Forthcoming in *Philosophical Psychology*

Visual Stuff and Active Vision[†]

Wayne Wright
Washington University in St. Louis
www.mindstuff.net
wrightwt@mindstuff.net

Abstract: This paper examines the status of unattended visual stimuli in the light of recent work on the role of attention in visual perception. Although the question of whether attention is required for visual experience seems very interesting, this paper argues that there currently is no good reason to take a stand on the issue. Moreover, it is argued that much of the allure of that question stems from a continued attachment to the defective 'inner picture view' of experience and a mistaken notion that the ultimate goal of vision is to produce visual experience. The paper considers a promising general account of the content and structure of vision and presents reasons for not taking that account to be committed to any substantive claims about the experiential status of unattended visual stimuli. Also addressed are the active nature of vision and the role of vision in enabling our ecological success. These considerations highlight that visual experience is not the whole of vision and that a much more important question about unattended visual stimuli than whether they are consciously experienced is what contribution they make to how we interact with the world.

Recent work on the role of attention in vision challenges the idea that our visual experience is as detailed, stable, and coherent as it introspectively seems to be. One way of interpreting the results of these studies is that attention is required for visual experience, either of changes from one moment of view to the next or of features in an isolated visual episode. In the absence of attention, the result is either change blindness (CB) or inattentional blindness (IB).[1] However, these findings might instead be interpreted in terms of an attention-induced lack of memory (e.g., inattentional amnesia, change amnesia): visual experience can occur in the absence of attention, and a failure of memory accounts for experimental subjects' inability to report the presence or

[†] I would like to thank William Bechtel and an anonymous referee of this journal for their useful comments and suggestions.

[1] For more on change blindness, see O'Regan (1992), Rensink (2000a), Simons (1996). For more on inattentional blindness, see Mack and Rock (1998a & 1998b), O'Regan and Noë (2001), Simons (1999 & 2000), Simons and Chabris (1999). Several of these references discuss both topics.

change of unattended stimuli. Jeremy Wolfe, while suggesting that unattended stimuli may be seen but not put into memory, comments that deciding the matter "rides on the definition of what it means to consciously perceive something" (1999, p.89). I want to make a similar, but crucially different claim. Namely, that treating visual experience as a solitary thing unto itself—rather than as an element in a complicated web of visual, cognitive, and behavioral phenomena—leads to all manner of confusion. Considered from the perspective of <u>active vision</u>—a framework for vision research that rejects the notion that the fundamental purpose of vision is to construct picture-like mental representations of the scene before the eyes—the question of whether unattended stimuli are consciously seen becomes rather unimportant, and we should neither expect nor demand an answer to it. However, given the active vision paradigm's emphasis on gaze shifts and other activity initiated by the perceiving subject, and the crucial role of such activity in meeting the informational needs of the perceiving creature (particularly with respect to decision making and behavior guidance), there is much useful research to be done on the effects of unattended visual stimuli on a variety of mental processes.

1 The inner picture view

A slightly more formal way to describe the introspective sense of complete detail, is that the contents of our visual experiences are structured entities that contain a considerable amount of metric-preserving information about the stimulus scene, and we are free to attend to as much (within empirical limits) or as little of that information as we choose. Call this the <u>inner picture</u> view. It is not as though only "the folk" take the inner picture view seriously; cf. Levin (2002) and Levin et al (2000) for research on people's beliefs about what they see that suggests that they radically overestimate their abilities for visual detection and recognition. In some form or other, it has often been the operating understanding of visual experience in philosophy and cognitive science. For example, Michael Tye claims that the module relevant to visual experience "is the one that has as its output a unified representation of the entire visual scene" (1995, p.138); cf. also Feldman (1985, pp. 265-266) and Trehub (1991, p.55) for endorsements of the idea that visual experience involves a rich and coherent representation of the entire visual scene.

As I will discuss in sections 2 and 3, studies of attention and vision show that in its most extreme or naïve form, the inner picture view is mistaken. The further issue of whether a more modest form of the inner picture view can be maintained will be addressed in section 5. Section 4 discusses a promising general account of the content and structure of vision and presents some reasons for not taking that account to be committed to any substantive claims about the experiential status of unattended visual stimuli. Section 6 addresses the active nature of vision and the role of vision in enabling our ecological success, which leads to the realizations that visual experience is not the whole of vision and that a much more important question about unattended visual stimuli than whether they are consciously experienced is what contribution they make to how we interact with the world. I conclude in section 7.

2 Change blindness

Building on the inner picture view, one might think that a perceiver is able to detect changes in successive views of the same scene by comparing the detail of her current visual episode with that of the moment just before. Call the idea that a detailed, metric-preserving representation of a scene survives or is constructed across views of that scene the <u>persistent picture</u> view. Recent studies of CB undermine this view. Change is typically perceived through the detection of associated "flickering" in the visual scene that attracts attention. However, flickering can go undetected when a perceiver is distracted by a disturbance in the scene simultaneous with the local change, with the result that the perceiver does not notice the change. Even prominently located and large-scale changes—the kind that would usually be easy to detect—go undetected in CB studies.[2]

Consider the following description of an experiment that demonstrates the effects of CB (O'Regan et al, 1999). Naïve subjects viewed a video display in which a sequence of three images was shown: the original, the distractor, and the modified original. The first and last are shown for approximately 3 s, while the middle is displayed for 80 ms. The original image is of a car on the road, with a solid traffic line to its right. The distractor image is just like the original, but with 'mudsplashes' (small black-and-white patterned rectangles) occluding small parts of the image, but not any part of the traffic line. The modified image is also just like the first, except now the traffic line is broken instead of solid. The brief and minor disturbance created by the mudsplashes has a profound effect on perceivers' ability to detect the change in the traffic lines: without knowledge of what changes between the first and last images, one does not notice the difference between them. Once one learns what to look for, however, the change becomes prominent. In other tests using image sequences of a similar sort, subjects sometimes noticed the change after seeing the images cycled through several times, and there were also cases when they never noticed the change during the 40 s interval that each sequence was cycled.

These results are incompatible with the persistent picture view, which would have it that the distractor should not negate one's ability to detect change. The information from the

---

[2] The reader is encouraged to go to Daniel Simons' Change Detection Database at http://viscog.beckman.uiuc.edu/change/demolinks.shtml and Kevin O'Regan's webpage at http://nivea.psycho.univ-paris5.fr to view some of the available demos. Note that, because they more clearly tie into the examination of the inner picture view, I am mainly discussing here cases of "fast" changes that go undetected due to a global distraction. Simons, for example, has several examples of blindness when change occurs slowly and without any distracting flicker.

immediately prior view should still be accessible, and a comparison between the stored representation of the before state and the representation of the after state should reveal what has changed. The fact that CB goes away immediately after one is told (or manages to discover) what changes highlights the role of attention in detecting change. Attention marks certain elements of the scene and the rest appear to not be encoded in memory. Those elements of the original view that have been indexed are loaded into visual short-term memory, persist through the distractor event, and are available for comparison to the modified view. If a richly detailed, coherent representation is constructed for every visual episode, CB makes clear that very little of that information survives even a short while.

