases in importance. This kind of gain has a similar effect as decreasing the internal noise. If internal noise were zero, there would be no benefit at all in raising the volume on both the signal and the noise. The benefit of raising the volume however dwindles away as external noise increases since the increase in volume increases the effects of external noise too. This is indicated by the lowered threshold on the bottom left of (a) where the advantage in lowering the threshold decreases as external noise rises.

A different picture emerges from the model of the bottom right (b) where the benefit of tuning is greatest when external noise is greatest. (Note that if there is no external noise, tuning is of no benefit.) Thus the benefit should increase as noise increases as pictured in the bottom right (b). These models were tested by a procedure somewhat like that in Figure 5, except using voluntary attention. A line indicated where the subjects were supposed to attend and then a tone indicated that the stimulus was about to appear. Subjects could be cued to one of 4 locations where their task was to report the direction of motion of a stimulus. Sometimes there was a tone but no cue. The result was unequivocal: a pattern like that of the bottom left of Figure 14, indicating an effect of gain but no tuning. "The data showed that spatial attention yielded benefits strictly with low external noise, and no benefits with high external noise" (Ling et al. 2009, p. 1201). They also used the same setup with feature-based attention in which the subjects were cued with an indicator of what the direction of the stimulus would be. In this version, there was both tuning and gain, showing a hybrid of the patterns of a and b in Figure 14.

Again, spatial attention does not appear to narrow representational precision, contrary to the representationist position. This is graphically shown in the tuning model of Figure 14: suppression of values outside the selected value directly reduces precision. This is what does not happen with spatial attention.

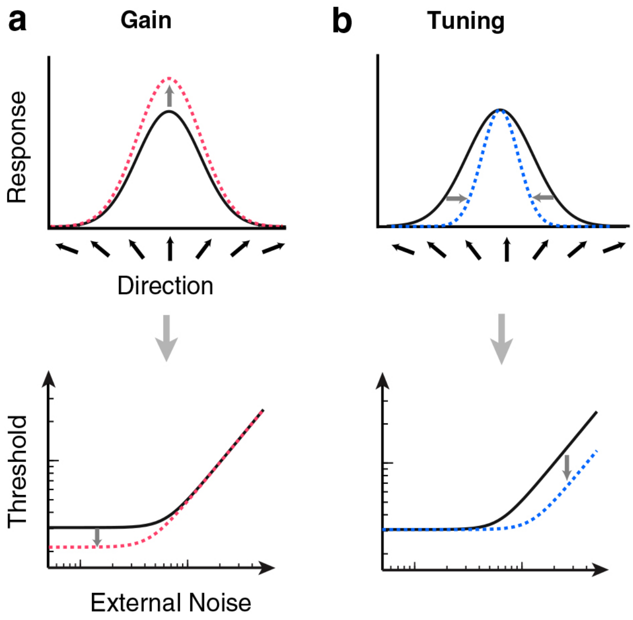

Figure 14: Two models of how attention boosts performance. According to the gain model indicated in the top left (a), the boost derives from increasing the firing of all directional feature detectors. The arrows along the x-axis indicate receptors for motion in different directions. The dotted lines represent the change due to attention (as compared with the solid lines). The tuning model at the top right (b) says performance is boosted by sharpening the response, decreasing the range of the intervalic content, as indicated by the narrowed shape of the dotted line. See the text for an explanation of the bottom diagrams. From Ling et al. (2009). With permission of *Vision Research*.

But why does this result concern representational precision rather than phenomenal precision? I considered an analog of this question concerning peripheral vision in section 2. There I noted that the anatomical asymmetries that are the probable basis of the inhomogeneities discussed are bound to affect unconscious perception in the same way as conscious perception. And a similar point applies here. The narrowing of receptive fields that is the main underlying mechanism of the attentional effects concerns perception *simpliciter* rather than conscious perception *per se*. As I mentioned earlier, spatial attention operates in unconscious perception in a similar manner to conscious perception (Chica et al. 2011; Kentridge et al. 2008; Norman et al. 2013).

The dimensions used in both of the experiments described are “metathetic” as opposed to “prothetic” (Stevens & Galanter 1957). Prothetic dimensions have a zero point and intrinsic directionality, whereas metathetic dimensions have neither (Fuller & Carrasco 2006). Color saturation is prothetic because there is a zero point—achromaticity—and colors are more or less saturated. Hues are metathetic. At least for primary hues such as red and green, neither has more of any hue. Carrasco’s work shows that the attentional effects involved in increasing size, speed, flicker rate and the like work for prothetic dimensions like color saturation but not metathetic dimensions like hue (Fuller & Carrasco 2006). And that fact leads to the question of whether the conclusion that attention does not change precision depends on the magnitude tested being metathetic.

The studies on prothetic dimensions are not as easy to interpret as the ones I just described. One reason is that for metathetic dimensions, the psychological meaning of a difference is roughly the same throughout the dimension. A 90° shift in direction has roughly the same significance independently of the starting direction. But for prothetic magnitudes that is dramatically not so. A one inch change in a length of .01 inch has a different psychological significance than a one inch change in a length of one mile. Baldassi & Verghese (2005) give some evidence that spatial attention does not change the intervalic range of detection of contrast—a metathetic magnitude—though feature-based attention does narrow intervalic range.

The review I mentioned (Ling et al. 2014) surveys many different studies on this issue, concluding (references removed):

> By and large, studies using psychophysical techniques to assess selectivity have converged on results that square quite nicely with the neurophysiological results…: feature-based attention to an item selectively changes psychophysical tuning curves…, while directing spatial attention to that item leaves behavioral feature tuning untouched…

I mentioned that increasing precision normally involves suppression of responses outside the expected range. There is no reason for spatial attention to increase the precision of anything else other than spatial area. In particular, why would spatial attention suppress some directions of motion and not others? However if attention is directed towards motion in a certain direction (feature based attention) then increasing precision does make sense. The point applies equally to prothetic as to metathetic dimensions. Why should spatial attention tune for some values but not others of contrast or gap size given that tuning involves suppression of some range of contrasts or gap sizes. So there is good reason to expect these results to apply to prothetic dimensions.

