Attention and Representational Precision

Azenet L. Lopez

Abstract: Visual experiences often feel crisper, sharper or more vivid when one pays attention to the seen object. According to some representationalist theories of perception, these felt effects occur because attentive experiences represent more *determinate* or *precise* properties than their inattentive counterparts: a color experience represents *vermillion* rather than *red* if the color is perceived with attention rather than without it. Recently, this idea has been expressed in terms of ranges of feature values represented, so that attentive experiences shall represent narrower ranges than inattentive experiences. However, the latter claim seems to run counter to evidence, recently discussed by Block (2015), that neither spatial attention nor attention as deployed in the attentional blink paradigm alters the ranges of properties represented in visual perception. In this chapter, I clarify what these findings in fact tell us about the representational contents of visual perception and argue that they do not refute the representational account of the attentional effects.

1. Introduction

The way we distribute our attention undoubtedly affects how we perceive the world around us. As a process of selection and prioritization of information (Posner 1980; Desimone & Duncan 1995; Carrasco 2011, 2018), attention often delimits and shapes the contents of our experience. For perceptual experience, this involves defining which properties we perceive, by modulating the information impinging our sensory channels. Such effects of attention on information processing can also be traced up to the way this information is consciously experienced. If you focus your attention on the red mug on your desk while looking at it directly, you enjoy a crisp and sharp experience of its shape and color. You see a small dent in its rand and its specific shade of red. If you distribute your attention differently, for instance by switching your gaze and your attention to the computer screen while the mug remains unattended in your visual periphery, your experience undergoes some changes. The crispness and sharpness of the perceived properties appears to diminish. You may now fail to see the dent in the mug's rand, and you may no longer be able to distinguish whether it is vermilion or maroon. The effects of attention on visual perception have been quantified in several studies showing that attention alone can make objects appear brighter, larger or more colorful (Carrasco 2018; Carrasco & Barbot 2019).

¹ Though the selection view of attention is still widely accepted, recent analyses instead emphasize signal amplification (Fazekas & Nanay 2021) or prioritization (Watzl 2017, Jennings 2020) as the distinctive features of attention.

There is an ongoing debate about the possibility to exhaustively explain these attentional effects in terms of the representational contents of experience, as some representationalist theories of perception entail.² Representational contents, on a widespread characterization, are the properties an experience attributes to an object, such that if the object has these properties, the experience is veridical. Your visual experience attributes to the mug the property *red*; this is a veridical experience if the mug is in fact red. Hence, *red* is part of the representational content of your experience.³ Representationalists contend that attentive experiences have more determinate contents than inattentive experiences, by attributing more determinate properties (Nanay 2010; Stazicker 2011; Brogaard 2015). Accordingly, attentive experiences attribute properties with more specific and delimited conditions of instantiation and thus have more constraining veridicality conditions. For instance, *vermillion* is more determinate than *red*. An experience attributing *vermillion* is veridical in a much more limited range of cases than an experience attributing *red*: there are only a few shades of vermillion, but there are many more shades of red.

Such representationalist accounts of the attentional effects run counter to empirical evidence suggesting that attention does not affect the ranges of visually represented properties. One piece of evidence is that spatial attention indiscriminately boosts the signal of any stimulus property present at the attended location and does not really *suppress* any feature values at that location. This is problematic for representationalism, since most of the demonstrated effects of attention on visual appearances focus on spatial attention, and since narrowing down the ranges of represented properties plausibly requires suppression of some property values. A second piece of evidence is that visual representations processed during the attentional blink, an effect that arguably deprives these representations of attention, seem to represent the same ranges of values than representations processed with full attention. These findings have been recently cited by Block (2015), who defends two claims:

- (1) Spatial attention increases spatial resolution but not representational precision.
- (2) Some visual representations are equally precise with and without attention.

My aim in this chapter is to show that the evidence does not really support claims 1 and 2. Elsewhere (Lopez & Simsova 2023), I have argued that current evidence on the effects of spatial attention on visual appearances neither supports nor refutes the claim that attention narrows down represented ranges of properties for the specific case of color. Here, I will argue that the evidence cited by Block also fails to refute the more general representationalist claim, namely:

(3) Visual perception represents more precise properties with (than without) attention.

The chapter is structured as follows. In section 2, I argue that increased spatial resolution, a well demonstrated effect of spatial attention, may indeed contribute to increasing the precision of visual experience, by narrowing down ranges of represented properties like edges and shapes. In section 3, I argue that the absence of feature suppression in the mechanisms of spatial attention is no obstacle for it to increase representational precision in the mentioned sense. Finally, in section 4 I argue that the attentional blink findings have no direct implications about what properties are represented in a visual experience and are in fact compatible with the representationalist claim. I conclude in section 5 by

² Specifically, by *pure representationalism* about perception (in the sense of Chalmers 2004): the view that all perceptual properties are (grounded in, supervenient on, or fully explained by) properties of representational contents.

