When is Perception Conscious?[1]

Jesse J. Prinz

Once upon a time, people thought that all perception was conscious. Indeed, it was widely believed that all mental states are conscious, so the problem of explaining consciousness collapses into the problem of explaining mentality. But things have changed. Most people now believe that a lot goes on unconsciously. Indeed, some people believe that mental states that are not perceptual in nature are never conscious. That's a matter of controversy. Less controversial is the claim that perceptual states are conscious some of the time, but not all of the time. This raises a question. When are perceptual states conscious? A theory of consciousness is, in large part, an answer to that question. In this chapter, I will offer a few critical remarks on one answer that has been popular in philosophy, and then I will offer a defense of another answer that has emerged out of cognitive science. To avoid undue suspense, the answer that I favor is that perceptual states become conscious when and only when the perceiver is attending.

1. The Conscious/Unconscious Divide

1.1 Unconscious Perception

In case there is any doubt, it will help to briefly review some reasons for thinking that perception can occur unconsciously. Evidence comes in various forms. I will focus here and throughout on the visual modality, because vision has been most thoroughly studied. One line of evidence comes from brain injury. Consider blindsight. People with lesions in the primary visual cortex (V1) are said to be cortically blind: they claim not to see objects presented in the affected portion of the visual field. In cases of global V1 damage, these individuals cannot spontaneously detect any changes made available to the eyes, even if they are quite dramatic. In one recent case report, for example, a subject with global V1 damage could not tell when the bright lights in a room were turned on and off (Hamm et al. 2003). But, in some cases of cortical blindness, there is residual visual ability: when asked to guess where an object is located, these individuals do quite well, even though they think they do not see the objects they are pointing at (Weiskrantz 1986).

Blindsight, as it is called, is a case of visual perception without conscious experience of what is being perceived. It counts as perception because people with blindsight are receiving information through a sensory transducer (the ideas) and responding to it. Some people with blindsight even retain some ability to recognize objects in the blind field. De Gelder et al. (2005) presented a blindsight subject with

[1] In writing this chapter, I had helpful discussions with Felipe De Brigard, Dave Chalmers, Anya Farennikova (who also corrected my English), Bill Lycan, and Bence Nanay (without whom this chapter would not exist). I also benefitted considerably from audience feedback at the 2008 SPAWN conference at Syracuse University.

emotional facial expressions and these influenced subsequent emotion tasks, even though he showed no sign of having consciously perceived the faces. Hamm et al. (2003) were able to condition a person with blindsight to have an aversive response to pictures of airplanes even though he reported no conscious awareness of the airplane pictures or control pictures when they were being presented.

Unconscious perception can also be demonstrated in individuals with intact brains. The most widely practiced method for doing this is backwards masking. Subjects are briefly presented with a visual stimulus, which is followed by a second stimulus (the "mask") that prevents them from consciously seeing the first. If the first stimulus is presented for a long enough duration (say, 200 milliseconds) subjects can consciously see it, identify it, and spontaneously report on their experience. At shorter durations (say, 50 milliseconds), the stimulus cannot be identified, but subjects are confident that they saw *something* flash before their eyes. At even shorter durations (say, 16 milliseconds) subjects do not report seeing a first stimulus, prior to the mask. Indeed, if they are given a sequence of trials in which some show a stimulus before the mask and others don't, they are at chance in guessing which kind of trial they are on. They have no idea that there has been a first stimulus. Nevertheless, the masked stimuli are perceived. They influence information processing in measurable ways.

Unconscious perception is shown to occur in masking studies using a family of methods called priming. A masked word can influence answers given on a word completion task or facilitate memory access to associated words. A masked color can facilitate detection of the same color. In some cases, the influence is dramatic and wide ranging. For example, Winkielman et al. (2005) presented subjects with masked facial expressions, which were either angry or happy, and then asked them to evaluate a soft drink. As after unconsciously seeing happy faces, as compared to angry faces, subjects said the soft drink was more delicious, they were willing to pay more for it, and they poured more of it into their cups. Clearly they were recognizing the emotional valence of the face, but they were doing so unconsciously.

1.2 Which Perceptions Can Be Conscious?

Before we can figure out how perceptions become conscious, we need to figure out which perceptions are even candidates for consciousness. Sensory systems contain many subsystems and some of these never produce conscious representations. We must figure out *where* in perceptual processing consciousness arises, and then figure *when* representations in these privileged subsystems become conscious.

Like all sensory systems, the visual system is hierarchical (Marr 1982). It begins with cell populations that operate relatively independently of each other and respond to very discrete features of a stimulus, such as a small edge or bit of color. At this level (associated with V1), the visual system is not very interested in how these edges and colors hang together, or how foreground differs from background. This is low-level vision. It corresponds to what Marr calls "the primal sketch."

At the next stage of processing (the intermediate-level), edges are connected together, to form coherent, bounded contours. In addition, figure is separated from ground, and binocular information is more thoroughly integrated to reveal depth.