Let me return to the issue of peripheral perception as compared with foveal perception. I mentioned the experiment that shows that discrimination of contrast in the periphery is as good as in the fovea. But there is an additional fact about peripheral vision, a phenomenon of “crowding” in which things lose the quality of ‘‘form…without losing crispness…’’ (Lettvin 1976). We can ignore crowding for the purposes discussed here so long as we confine ourselves to perception of what the visual system treats as single objects. For more on this, see Block (2012, 2013).

To conclude, there is evidence that attended and foveal perception can be greater in phenomenological precision without being greater in representational precision, contrary to representationism. In direct realist terms, there is evidence that attended and foveal perception can be greater in phenomenological precision without involving awareness of more precise environmental properties.

13 Abstraction and indeterminacy

The purpose of this section is to argue that the 6% difference between the attended 22% and unattended 28% patches underestimates the effect of attention. The reader who is willing to take that on faith can skip to the conclusion.

I argued that since the attended 22% and unattended 28% patches look the same when seen in peripheral vision but look determinately different from one another when seen foveally and attentively, we can conclude that the precision of the phenomenal and therefore representational content of the attended foveal percepts must be greater than that of the prior percepts —if representationism is true and all the mentioned percepts are veridical. The reader may not be convinced however that the 22% and 28% patches do look determinately different when seen foveally and attentively. Perhaps the sense that they look different is a matter of an ability to discriminate rather than an appreciation of appearances that are determinately different.

The Carrasco lab experiments reported so far use stimuli that are 4° from the fixation point. But you might have noticed that when you fixated the bottom left square in Figure 9, you could also see the 28% patch to the far right. And some of the Carrasco lab's experiments have been done with 9° angle of separation.

If one combines the two different angles of separation as in Figure 15 an attended 16% patch looks the same in contrast as an unattended 28% patch, a larger difference than mentioned earlier for this absolute level of contrast. (The differences produced by attention increase with absolute level.) Of course the logic of the case is the same as before. I introduce it because I think it is easier to be sure that the patches in Figure 15 look determinately different when foveated and attended.

Of course there is a difference between the relations between the perceiver and the two patches—in the different angles of separation from the fixation point. Does that ruin the case for my purposes? Note that there was a difference in the relations between the perceiver and the two patches in the experiments of Figure 6 and Figure 7, namely one was on the left of fixation and the other was on the right. Why would there be a difference in relevance between left/right and number of degrees of peripherality?

Recall that imprecise contents were introduced in the first place via the following reasoning. An attended 22% patch looks the same as an unattended 28% patch. But both percepts with that same contrast phenomenology are veridical. In order for percepts with that phenomenology of contrast to be grounded in the representation of contrast, the imprecision of the representational content has to be at a minimum 22%-28% (inclusive of 22% and 28%). Suppose the representationist had said "No no, those phenomenologies are different since one is leftish and one is rightish so there can be no legitimate demand for a representational content in virtue of which they have the same phenomenology. That argument would look silly and be silly because we have an appreciation of how contrast looks independently of which side it is on. We can easily abstract the percept of contrast from a total percept of contrast on the left or contrast on the right. The sense of 'abstract' here is a question of appreciation of the phenomenology of contrast independently of perceived location: I speak of abstraction because location is abstracted away from."

I suggest that the same reasoning applies to Figure 15 even though the difference in peripherality is causally implicated in producing the apparent contrast. The point is that we have an appreciation of that contrasty look independently of degree of peripherality and can appreciate that the two patches look the same in contrast when I am attending to the one on the left. The point is that a 16% patch can look the same in contrast as a 28% patch with the right distribution of attention and we need a representational account of what it is in virtue of which these apparent contrasts are the same. And with respect to that issue there is nothing illicit about comparing 4° with 9°.

The issue of abstraction I just mentioned comes up often in discussions of problem cases for representationism. Consider the phenomenal difference in seeing the round rim of a drinking glass and feeling it with one's hand. Both are percepts of one property, circularity, but the phenomenology is different. How can representationists cope with this case? Michael Tye (1995, p. 157; 2000, p. 93-95) has noted