³ I intend to remain neutral regarding the nature of content. The claims made here about range content should apply to, e.g., propositional or imagistic treatments of content.

briefly situating the present discussion within Block's more overarching argument against representationalism.

2. Spatial resolution, precision and range content

Representationalists often invoke the well-documented effects of attention in increasing spatial resolution in support of the claim that attentive experiences represent more determinate properties (Nanay 2010; Stazicker 2011). Behavioral evidence shows that directing attention to a spatial location improves ability to detect gaps (Yeshurun & Carrasco 1999; Montagna et al. 2009; Gobell & Carrasco 2005) or segregate fine-grained textures at that location (Yeshurun & Carrasco 1998; Yeshurun et al. 2008). Relatedly, neuroscience studies show that spatial attention alters neuronal receptive fields by shrinking them around the location of the attended stimulus (Reynolds et al. 2000; Martinez-Trujillo & Treue 2002, 2004; Moran & Desimone 1985) or by shifting them to attended locations (Anton-Erxleben & Carrasco 2013; Anton-Erxleben et al. 2007).

Block accepts that increasing spatial resolution is a way of making experiences *more determinate*. *Precision*, however, underwrites a different kind of determinacy (see Figure 1). As noted above, precision is a measure of the *range of values* represented within a property dimension (hue, length, etc.). A representation is more precise if it attributes a narrower range of values of a given property. For example, representing someone's height as *between 175 and 180cm* attributes to this person a range of height values that is narrower than the range attributed by *between 170 and 200cm*; hence, the former representation is more precise than the latter. The notion of *intervallic* or *range content* (Block 2015; Boone 2020) helps expressing this idea: *between 175 and 180cm* is a narrower range content than *between 170 and 200cm*. Spatial resolution, in turn, is not a measure of ranges of values but of amount of *spatial detail*. A representation has higher spatial resolution if it represents more spatial detail from a location in space, for instance, "two dots instead of one" (Block 2015: 35) or a dotted line instead of a solid line (Block 2015: 15).




Figure 1. Precision as a species of determinacy.

These distinctions show how attention could make representations more determinate in one sense, without making them more determinate in another.⁴ Greater spatial detail can come apart from narrower range contents. To make the point more intuitive, consider representations R1 and R2 in Figure 2. R2 has higher spatial resolution than R1 because it represents the dots and the gaps between the dots, while R1 misses this detail. However, R1 and R2 do not obviously differ in range contents. The contents *solid line* and *dotted line* are not intervallic, so there is no question as to which attributes a narrower range of values. Furthermore, veridicality conditions seem equally constraining for both representations. R1 is veridical if and only if the stimulus is a solid line, and R2 is veridical if and only if the stimulus is a dotted line.




Figure 2. Representations with low and high spatial resolution

In these ways, there is conceptual room for Block's first claim, which we can now cash out as follows:

(1.1) Spatial attention increases represented spatial detail but does not narrow down represented ranges of values.

Now, one can dispute this claim by noting that there are some ways how increasing the amount of represented spatial detail *can* plausibly narrow down ranges of represented properties, especially, properties involving spatial locations. Consider again R1 and R2 in Figure 2. Plausibly, these also represent luminance differences between two points. Say these two points are represented by coordinates X_1 , Y_1 and X_{10} , Y_{10} . The content of R1 then includes *luminance difference from* X_1 , Y_1 to X_{10} , Y_{10} . In turn, the content of R2 includes *luminance differences at* X_1 , Y_1 ; X_3 , Y_3 ; ... X_{10} , Y_{10} , where each pair of coordinates signals the location of a dot. Understood in this way, the contents of R2 do afford more constraining veridicality conditions than the contents of R1. R2 is veridical only if there are luminance differences at specific individual points. Contrastingly, R1 only requires that there are luminance differences between X_1 , Y_1 and X_{10} , Y_{10} . This condition is met in a broader range of cases, including but not limited to the case that makes R2 veridical (for instance, R1 could still be veridical if one intermediate dot was missing or slightly displaced). In this respect, R2 attributes a narrower range of properties than R1.

⁴ About the effects of attention on response variability, see Arazi (2019).