Marr calls this the 2.5 dimension sketch; it is 2D like a painting, but also includes depth, like a 3D movie, hence the extra half dimension. Neuroscientists associate this level with a family of structures in extra-striate cortex, labeled V2-V5. In these structures, the visual system shows sensitivity to global context effects: color sensitive cells fire in a way that is sensitive to the color in other areas, spanning over distances that are larger than the receptive fields of cells in early vision. Illusory contours are also registered, completing shapes that are only partially present in the stimulus.

The intermediate level is not the last stage of visual processing. Representations there are very specific to vantage point. Objects are represented form a specific point of view and they occupy a specific location and size in the visual field. This level of specificity is not ideal for object recognition. To recognize objects, it is useful to abstract away from specific vantage points and discern the underlying structure of objects (what Marr calls 3 dimensional models). Cells in high-level visual areas (in inferior temporal cortex) do just that. They abstract away from location, visual size, and, to a considerable degree, viewing angle. In some sense, cells at these different stages are representing different things (a vantage-point specific representation encodes different features than a vantage-point invariant representation), but the crucial difference is in the degree of abstraction (the same kinds of stimulus dimensions are represented in ways that are less and less like the pattern of stimulation on the retina).

It was against the background of this picture, that Ray Jackendoff (1987) put forward his Intermediate-Level Theory of Consciousness. Of the three levels just described, it is the intermediate level that best corresponds to what people report in conscious experience. If I see a walrus in the zoo, I experience it as a bounded whole, separate from its environment, occupying a specific region of space, and oriented from a specific point of view. In this respect, visual experiences are a little bit like realist paintings. They are not like Picasso's cubist portraits, simultaneously representing multiple perspectives, as we find in high-level vision, and nor are they like Seurat's pointillist paintings viewed from an inch away, obscuring the boundaries and decomposing the surfaces.

I have argued elsewhere, on neurobiological grounds, that Jackendoff's conjecture is right (Prinz 2000; see also Koch and Braun 1996). Visual consciousness arises at an intermediate level of processing, and other visual representations are always unconscious. We know, however, that mere activation of visual cells at the intermediate level is not sufficient for consciousness. That is because there is such a thing as unconscious perception. As we have seen, unconscious perception often involves the unconscious recognition of objects. Priming studies show that people can respond in content-specific ways to unconsciously perceived faces, spiders, airplanes, words, numbers, and everyday objects, such as desk lamps. This suggests that cases of unconscious vision involve processing through the entire visual hierarchy (see, e.g., Bar 2000). For many objects recognition is achieved only in high-level visual centers, and this, in turn, requires prior activation in intermediate- and low-levels of processing. This conclusion, confirmed by neuroimaging, shows that mere activation of intermediate-level visual representations is not sufficient for consciousness.

All this confirms the common observation that theories of consciousness need two parts: an account of qualitative character of conscious states (what kinds of things are we conscious of) and an account of how states become conscious. The intermediate-level hypothesis is a contribution to a theory of the first part. But what about the second part? We need a theory of how intermediate-level representations become conscious. What is the difference between conditions under which these representations are conscious and when they are not?

2. The AIR Theory

2.1 Attention As the Key to Consciousness

To empirically identify the factor that is responsible for conscious experience, one needs to determine what variable co-vary with consciousness: what variables allow consciousness when present and prevent consciousness when absent? We have already seen one variable that seems to make a difference: time. If a stimulus is presented too quickly, it won't be experienced. But time itself probably isn't the key. There are, as we will see in a moment, conditions under which stimuli presented for reasonably long durations do not get consciously experienced. It is more likely that time plays an indirect role. Stimuli that are presented for long enough durations become candidates for other processes to act on them, and one of these other processes is responsible for rendering perceptual states conscious. But which one? Happily, recent research points to one very plausible answer. There is increasing support for the conjecture that attention is the key. We are conscious when and only when we attend.

It's a commonplace that we can fail to see things when we're not paying attention. But platitudes are often false. To put this one to the test, vision veterans Mack and Rock (1998) developed an experimental paradigm in which subjects are presented with shapes or words while performing a concurrent task that demands a lot of attention. In particular, subjects are asked to determine which of two intersecting lines is longer—hard work if line lengths are similar in length. While doing this, an unexpected polygon, face, or word flashes in the center of the visual field for 200 milliseconds, which is normally well above the conscious threshold. In these studies, about 25% of observers fail to notice the unexpected stimulus. It's not just that they see something and can't identify it. They seem to see nothing at all. Immediately after the critical trial, subjects are a series of increasingly leading questions to see if they experienced anything other than the intersecting lines. In many cases, the answer is no; the subjects simply don't see an object presented in clear view.

This is called inattentional blindness, and it has now been replicated many times. The most dramatic case may be a study by Most et al. (2005). They have subjects count how many times animated black or white letters bounce against the side of a computer screen. While this is going on a red cross travels across the center of screen, taking a full five seconds to reach the other side. The cross differs from the other letters in color, luminance, shape, and trajectory, but it is still missed by 28% of the subjects.