3 Inattentional blindness

IB is the failure to report, and perhaps to detect, the presence of visual stimuli that occur in plain view due to a lack of attention.[3] When IB subjects are confronted with what they claim they did not notice during the trials, they have no problem seeing what they missed and are shocked to learn that they did so. Arien Mack and Irvin Rock conclude from their research on IB that "there seems to be <u>no conscious perception without attention</u>" (1998a, p.ix; their emphasis). According to them, although we may have the introspective sense that our visual experience involves a robust inner model of the stimulus scene, we experience only those features that we attend to.

In Mack and Rock's experiments, subjects were engaged in tasks that demanded a high degree of attention, such as looking at a cross and trying to determine which arm is longer. After several trials of this task, an unexpected stimulus (e.g., a small, black square) was presented in one of the cross' quadrants. Focused on their task, approximately 25% of perceivers reported not

[3] For the moment, I am lumping together inattentional blindness and inattentional amnesia as IB.

noticing the unexpected object when it appeared parafoveally and the cross was presented at fixation. Even more surprising and revealing is that when the cross was presented parafoveally and the unexpected stimulus appeared at fixation, approximately 75% of the subjects reported not perceiving the target object. In subsequent trials that included priming to look for the target object, almost every subject reported noticing the target object. Based on the evidence they gathered, Mack and Rock argued that subjects did not consciously perceive those items that they did not attend to and that attention is necessary for visual stimuli to become conscious.

A well-known study with results consistent with Mack and Rock's research had subjects watch a 75 s videotape of a basketball-like game played in a small area, while tasked with counting the number of bounce passes and chest passes one of the two teams makes (Simons and Chabris, 1999, which follows the experiments of Neisser, 1979). In one condition, in which separate video recordings of the two teams (one wearing black and the other wearing white) were superimposed, giving the players a 'ghostly' or transparent appearance, 75% of the test subjects observing the black team and 92% of those observing the white team did not notice a person in a gorilla suit stroll across the court from left to right over a span of 5 s. At points during its path across the court, the gorilla 'passed through' the players and the ball, both of which the subjects were supposed to be attending to, yet went unnoticed. Due to the transparent quality of the figures in the video, the gorilla was occupying the same retinal location as the features of the scene that were being attended to, suggesting that the subjects' attention was on something (or some feature) rather than a particular location in the visual field. In another test condition, filmed in live action with a single camera (giving the actors a normal, opaque appearance) roughly half the test subjects failed to notice that the gorilla stepped into the middle of the visual display, thumped its chest, and strolled off the court, all of which took 9 s and had the gorilla both occluding and occluded by the players.

The data on CB make an imposing case for the requirement of attention to experience change. Less clear is the status of unattended features in an isolated view, as one possible explanation of CB is that all the detail of the visual scene is encoded for the before and after views, but the lack of attention leads to a failure of memory or of comparison. Likewise, the failure to report stimuli in the IB experiments could be an effect of limits on memory (or categorization), not on experienced detail, as subjects are not asked whether they saw the unexpected stimuli until after the trial (Wolfe, 1999 and Simons, 2000). There certainly is room to question Mack and Rock's strong interpretation of the connection between attention and visual experience, the chief issues being that we cannot directly observe the experiences of others, the powerful introspective sense of the phenomenological presence of visual stimuli outside the scope of attention, and the effects of unattended stimuli on subsequent behavior. Regardless, it is clear that we have good reason to reject strong versions of the inner picture view. For one thing, from the standpoint of evolutionary engineering, it would stand to good reason that our visual system would not create—especially if it is not going to store or exploit—richly detailed representations for features that are not of interest to the perceiver, with what is of interest being determined by what is attended to.[4] The steady diminishment of visual acuity and color sensitivity away from the fovea should itself be sufficient to undermine the strongest forms of the inner picture view. More closely tied to the results of the CB and IB studies is that given the salience and prolonged duration of (or repeated exposure to) the unexpected stimuli in some of

[4] Rojer & Schwartz (1990) discuss the prohibitively high computational costs of high-fidelity representations of the detail present in the visual scene; cf. also Findlay & Gilchrist (2003, p.5). Diamond (1992 & 1993) examines nature's reluctance to over-engineer, focusing on studies of the quantitative match between biological capacities and natural loads.

these experiments, the subjects' failures to report them becomes extremely difficult to comprehend if the visual scene were represented in vivid detail.

I do not mean to rule out that at least some features of the unexpected stimuli are coarsely represented in visual experience in the absence of attention. In fact, I don't think we can exclude that possibility, although it is also unclear whether we should embrace it. The idea that only attended contents are represented in coherent detail while unattended features can be present but 'worked up' to a considerably lesser degree is consistent with the results of the study conducted by Rees et al (1999) on inattentional blindness for ignored words. The subjects' performance and functional imaging data might be taken to suggest that while they "were blind to those properties that distinguish words from random strings of consonants" (Rees et al, 1999, p.232), the ignored words (which were superimposed with the pictures the subjects were attending to in their detection task) had <u>some</u> presence in the subjects' visual experience. The attentional demands of the task the subjects were engaged in while the words were presented thwarted further processing in both language areas (e.g., left prefrontal cortex) and areas of visual cortex, but that does not dictate that some features of the words were not consciously present.

Also relevant is Haines' (1991) discussion of airplane landings in flight simulators equipped with head-up displays, which presented standard cockpit information on the windshield of the plane. The displayed information, which the pilots were attending to while conducting their landings, was retinally co-located with features of the scene outside the aircraft. Two out of the eight commercial airline pilots in the study failed to notice that another airliner was sitting on the runway they were approaching. Those pilots landed their planes 'through' the parked airplane, completely unaware of their blunder. Following up on Rees et al's (1999) conjecture that a less attention-hungry detection task might have allowed some linguistic processing of the words their subjects ignored, if the attentional demands of the pilots' tasks in conducting their landings had been higher (or high more continuously), likely more of the pilots would have committed the same error. The fact that subjects can fail to register significant features of the visual scene that they have fixated—the unexpected stimuli presented at fixation while the attended cross is presented parafoveally, the ghostly gorilla that 'passes through' the basketball players, the parked jet that is retinally co-located with the HUD flight information—undercuts the claim that our experiences involve high-fidelity representations and establishes that the impression that "all of the information in the visual environment is potentially available for attentive processing" is mistaken (Simons and Chabris, 1999, p.1060).

4 Preattentive processing, memory, and visual stuff

What does the IB and CB research tell us about the structure and content of visual experience? One safe conclusion is that whatever detail and structure is included in visual experience, all that survives from moment-to-moment is that which receives attention. We simply do not build up in memory a fused conscious representation of the scene before us, even one of short duration; cf. Irwin et al (1983), Irwin (1991 & 1996).[5] However, we are not in a position to deny that something beyond the scope of attention is consciously present. Rejecting that possibility requires dealing with the overwhelming phenomenological sense that there is more to visual experience than what is the focus of attention, the impressive body of research on preattentive

[5] There is considerable evidence that while a fused (e.g., detailed, coherent) representation is not formed across fixations, some information about the scene affects subsequent visual activity; cf. Pollatsek et al (1984) and Germeys et al (2002). This information need not be taken to be conscious, despite its effects on visual activity, although some of it might be consciously accessible. This is consistent with the broader understanding of vision I am advocating in this paper and fits with the 'visual stuff' account described below.