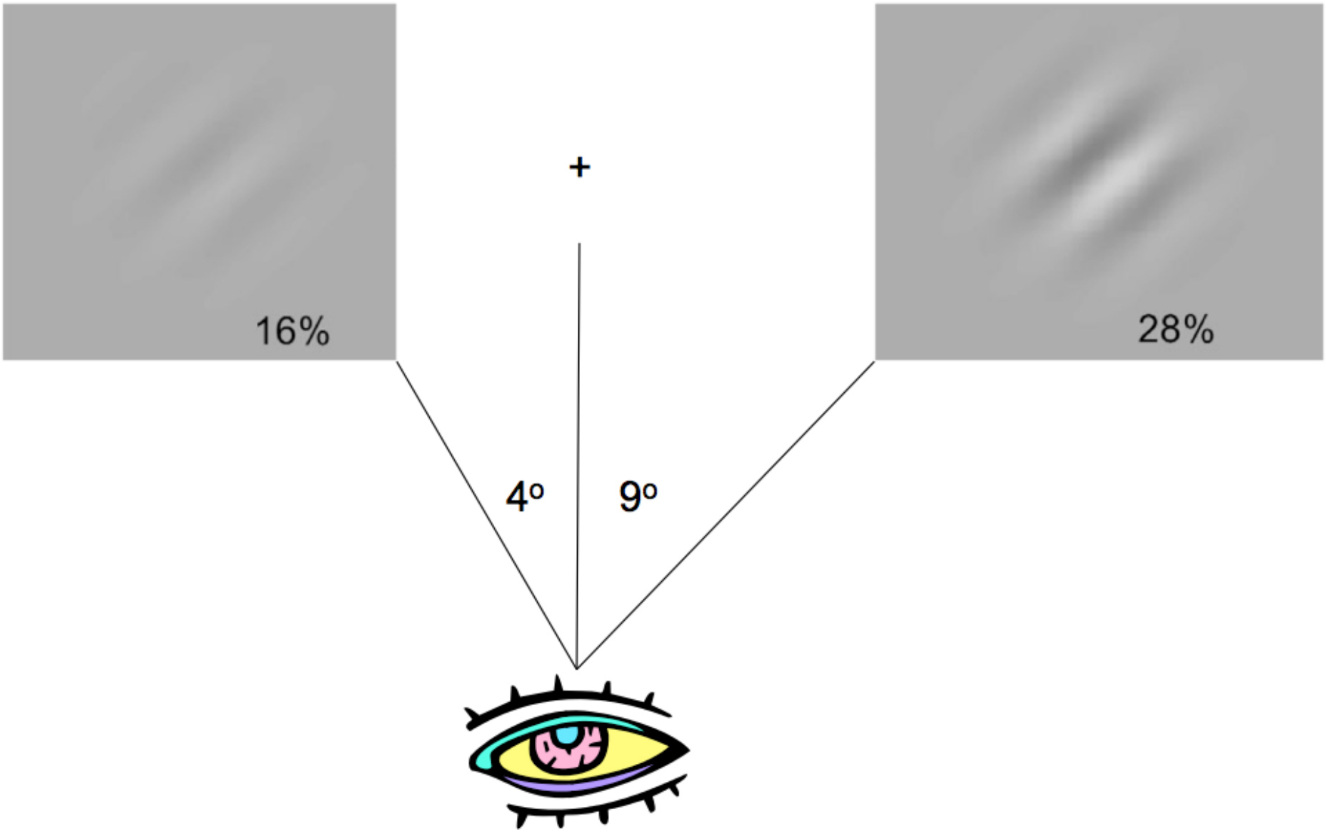

Figure 15: If you fixate at the "+" sign and attend to the left patch, it should look approximately equal in contrast to the right patch. My thanks to Jared Abrams for help in constructing this figure.

that the "total" percepts involve representation of different properties. These "collateral" properties might be shininess for the visual percept of the circularity and temperature for the tactile experience. The difference between the percepts can be blamed on the perceptions of these different properties. That is, what we are visually representing is circularity-&-shininess and what one is tactually representing is circularity-&-coldness. Can one abstract the visual impression of circularity from the total visual percept? Can one abstract the tactile impression of circularity from the total tactile percept? Tye says he cannot make sense of such abstraction. However, our ability to abstract shape from location on the right vs the left suggests the Principle of *Spatial* Abstraction: perceptual placing of a feature at a location can be abstracted from the perception of the location. I have a visual appreciation of the color of an object even as it moves, changing location. To the extent that this principle is accepted it licenses the use of Figure 15 in the premise that the contrast percepts are determinately different.

14 Conclusion

I can now summarize the overall argument. First, the short version. The 22% patch and the 28% patch look different when foveated and attended one after the other. However, fixating in between them and attending to the 22% patch, they look the same. How can this be explained representationally without supposing that the precision of attentive foveal vision is narrower than that of inattentive peripheral vision? As before, this is a burden of proof argument that does not explicitly utilize the idea of phenomenal precision.

And as before, here is the long version:

1. The attended 22% patch and the unattended 28% patch, being the same in contrast-phenomenology are the same in contrast-representational contents.

2. Both are veridical.
3. The contrast attributed by vision to the two patches has a minimum span of 22%-28%.
4. Attended and foveal percepts of 22% and 28% (seen sequentially) are determinately different in phenomenology.
5. Phenomenal precision principle: the phenomenal precision of the percepts of the patches seen attended and foveally is narrower than the phenomenal precision of at least one of the percepts seen in the periphery with only one attended. And it is plausible to suppose it is the unattended percept that has the wider precision.
6. So the phenomenal precision of the attended foveal percepts must be narrower than at least one of the peripheral percepts (probably the unattended one).
7. Representationism requires that a difference in phenomenal precision be grounded in a commensurate difference in representational precision.
8. So representationism requires that the precision of the foveal attended percepts be narrower than at least one of peripheral percepts. We have already seen that peripherality pre se probably does not decrease precision so if precision is decreased, it probably is due to withdrawal of attention. But empirical results suggest that withdrawal of attention does not decrease precision.
9. Conclusion: there is some reason to think that the phenomenology of perception is not grounded in its representational content.

Thus, for the perception of some properties, we have reason to believe that the representational content of perception neither grounds nor is grounded by the phenomenology of perception.

I argued that an attended .20° gap looks the same in respect of size as an unattended .23° gap. The comparative percept—the gaps looking the same—is illusory. But what about the percepts of each gap, considered separately? I argued that we would need a good reason to suppose that one but not the other is illusory and that the view that that both are illusory would undermine the notion of representational content altogether. I said that both are (or rather can be in normal circumstances) veridical. A similar point applies to the version of the experiments involving contrast in which an attended 22% patch looks the same in contrast as an unattended 28% patch. If the two patches look the same and if looking the same is a matter of sameness in representational content, and if the percepts are veridical, the size properties the patches are represented as having must be intervalic. And the interval—an index of precision—must be wide enough to encompass both patches. So the representational content has to have a precision range of 6%. And further considerations I mentioned suggest a range of 12%. The phenomenal precision principle says if percepts of 22% and 28% are phenomenally the same with one unattended in peripheral vision but determinately different when attended and foveal, then the attended and foveal percepts must have a narrower phenomenal precision than at least one of the peripheral percepts. The 22% and 28% patches do look determinately different if foveated and attended. So the attended and foveal percepts must have a narrower phenomenal precision than one of the peripheral percepts. The only way that this can happen on the representationist point of view is if one of the peripheral representational content is less precise than the foveal attended content. But experimental results that I cited suggest that may not be true. It may not be true of foveal vs peripheral vision independently of attention, and it may not be true for attention independently of foveal vs peripheral perception.