Inattentional blindness suggests that attention is necessary for consciousness. This conjecture gains further support from other experimental paradigms. For example, Macdonald and Lavie (2008) have demonstrated what they call "load-induced blindness." Subjects are asked to search for a target letter in a group of letters, and, while performing this task, a meaningless shape is flashed. Subjects know what the shape will look like in advance but only 37% detect it when they are given six letters for the search task, which introduces a large attentional load. In addition, there is a phenomenon called the attentional blink, in which subjects fail to see the second of two targets in a rapid series of letters or numerals (Raymond et al. 1992). The first target captures attention briefly, and the second goes unseen if presented shortly thereafter. Researchers have recently discovered a related phenomenon called the emotional blink (Arnell et al. 2007). If subjects are asked to look for the name of a color in a series of words, they will fail to see it if it is displayed shortly after an emotionally charged word that captures attention, such as "orgasm."

These behavioral experiments on healthy subjects re-confirm something that has been well known in neurology for a long time: damage to attention centers of the brain disrupts consciousness. The most familiar case of this is visual neglect. People who sustain injuries to attention centers in right inferior parietal cortex seem to be blind in the contralesional region. They are oblivious to stimuli presented on the left (or the left side of object), despite the fact that their visual systems are intact and responsive to those stimuli. When asked to compare two objects that differ only on the left, subjects with neglect report that the objects are identical, even though there is sometimes evidence for unconscious processing of the invisible features (Marshall and Halligan 1988). Neglect can also affect perception in other sense modalities. Some patients become oblivious to the left sides of their bodies, even insisting that their left limbs are not their own. In auditory neglect, words heard in the left go unnoticed during dichotic listening tasks.

The evidence strongly suggests that attention is necessary for consciousness. Attention may also be sufficient. When we attend to things, they are consciously experienced. This can be demonstrated, for example, in cases of visual "pop-out" where objects capture attention. If you see an array of objects in which one differs from the others in color, shape, position, or orientation (Figure 1). Objects that pop-out this way are nearly impossible to ignore, they are experienced whether we are looking for them or not. In fact, pop-out can cause us to experience stimuli that are presented under conditions that ordinarily preclude conscious detection. This happens with stimuli that have personal significance. The classic example is the cocktail party effect: if your name is mentioned by someone across the room at a crowded cocktail party, you will hear it, even if the conversation there had been inaudible. Mack and Rock (1998) found a visual analogue for the cocktail party effect in their inattentional blindness studies. Subjects often fail to notice surprise stimuli when they are focusing attention on the central task, but they always notice the surprise stimulus if it is their own name. Mack et al. (2002) show that subjects also can't fail to notice their own names in attention blink and attentional load paradigms.

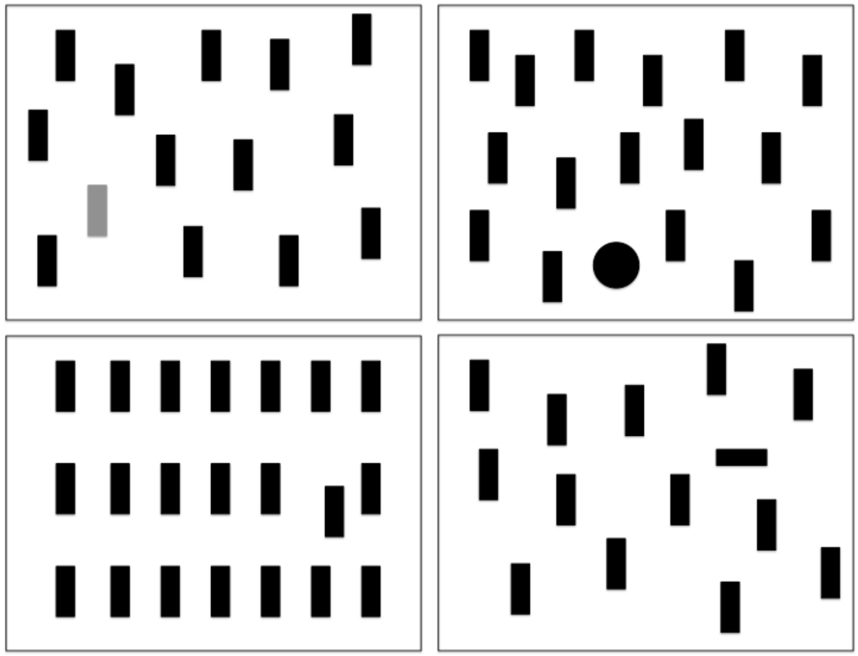

Figure 1

The cases I have been discussing involve the capture of attention by a stimulus; they are bottom up. The impact of attention on consciousness can also be demonstrated in top-down cases. If you search for an object in an array, you use stored object representations as a filter so that the sought object comes under the attentional spotlight more readily. For example, look at Figure 2, and try to find the circle. When you do this, the circle is more quickly identified and it enters consciousness. Objects similar to the circle, like the crescent become more vivid in consciousness too. If you search for the hexagon, it will become salient, as will the octagon. If you search for the triangle, you might spot the trapezoid or the diamond. In each case, search draws attention to a class of shapes, and they are consciously experienced as a result.