vision (cf. Treisman, 1986), and the evidence that unattended features can influence conscious-level activity (e.g., syntactic disambiguation, word-stem completion, behavior guidance, expectations of subsequent stimuli). In the light of the research on attention, memory, and vision, a general picture of how to understand visual perception has attracted some support:[6]

1. All regions of the visual field contain 'visual stuff' (Wolfe, 1999) or 'proto-objects' (Rensink, 2000b). Visual stuff is the product of preattentive processing and includes some coherent feature detail, but the coherence is only local; i.e., the visual stuff consists of groupings of features in small regions of visual space and does not involve any broader-range groupings of the local groupings.
2. When a batch of available visual stuff receives attention, it is grouped together in a spatially coherent, temporally durable representation. This representation can be placed into memory and can serve as the target of other mental activity; e.g., decision-making, object recognition.
3. Besides having only local coherence, unattended visual stuff has no life beyond the present. Visual stuff is constantly updated and is not encoded in memory.
4. When attention is removed, the representation in (2) falls apart and its elements revert to 'mere' visual stuff.
5. Pre-attentive processing provides information about structural aspects of the scene that can be used to guide attention to locations or objects that are relevant to the task at hand. This information includes the gist and spatial layout of the scene (Rensink, 2000b, p.1476; Wolfe, 1994 & 1998).

[6] Most notably, Rensink (2000b), Wolfe (1999), Wolfe & Bennett (1997), Wolfe et al (2000). Hollingworth (in press) offers a critique of this approach to visual representation, particularly in regard to the effects of withdrawing attention.

This promising account rules out perceiving change without attention, but by itself does not address the issue of whether unattended features of a visual scene are consciously perceived. Wolfe contends that conscious visual experience consists of both all the unattended visual stuff and the representation generated by focusing attention on some of that visual stuff (1999, p.74), whereas Rensink is less sanguine about the inclusion of unattended visual stuff in visual experience (2000b, p.1475). My suggestion is that there is a legitimate sense in which 'mere' visual stuff is not conscious. This still doesn't resolve for us the difficult matter of whether unattended features are conscious in a distinctively visual way (i.e., involving visual qualia), but I will later argue that taking that question very seriously is mistaken, as it is significant only from the standpoint of a fundamental misunderstanding of the role of vision in our lives.

On one common understanding of what it is for a perceptual content to be conscious, it must be poised to directly impact activity such as deliberation and the direction of intentional action (Tye, 2000, p.62). In being poised, however, the content doesn't have to make an actual impact on mental processes. 'Directly' is a slippery term, but presumably it is tied up with the subject's ability to selectively manipulate that content and grasp its role in the processes in which it figures. This understanding of consciousness is basically a requirement of an appropriate functional pedigree; i.e., having a role in mediating inputs and outputs in certain ways. By this standard, mere visual stuff is not conscious. Unattended visual stuff is too volatile—constantly regenerated and exhibiting only local structure—to even be poised to, much less actually make, such a direct impact. Furthermore, unattended visual stuff will not carry over into an accessible representation built-up across a series of fixations and so cannot attain poise in a 'secondary' way. Thus unattended contents are not available to be manipulated in the kinds of ways characteristic of this notion of consciousness.

Of course, one could argue that visual stuff is poised in the sense that were it to receive attention, it would take on a temporally durable and spatially coherent nature sufficient to make a direct impact on, e.g., the belief/desire system. However, such a maneuver stretches the notion of poise to vacuity and those opting to understand poise so liberally would take on the burden of marking out some useful way of distinguishing between contents that are poised and those that are not. Additionally, it ignores the important similarities between IB and conditions such as blindsight and visual neglect (Mack, 2002). For example, like blindsighters, IB subjects cannot report the presence of certain stimuli in their visual field and do not actively engage information about the features they claim to not visually experience, yet further studies show that information about the blind area can influence the guidance of behavior and expectations of subsequent stimuli. If we take seriously blindsighters' reports that they do not have visual experiences in their blind fields, why do we not do the same for the IB test subjects regarding their reports that they do not see the unexpected stimuli they are presented? Although visual neglect and blindsight are neural deficit conditions and IB occurs as part of normal visual experience, there seem to be compelling reasons to conclude that all three preclude certain contents from directly impacting other mental processing.

Relevant to this discussion is Wolfe's treatment of the question, "Can one be said to consciously see something if one cannot remember seeing it an instant later?" (Wolfe, 1999, p.89). Recall that Wolfe takes the conscious visual representation to consist of all the available visual stuff, both attended and unattended, and he takes attention to be the gateway to memory. So, for him, the question is either about whether fleeting awareness is sufficient for perception or about the support for his inclusion of unattended visual stuff in the content of visual experience. Wolfe responds to his own question by saying that a failure of memory does not entail a failure of conscious experience (ibid, pp.89-80), so it looks like he has the latter concern in mind. To make his case, Wolfe uses a hypothetical example in which he shows a subject a picture and immediately thereafter gives the subject an electroconvulsive shock, thereby preventing the subject from forming a memory of seeing the picture. Wolfe contends that it would be a mistake to conclude from the subject's failure of memory that she did not consciously experience the picture. I don't think the matter is so simple. For one thing, it looks like the kind of conscious experience Wolfe is appealing to has to do with 'bare phenomenology', rather than the 'poised' notion of conscious experience. If the requirement of poise were in play, the unattended visual stuff of the subject's visual episode would be ruled out independently of the effects of the subsequent electroconvulsive shock. The only difficult issue would concern the details of what counterfactual conditions we would want to build into the notion of poise to deal with situations in which external factors interfere with the attended content standing ready to directly impact mental processes.

So, if Wolfe's example is supposed to establish anything about the possibility of conscious experience in the face of a lack of memory, the kind of conscious experience at issue has to be intimately tied up with the visual qualia of the subject's experience. Does it make sense to talk about regions of the visual field exhibiting qualitative features in the absence of any memory of those features just a moment later? If one thinks that a content must be poised in order to be phenomenally consciousness (as Tye and several others do), then the answer is the one given above—the unattended contents are not conscious and we may very well count the attended parts of the picture as having been consciously seen.[7] However, one might doubt that

[7] Important to note is that Tye explicitly claims that there can be phenomenal consciousness in the absence of attention (Tye, 2000, p.14). However, I have argued elsewhere (Wright, 2005), on grounds

the poised notion of consciousness exhausts the phenomenal aspect of consciousness, perhaps drawing on Block's (1995) distinction between access-consciousness and phenomenal-consciousness. While I am not persuaded that a convincing case has been made that phenomenal consciousness and access consciousness can come apart, I also don't see much reason to insist that they must always go together.[8] At this point, I am going to show my hand and say that on the 'bare phenomenology' reading, we just don't know the answer to Wolfe's question and we may never be in a position to settle the matter. Unless a future neuroscience enables us to read phenomenology off of brain activity, there will always be obstacles to characterizing those aspects of experience that subjects do not report.[9] Perhaps experimental methods can be devised to further chip away at the domain of what evades report, but there is good reason for pessimism about getting a non-speculative complete answer to questions about the phenomenological status of unattended (or otherwise unremembered) elements of the visual scene. In section 6, I will explain why this pessimism should not cause us much worry.