Note that construing the content of R1 as attributing a *continuous* luminance difference from X_1 , Y_1 to X_{10} , Y_{10} will not help Block's claim. If this is R1's content, then R1 is not veridical. But this implies that most of ordinary perception is nonveridical, at least with respect to properties depending on spatial resolution, since most of ordinary perception occurs in the visual periphery and consequently has low resolution (Cavanagh et al 1999; Pelli 2008; see also Strasburger et al. 2011). This is a result that Block explicitly wants to avoid, as he contends that (i) perception must be mostly veridical, and (ii) there is no principled way of deciding whether attention increases or decreases veridicality. Though the present discussion is largely independent from issues about veridicality, I shall briefly come back to this point in the final section.

Shape representations can be handled in a similar way. Consider R3 and R4 in Figure 3. R4 has higher spatial resolution, in virtue of which it represents a more specific shape than R3. Call these shapes S₃ and S₄. As before, one could well argue that S₃ and S₄ are not intervallic contents, as each is a maximally specific shape (though see Boone 2020 for argument that perception rarely represents maximally specific properties). However, R3 and R4 also represent the boundaries of a region in bidimensional space. In the lower resolution R3, these boundaries are less specific and thus are more permissive about the occupied coordinates that would make R3 veridical. Though Block is careful to distinguish fuzziness of borders as a matter of *vagueness* rather than imprecision (Block 2015: 4, n.2), it is unclear whether a representation with fuzzier *spatial* boundaries could represent narrower ranges of spatial properties than another representation with sharper spatial boundaries.⁵




Figure 3. Shape representations with low and high spatial resolution

⁵ See Vance (this volume), for a lengthier discussion of vagueness in perceptual experience.

To be sure, I do not deny that spatial resolution and representational precision are two different *concepts*, which might be usefully kept separate. My point is rather that these two properties might not really come apart in our visual representations, in the way required to refute the representationalist claim. One final, empirical argument to think that they may come apart draws on (i) the demonstrated variations of spatial resolution across the visual field, which is lower at the visual periphery and higher near the fovea (Cavanagh et al 1999; see also Strasburger et al. 2011) and (ii) evidence that contrast discrimination does not vary across the visual field, as subjects can judge whether a patch has higher or lower contrast than neighboring patches with equal accuracy at the fovea and at the periphery (Hess & Field 1993). Indeed, Block cites the latter findings as evidence that foveal and peripheral vision do not differ in representational precision, thus effectively associating representational precision with contrast discriminability. However, pinning down representational precision in this way defeats the purpose of distinguishing it from spatial resolution, as spatial attention is also well-known to improve contrast discriminations (Carrasco et al. 2000; Huang & Dobkins 2005; see also Brogaard 2015).

3. Spatial attention and featural precision

We come now to the tentative evidence that representational precision or breadth of range contents does not differ between attended and unattended representations. The first relevant finding concerns spatial attention, that is, attention directed at a spatial location (rather than an object or a feature). A well-known mechanism of spatial attention is *response gain*, the boosting of neuronal responses to stimuli present at the attended location (McAdams & Maunsell 1999; Martinez-Trujillo & Treue 2002, 2004; Reynolds & Heeger 2009). Notably, though spatial attention suppresses information at unattended locations (Datta and DeYoe 2009), it does not suppress any *feature values* at the attended location.

In this respect, spatial attention contrasts with feature attention, which boosts neural responses to the selected feature values while suppressing responses to different ones, depending on how dissimilar they are to the selected one, and regardless of spatial location (Theeuwes 2013). Consistently with these different neural mechanisms, a behavioral study by Ling et al. (2009) shows that only feature attention, but not spatial attention improves discrimination of the predominant motion direction in a dot array when noise is high, that is, when a greater proportion of the dots is moving randomly. Spatial attention indiscriminately boosts responses to all motion directions, so that if more dots are moving randomly this boosted noise swamps the signal of dots moving together. Conversely, feature attention does improve discrimination because it only boosts responses to specific motion directions, while suppressing irrelevant values.

On these bases, Block argues that though *feature* attention narrows down the ranges of features represented, *spatial* attention does not. Thus, the claim is:

(1.2) Spatial attention does not narrow down the ranges of represented features.

This claim is problematic for representationalists, since most psychophysical paradigms assessing the effects of attention on visual perception focus on manipulations of spatial attention (see reviews in

⁶ That said, it is also possible to conceptualize spatial resolution as a species of precision, concerned with ranges of properties underwritten by amount of spatial detail. In such view, precision could remain a species of determinacy, but it could become interchangeable with determinacy (as assumed in some representationalist views, e.g., Brogaard 2015). In any case, here we do not need to further specify the relation between precision and determinacy, beyond the points made in the main text. For arguments that precision and determinacy are two separate properties, see Boone (2020) and Lee (2021).