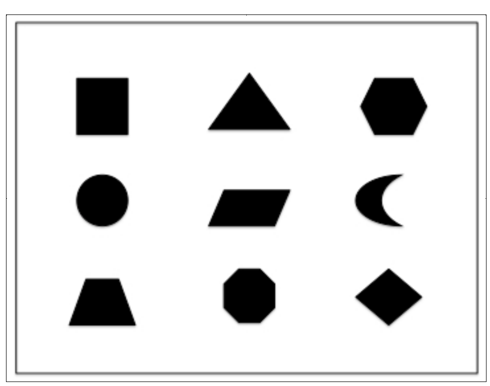

Figure 2

In sum, it seems that attention is both necessary and sufficient for making intermediate-level representations conscious. Conscious states are attended intermediate-level representations, or AIRs (Prinz 2000; 2005).

2.2 What Is Attention?

I have been claiming that consciousness arises when, and only when, we attend. But, as stated, this may seem unsatisfying because I have not yet said what attention is. By way of elucidation, let me consider four objections that will lead to a substantive theory of attention, and more precise statement of the AIR theory.

First, one might object that the theory is circular. I say perception becomes conscious when we attend. But, one might think that attention should be defined in terms of consciousness. An attended stimulus, on this view, is by definition one when are conscious of.

I do not think the conceptual link between consciousness and attention is so tight. Indeed, we will see that some researchers claim that they are dissociable. In any case, I think attention can be defined without reference to consciousness. Empirically, there is a close link between attention and working memory. "Working memory" refers to the storage buffers that are used to temporarily retain information and use it for controlled cognitive processes. It is what we use when we hold a phone number in our heads, calculate in our heads, or keep track of which plates have been filled when doling out potatoes at a dinner party. Working memory is the mind's scratch pad. In perception, working memory is where we jot down the contents of our senses in order to report on them, reflect on them, act on them, or just keep them in mind for brief intervals. And it couldn't do this without attention. Attention is a change in the way perceptual representations are processed that allows them to send signals to working memory. Put more simply, attention is what makes information *available* to working memory. In a slogan, attention is availability.

This is not a bit conceptual analysis, but rather an *a posteriori* identity, supported by empirical evidence. For example, Rock and Gutman (1981) showed that, when attending to one of two overlapping shapes, subjects can remember the attended one, but not the unattended one. When attention is spread thin, fewer things can get into working memory (Sperling 1960). And, when working memory is occupied, attention is limited. Fougnie and Marois (2007) have shown that inattentional blindness can be induced by giving people a heavy working memory load. Neuroimaging studies show that when compared to unconsciously perceived stimuli, conscious perception is associated with co-activation of parietal attention structures and dorsolateral prefrontal working memory structures (Rees et al. 2002).

In response to this *a posteriori* definition of attention, critics might advance a second objection. Surely, they'll say, attention cannot be identified with any one thing. It is very tempting to say attention is not a coherent category. After all, we use the word to refer to a wide range of different mental phenomena. These include vigilance ("Pay attention!"), monitoring ("Maintain attention on this spot here"), tracking ("Attend to the ball!"), pop-out ("The stimulus captured my attention"), focus ("Attend to the fine details"), and selection ("I attended to Xs and ignored Ys"). I already noted that attention can be top down or bottom up. Why think that all these phenomena have a common essence? Attention seems to be mixed bag of independent mechanisms. If so, the claim that consciousness arises when we attend is unpromising, because attention is not one thing, but many.

Here again, I think the working memory account can do some work. It is true that there are many ways to control attention. One can monitor a whole scene, track an object, or select one object over another. Attention can be controlled by a perceiver's plans or by features of the stimuli perceived. But the variability of control should not dupe us into thinking attention itself is varied. In each of these cases, attention coincides with availability. If you monitor an object in the scene in front of you, its features and changes become available to working memory, and if you allow things to pop out, those things will become available to working memory. It seems the impact of attention is more or less the same regardless of the source. There may be uses of the term "attention" outside the perceptual domain that involve different mechanisms, but perceptual attention, at lease, seems to involve a single process.

One important consequence of the single process view is that attention can retain its identity as a process even when applied very differently. Attention can be highly focused, as when we study discrete features of an object in our field of view, or quite diffuse, as when we scan a scene or passively take in a vista. Stage light metaphors are useful here. Attention can act like a floodlight or a spotlight, and there can even be an attentive spotlight (point of focus) within a diffusely lit scene. Attention also seems to have a limited capacity; some things go unlit. As focus increased, there may be less attention left for diffuse monitoring.

Talk of diffuse attention raises a third objection. Working memory has a highly limited capacity. It seems we can only store about four items at once (Cowan 1999). And what we store seems quite coarse grained. For example, recognition of color patches after short temporal delays is very limited. So it looks like working memory stored small numbers of relatively abstract chunks. But attention, I have just said, can be diffuse, covering large areas, and it can also be highly detailed. We can attend to specific colors and shapes, which have no hope of getting stored in working memory. The hypothesis that attention is a mechanism that makes information available to working memory seems hopeless when these facts are considered.

The difficulty can be addressed if we distinguish availability and encoding. A perceptual stimulus is available to working memory when it is processed in a way that would allow for working memory encoding, but many things are available and never get encoded. If you glance at a scene, you could store information about many parts of it, but you don't. Most things are forgotten instantly, because they are not stored in working memory. Evidence for this comes from studies of change blindness. When subjects are consecutively presented with two slightly different images, they regularly fail to detect how the second differs from the first (Simons and Levin 1997). This is happens even if they are looking for the change and see the two images repeatedly. Presumably, they can experience each image in its entirely, but they cannot store the details from moment to moment. When a change is detected, it is usually because the changing feature was, by chance, encoded. The changing feature could have been encoded at any point—it is always *available* for encoding—but typically appears without being encoded numerous times before subjects notice the change.