In the section on inhomogeneities of the visual field, I mentioned a route to the same conclusion based on introspection. And I will update that point to include attention. The more introspective route is this: it is natural to feel that the phenomenology of seeing the contrast between the lines and spaces on a piece of lined paper attentively and foveally differs in precision from seeing the same lines inattentively and peripherally. The foveal attentive percept seems more "crisp" than the inattentive peripheral percept. As we have seen, location is indeed represented more precisely but the same is not true for other properties such as hue or contrast. If this intuitive judgment is correct, there is introspective evidence for a discrepancy between the precision of phenomenology and the precision of representational content.

As I mentioned at the outset, the phenomenal precision principle needs more clarification and justification. It depends on notions of overlapping and of determinately different phenomenologies that are not as clear as one would like. My rationale is that if any advance in understanding of the phenomenology of perception is possible, it will have to start with underdeveloped ideas. I believe that there is enough in these ideas to give some credence to the conclusion. A second issue is whether the percepts that I say are determinately different in phenomenology really are.

The reader will have noticed that for the experimental results I have discussed it can often be difficult to figure out what aspects of the results concerned visual phenomenology and what aspects concern visual representation. As I mentioned earlier we have a real science of perception but very little science of the phenomenology of perception. If we are ever to turn what we know about perception into a scientific approach to the phenomenology of perception, we have no alternative but to start with some vague intuitive notions and proceed from there.

Although there are some loose ends, I think I have said enough to suggest a disconnect between the representational content of perception and what it is like to perceive.

Acknowledgements

I am grateful to Worth Boone, Tyler Burge, Marisa Carrasco, Jeremy Goodman, Eric Mandelbaum, John Morrison, Susanna Siegel, James Stazicker, Thomas Metzinger, Jennifer Windt and two anonymous reviewers for the Open MIND Project for comments on an earlier draft. And I am especially grateful to Jeremy Goodman and James Stazicker for discussion of these topics.

Glossary

Acuity	Also known as spatial resolution-- is the ability to resolve elements of stimuli. Common measures in the case of vision are the extent to which the subject can distinguish one dot from two dots, detect a gap between two figures, determine whether a rotating figure is rotating clockwise rather than counter-clockwise, ascertain whether two line segments are co-linear, distinguish a dotted from a solid line or detect which side of a Landolt Square a gap is on.
Attention	William James (1890, p. 404) famously said attention "…is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others." Except for the exclusion of unconscious attention, most scientists would accept something like that characterization today. Spatial attention is attention directed to portion of environmental space and is distinct from attention to a thing or a property.
Content	See representational content.
Contrast	Contrast in an environmental layout is often defined as the average difference in luminance between light and dark areas. (Luminance is the amount of light reflected.) More specifically, it is the luminance difference between the lightest and darkest areas divided by the sum of those luminances. There are alternative ways of defining the notion but the differences won't matter here.
Determinately different	For items to look determinately different in contrast, their contrast phenomenologies cannot be almost completely overlapping. I noted that this notion makes sense from a representationist perspective. I said that if one patch is represented as 10%-30% in contrast and another patch as 10.5%-30.5% the representationist would need a good reason to deny that the phenomenologies almost completely overlap. Given that representationism is committed to phenomenal precision and phenomenal overlap, it is legitimate to assume them in an argument against representationism.
Diaphanousness	G. E. Moore (1903) famously said "… the moment we try to fix our attention upon consciousness and to see what, distinctly, it is, it seems to vanish: it seems as if we had before us a mere emptiness. When we try to introspect the sensation of blue, all we can see is the blue; the other element is as if it were diaphanous …"
Direct realism	The view that the phenomenal character of perceptual experience is grounded in direct awareness of objects and properties in the world.
Endogenous attention	Endogenous attention is voluntary—what people often mean by "paying attention".
Exogenous attention	Exogenous spatial attention is attention that is attracted, automatically by a highly visible change. It is sometimes referred to as "transient" attention, whereas endogenous spatial attention is "sustained". Exogenous spatial attention peaks by 120 ms after the cue, whereas endogenous spatial attention requires at least 300 ms to peak and has no known upper temporal limit.
Fixation	To fixate a thing or area of space is to point your eyes at it.
Fovea	The fovea is the high density center of the retina. Foveal vision is the only vision that can be 20/20. If you hold your hand at arm's length, your foveal perception encompasses about double the width of your thumb.

Gabor patches	The fuzzy (actually sinusoidal) grids in Figure 1 and other figures.
Grounding	phenomenology is grounded in representational content just in case it is in virtue of the representational content of an experience that it has the phenomenology it has.
Identity formulation of representationism	What it is for an experience to have a certain phenomenal character is for it to have a certain representational content.
Landolt Square	See Figure 2.
Phenomenal precision principle	(one form) If two things look the same in peripheral vision but determinately different in foveal vision, then the phenomenal precision of foveal vision is narrower than that of peripheral vision.
Phenomenal precision	As with everything phenomenal, nothing like a definition is possible. The best you can do is use words to point to a phenomenon that the reader has to experience from the first person point of view. The experience of a color as red is less precise than the experience of a color as crimson. According to representationism, phenomenal precision is just the phenomenology of the precision of representational content. We experience a percept with representational content of 10%-20% as having more precision than we experience a percept with representational content 10%-30%. For a direct realist, phenomenal precision is just the precision of the properties we are directly aware of. We can be directly aware of properties with different precisions, for example, crimson, or alternatively red. Similarly we can be directly aware of a 10%-20% contrast property and also a 10%-30% contrast property and the difference constitutes a phenomenal precision difference.
Prothetic vs metathetic	Prothetic dimensions have a zero point and intrinsic directionality, whereas metathetic dimensions have neither.
Receptive field	In vision, the receptive field of a neuron is the area of space that a neuron responds to. In tactile perception the receptive field of a neuron is often gauged physiologically—the field of sensory receptors that feed to that neuron.
Representational content	Condition of veridicality. A simple percept consists of a representation of an environmental property and a singular element that picks out an individual item (Burge 2010). The representational content is satisfied when the referent of the singular element has the property represented by the property-representation.
Representational Precision	The precision of a representation is a matter of the intervalic range. For example, the precision of a representation of contrast of 10%-20% is narrower than a representation of 10%-30%. Precision in the sense used here is not a matter of indeterminacy of interval borders.
Spatial frequency	A measure of how closely spaced light and dark areas are. One could think of it with regard to the Gabor patches as a matter of stripe density.
Supervenience formulation of representationism	If phenomenology supervenes on representational content, there can be no difference in the phenomenology of perception without a difference in its representational content.
Veridicality	The veridicality of the most basic percept representations is a matter of the item referred to by the singular element having the property represented by the property representation.