Before concluding this section, I want to note a few other reasons why one should not rush to embrace Wolfe's inattentional amnesia account (while also granting that it might be correct). For one thing, should it turn out that the priming effects of unattended stimuli require that they receive some representation in memory (a claim I will not attempt to argue for here and that I have no firm opinion on), the diagnosis of amnesia becomes hard to maintain (Mack, 2003, p.181). In that case, even if it turns out that unattended stimuli are visually experienced, their

---

similar to those provided just before, that the CB and IB research, together with the requirement of poise, demonstrate that Tye is committed to unattended contents lacking phenomenal consciousness.

[8] The recognition of this latter point accounts in part for my more tempered claims here about attention and phenomenal consciousness than those developed in Wright (2005).

[9] As Dennett (1991 & 2002) often points out, there is also a serious question of how to interpret the reports subjects do make about their conscious experience. Wright (in press) discusses some of the methodology employed in recovering first-person data about conscious experience on the basis of third-person data.

priming effects would suggest that whatever it is that prevents subjects from reporting them cannot be a failure of memory (although it could be a matter of the inaccessibility of the memory once it is encoded). A link between priming and memory would undercut Wolfe's appeal to Moore & Egeth's (1997) study that demonstrated that unattended stimuli can affect perceptual judgments about Ponzo illusions (Wolfe et al, 2000, p.711). Furthermore, the attempt to recharacterize change blindness as amnesia for pre-change contents, which is then extended to recharacterize IB in terms of amnesia, does not seem strongly supported:

> [A visual] transient supports detection of the change. It does not give the viewer information about the original contents of the display. Brawn, Snowden, and Wolfe (1999) had participants view a field of red and green dots. At one moment, a transient summoned attention to one dot. At that same moment, the dot either did or did not change color. It was then quickly removed and masked. Participants had little trouble identifying the color of the dot in the interval between cue and mask. However, they were at chance level in regard to naming that color before the attention-summoning cue (Brawn et al, 1999). Change blindness is a catchy misnomer. As a general rule, the visual system is not blind to change. Rather, it is amnesic for the contents of the stimulus before the change. (Wolfe et al, 2000, p.710)

One point of easy agreement with this passage is that the visual system typically is not blind to large-scale changes that occur right before the eyes in normal situations; e.g., situations that don't involve global disturbances like those found in CB demonstrations. However, I fail to see how the inability of subjects to identify the dot's color prior to having their attention drawn to it and their ability to identify the dot's color after attention had been drawn to it (and the dot subsequently masked) supports the claim that the dot's color before cueing was consciously perceived and quickly forgotten. This phenomenon is consistent with the diagnosis of amnesia, but it also is consistent with the diagnosis of blindness and thus does not militate on behalf of Wolfe's position. Since flicker-detection is an automatic mechanism of the visual system, we need not infer from the fact that the transient captured attention that subjects had any awareness of the dot before the cue (O'Regan & Noë, 2001, p. 947). Wolfe et al might be right in taking change blindness and inattentional blindness to be misnomers, but at this point it's unclear why one should side with them on this issue.

5 Accommodating (some of) the inner picture view

Despite the meager prospects for resolving whether unattended features are visually experienced, headway can be made in reconciling the introspective sense of phenomenological richness with subjects' failures to report unattended stimuli. Recalling the suggestion by Rees et al that the lack of further processing of the unattended words shown to their subjects may be attributable to the exceptionally high attentional demands of the detection task they were engaged in, the IB effects are demonstrated in the laboratory under conditions in which attention is focused on a single task and are discovered in everyday life in situations in which people are engrossed in some activity. It is natural to think that we don't normally walk around in states of such highly focused attention. Rather, our attention is usually much more widely distributed. In that case, "the sense we have of seeing the entire scene may derive from the many occasions in which our attention is distributed broadly or is divided and is not exclusively focused on a single object" (Mack, 2002, p.105).

This explanation of the robust sense of experienced visual detail fits with the 'visual stuff' account. Since attention is required to create spatially coherent and temporally durable representations out of visual stuff, spreading one's attention across the visual scene should have the effect of a more varied sampling of features being rendered into representations that are available to further stages of processing. Hence there is something importantly right about the everyday conception of visual experience. However, there still is room for CB and IB, even with our attention being widely allotted in ordinary situations. One simple way to acknowledge that we can still fall prey to CB or IB is that distributing attention across the visual scene does not mean that all of the available visual stuff is bound together in stable representations. There is only so much attention to go around, so it is likely that there will always be plenty of visual stuff that goes unattended. Furthermore, as attention gets spread out, we should expect the representations to become increasingly fragile, along both temporal and spatial dimensions. The fewer attentional resources devoted to each coherent object, the more coarse-grained their representation becomes and the more likely it becomes that non-significant features will come to be ignored. Additionally, the degradation of information pickup away from the fovea and the preferential concentration of processing resources on foveal information should result in lower-fidelity representations being generated by the direction of attention to visual stuff in parafoveal and peripheral regions, which should diminish the usefulness of such representations in further processing.

Mack's explanation has the virtues of meshing nicely with the compelling visual stuff approach and accommodating much of what we'd be led to say about our visual experiences on the basis of introspection. However, it does have its limitations, particularly regarding how many items we can attend to in each view. Largely on the basis of the limits on attention, Rensink (2000b, p.1475) claims that "there is usually only one object in play at any one time … tasks involving more than a few objects can be handled by rapidly switching attention between the objects”; cf. also Ballard (1991) and Findlay & Gilchrist (2003, pp. 145-149). Of course, the “one object in play” need not neatly correspond to a single object of the physical scene in view, as it might be a creation of the visual system: an amalgam of diverse features from across the scene that have been grouped together by attention. Rensink suggests that the limits on attention revealed by research on change blindness and visual search should lead us to consider that attended visual stuff is pooled together in a single <u>nexus</u>, which is to be identified with the “one object in play” (Rensink, 2000b, p.1483). Such an entity might be advantageous in handling the binding problem; i.e., the problem of how features processed by many different visual subsystems get correctly mapped onto objects (Palmer, 1999, p.556-557). Instead of dealing with the challenge of binding together a motley bunch of features into multiple coherent objects across the visual field, the visual system has a much more manageable task of correctly mapping features to one object.

Note that on this proposal, according to which the single nexus is phenomenologically prominent, the phenomenological status of mere visual stuff is still left unsettled. Nonetheless, we still have a satisfying explanation of much, if not all, of the seductive power of the inner picture and persistent picture views. For one thing, the ability to distribute attention leads to features from across the visual scene figuring in the coherent representation that is so central to—if not exhausting of—experience, even if the upper limit on how widely attention can be spread is rather modest. Of course, as attention is distributed, we lose our ability to notice things, as the resulting representation will not include specific detail for all features of the independent physical objects that are attended to. The oft-noted ease and quickness with which one can shift attention to other parts of the same visual scene also fosters the belief that each view is itself highly detailed; this point is made in several of the papers included in Noë (2002). So, in everyday experience, unlike the special situations in which IB and CB are most apparent, detail from across the visual scene is typically captured by one's attention and one can always quickly turn attention to other features that may not have been included in the coherent representation from just a moment ago. In that case, it is quite easy to see why subjects would so readily come to the inner and persistent picture views. For all they can tell on the basis of untutored (or mistutored, perhaps, in the case of those familiar with traditional theories of vision) introspection, all the detail of the visual scene is right there in their experience. The relatively slow speed at which visual information is processed into conscious experience and the relatively fast speed at which saccades carry the eye across the visual scene together can readily lead to introspectively confusing lots of detail being easily accessible for lots of detail being actually present (Ballard, 2002; Hayhoe et al, 1998). The time needed to generate visual experience also supports the caution urged earlier regarding Wolfe's example of the subject who is electroconvulsively shocked immediately after being shown a picture. Another potential contributor to the appeal of the inner picture view is visually acquired information (e.g., the gist and spatial layout of the scene) that is consciously present and available to guide the subject's visual activity, but that is not accompanied by visual qualia. Not realizing (or perhaps unable to realize, due to limits on introspection) that there is more to vision than visual experience, subjects might confound such information with that which is present in their visual experience, further reinforcing their sense of rich and coherent experience of the entire visual scene.