Carrasco 2018; Carrasco & Barbot 2019). These paradigms show that cueing subjects' attention to a stimulus' location makes the stimulus appear larger, darker (or brighter, depending on presentation conditions⁷), more saturated, etc. Indeed, Block discusses at length the prominent effects of spatial attention on apparent contrast (Block 2010; 2015). Spatial attention thus induces phenomenal changes, which according to representationalism must be explained by changes in represented properties.

In response to Block, Boone (2020: 13) recently noted that spatial attention plausibly does increase the precision of represented features under certain conditions, specifically, when external noise is low. Indeed, Ling et al (2009) also found that spatial attention improves motion discrimination when a smaller proportion of dots are moving randomly. This suggests that spatial attention can also increase *signal to noise ratio*, though in a different way and under different conditions than feature attention. After all, one can increase this ratio either by reducing noise *or by strengthening the signal*. In this way, both spatial and feature attention increase representational precision, though in different (and, may I add, complementary) ways.

Boone makes a plausible point. Indeed, increases of signal to noise ratio plausibly underpin representations of more precise properties (compare how reducing static noise in a radio-transmitted message helps conveying the specific words spoken). And indeed, psychophysical studies on the effects of spatial attention on appearances typically present their stimuli under low noise conditions. Notably, the extensively discussed studies on apparent contrast (Carrasco et al. 2004) simply present one Gabor patch on each side of fixation; thus, representationalists could safely rely on these studies to support their claims about attentional differences in content.

However, Boone's argument must be defended from a possible objection. Block seems to take range contents to be determined by the width of a tuning curve, which becomes narrower with feature attention but not with spatial attention. Figure 4 roughly illustrates this idea. Represented values are the values "covered" by this wider curve portion. One could then say that range contents do not change if the curve width does not change —indeed, we will see that Block makes this point in the argument discussed in next section. Moreover, the curve width will not change unless there is some active suppression of some of the values at its basis. Since spatial attention does not suppress any values, precision in the sense at stake does not change with spatial attention. Changes in the signal to noise ratio shall then underpin a different property — a different kind of precision even, but not range content.

⁷ See discussion in Lopez & Simsova (2023).

⁸ External noise is noise in the environment, as manipulated in random dot cinematograms like the one from Ling et al. (2009). *Internal* noise, which is constituted by variations intrinsic to the perceptual system. Notably, Block acknowledges that spatial attention can swamp internal noise, but he does not discuss how suppression of internal noise might be relevant for representational precision.

⁹ Feature attention does not have a significant effect when external noise is low. The offered explanation is that suppression of irrelevant values is not of great benefit when there are already few of these.

¹⁰ This figure is admittedly very schematic. Figures 13 and 14 in Block (2015) are a more exact illustration.




Figure 4. Spatial attention, feature attention and range contents. 11

In response, one may note that the objection trades on some unclarities and ambiguities in the notion of range content. Consider Block's paradigmatic examples, between 175 and 180cm and between 170 and 200cm. The former is more precise because it attributes a smaller range of height values than the latter. Following the rationale in the present argument, the former content should be associated with a narrower tuning curve than the latter. But building tuning curves for these contents requires introducing an additional dimension, which shall determine the height of the curve at different value points (171cm, 172cm...). Since our current example is a priori and hence does not involve any actual measurements, one could think of the additional dimension as expected signal strength for each value, or probability that the value is instantiated. Block is not keen to a probabilistic treatment of range contents (Block 2015: 28), so this might already be a stumbling block for his argument. In other words, if one wants to say that between 175 and 180cm is a range content but not necessarily a probabilistic content, then, absent a plausible specification of what determines the height of the relevant tuning curve, one might need to avoid pinning range contents breadth onto tuning curve width. But then, the reason to think that spatial attention leaves range content unaffected is lost.

Now say, for the sake of the argument, that one accepts the probabilistic treatment to retain the link between range contents and tuning curve width. The most straightforward way to plot the curve for, say, between 175 and 180cm is as a normal distribution, assigning greater probability values to intermediate centimeter values (e.g., 177cm, 178cm) and smaller probability values to extreme centimeter values (e.g., 175cm, 180cm). Though this would yield the desired results (between 175 and 180cm has fewer intermediate values than between 170 and 200cm), it does not seem to capture the properties that the relevant representations would attribute. If my visual experience represents my

¹¹ I thank Eliska Simsova for her help in improving this figure.