Working memory encoding is not simple copying. When a perceived item is stored in working memory, working memory does not encode a duplicate of the item. Indeed, the duplication metaphor is false on two counts. First, as already seen, working memory stores information in a coarser code than what we experience in perception. Second, working memory may store information in form of procedural knowledge, rather than be re-representing features of perception. More specifically, working memory encodings can be thought of as commands for maintaining activity in perceptual systems rather than reproducing copies of perceptual states. Putting these two points together, we might imagine the following picture. When an intermediate-level perceptual representation is encoded, that means a high-level perceptual representation that captures the gist of the intermediate-level representation sends a signal to working memory systems, which, in turn, uses this representation as a set of instructions for maintaining perceptual states during a temporal delay. The high-level representation can be used to generate a mental image of the stored item by projecting back into intermediate-level areas, but, because it is course grained, the resulting image will be indeterminate and unstable (for more details, see Prinz 2007).

If this picture is right, then consciousness does not require working memory encoding. It just requires availability for encoding. Available representations can be coarse-coded and stored, but consciousness does not depend on that. The picture in borne out by various psychological studies; I will mention three. It finds support in the fact that color discrimination far outstrips recall, suggesting that we can experience many more colors than we can encode (Halsey and Chapanis 1951). Or consider Sperling's (1960) widely discussed experiments with letter arrays. When subjects are briefly presented with a three by three array, they can recall three or four letters, but not more. Nevertheless, they seem to experience the entire array, and, absolutely any of the letters could be stored if it is cued after the stimulus is removed. Thus, each letter is available for encoding and consciously experienced, but only a small handful are encoded. The claim that encoding in not necessary for consciousness can also be confirmed by neuroiming results. Working memory is located in lateral frontal cortex, but there structures are not always operative during conscious perceptual tasks. Lateral frontal cortex is more or less dormant, for example, when we watch movies (Goldberg et al. 2001). This suggests that encoding doesn't matter. Conscious arises with mere availability.

This raises one final objection. Availability is a dispositional property. Consciousness, on the other hand, cannot be dispositional. Having an experiential quality is an occurrent property of a perceptual state, not merely something that the state *could* do. So, if attention is availability, attention cannot be what makes perception conscious.

The slogan, attention is availability, is a bit of loose talk. It would be more accurate to say, attention is the categorical basis of availability. This is what I implied above when I introduced the proposal. Attention is a process, I said, that makes information available to working memory. It's an actual change that takes place in perceptual representations in virtue of which they become candidates for encoding.

But what is this process? What is the categorical basis for availability? Here the science is still underway, but it is possible to advance an empirically driven guess (see Prinz, forthcoming, for more discussion). To see how perceptual representations become available for working memory encoding, it is necessary to descend to the level of computational neuroscience. The mechanisms of availability cannot be adequately described using psychological vocabulary. What accounts for availability at the neuronal levels seems to be something like this. The various structures that control levels of attention seem to work by increasing activity in interneurons, which are inhibitory cells that modulate activity in pyramidal cells, which respond to the features of perceptual stimuli. By inhibiting pyramidal cells, interneurons cause them to synchronize their activity (Sohal and Huguenard 2005). Synchrony can be measured in the axon potentials and in the local field potentials around dendrites, though it's not fully clear which, if either, form of synchrony is more important. They key thing is that neural synchrony seems to be a very good candidate for availability. Neurons that are in sync, speak as one voice and can be heard above the din (Salinas and Sejnowski 2001). This may turn out to be mistaken, but it gives some idea of what a theory of availability should look like, and it currently enjoys some empirical support.

In this section, I've considered four objections that have lead, in turn, to helpful insights about the nature of attention. Attention is the categorical basis of availability. Attention can be controlled by different sources and can therefore be diffuse or focused, top-down or bottom-up. Attention can occur without encoding, and has a finer grain. And, at the neurocomputational level, attention may be identified with processes such as interneuron inhibition and pyramidal synchronization, which make neural signals available for downstream propagation.

The AIR theory of consciousness states that consciousness arises when and only when we have attended intermediate-level representations. The four objections that I have been considering allow us to make this a bit more precise. We can now say that, attention is a process that:

- makes perceptual states available to working memory
- qualifies as a natural kind, even though it can be initiated by different control structures and allocated with varying degrees of diffuseness
- does not necessarily result in perceptual states being encoded in working memory
- is physically realized by neural synchronization and interneuron inhibition

Thus, the AIR theory can be elaborated as the view that that consciousness arises when intermediate-level representations become available for working memory encoding as a result of specific neural process that can operate under multiple sources of influence.

2.3 AIR Compared to Other Theories

AIR is not the only game in town. There are other theories of consciousness in the literature. This is not the place to critically review these, but it will help to understand the AIR theory, if I compare and contrast it to some of the leading alternative accounts.