6 Vision and action

Gerd Gigerenzer and his colleagues have argued that our decision-making processes do not involve robust informational models of the world. Rather, given the real-world limitations confronting us, we often make use of 'fast and frugal' heuristics that exploit relatively stable features of our environment in order to quickly make decisions in complex, underinformed situations (Gigerenzer et al, 1999; cf. also Simon 1982 and 1991). On this view, our decisions are far from ideally rational and our behavior has been designed with an eye toward producing results that are 'good enough' at satisfying our goals. The points that motivate this conclusion about decision-making also raise the question of whether our perceptual systems make widespread use of heuristics (Gigerenzer et al, 1999, p.363), which would fit well with the findings of CB and IB research. Considering the computational, metabolic, and architectural burdens that come along with constructing high-fidelity representations of the visual scene at every moment of view, it is natural to suppose that our visual systems are the product of engineering that emphasizes satisficing in acquiring information about the environment on the basis of phototransduction, at the expense of carrying out that task in an ideal fashion (Ramachandran, 1990). Furthermore, cumulative selection can create slippage between the model of the world the visual system creates and the world itself. As trial-and-error processes acted on the perceptual systems of our antecedents, maintaining those variants that worked and discarding the failures, new engineering problems could be solved only by slight tinkering with the materials and structures already on-hand. This places significant limits on the fidelity of information that can be visually acquired and made available for further processing.

Despite all that can be said on behalf of the idea that informational impoverishment is prevalent, many of our most important tasks are informationally demanding. In circumstances in which useful information is easy to come by and error has a low cost, we can get by with decisions that take the first acceptable choice out of a short list of options, because it does not pay (and it may be counter-productive) to consider a broad range of options. Other situations are not like this. Predation is a good example, as the stakes are high and success is strongly sensitive to the amount and quality of information on hand. In established predatory relations, prey are typically equipped with appropriate defenses; e.g., armaments, camouflage, heightened alertness. These defenses make it harder for a predator to succeed in the hunt, because they can either make relevant information difficult to acquire or make the costs of poorly planned actions punishing. Predators come outfitted in ways that create similar problems for their prey. Besides the immediate risks of predatory conflicts, managing opportunity costs is important, as valuable resources can be wasted by fleeing prime foraging territory due to a false-positive in predator detection or by engaging in an unrewarding hunt. Poorly informed decisions will not do here, although the ecological profiles of both hunter and hunted can make a big difference to what counts as poor. The parties involved benefit from making it harder for their opponent to acquire relevant information (e.g., misinforming, giving away little veridical information), increasing their own abilities to extract veridical information about the scenario, and capitalizing on their opponent's limitations and failures. As the informational loads on appropriate behavior and the costs of inappropriate behavior both increase, it becomes less likely that 'fast and frugal' heuristics will suit a creature's needs and more likely that creatures employing informationally rich models of the world will do better than those that don't (Sterelny, 2003).

Most complex creatures' lives are littered with situations that are more like predation than they are like buying a new pair of shoes. Opponents are <u>everywhere</u>. Predators, prey, potential mates, competitors for potential mates, and competitors for limited material resources are just a few examples that come to mind. These rivals benefit by making it more difficult for their adversaries to acquire information that leads to successful action and by making the costs of failure high. There just isn't much of a second-prize in predatory conflicts or in mating competitions. Creatures that have a better model of the actual state of the world and who are better able to exploit that model will consistently win out over those that do not. Thus it's reasonable to suppose that our mechanisms for extracting information from the environment would be designed in such a way that they meet the high informational loads needed for us to act appropriately in such critical tasks. In that case, it's difficult to see how 'fast and frugal' heuristics could be the entire story of our mental lives and it may be implausible to suppose that they constitute a significant part of our mental lives.

So, granting that CB and IB establish something significant about the informational limitations of everyday visual experience, how do we meet the high informational loads of the many crucial tasks in which vision plays a key role? There are two related responses to this question. The first is a continuation of a point made in the previous section. Although any moment of view has a quite limited content that is available for further use, plenty of information is easily and quickly available to the perceiver, as it's right there in her environment. The perceiver need only actively go out and gather it. This gathering can include saccading across the visual field, tilting one's head, moving to a different vantage point, and many other activities. This is closely linked to Kevin O'Regan's (1992) now-familiar idea of the external world as a memory store. Instead of going to the trouble of building up a detailed representation of the scene before the eyes, our visual systems save their most valuable resources for recovery and processing of those features that best suit the perceiver's immediate needs, while relying on assumptions about the relative stability of the external environment to obviate the need to create an exhaustive, resource-intensive representation for every view. In doing this, the visual system can also accommodate the acquisition of information about some features that have not been attended to, but which can unconsciously influence subsequent beliefs or actions, and which can draw the perceiver's attention, thereby bringing them to the forefront of consciousness.

This conception of seeing as an active process is gaining a number of supporters and is proving successful in addressing questions that stymied previous efforts, but it is important to recognize how strongly it runs against the current of established vision science.[10] For the longest time, vision was studied as a passive process, one with the goal of recovering from confounding retinal input a faithful mental representation of the stimulus scene. This approach feeds (and is fed by) the inner picture view and treats vision more-or-less in isolation, as a faculty largely independent of the cognitive and behavioral capacities of the sighted creature; 'pure vision', as Churchland et al (1994) termed it. On the active understanding, seeing is something that a creature does and that can involve a whole range of bodily movements in an effort to acquire information through phototransduction. In jettisoning the passive conception of vision, one need not dispense altogether with the central feature in passive theories, mental representations, although some have done so; cf. O'Regan & Noë (2001). At a minimum, active approaches emphasize that visual processing is driven by the interests and needs of the perceiver, with what is of interest or need being both subject to change quite quickly and fixated by the subject by means of exploration of her surroundings. There is room here for mental representations, although those representations will not be 'snapshots' of the stimulus scene. The active approach gives a start to understanding how it is that vision makes its valuable contribution to meeting the informational loads of the tasks we engage in. Doing away with the inner-picture view does not entail that vision is informationally poor, it simply means that perceivers are making a trade-off in how they visually acquire the information they need. Instead of paying exorbitant costs to

[10] Findlay & Gilchrist (2003) offer a thorough discussion of the motivations for and successes of the active vision research program.

generate and store high quality representations of the entire scene before the eyes, the visual system thriftily makes the most of its resources by extending information acquisition through time (although still at quite short intervals) and using the body to which it is harnessed to get the eyes—the fovea in particular (Findlay & Gilchrist, 2003, pp.5, 162-167)—in a position to gather the information needed at that moment.