¹² I changed the metric unit from inches to centimeters, but nothing in the relevant arguments depends on metric unit choice.

colleague's height as between 175 and 180cm, then for all my experience tells me it is equally likely that my colleague is 175cm tall, 176cm tall, etc. Now, the relevant curve need not be a normal distribution. But then, how should one distribute the probability values? Here is a tentative answer: assign equal probabilities to each point in the range, so that their sum equals 1. This will also make between 175 and 180cm a narrower range content than between 170 and 200cm, but then it is unclear how the difference is still a matter of the width of a curve. And if the difference is not about the width of a curve, then once again the reason to think that spatial attention leaves range content unaffected is lost.

Thus, to press the point that spatial attention does not affect range content because it does not suppress feature values, one must be clearer about how the values at the basis of a tuning curve relate to the properties represented in perception. Here is one last reason for thinking that the two might not line up. Say there is a bunch of red, orange, and green chips spilled on your desk (see Figure 4). The chips are in your visual field, but you are not paying attention to them in any particular way. Call your visual representation of this portion of the desk R_U (for "unattended"). According to Block, R_U represents a range of three hue values: red, orange, and green. If you now direct your attention to the location of the chips, hue (and other) information from surrounding portions of your visual field will be suppressed, while hue information at this location will be boosted, regardless of hue value. Call this new representation R_S. According to Block, R_S still represents the same range of hue values as R_U, namely, red, orange, and green. But if you instead direct your attention to the red things on the desk, this range will shrink, because red hues will be boosted everywhere, while green hues will be suppressed everywhere. Call the resulting representation R_F. Because of the green hue suppression, R_F is supposed to represent a narrower range of features than both R_S and R_U. Now, say that because of this R_F has narrower range content than R_S. At the outset, we linked range content breadth with permissiveness of veridicality conditions, so that the broader the former, the more permissive the latter. But R_F and R_S illustrate the opposite pattern: R_S is supposed to have broader range content than R_F, yet R_F seems to have more permissive veridicality conditions than R_S. R_S is veridical just in case there are red, orange and green things on the desk, but not if there are only red and green or red and orange things. Contrastingly, for R_F to be veridical it suffices that there are red and maybe orange things on the desk – as orange is similar to red, it is only partially suppressed. See Figure 5.

Worldly condition	R _F veridical	R _s veridical
	Yes	Yes
	Yes	No
	Yes	No

Figure 5. Veridicality conditions of representations with spatial and feature attention.

4. Precise inattentive representations

Let us now consider the second piece of evidence against representationalism. Here, the focus shifts from the effects of spatial attention to the effects of attention more generally understood as a perceptual

and cognitive processing resource.¹³ Asplund et al (2014) compared perception of a stimulus' color with and without attention.¹⁴ They manipulated attention with an attentional blink paradigm, which renders a target (T2) "unattended" by presenting it within a 200-400ms window after another target (T1; see Raymond et al. 1992). This is an interval where attention is thought to be exhausted by T1 processing, as suggested by significant drops in T2 identification performance (see Figure 6). Subjects reported T2 color by picking it from a wheel with 180 hue values. Though, as expected, more *correct* responses were given for attended than for unattended targets (which were at chance), attended and unattended targets did not differ in how far off *incorrect* reports were from the correct values. Asplund et al concluded that attentive and inattentive perception are equally precise.¹⁵




Figure 6. Schematic illustration of the attentional blink effect

Block invokes these results to support his second claim, namely, that attention does not make a difference to representational precision (Block 2015: 36). If cashed out in terms of range content, the claim becomes:

(2.1) Visual representations attribute equal ranges of values with and without attention.

One initial worry about the extent to which Asplund et al's results support this claim is that the attentional blink might not cleanly separate unattended from attended targets. The unattended condition was constituted by T2 presentations within the mentioned 200-400ms window, dubbed 'lag 2'. ¹⁶ In turn, the attended condition was given by T2 presentations at a longer lag, dubbed 'lag 8', at which T2

¹³ Or capacity; this difference does not matter for the present discussion.

¹⁴ Boone (2020) also discusses this study, but he focuses on explaining how accuracy differences in between attended and unattended can readily explain attentional changes in phenomenology. He does not challenge the claim that attention leaves precision intact.

¹⁵ Specifically, their claim is that perception does not gradually become more precise as the amount of available attention increases. These details are not relevant for the present discussion, but the interested reader can find further elaboration in Lopez (2023).