The most prominent class of alternatives comprises the Higher-Order Representation (HOR) theories of consciousness. According to these, a mental state becomes conscious when and only when it is represented (or disposed to be represented) by another mental state. That other mental state is, on some versions, a higher-order thought (HOT), or, on other versions, a higher-order perception (HOP). Proponents of HOR theories often defend their approach by armchair arguments. One popular move goes roughly like this (Lycan 2001; see also Rosenthal 1997):

P1. A conscious state is a state of which a subject is conscious
P2. To be conscious *of* a state a subject must represent it
C . Therefore a state is conscious only if it is represented by a subject

I reject both premises. *Defending* my dissent is beyond the scope of this chapter, but *declaring* my rejection will help clarify how AIR and HOR theories differ. I reject P1, because I think consciousness is a first-order affair. States are conscious when they become available to working memory, not when a subject becomes conscious *of* them. This is one reason why the distinction between availability and encoding is so very important. As emphasized earlier, conscious states need not be encoded. When there is no encoding, there can be conscious states, but we are not conscious of those states, contrary to what HOR theorists maintain.

I also have misgivings about P2. I think the locution "conscious of" deceptively implies that we have a representation of the states of which we are conscious. But an alternative analysis would say a subject is conscious of a mental perceptual when that state is encoded in working memory. Encoding may explain consciousness of (also sometimes called "awareness"). But encoding, I submit, is not meta-representational. A working memory encoding is not a higher-order representation of a perceptual state, but rather a mechanism that can serve to maintain a perceptual state or re-activate it after a brief interval. As, I said above, working memory encodings can be thought of as a kind of procedural knowledge (see also Prinz 2007). If I am right, then a subject becomes conscious of a perceptual state without representing it, contrary to P2. This distinguishes the AIR theory from HOR theories.

The point about availability vs. encoding is also crucial for distinguishing AIR from another approach to consciousness that enjoys considerable support: the Global Workspace theory (Baars 1988; Dehaene and Nacchache 2001). Proponents say that consciousness arises when perceptual states are broadcast to more central cognitive mechanism, involved in reporting, deliberation, and executive control. This looks a lot like the AIR theory at first glance, but there is a crucial difference. On the global workspace theory, consciousness requires working memory encoding. Dehaene and Nacchache are quite explicit about this, but the evidence from their own research suggests that consciousness involves availability rather than encoding. For example, in masked priming studies, they have shown that under some short latencies, stimuli can be consciously detected, but not recalled, suggesting that they are available but not encoded (Kouider et al., in press).

AIR theory is not the only theory, however, to identify consciousness with availability. Consider, for example Kirk's (1994) account according to which perceptual states become conscious when they become available for interpretation, assessment, decision, and control processes (p. 146). Or consider Tye's (1995) thesis that conscious states are PANICs: poised, abstract, non-conceptual, intentional contents. The last three criteria (abstract, non-conceptual, intentional contents) constitute Tye's account of which mental states are candidates for consciousness, and will not concern me here. The key thing is the term "poised," which is supposed to characterize how these states become conscious. More specifically, Tye thinks consciousness arises when perceptual states are available to make an impact on beliefs and desires. Thus, both Kirk and Tye invoke availability, which strikes me as exactly right. The problem is neither author offers an adequate account of what availability consists in. What is the mechanism that makes perceptual states available? The AIR theory answers that question. Attention is the essence of consciousness. The theory is not so much an alternative to these order accounts, as an implementation. Kirk and Tye are too vague, and consequently incomplete and overly liberal. They identify they key symptom of consciousness without identifying the mechanism that makes this symptom possible. By analogy, one might say that Kirk and Tye offer something like the early theory of genes (units of inheritance) where AIR gives us the double helix.

In summary, then, the AIR theory resembles other theories of consciousness in various respects but also differs in important ways. It is a first-order account that emphasizes availability, rather than encoding in working memory. Of course, the major difference can be stated more directly. None of these other accounts say that attention is the essence of consciousness. This is the chief claim that I have been trying to defend. The link between attention and consciousness enjoys considerable empirical support, which has not been fully recognized by defenders of these other theories.

That said, the AIR theory also faces objections. The strong evidence linking consciousness to attention has come under attack recently on both empirical and methodological grounds. I consider what I take to be the most pressing concerns in the next section.

3. Objections

3.1 Attention is Not Necessary

I have claimed that attention is necessary and sufficient for making intermediate-level perceptual states conscious, but this claim can be challenged. Within cognitive neuroscience, it has become popular to argue that consciousness and attention are dissociable. Much of the evidence is reviewed by Koch and Tsuchiya (2006). Koch used to think attention and consciousness were closely linked (Crick and Koch 1990), but now thinks they are independent mechanisms. My goal here is not to assess each experiment that has been described as supporting this conclusion, but rather to consider a few representative examples. In showing how that AIR theory can handle these apparent counter-examples, I am hoping that readers will be able to extrapolate or devise replies to others (see also Prinz, forthcoming; and De Brigard and Prinz, forthcoming). I will begin with two studies that purport to show that attention is not necessary for consciousness.