This observation also suggests an interesting point of difference between Wolfe's understanding of the role of preattentive visual stuff and the role of visual stuff on an active conception of vision. On Wolfe's Guided Search model, there is a systematic scanning of feature maps (roughly, independent representations of basic visual features from across the visual field that are generated preattentively and in parallel) by an internal 'spotlight' of attention to identify relevant targets, such as unusual or sought-after items (Wolfe, 1994, pp.204-209). This account does not include a role for eye movements, nor does it accommodate the decrease in visual processing as eccentricity from the fovea increases (Findlay & Gilchrist, 2003, p.109). The latter point is especially significant, considering that attempts at detection in circumstances of impoverished visual resolution are likely to be severely compromised. Such circumstances are the norm outside the fovea and near parafovea. Within the framework of active vision, preattentively processed information (e.g., gist, spatial layout) can be used to guide eye movements to facilitate identification of relevant items in the scene by foveating them. Making those items the target of such movements results in them achieving higher acuity and receiving a greater concentration of processing resources, which are beneficial, if not critical, to success in a great many real-world tasks.

The second response to the question posed above is that a proper understanding of vision's role in facilitating ecologically appropriate behavior requires recognizing that visual experience is not the whole of vision. Additionally, when inquiring into the nature of visual experience, some of the most important questions we might ask have to do with what need there is for us to have the kinds of experiences we do. These points play a central role in the work of Mel Goodale and his colleagues (Milner & Goodale, 1995; Goodale & Humphrey, 1998; Goodale & Milner, 2004). The phenomenology of our visual experience is so gripping that it is easy to lose track of the fact that we only have it because it must have made some positive contribution to the differential reproductive success of our ancestors. Apart from those inclined to epiphenomenalism, one should expect that we would not be subject to conscious perceptual states unless those states were exposed to selection.[11] Since hidden, subjective states considered in isolation are not good candidates for being direct objects of selection, it figures that conscious experience was engineered by selection on the basis of effects it has on the kinds of behavior we engage in; it is rather easy to understand how behavior is selected for or against. We also have a fairly good idea of how conscious experience could affect behavior; e.g., the creature over yonder looks to be moving toward me with great speed and, based on a comparison of its visual characteristics to what I have learned from past interactions, I recognize it as a predator, so I take off for a known hiding place. Much of what we do seems to be guided by the way the world appears in visual experience. It turns out, however, that this understanding of the connection between vision and action, while not incorrect, is quite incomplete. In fact, not only can unattended visual stuff influence our behavior (as noted earlier), but Goodale et al suggest that it is <u>unconscious</u> visual information that plays the dominant role in closely guiding our actions.

[11] In focusing on selection, I am not claiming that other factors (e.g., evolutionary drift) are irrelevant to the evolution of conscious experience. It is compelling to think, though, that selection has the most important role in the evolution of our perceptual equipment.

Goodale et al have studied subjects with visual deficits and concluded that instead of there being a single, general-purpose visual system that both creates an 'inner model' of the external world and guides action, there are two interacting visual systems (Goodale & Milner, 2004, pp.48-49, 101-103). One of these systems consists of projections from primary visual cortex (V1) to the posterior parietal cortex and the other system consists of projections from V1 to the inferotemporal cortex; the former is the dorsal stream and the latter is the ventral stream. Goodale et al suggest that "the ventral stream plays the major role in constructing the perceptual representation of the world and the objects in it, while the dorsal stream mediates the visual control of actions directed at those objects" (Goodale & Humphrey, 1998, p.186). Activity in the ventral stream is associated with conscious visual perception and is a close fit for the commonsense understanding of the visual system as creating a model of the external world. Activity in the dorsal stream is not consciously accessible and it feeds into motor activity by providing information about, for example, form, size, and location that can be used to guide activities such as reaching, ducking, and leaping. The separate streams are not processing different information. Rather, each stream makes its own particular use of the information projected from V1. Conscious visual information is used to select targets of action, while unconscious visuomotor information is used to finely guide motor behavior.[12] One way to think about the difference is that ventral stream activity is 'epistemically charged' while dorsal stream

---

[12] The division between the two systems may not be as tidy as Milner & Goodale suppose. Jacob & Jeannerod (2003, Chapter 4) discuss evidence that conscious visual information can directly affect, although minimally, motor activity. As they note, this evidence does not undermine either the dual systems hypothesis or the related claim that the strong introspective sense of experienced-based control of action is mistaken; cf. also Clark (2001). It simply militates on behalf of a weaker form of the dual systems model than that advocated by Milner & Goodale. Kanwisher (2001) presents an (admittedly speculative) challenge to the notion that only ventral stream activity is involved in visual experience, arguing that conscious visual perception is grounded in binding and global availability of activity in both streams.

activity deals with body-centered representations of more-or-less objective features of the world like shape, motion, and location.

As Goodale et al note, the specialized functions of the two visual streams and studies of visual deficits in non-human animals reinforce the point that brains have been shaped to produce action that suits the ecological needs of the creature in whose skull the brain rests (Goodale & Milner, 2004, p.55). It is likely that the action-guiding dorsal stream is evolutionarily quite old and that the visual experience generating ventral system is a much more recent development. Creatures occupying relatively simple ecological niches can benefit from a system that drives hardwired reactions to crude information acquired through phototransduction. No inner model of the world is needed in such cases. Such creatures were among our ancestors and elements of our visuomotor systems have certainly been inherited from them. Those of our ancestors that were more flexible in their reactions to sensory inputs had an advantage in creating and occupying new, increasingly complex ecological niches (with respect to, inter alia, reproduction, energy acquisition, and energy use), thus increasing their chances of reproduction and survival. The value of an inner model of the world in aiding such flexibility is readily seen. Once proto-forms of systems that enabled flexible behavior were in place, an arms-race took off in which adaptive advantage was conferred upon those creatures who were in a better position to engage in novel (and appropriate) behavior in response to environmental conditions. As variants of those systems spread, a feedback loop developed in which creatures were both shaping and being shaped by their environments, which only increased the value of flexible behavior.

It sounds banal to say that vision is closely connected to action and that visual experience does not exist for our phenomenological pleasure. However, I submit that the strong pull felt by some to decide whether unattended visual stuff is consciously seen is evidence that the inner picture view and a misunderstanding of the ultimate goal of vision as producing visual experience both still exert a powerful influence on much theorizing about vision. From the standpoint of an approach grounded in the close connection between vision and action, the question of the phenomenological status of unattended visual stuff should be more of a sideline curiosity than a pressing question. We know that unattended visual stuff doesn't directly contribute to intentional action and decision making (the subjects' reports and behavior give us that), and we know that subjects have not found that visual stuff to be of interest (otherwise it would have received attention and been coherently represented). We also know, however, that unattended visual stimuli can have an effect of some sort on perceivers' subsequent activity. It is the further effects of unattended visual stimuli that are in need of careful study as part of active vision, since those effects fit with the guiding principles of the entire approach. They do something for us as far as affecting our interactions with others and our environment (Simons, 1999, p.168), and they also play a role in directing our attention around the visual scene to gather more information (O'Regan & Noë, 2001, p.947; Rensink, 2000b, p.1476). Besides characterizing the effects of unattended visual stimuli, understanding whether those effects are attributable to 'informational leakiness' within the visual system (which might explain the effects of unattended visual stimuli either as the result of information leaked by one subsystem being co-opted by other subsystems for various purposes or as a mere side-effect) or a more streamlined design (that reflects a need to reduce the burdens on conscious-level decision-making and behavior-guidance) is also an important matter that can be investigated without worrying about whether unattended visual stimuli are experienced. This is <u>not</u> to downgrade the scientific status of or eliminate altogether phenomenal experience. Rather, it is to say that while visual experiences mediate some inputs and outputs, other visual phenomena also mediate inputs and outputs. Those other visual phenomena may be either unconscious or conscious in a way that is not qualitatively visual. The key point here is that as far as our ecological success is concerned, there is much more to vision than visual experience, despite the strong introspective sense that the latter exhausts the former.