¹⁶ Instead of "unattended" and "attended", one could also think of these more cautiously as conditions of less and more attention, respectively. The claims made here also apply if attention is treated as gradable.

identification returns to baseline and thus attention is thought to be available again (duration unspecified in the paper, but typically around 800ms). However, Asplund et al observed that response pools at each of these lags could be analyzed as a mixture of two components: a set of randomly scattered responses, plausibly corresponding to guesses in trials where subjects did not see T2, and another set with responses clustering around the correct value, plausibly corresponding to trials where subjects did see T2. ¹⁷ Unsurprisingly, there were more random responses at lag 2 than lag 8, which is consistent with the idea that attention is not available at lag 2, as the absence of attention should make T2 harder to see (Hutchinson 2019). But then, it is unclear what one should make of the targets in *the other set* of lag-2 responses, where subjects presumably did see T2. By the previous reasoning, it is possible that some attention was in fact available for T2 in these trials. Hence, there are two available accounts of what is happening in these trials: either subjects see T2 without attention, or they see T2 *because their attention is not fully exhausted*. But if it is the latter, both lag 2 and lag 8 provide *attended* conditions.

Admittedly, this worry hinges on general difficulties with demonstrating total absence of attention (Cohen & Dennett 2011). However, a more direct issue concerns the measure of precision at work in Asplund et al.'s study. They understand precision in terms of *the non-random components* in lag-2 and lag-8 response pools. Since the random component is thought to reflect trials were T2 is not seen, these trials might not involve a T2 representation, and thus might be irrelevant for the precision of T2 representations. To be sure, T2 representations that are too fleeting or too degraded to raise performance above chance cannot be fully ruled out. The random response subset could involve such representations, in which case these responses would indeed be relevant for representational precision and should not be excluded from the measurement. However, since chance performance is typically accepted in cognitive psychology as an indicator of absence of target information or even perception, I shall not dispute Asplund et al's measure choice on these grounds. Instead, I will emphasize a dissociation between their notion of precision and the one at stake in claims about range content.

Asplund et al compared the precision of unattended and attended perception by comparing how lag-2 and lag-8 responses distributed in the color wheel. As mentioned above, they focused on the subsets of non-random responses within each of these pools, though for simplicity I will simply refer to lag-2 and lag-8 response pools. Precision was defined as a measure of how tightly the responses in each pool clustered together. ¹⁸ "Tightness of clustering" consists in the error range within each pool, which could go from 0 (no error/maximal accuracy) to 180 degrees off the mark (maximal inaccuracy). To illustrate, suppose you have two pools of ten responses each (see Figure 6). In the first pool, all responses are off by 45 degrees. In the second pool, six responses are off by 45 degrees and four more are off by 90 degrees. The first response pool is more tightly clustered than the second; accordingly, responses are more precise in the former than in the latter.

¹⁷ Specifically, they modeled their results as a mixture of these two distributions, i.e., a random one for guesses and a non-random one for non-guesses. This model fitted their findings better than alternative models, some of which incorporated variations in precision at different lags.

¹⁸ More formally, precision was indexed by the standard deviation (σ) of these response distributions. This measure was virtually indistinguishable between both response pools, though notably, σ is slightly greater for unattended than for attended targets (σ = 20.4 versus 20.6).




Figure 7. Response error distributions and range contents

In Asplund et al's study, lag-2 and lag-8 response pools were equally spread out: in both, responses were off for as much as 180 degrees in either direction (Asplund et al 2014: 10, Fig. 3). More importantly, the bulk of off responses in both pools concentrated around 90 degrees, so both pools clustered together equally tightly –the width of the bell-shaped error distribution curve was roughly the same in both (Block 2015: 38, Fig. 12). If lag-2 targets were represented with less precision than lag-8 targets, then this curve width should be broader for lag-2 than for lag-8 responses –for instance, with the bulk of lag-2 off responses concentrated around 120 rather than 90 degrees.

Now, here is why these results might lack direct bearing on the issue of range contents. In the last section, we saw how the width of tuning and probability distribution curves might come apart from range content breadth. Similar considerations apply to the error distributions discussed here. Take our toy example from Figure 7. Though responses cluster together more tightly in Pool 1 than in Pool 2, it is not clear that the former "represents" a narrower range of values than the latter. At any rate, even if Pools 1 and 2 are at all *representations*, it is doubtful that we should count them as *visual* representations. They are more plausibly *sets* of visual representations, or representations of sets of visual representations. But then, the predicated precision is a property of these sets, not of the individual representations within them. Hence, equal error distributions for attended versus unattended target colors are in principle compatible with a difference in the value ranges represented by individual attentive and inattentive percepts.