First consider a study by Li et al. (2002). Here subjects are presented with the attention-demanding task of finding a rotated T in a group of rotated Ls. While performing the task, an image is briefly presented in the periphery and subjects are asked to make a judgment about the image, such as whether it contained an animal or whether it contained a vehicle. For some of these discriminations, subjects perform exceptionally well. This leads the Koch and Tsuchiya (2006) to describe the study as evidence for conscious perception in the near absence of attention.

This counter-example is problematic in various ways. First, "near absence" is not absence. Subjects may be allocating some attention to the periphery. In fact, if they were not we have antecedent reason to think they would not be able to freely report on what they saw, because this is precisely what change blindness studies establish. Second, the flashed stimuli, which are complex, richly colored, and high contrast, may capture attention, unlike the small shapes used in inattentional blindness studies. Performance was poor for meaningless color patterns or letters, and excellent for meaningful objects, which may be especially effective as attention lures. Third, subjects had 10 hours of training, going through 12,000 trials before being tested, which may have reduced the attentional demands of the central task. Forth, we don't even know for sure that subjects are conscious of the peripheral stimuli. Success at forced choice guessing after extensive training is not necessarily a measure of conscious awareness. In sum, the research by Li et al. (2002) does not provide strong evidence for consciousness without attention

Turn now to another study purporting to show that there can be consciousness without attention. Landman et al. (2003) devised a clever experiment that combines change blindness with Sperling's array paradigm. Subjects see an array of rectangles, which are either horizontal or vertical, followed by a gray screen, and then a second array. In the second array, one rectangle has changed its orientation, but subjects are usually incapable of detecting the change. However, on some trials, subjects see a cue during the gray screen pointing to where a particular rectangle has been located in the original array. When this is done, subjects can reliably report whether than rectangle changed when the second array is presented. Thus subjects have the potential to detect every change if attention is directed. The crucial finding is that the attention cue can come after the stimulus. The authors interpret this as showing that the stimulus has been consciously perceived when the display was originally presented, but unattended (see also, Lamme 2003; Block 2007). This interpretation rests on two conditional assumptions. If the stimulus had been attended, it would have been reportable. If it had been unconscious, the presentation of a post-display cue shouldn't be effective.

Both of these conditional assumptions can be questioned. First, consider the claim that unconscious stimuli cannot be cued after they are removed. Sperling's original study proved that visual stimuli produce iconic memories: rapidly fading traces in the visual system. For all we know from this research, iconic memories can be produced by unconsciously presented stimuli. If so, a cue presented during the period in which an unconsciously induced iconic memory is fading may serve to bring that iconic trace into consciousness. On the AIR theory, this would be readily explained. The cue brings attention to a visual trace, and the trace becomes conscious thereby. In the Landman et al. study, this may be exactly what takes place. The rectangles used in their displays are depicted with a noise gradient against a noisy background, and it's far from obvious that every rectangle is consciously perceived.

But suppose these rectangles in the original display are consciously perceived. There is still a possibility that the second conditional assumption is false. To argue that this study illustrates consciousness without attention, Landman et al. must say that, if a rectangle had been attended, it would have been reportable. But I reject that assumption. Reportability arises when attended stimuli are encoded in working memory, but I have already argued that attention outstrips working memory. So Landman et al. do not provide evidence for thinking their stimuli are not attended. It's quite plausible that subject try to attend to the whole display, since they know they will be tested on it. And the allocation of diffuse attention may bring each rectangle into consciousness. If so, the study established only that consciousness can arise without encoding, which is consistent with the AIR theory. In summary, Landman et al. do not establish consciousness without attention because they neither establish that the stimuli in question are conscious nor that they are unattended.

These studies are illustrative of the best recent attempts to show that consciousness can arise without attention, and they have been emphasized by leading defenders of that dissociation. But the studies fail to establish any such thing.

3.2 Attention is Not Sufficient

Granting that attention is necessary consciousness, critics may still object that it is not sufficient. Recent experiments have been constructed to establish that attention can occur in the absence of consciousness. These experiments pose a threat to the AIR theory. Fortunately for AIR, they don't provide compelling evidence for what they seek to establish.

Let's begin with a study by Kentridge et al. (2008), which extends earlier work by the first author. The study combines "metacontrast masking" with attentional cueing. Metacontrast masking is a special kind of masking. In typical masking studies, the mask is a stimulus that occupies the same region of space that was occupied by the first stimulus. In metacontrast masking, subjects see a mask that occupies the area *surrounding* the space in which the first stimulus had been presented. Just like masks that overlap the region that was occupied by a previous stimulus, these surrounding masks also prevent the previous stimulus from being consciously perceived. In the present study, subjects see colored disks followed by colored rings, and the rings mask the disks, resulting in unconscious perception. Their task is to detect the ring as quickly as possible. On some trials the disk is preceded by an arrow that serves a lure for attention. The arrow does not bring the disk into consciousness, but it does have a significant affect: if the disk is the same color as the ring, it facilitates ring detection, but this happen only when the disk is preceded by an arrow. Thus the arrow seems to give this unconscious stimulus the power to exert priming influence. The authors interpret this as showing that the unconscious stimulus has been enhanced by attention. Thus, it seems to be a case of attention without consciousness.