To further clarify, the insignificance of the phenomenological status of unattended visual stuff is not at all an eliminativist claim. Rather, it is a methodological point driven by the currently most promising scientific conception of vision, active vision. Active vision research is centered around the role of eye movements (including movements of the rest of the body to move the eye) in sampling the environment and the ways in which the sampled information is processed to—consciously and unconsciously—support the activities of the perceiving creature. In its current state, this approach has nothing to say that would help us decide whether our visual experiences are anywhere near as heavy with unattended detail as they introspectively seem to be. Furthermore, research on active vision can profitably continue without addressing that issue. As was discussed in section 5 and earlier in this section, among the many advantages of this re-thinking of the nature of vision is that it provides some useful explanation for why our experiences would seem so richly detailed, even if that impression is illusory. Thus we are able to meet one of Wolfe's stated motivations for preferring the diagnosis of amnesia over the diagnosis of blindness—to reconcile the introspective sense of rich phenomenological detail with the results of research on the role of attention in vision (Wolfe et al, 2000, p.694)—without having to force a decision between the two alternatives.

Of course, future research may offer better guidance on the issue by underwriting certain inferences from the empirical data on vision, memory, and attention. For example, research on two phenomena closely connected to attention and capacity limitations within the visual system—the attentional blink (AB) and repetition blindness (RB)—might help reveal some of the connections between attention, priming, and visual experience.[13] AB is an impairment in detecting or reporting a second target in a rapid serial visual presentation (RSVP) when it follows the first within approximately half a second. While subjects cannot report the second stimulus in AB, the unreported stimulus can have a priming effect on subsequent stimuli, such as enhancing detection of subsequent stimuli with relevantly similar characteristics. RB is an impairment in detecting or reporting repetitive instances of a stimulus in RSVP, even when a different stimulus is presented between the repeated instances of the same kind of stimulus. The strength of the RB effect is highly dependent on characteristics of the interval between the first and repeat instances of the stimulus; viz., the number of different stimuli presented in the interval, the time of that interval (Kanwisher, 1987). Research on these phenomena might eventually be able to tell us how much visual processing the unreported stimuli receive, what capacity limitations and mechanisms are relevant to the inability to report, why priming should occur in some cases but inhibition in others, and so forth. It is not unrealistic to expect that such discoveries would make some claims about the phenomenological status of unattended stimuli more plausible than others, and would help us better understand why we are conscious of some visual stimuli, primed or otherwise affected by certain stimuli that (at a minimum) are not available for report, and completely oblivious to and unaffected by some stimuli.

Right now, however, it is reasonable to maintain that substantive claims about the phenomenological status of unattended visual stimuli are highly speculative. While I don't claim that such a speculative enterprise has no value (after all, we seem to be very curious about the

[13] Chun (1997) studies the relationship between these two phenomena. Kanwisher (1987) introduced RB, and Shapiro et al (1997) investigate what happens to the undetected target in AB.

possibility that our experience is significantly different than we take it to be), I do want to press the epistemological point that whatever justification that enterprise may have must come from outside the active vision research program as it currently stands. This is significant because without proper justification, it is unclear why one would be motivated to make claims about whether unattended visual stuff is consciously seen, much less how we are to evaluate such claims.

7 Conclusion

In sum, the position advanced here is that what prima facie appears to be a deeply interesting question regarding the phenomenology of visual experience currently does not admit of a definitive answer, but that is actually of much less consequence than we might initially suppose. Unattended visual stuff might very well have a qualitative presence in visual experience, but then again, it might not. Since we do not have direct access to the minds of experimental subjects, we have to rely on the commitments of our best scientific theories to guide the inferences we make to mental phenomena on the basis of analysis of experimental data. As near as I can tell, there is no decisive evidence that forces us to take a side on this issue and nothing in the active vision research program right now supports any inferences about the presence or absence of unattended visual stuff in visual experience. This agnosticism is provisional. As vision research advances, any number of developments might lead us to be able to take a position on the issue, including that the matter cannot be settled at all by empirical investigation. Regardless, the difficult nature of this issue should not be taken as a weakness in the scientific study of vision, as the realizations that vision is more than visual experience and that there is much to be learned about the effects of unattended visual stimuli on perceptual, cognitive and behavioral activity will both be of great benefit to investigation of the structure of the mind.

8 References

Aglioti, S., J.F. DeSouza, & M.A. Goodale. 1995. "Size-contrast illusions deceive the eye but not the hand." Current Biology, *5,* 679-685.

Ballard, D. 1991. "Animate Vision." Artificial Intelligence, 48, 57-86.

_____. 2002. "Our Perception of the World Has To Be an Illusion." In Noë 2002.

Ballard, D.H., M.H. Hayhoe, P.K. Pook, & R.P.N. Rao. 1997. "Deictic codes for the embodiment of cognition." Behavioral and Brain Sciences, 20, 723-767.

Block, N. 1995. "On a confusion about a function of consciousness." Behavioral and Brain Sciences, 18, 227-247.

Brawn, P., R. Snowden, & J. Wolfe. 1999. "The minimal conditions for 'change blindness': What is hides what was." Investigative Ophthalmology and Visual Science, S49.

Chun, M. 1997. "Types and Tokens in Visual Processing: A Double Dissociation Between the Attentional Blink and Repetition Blindness." Journal of experimental psychology: Human perception and performance, 213, 738-755.

Churchland, P.S., V.S. Ramachandran, & T.J. Sejnowski. 1994. "A critique of pure vision." In C. Koch & J. Davis (eds.) Large scale neuronal theories of the brain. Cambridge, MA: MIT Press.

Clark, A. 2001. "Visual Experience and Motor Action: Are the Bonds Too Tight?" Philosophical Review, 110, 495-519.

Dennett, D. 1988. "Quining Qualia." In A. Marcel & E. Bisiach (eds.) Consciousness and Contemporary Science. Oxford: Oxford University Press.

_________. 1991. Consciousness Explained. Boston: Little, Brown, & Co.

_________. 2002. "How Could I Be Wrong? How Wrong Could I Be?" In Noë 2002.

Diamond, J. 1992. "The red flag of optimality." Nature, 355, 204-206.

_____. 1993. "Quantitative design of life." Nature, 366, 405-406.

Feldman, J. 1985. "Four frames suffice: a provisional model of vision and space." Behavioral and Brain Sciences, 8, 265-289.

Findlay, J. & I. Gilchrist. 2003. Active Vision. Oxford: Oxford University Press.

Germeys, F., P. De Graef, K. Verfaillie. 2002. "Transaccadic Perception of Saccade Target and Flanker Objects." Journal of Experimental Psychology: Human Perception and Performance, 28, 868-883.