Block acknowledges that the study by Asplund et al does not measure the precision of individual percepts (though Asplund et al might not acknowledge this themselves; see Block 2015: 37). Accordingly, it is expected that one cannot directly read the results in terms of range contents of individual percepts. Thus, the relevant response distributions are not supposed to be themselves representations with equally broad range content. Still, Block notes that the precision of a set of responses can plausibly reflect the precision of the individual representations in it. Block does not

elaborate on how, but here is a possible way. If individual unattended targets were represented more imprecisely, it would be reasonable to expect this response pool to be more scattered. Say a target is vermillion and say that when you see it at lag 2, your percept represents it as reddish. Suppose this reddish percept covers only two values: vermillion and pumpkin orange. You can only pick one value to report your percept, and since your reddish representation is indeterminate between vermillion and pumpkin orange, you are equally likely to pick either. Thus, if you see a vermillion target at lag 2 ten times, you are likely to pick vermillion five times and pumpkin orange five times. In this case, your range of incorrect responses will span two degrees (for this is the portion of the wheel occupied by each hue value). Now suppose your reddish percept is more imprecise, covering four values instead of two. In that case, you are likely to pick the correct value about two to three times, and your range of incorrect responses will be six to eight degrees. In this way, increasing the imprecision of individual responses makes response distributions more scattered. Accordingly, equal error ranges in two response sets indicate that the percepts associated with individual responses in each set represent equal (or similar) ranges of values.

Though this proposal has initial plausibility, more elucidation is needed on how exactly response pool precision reflects individual response precision. For it seems prima facie possible that a set of precise individual responses is widely scattered, and it seems likewise possible that a set of imprecise individual responses is tightly clustered together. Similarly, a given response distribution could be obtained from responses with different individual precisions. Moreover, more studies with different attentional manipulations and response measures are needed for assessing how attention affects the precision of individual percepts.

5. Concluding remarks

I have argued that evidence on the mechanisms of spatial attention and on visual perception during the attentional blink does not support the claims that attention does not affect representational precision, understood as a measure of the ranges of values attributed by a perceptual experience within a property dimension. In other words, I have argued that this evidence is compatible with the claim that attentive experiences represent narrower ranges of properties than inattentive experiences. However, along the way I have also uncovered ambiguities and unclarities in the concept of range content. Subsequent arguments for or against representationalism should first clarify these issues. Finally, I have also identified some potential directions for future empirical investigations.

In the foregoing, I have focused on Block's claims about the precision of attentive and inattentive experiences, and I have bypassed the more intricate and overarching argument against representationalism of which these claims are part. In this argument, considerations about veridicality play a center role in constraining representational content. I will not attempt to make justice to this argument here, but two key ideas are that we must assume perception to be largely veridical, and that there is no principled way of deciding whether attention increases or decreases veridicality. Attributing broad range contents to both attended and unattended percepts is one way of handling this difficulty, as these broad ranges can include the actual properties of the object (e.g., its actual contrast value), while at the same time accommodating the puzzling case where two stimuli with different properties (e.g., different contrast values) look indistinguishable when only one of them is attended.

Here, I have only mentioned veridicality because of its role in spelling out representational content. I have not made any claims about the general veridicality of perception or about whether

attentive experiences are more or less veridical than inattentive experiences. ¹⁹ A positive defense of the representationalist claim that attentive experiences have narrower range contents than inattentive experiences (which I also have not offered here) would plausibly have to say more about these issues.

References

Anton-Erxleben, K. and Carrasco, M. (2013). Attentional enhancement of spatial resolution: Linking behavioural and neurophysiological evidence. *Nature Reviews Neuroscience* 14,18-200.

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

Arazi, A., Yeshurun, Y., & Dinstein, I. (2019). Neural Variability Is Quenched by Attention. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, *39*(30), 5975–5985. https://doi.org/10.1523/JNEUROSCI.0355-19.2019

Asplund, C. L., Fougnie, D., Zughni, S., Martin, J. W. and Marois, R. (2014). The attentional blink reveals the probabilistic nature of discrete conscious perception. Psychol Sci. 2014 Mar;25(3):824-31. doi: 10.1177/0956797613513810.

Block, N. (2015). The puzzle of perceptual precision. In T. Metzinger & J. M. Windt (Eds.), *Open mind* (Vol. 5). MIND Group.

Block, N. (2010). Attention and mental paint. *Philosophical Issues*, 20(1), 23–63. https://doi.org/10.1111/j.1533-6077.2010.00177.x

Boone, T. (2020). Range content, attention, and the precision of representation. *Philosophical Psychology* 33 (8):1141-1161.

Brogaard, B. (2015). Type 2 blindsight and the nature of visual experience. *Conscious Cogn.* Mar;32:92-103. doi: 10.1016/j.concog.2014.09.017.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Res.* July 1; 51(13): 1484–1525. doi:10.1016/j.visres.2011.04.012.

Carrasco M. (2018). How visual spatial attention alters perception. *Cognitive processing*, 19(Suppl 1), 77–88. https://doi.org/10.1007/s10339-018-0883-4

Carrasco, M., & Barbot, A. (2019). Spatial attention alters visual appearance. *Current Opinion in Psychology*, 29, 56–64. https://doi.org/10.1016/j.copsyc.2018.10.010

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

¹⁹ Note that this is not an issue of how constraining or permissive the veridicality conditions associated with these experiences are.