References

Abrams, J., Nizam, A. & Carrasco, M. (2012). Isoeccentric locations are not equivalent: The extent of the vertical meridian asymmetry. *Vision Research*, *52* (1), 70-78. 10.1016/j.visres.2011.10.016

Afraz, A., Pashkam, M. & Cavanagh, P. (2010). Spatial heterogeneity in the perception of face and form attributes. *Current Biology*, *20* (23), 2112-2116. 10.1016/j.cub.2010.11.017

Alvarez, G. & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science*, *15* (2), 106-111. 10.1111/j.0963-7214.2004.01502006.x

Anton-Erxleben, K. & Carrasco, M. (2013). Attentional enhancement of spatial resolution: Linking behavioural and neurophysiological evidence. *Nature Reviews Neuroscience*, *14*, 188-200. 10.1038/nrn3443

Anton-Erxleben, K., Henrich, C. & Treue, S. (2007). Attention changes perceived size of moving visual patterns. *Journal of Vision*, *7* (11), 1-9. 10.1167/7.11.5

Anton-Erxleben, K., Abrams, J. & Carrasco, M. (2010). Evaluating comparative and equality judgments in contrast perception: Attention alters appearance. *Journal of Vision*, *10* (11), 1-22. 10.1167/10.11.6

——— (2011). Equality judgments cannot distinguish between attentioneffects on appearance and criterion: A reply to Schneider. *Journal of Vision*, *11* (13), 1-8. 10.1167/11.13.8

Anton-Erxleben, K., Herrmann, K. & Carrasco, M. (2013). Independent effects of adaptation and attention on perceived speed. *Psychological Science*, *24* (2), 150-159. 10.1177/0956797612449178

Aristotle (1955). On dreams. In W. D. Ross (Ed.) *Aristotle. Parva naturalia.* Oxford, UK: Clarendon Press.

Asplund, C., Fougnie, D., Zughini, S., Martin, J. W. & Marois, R. (2014). The attentional blink reveals the probabilistic nature of discrete conscious perception. *Psychological Science*, *25* (3), 824-831. 10.1177/0956797613513810

Baldassi, S. & Verghese, P. (2005). Attention to locations and features: Different top-down modulation of detector weights. *Journal of Vision*, *5* (6), 556-570. 10.1167/5.6.7

Barbot, A., Landy, M. & Carrasco, M. (2012). Differential effects of exogenous and endogenous attention on 2nd-order texture contrast sensitivity. *Journal of Vision*, *12* (8), 1-15. 10.1167/5.6.7

Barner, D., Wood, J., Hauser, M. & Carey, S. (2008). Evidence for a non-linguistic distinction between singular and plural sets in rhesus monkeys. *Cognition*, *107* (2), 603-622. 10.1016/j.cognition.2007.11.010

Bayne, T. (2014). Gist! *Paper presented at NYU October 29, 2014.*

Block, N. (2007). Overflow, access and attention. *Behavioral and Brain Sciences*, *30* (5-6), 530-542. 10.1017/S0140525X07003044

——— (2010). Attention and mental paint. *Philosophical Issues: A Supplement to Noûs*, *20* (1), 23-63. 10.1111/j.1533-6077.2010.00177.x

——— (2012). The grain of vision and the grain of attention. *Thought*, *1* (3), 170-184. 10.1002/tht3.28

——— (2013). Seeing and windows of integration. *Thought*, *2*, 29-39. 10.1002/tht.62

——— (2014a). The Canberra plan neglects ground. In T. Horgan, M. H. Sabates & D. Sosa (Eds.) *Qualia and mental causation in a physical world: Themes from the philosophy of Jaegwon Kim.* Cambridge, MA: Cambridge University Press.

——— (2014b). Seeing-as in the light of vision science. *Philosophy and Phenomenological Research*, *89* (3), 560-572. 10.1111/phpr.12135

Boone, W. (2013). Range content, attention and the precision of representation. *Presentation at the Society for Philosophy and Psychology, June 15, 2013, Brown University*

Botta, F., Lupiáñez, J. & Chica, A. (2014). When endogenous spatial attention improves conscious perception: Effects of alerting and bottom-up activation. *Consciousness and Cognition*, *23*, 63-73. 10.1016/j.concog.2013.12.003

Burge, T. (2005). Disjunctivism and perceptual psychology. *Philosophical Topics*, *33* (1), 1-78. 10.5840/philtopics20053311

——— (2010). *Origins of objectivity.* Oxford, UK: Oxford University Press.

Cameron, E., Tai, J. & Carrasco, M. (2002). Covert attention affects the psychometric function of contrast sensitivity. *Vision Research*, *42* (8), 949-967. 10.1016/S0042-6989(02)00039-1

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51* (13), 1484-1525. 10.1016/j.visres.2011.04.012

Carrasco, M., Talgar, C. & Cameron, E. (2001). Characterizing visual performance fields: Effects of transient covert attention, spatial frequency, eccentricity, task and set size. *Spatial Vision*, *15* (1), 61-75. 10.1163/15685680152692015

Carrasco, M., Williams, P. E. & Yeshurun, Y. (2002). Covert attention increases spatial resolution with or without masks: Support for signal enhancement. *Journal of Vision*, *2* (6), 467-479. 10.1167/2.6.4

Carrasco, M., Ling, S. & Read, S. (2004). Attention alters appearance. *Nature Neuroscience*, *7* (3), 308-313. 10.1038/nn1194

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 and Psychophysics*, *70* (7), 1151-1164. 10.3758/PP.70.7.1151

Chalmers, D. (2004). The representational character of experience. In B. Leiter (Ed.) *The future for philosophy* (pp. 153-181). Oxford, UK: Oxford University Press.