Gigerenzer, G., P. Todd, & the ABC Research Group. 1999. Simple Heuristics that Make Us Smart. New York: Oxford University Press.

Goodale, M. & G.K. Humphrey. 1998. "The objects of action and perception," Cognition, 67, 181-207.

Goodale, M. & A.D. Milner. 2004. Sight Unseen. New York: Oxford University Press.

Haines, R. 1991. "A breakdown in simultaneous information processing." In L. Stark & G. Obrecht (eds.), 4th international Symposium on Presbyopia. New York: Plenum.

Hayhoe, M, D. Bensinger, & D. Ballard. 1998. "Task Constraints in Working Memory." Vision

Research, 38, 125-137.
Hollingworth, A. In press. "Visual Memory for Natural Scenes: Evidence from Change Detection and Visual Search." Visual Cognition.
Irwin, D. 1991. "Information integration across saccadic eye movements." Cognitive Psychology, 23, 420-456.
_______. 1996. "Integrating information across saccadic eye movements." Current Directions in Psychological Science, 5, 94-100.
Irwin, D., S. Yantis, & J. Jonedies. 1983. Evidence against visual integration across saccadic eye movements." Perception and Psychophysics, 34, 49-57.
Kanwisher, N. 1987. "Repetition Blindness: Type recognition without token individuation." Cognition, 27, 117-143.
_______. 2001. "Neural events and perceptual awareness." Cognition, 79, 89-113.
Jacob, P. & M. Jeannerod. 2003. Ways of Seeing. Oxford: Oxford University Press.
Levin, D. 2002. "Change blindness blindness as visual metacognition." In Noë 2002.
Levin, D., N. Momen, S. Drivdahl, & D. Simons. 2000. "Change blindness blindness: The metacognitive error of overestimating change-detection ability." Visual Cognition, 7, 397-412
Mack. A. 2002. "Is the Visual World a Grand Illusion? A Response." In Noë 2002.
_______. 2003. "Inattentional blindness: looking without seeing." Current Directions in Psychological Science, 12, 180-184.
Mack, A. & I. Rock. 1998a. Inattentional Blindness. Cambridge, MA: MIT Press.
________. 1998b. "Inattentional blindness: perception without attention." In R.D. Wright (Ed.), Visual Attention. New York: Oxford University Press.
Milner, A.D. & M.A. Goodale. 1995. The Visual Brain in Action. Oxford: Oxford University Press.
Moore, C. & H. Egeth. 1997. "Perception without attention: Evidence of grouping under conditions of inattention." Journal of Experimental Psychology: Human Perception and Performance, 23, 339-352.
Neisser, U. 1979. "The control of information pickup in selective looking." In A.D. Pick (ed.) Perception and its Development: A Tribute to Eleanor J. Gibson. Hillsdale, NJ: Lawrence Erlbaum Associates.
Noë, A. (ed.) 2002. Is the Visual World a Grand Illusion? Special issue of Journal of Consciousness Studies.
O'Regan, J. K. 1992. "Solving the 'real' mysteries of visual perception: the world as an outside memory." Canadian Journal of Psychology, 46, 461-488.
O'Regan, J. K., R.A. Rensink, & J.J. Clark. 1999. "Change-blindness as a result of 'mudsplashes'." Nature, 398, 34.
O'Regan, J. K. & A. Noë. 2001. "A sensorimotor account of vision and visual consciousness." Behavioral and Brain Sciences, 24, 955-975.
Palmer, S. 1999. Vision Science: From Photons to Phenomenology. Cambridge, MA: MIT Press.
Pollatsek, A., K. Rayner, W. Collins. 1984. "Integrating pictorial information across eye movements." Journal of Experimental Psychology: General, 113, 426-442.
Ramachandran, V.S. 1990. "Interactions between motion, depth, color, and form: the utilitarian theory of perception." In C. Blakemore (ed.) Vision: Coding and Efficiency. New York: Cambridge University Press.
Rees, G., C. Russell, C. Frith, J. Driver. 1999. "Inattentional Amnesia for Fixed but Ignored

Words." Science, 286, 2504-2507.
Rensink, R. 2000a. "When good perceivers go bad: change blindness, inattentional blindness, and visual experience," Psyche, 6.
_________. 2000b. "Seeing, sensing, and scrutinizing," Vision Research, 40: 1469-1487.
Rensink, R. A., J.K. O'Regan, & J.J. Clark. 1997. "To see or not to see: the need for attention to perceive changes in scenes," Psychological Science, 8, 368-373.
Rojer, A. & E. Schwartz, E. 1990. "Design considerations for a space variant visual sensor with complex logarithmic geometry." In Proceedings of the International Conference on Pattern Recognition, Philadelphia, PA.
Shapiro, K. L., J. Driver, R. Ward, & R. Sorensen.1997. "Priming from the attentional blink: A failure to extract visual tokens but not visual types." Psychological Science, 8 , 95-100.
Simon, H.A. 1982. Models of Bounded Rationality. Cambridge, MA: MIT Press.
_________. 1991. "Cognitive architectures and rational analysis: Comment." In K. Van Lehn (ed.) Architectures for Intelligence. Hillsdale, NJ: Erlbaum.
Simons, D.J. 1996. "In sight, out of mind: when object representations fail," Psychological Science, 7, 301-305.
________. 1999. "To see but not to see," Journal of Mathematical Psychology, 43, 165-171.
________. 2000. "Attentional capture and inattentional blindness," Trends in Cognitive Sciences, 4, 147-155
Simons, D. J. & C.F. Chabris. 1999. "Gorillas in our midst: sustained inattentional blindness for dynamic events," Perception, 28, 1059-1074.
Sterelny, K. 2003. Thought in a Hostile World. Malden, MA: Blackwell Publishing.
Trehub, A. 1991. The Cognitive Brain. Cambridge, MA: MIT Press.
Treisman, A. 1986. "Properties, parts, and objects." In. K. Boff, L. Kaufmann, & J. Thomas (eds.) Handbook of Human Perception and Performance. New York: Wiley & Sons.
Tye, M. 1995. Ten Problems of Consciousness. Cambridge, MA: MIT Press.
______. 2000. Consciousness, Color, and Content. Cambridge, MA: MIT Press.
Weiskrantz, L. 1990. Blindsight: A case study and implications. Oxford: Oxford University Press.
________. 1997. Consciousness lost and found: a neurophysiological exploration. Oxford: Oxford University Press.
Wolfe, J. 1994. "Guided Search 2.0: A revised model of visual search." Psychonomic Bulletin and Review, 1, 202-238.
______. 1998 " Visual Memory: What do you know about what you saw?" Current Biology, 8, R303-R304.
_______. 1999. "Inattentional Amnesia." In V. Coltheart (ed.) Fleeting Memories. Cambridge, MA: MIT Press.
Wolfe, J. & S. Bennett. 1997. "Preattentive Object Files: Shapeless Bundles of Basic Features," Vision Research, 37, 25-43.
Wolfe, J., N. Klemplen, K. Dahlen. 2000. "Postattentive Vision." Journal of Experimental Psychology: Human Perception and Performance, 26, 693-716.
Wright, W. 2005. "Distracted Drivers and Unattended Experience." Synthese.
________. In press. "Explanation and the hard problem." Philosophical Studies.