Carrasco, M, Penpeci-Talgar, C. & Eckstein, M. (2000). Spatial covert attention increases contrast sensitivity along the CSF: support for signal enhancement. *Vision Res.*; 40:1203–1215. [PubMed: 10788636]

Cavanagh, P., HE, S., & Intriligator, J. (1999). Attentional resolution: the grain and locus of visual awareness. In *Neuronal Bases And Psychological Aspects Of Consciousness* (pp. 41-52).

Chalmers, D. (2004). The representational character of experience. In *The Future for Philosophy*, ed. Brian Leiter. Oxford: Oxford University Press. pp. 153–81.

Cohen, Michael A. & Dennett, Daniel C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences* 15 (8):358--364.

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

Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18(1), 193–222.

Fazekas, Peter & Nanay, Bence (2021). Attention Is Amplification, Not Selection. *British Journal for the Philosophy of Science* 72 (1):299-324.

Gobel, J. and Carrasco, M. (2005). Attention alters the appearance of spatial frequency and gap effect. *Psychological Science*, *16*(8), 64-651.

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

Huang, L., and Dobkins, K.R. (2005). Attentional effects on contrast discrimination in humans: evidence for both contrast gain and response gain. *Vision Res.* Apr;45(9):1201-12. doi: 10.1016/j.visres.2004.10.024.

Hutchinson B. T. (2019). Toward a theory of consciousness: A review of the neural correlates of inattentional blindness. *Neuroscience and biobehavioral reviews*, *104*, 87–99. https://doi.org/10.1016/j.neubiorev.2019.06.003.

Jennings, Carolyn Dicey (2020). The Attending Mind. New York: Cambridge University Press.

Lee, Andrew Y. (2021). Modeling Mental Qualities. The Philosophical Review 130 (2):263-209

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

Lopez. A. (2023). Degrees of attention and degrees of consciousness. *Conscious and Unconscious Mentality: Examining their Nature, Similarities and Differences*. M. Polák, T. Marvan, & J. Hvorecký (eds.) Routledge.

Lopez, A., and Simsova, E. (2023). Enhanced but indeterminate? How attention colours our world. *Review of Philosophy and Psychology*. https://doi.org/10.1007/s13164-023-00697-7

Martinez-Trujillo, J.C., and Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr Biol*;14:744–751. [PubMed: 15120065]

Martinez-Trujllo J. C., and Treue, S. (2002). Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron*;35:365–370. [PubMed: 12160753]

McAdams, C. J., and Maunsell, J.H.R. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J Neurosci* 19:431–441.

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

Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, 229(4715), 782–84.

Nanay, B. (2010). Attention and perceptual content. Analysis 70 (2):263-270.

Pelli, D. G. (2008). Crowding: A cortical constraint on object recognition. *Current Opinion in Neurobiology* 2008 August; 18(4): 445 □ 451. doi:10.1016/j.conb.2008.09.008.

Posner, M. I. (1980). Orienting of attention. Quarterly Journal of Experimental Psychology, 32(FEB), 3e25.

Raymond, J. E., Shapiro, K. L. and Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: an attentional blink? *Journal of Experimental Psychology: Human Perception & Performance*; 18(3):849–60.

Reynolds, J. H., and D. J. Heeger (2009). The normalization model of attention. *Neuron* 2009 January 29; 61(2): $168 \square 185$. doi:10.1016/j.neuron.2009.01.002.

Reynolds, J.H., Pasternak, T. & Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron.*; 26:703–714. [PubMed: 10896165]

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

Strasburger, H., Rentschler, I. and Jüttner, M. (2011). Peripheral vision and pattern recognition: A review. *Journal of Vision*;11(5):13. doi: https://doi.org/10.1167/11.5.13.

Theeuwes, J. (2013). Feature- based attention: it is all bottom-up priming. *Philosophical Transactions of the Royal Society B* 368: 20130055. http://dx.doi.org/10.1098/rstb.2013.0055

Yeshurun, Y. and Carrasco, M. (1999). Spatial attention improves performance in spatial resolution tasks. *Vision Research*, 39(2), 293-306.

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

Yeshurun, Y., Montagna, B., and Carrasco, M. (2008). On the flexibility of sustained attention and its effects on a texture segmentation task. *Vision Research*, 48, 80–95.

Watzl, Sebastian (2017). *Structuring Mind. The Nature of Attention and How it Shapes Consciousness*. Oxford, UK: Oxford University Press.