——— (2006). Perception and the fall from Eden. In T. S. Gendler & J. Hawthorne (Eds.) *Perceptual Experience* (pp. 49-125). Oxford, UK: Oxford University Press.

Chica, A., Lasaponara, S., Lupiáñez, J., Doricchi, F. & Bartolomeo, P. (2010). Exogenous attention can capture perceptual consciousness: ERP and behavioural evidence. *NeuroImage*, *51* (3), 1205-1212. 10.1016/j.neuroimage.2010.03.002

Chica, A., Lasaponara, S., Chanes, L., Valero-Cabre, A., Doricchi, F., Lupiáñez, J. & Bartolomeo, P. (2011). Spatial attention and conscious perception: The role of endogenous and exogenous orienting. *Attention, Perception & Psychophysics*, *73* (4), 1065-1081. 10.3758/s13414-010-0082-6

Chica, A., Paz-Alonso, P., Valero-Cabre, A. & Bartolomeo, P. (2013). Neural bases of the interactions between spatial attention and conscious perception. *Cerebral Cortex*, *23* (6), 1269-1279. 10.1093/cercor/bhs087

Chirimuuta, M. & Tolhurst, D. (2005a). Accuracy of identification of grating contrast by human observers: Bayesian models of V1 contrast processing show correspondence between discrimination and identification performance. *Vision Research*, *45* (23), 2960-2971. 10.1016/j.visres.2005.06.021

——— (2005b). Does a Bayesian model of V1 contrast coding offer a neurophysiological account of human contrast discrimination? *Vision Research*, *45* (23), 2943-2959. 10.1016/j.visres.2005.06.022

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36* (3), 233-253. 10.1017/S0140525X12000477

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24* (1), 87-185. 10.1017/S0140525X01373922

Craig, J. & Johnson, K. (2000). The two-point threshold: Not a measure of tactile spatial resolution. *Current Directions in Psychological Science*, *9* (1), 29-31. 10.1111/1467-8721.00054

Cutrone, E., Heeger, D. J. & Carrasco, M. (in press). Attention enhances contrast appearance via increased input baseline of neural responses. *Journal of Vision.*

Datta, R. & DeYoe, E. (2009). I know where you are secretly attending! The topography of human visual attention revealed with fMRI. *Vision Research*, *49* (10), 1037-1044. 10.1167/10.7.9

David, S., Haydon, B., Mazer, J. & Gallant, J. (2008). Attention to stimulus features shifts spectral tuning of V4 neurons during natural vision. *Neuron*, *59* (3), 509-521. 10.1016/j.neuron.2008.07.001

Dehaene, S. (2014). *Consciousness and the brain: Deciphering how the brain codes our thoughts.* New York, NY: Viking.

Deubel, H., Irwin, D. & Schneider, W. X. (1999). The subjective direction of gaze shifts long before the saccade. In W. Becker, H. Deubel & T. Mergner (Eds.) *Current oculomotor research: Physiological and psychological aspects* (pp. 65-70). New York, NY: Plenum.

Emmanouil, T. & Magen, H. (2014). Neural evidence for sequential selection of object features. *Trends in Cognitive Sciences*, *18* (8), 390-391. 10.1016/j.tics.2014.04.007

Feigenson, L., Carey, S. & Hauser, M. (2002). The representations underlying infants' choice of more: Object files vs. analog magnitudes. *Psychological Science*, *13* (2), 150-156.

Feynman, R. P. (1988). *QED: The strange theory of light and matter.* Princeton, NJ: Princeton University Press.

Fine, K. (2012). Guide to ground. In F. Correia & B. Schnieder (Eds.) *Metaphysical grounding: Understanding the structure of reality* (pp. 37-80). Cambridge, UK: Cambridge University Press.

Fuller, S. & Carrasco, M. (2006). Exogenous attention and color perception: Performance and appearance of saturation and hue. *Vision Research*, *46* (23), 4032-4047. 10.1016/j.visres.2006.07.014

Gobell, J. & Carrasco, M. (2005). Attention alters the appearance of spatial frequency and gap effect. *Psychological Science*, *16* (8), 644-651. 10.1111/j.1467-9280.2005.01588.x

Goodale, M. A. & Murphy, K. (1997). Action and perception in the visual periphery. In P. Their & H.-O. Karnath (Eds.) *Parietal lobe contributions to orientation in 3-D space* (pp. 447-461). New York, NY: Springer.

Goodman, J. (2013). Inexact knowledge without improbable knowing. *Inquiry*, *56* (1), 30-53. 10.1080/0020174X.2013.775013

Gross, H. J., Pahl, M., Si, A., Zhu, H., Tautz, J. & Zhang, S. (2009). Number-based visual generalization in the honeybee. *PloS ONE*, *4* (1), 1-9. 10.1371/journal.pone.0004263

Hauser, M., Carey, S. & Hauser, L. (2000). Spontaneous number representation in semi-free ranging rhesus monkeys. *Proceedings of the Royal Society of London: Biological Sciences*, *267* (1445), 829-833. 10.1098/rspb.2000.1078

Herrmann, K., Montaser-Kouhsari, L., Carrasco, M. & Heeger, D. J. (2010). When size matters: Attention affects performance by contrast or response gain. *Nature Neuroscience*, *13* (12), 1554-1559. 10.1038/nn.2669

Hess, R. & Field, D. (1993). Is the increased spatial uncertainty in the normal periphery due to spatial undersampling or uncalibrated disarray? *Vision Research*, *33* (18), 2663-2670. 10.1016/0042-6989(93)90226-M

Hill, C. (2009). *Consciousness.* Cambridge, UK: Cambridge University Press.

Hohwy, J. (2013). *The predictive mind.* Oxford, UK: Oxford University Press.

Hohwy, J., Roepstorff, A. & Friston, K. (2008). Predictive coding explains binocular rivalry: An epistemological review. *Cognition*, *108* (3), 687-701. 10.1016/j.cognition.2008.05.010

Horgan, T. & Tienson, J. (2002). *Philosophy of mind: Classical and contemporary readings.* Oxford, UK: Oxford University Press.

Itti, L. & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, *40*, 1489-1506. 10.1.1.40.4697

James, W. (1890). *Principles of psychology.* New York, NY: Henry Holt.

Kentridge, R., Nijboer, T. & Heywood, C. (2008). Attended but unseen: Visual attention is not sufficient for visual awareness. *Neuropsychologia*, *46* (3), 864-869. 10.1016/j.neuropsychologia.2007.11.036

Kerzel, D., Zarian, L., Gauch, A. & Buetti, S. (2010). Large effects of peripheral cues on appearance correlate with low precision. *Journal of Vision*, *10* (11), 1-14. 10.1167/10.11.26

Kriegel, U. (2011). *The sources of intentionality.* Oxford, UK: Oxford University Press.

——— (2013). The phenomenal intentionality research program. In U. Kriegel (Ed.) *Phenomenal intentionality* (pp. 1-26). Oxford, UK: Oxford University Press.

Lettvin, J. Y. (1976). On seeing sidelong. *The Sciences*, *16* (4), 10-20. 10.1002/j.2326-1951.1976.tb01231.x

Levi, D. & Klein, D. (1996). Limitations on position coding imposed by undersampling and univariance. *Vision Research*, *36* (14), 2111-2120. 10.1016/0042-6989(95)00264-2

Levitin, D. (2005). Absolute pitch: Perception, coding, and controversies. *Trends in Cognitive Sciences*, *9* (1), 26-33. 10.1016/j.tics.2004.11.007

——— (2008). Absolute pitch: Both a curse and a blessing. In M. Klockars & M. Peltomaa (Eds.) *Music meets medicine. Proceedings of the Signe and Ane Gyllenberg Foundation* (pp. 124-132). Helsinki, Finland: Signe and Ane Gyllenberg Foundation.

Ling, S. & Carrasco, M. (2006). When sustained attention impairs perception. *Nature Neuroscience*, *9* (10), 1243-1245. 10.1038/nn1761

Ling, S., Liu, T. & Carrasco, M. (2009). How spatial and feature-based attention affect the gain and tuning of population responses. *Vision Research*, *49* (10), 1194-1204. 10.1016/j.visres.2008.05.025

Ling, S., Jehee, J. & Pestilli, F. (2014). A review of the mechanisms by which attentional feedback shapes visual selectivity. *Brain Structure and Function.* 10.1007/s00429-014-0818-5

Ma, W. J., Husain, M. & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, *17* (3), 347-356. 10.1038/nn.3655

Metzinger, T. (2003). *Being no one.* Cambridge, MA: MIT Press.

Montagna, B., Pestilli, F. & Carrasco, M. (2009). Attention trades off spatial acuity. *Vision Research*, *49* (7), 735-745.

Montaser-Kouhsari, L. & Carrasco, M. (2009). Perceptual asymmetries are preserved in short-term memory tasks. *Attention, Perception & Psychophysics*, *71* (8), 1782-1792. 10.3758/APP.71.8.1782

Moore, G. E. (1903). The refutation of idealism. *Mind*, *12* (48), 433-453.

Morrison, J. (2013). Anti-atomism about color representation. *Noûs.* 10.1111/nous.12018

Nanay, B. (2010). Attention and perceptual content. *Analysis*, *70* (2), 263-269. 10.1093/analys/anp165

Norman, L. J., Heywood, C. & Kentridge, R. (2013). Object-based attention without awareness. *Psychological*

Science, *24* (6), 836-843. 10.1177/0956797612461449

Pautz, A. (2010). Why explain visual experience in terms of content? In B. Nanay (Ed.) *Perceiving the world: New essays on perception.* New York, NY: Oxford Univeristy Press. 10.1111/j.1520-8583.2007.00134.x

Perkins, R. & Bayne, T. (2013). Representationalism and the problem of vagueness. *Philosophical Studies*, *162* (1), 71-86. 10.1007/s11098-012-9990-8

Phillips, I. B. (2011). Perception and iconic memory: What Sperling doesn't show. *Mind & Language*, *26* (4), 381-411. 10.1111/j.1468-0017.2011.01422.x

Prinz, J. J. (2012). *The conscious brain.* New York, NY: Oxford University Press.

Raffman, D. (1995). On the persistence of phenomenology. In T. Metzinger (Ed.) *Conscious experience* (pp. 293-308). Paderborn, Germany: Ferdinand Schöningh.

Ramachandran, V. S. & Hirstein, W. (1998). The perception of phantom limbs. *Brain*, *121* (Pt 9), 1603-1630. 10.1093/brain/121.9.1603

Reynolds, J. H. & Chelazzi, L. (2004). Attentional modulation of visual processing. *Annual Review of Neuroscience*, *27*, 611-647. 10.1146/annurev.neuro.26.041002.131039

Reynolds, J. H., Pasternak, T. & Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron*, *26* (3), 703-714. 10.1016/S0896-6273(00)81206-4

Reynolds, J. H. & Heeger, D. J. (2009). The normalization model of attention. *Neuron*, *61* (2), 168-185. 10.1016/j.neuron.2009.01.002

Rolfs, M. & Carrasco, M. (2012). Rapid simultaneous enhancement of visual sensitivity and perceived contrast during saccade preparation. *The Journal of Neuroscience*, *32* (40), 13744-13752. 10.1523/JNEUROSCI.2676-12.2012

Rolfs, M., Lawrence, B. & Carrasco, M. (2013). Reach preparation enhances visual performance and appearance. *Philosophical Transactions of the Royal Society B*, *368* (1628), 1-8. 10.1098/rstb.2013.0057

Schneider, K. A. (2006). Does attention alter appearance? *Perception and Psychophysics*, *68* (5), 800-814. 10.3758/BF03193703

——— (2011). Attention alters decision criteria but not appearance: A reanalysis of Anton-Erxleben, Abrams and Carrasco (2010). *Journal of Vision*, *11* (13), 1-8. 10.1167/11.13.7

Schneider, K. A. & Komlos, M. (2008). Attention biases decisions but does not alter appearance. *Journal of Vision*, *8* (15), 1-10. 10.1167/8.15.3

Schoenfeld, M. A., Hopf, J.-M., Merkel, C., Heinze, H.-J. & Hillyard, S. A. (2014). Object-based attention involves the sequential activation of feature-specific cortical modules. *Nature Neuroscience*, *17* (4), 619-624. 10.1038/nn.3656

Scholl, B. J., Noles, N. S., Pasheva, V. & Sussman, R. (2003). Talking on a cellular telephone dramatically increases 'sustained inattentional blindness'. *Journal of Vision*, *3* (9), 156-156. 10.1167/3.9.156

Sergent, C., Baillet, S. & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience*, *8* (10), 1391-1400. 10.1038/nn1549

Shoemaker, S. (2007). *Physical realization.* Oxford, UK: Oxford University Press.

Siegel, S. (2010). *The contents of visual experience.* Oxford, UK: Oxford University Press.

——— (2013). Are there edenic grounds of perceptual intentionality? *Analysis Reviews*, *73* (2), 329-344. 10.1093/analys/ans150

Sorenson, R. (2013). Vagueness. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy.* http://plato.stanford.edu/archives/win2013/entries/vagueness.

Speaks, J. (2010). Attention and intentionalism. *Philosophical Quarterly*, *60* (239), 325-342. 10.1111/j.1467-9213.2009.617.x

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, *74* (11), 1-29. 10.1037/h0093759

Stazicker, J. (2011a). Attention, visual consciousness and indeterminacy. *Mind & Language*, *26* (2), 156-184. 10.1111/j.1468-0017.2011.01414.x

——— (2011b). Attention, visual knowledge & psychophysics: Discriminating the determinable. *NYU Philosophy of Mind Discussion Group.*

——— (2013). Attending, knowing and detecting signals. In J. Taylor (Ed.) *Consciousness and attention workshop, as part of 'Philosophy and psychology: Integrating research across domains'.* Durham, UK: Durham University Press.

Stevens, S. S. & Galanter, E. H. (1957). Ratio scales and category scales for a dozen perceptual continua. *Journal of Experimental Psychology*, *54* (6), 377-411. 10.1037/h0043680

Stoljar, D. (2004). The argument from diaphanousness. In M. Escurdia, R. J. Stainton & C. D. Viger (Eds.) *Language, mind and world: Special issue of the Canadian Journal of Philosophy (Vol. 30)* (pp. 341-390). Edmonton, Canada: University of Alberta Press.

Strasburger, H., Rentschler, I. & Jütner, M. (2011). Peripheral vision and pattern recognition: A review. *Journal of Vision*, *11* (5), 1-82. 10.1167/11.5.13

Suchow, J., Fougnie, D., Brady, T. F. & Alvarez, G. (2014). Terms of the debate on the format and structure of visual memory. *Attention, Perception & Psychophysics*, *76* (7), 2071-2079. 10.3758/s13414-014-0690-7

Tong, J., Mao, O. & Goldreich, D. (2013). Two-point orientation discrimination versus the traditional two-point test for tactile spatial acuity assessment. *Frontiers in Human Neuroscience*, *7* (579), 1-11. 10.3389/fnhum.2013.00579

Treue, S. (2004). Perceptual enhancement of contrast by attention. *Trends in Cognitive Sciences*, *8* (10), 435-437. 10.1016/j.tics.2004.08.001

Turatto, M., Vescovi, M. & Valsecchi, M. (2007). Attention makes moving objects be perceived to move faster. *Vision Research*, *47* (2), 166-178. 10.1016/j.visres.2006.10.002

Tye, M. (1995). *Ten problems of consciousness.* Cambridge, MA: MIT Press.

——— (2000). *Consciousness, color, and content.* Cambridge, MA: MIT Press.

——— (2009). *Consciousness revisited.* Cambridge, MA: MIT Press.

Valsecchi, M., Vescovi, M. & Turatto, M. (2010). Are the effects of attention on speed judgments genuinely perceptual? *Attention, Perception & Psychophysics*, *72* (3), 637-650. 10.3758/APP.72.3.637

van den Berg, R., Shin, H., Chou, W., George, R. & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences, USA*, *109* (22), 8780-8785. 10.1073/pnas.1117465109

Watzl, S. (forthcoming). Can intentionalism explain how attention affects appearances? In A. Pautz & D. Stoljar (Eds.) *Themes from Block*. Cambridge, MA: MIT Press.

Wilkinson, S. (2014). Book review: The predictive mind. *Analysis*

Wu, W. (2014). *Attention.* New York, NY: Routledge.

Yeshurun, Y. & Carrasco, M. (1998). Attention improves or impairs visual performance by enhancing spatial resolution. *Nature*, *396* (6706), 72-75. 10.1038/23936

——— (1999). Spatial attenttion improves performance in spatial resolution tasks. *Vision Research*, *39* (2), 293-306. 10.1016/S0042-6989(98)00114-X

Zhang, S., Xu, M., Kamigaki, T., Do, J., Chang, W.-C., Jenvay, S., Miyamichi, K., Luo, L. & Dan, Y. (2014). Long-range and local circuits for top-down modulation of visual cortex processing. *Science*, *345* (6197), 660-665. 10.1126/science.1254126