As compelling as this may seem, it is not a counter-example to the AIR theory. In fact, there is reason to think that attention *cannot* be responsible for the observed effect. The disk is presented very briefly and followed by a ring in the surrounding area. According to the leading interpretation of metacontrast masking, the ring is able to mask the disk precisely because it draws attention away from the region of space in which the disk is located (Enns and Dilollo 2000). Attention does not have time to enhance the representation of disk, and it cannot enhance the iconic trace of the disk because the ring draws attention away. How, then, can we explain why the subjects in the Kentridge et al. study show improved performance when cued? There are several possibilities, but the simplest answer is that cueing affects where subject look. More specifically, the arrow may elicit saccadic eye movements to the region in which the disk will be presented, and that could result in a more accurate representation of the disk. Vision scientists sometimes refer to saccadic eye movements as "overt shifts of attention," but the term is misleading, because attention and eye movements are dissociable; we can look is one place while attending to another (Posner 1980). Direction of gaze can enhance vision even when attention is absent. Foveal vision has a higher concentration of color receptive cells, and results in more saturated representations. If the cue leads to a more saturated representation, that could explain why color-based priming is found in the cued condition. This interpretation of the results appeals to a well-known process, and, for reasons just mentioned, is more plausible than the claim that attention is involved.

Turn now to a second study that attempts to establish attention in the absence of consciousness. Jiang et al. (2006) use a technique called interoccular suppression in which a low-contrast stimulus presented to one eye is masked by a higher contrast stimulus presented to the other eye. Cleverly, Jiang et al. used nude photographs as the masked stimuli assuming that these would attract attention, despite the fact that they could not be consciously perceived. Sure enough, when the interoccular displays were taken away, subjects were asked to detect a target and target detection was superior in the location that had been occupied by the nude. The effect worked only when the nude was a member of the subject's preferred sex. The study has two features that make it a very powerful response to the AIR theory. First, the method allows the stimulus to be presented for an extended duration, which means attention has enough time to act on the stimulus representation. Second, the stimulus itself serves as the attention cue, so it seems especially plausible that the representation of the stimulus is modulated by attention. This looks like a case of an unconscious AIR.

But, once again, an alternative interpretation is available. Perhaps the nude attracts saccades and not attention. This would facilitate target detection when the target is displayed in the region that was occupied by the nude. And, once again, there is reason to prefer this interpretation over the one given by Jiang et al. They say that attention is being allocated to the unconscious stimulus, but the experimental set up seems to preclude that possibility. In interoccular suppression the rival stimulus is an attention lure. Suppression results from the fact that the high contrast stimulus is able to attract attention away from the low contrast stimulus. In support of this, there is evidence that binocular rivalry, of which interoccular suppression is a special case, involves the same mechanisms as selective attention (Mitchell et al. 2004). Moreover, if the nude were being attentionally enhanced, we might expect to see increased activation in the ventral stream, where object representations are processed. Jiang et al. did not measure brain response, but fMRI studies of interoccular suppression suggest that suppressed stimuli are not in fact associated with increased ventral processing (Fang and He 2005). The increases are observed in the dorsal stream, which plays a role in saccadic eye movements and spatial perception. Thus, the existing evidence offers better support for my interpretation than for the interpretation given by Jiang et al.

The two studies I have been discussing fail to establish that there can be unconscious AIRs. I think these studies are the best efforts to establish that attention is insufficient for consciousness, and, if I am right, they do not succeed. So I conclude that current empirical evidence provides strong support for the claim that attention is necessary and sufficient for making perceptual representations conscious, and no convincing evidence to reject that claim.

3.3 Alternative Interpretations of the Evidence

I might rest my case here, but there is one more challenge that deserves attention. This one was put to me forcefully in a commentary that David Chalmers delivered in response to a related paper of mine (Chalmers 2008). I will not do justice to all the moves that Chalmers makes, but I will address the main thrust of his critique. Chalmers attacks the evidence offered in support of the view that attention is necessary for consciousness. In this respect, his critique echoes the theme of the studies discussed in section 3.1, but it deserves to be treated separately because it raises very different issues. Chalmers is not offering counter-examples to the necessity thesis; he is challenging the interpretations of experiments that were used to support the thesis in first place. If these challenges succeed, the AIR theory cannot even get off the ground.

I offered two main sources of evidence for the conclusion that consciousness requires attention. Healthy subjects fail to report stimuli when attention is divided (inattentional blindness) and damage to attention centers leads to neglect. Chalmers argues that both sources of evidence are open to an alternative interpretation. Subjects may have conscious experiences that they simply fail to report. In neglect, subjects may be incapable of reporting stimuli in their blind fields precisely because they cannot attend. Perhaps attention is required for reporting, but not for experience itself. If so, neglect provides no evidence for the conclusion that consciousness requires attention. It discussing inattentional blindness, Chalmers notes that there is an alternative interpretation of these studies: subject may suffer from inattentional amnesia (Wolfe 1999). They may consciously experience the surprise stimuli and instantly forget them. We know that visual memory isn't very good, so the subjects' retrospective reports are not necessarily reliable.

Let me reply to these two alternative interpretations in turn. First consider neglect. The evidence that people with neglect do not experience things on the left is not restricted to self-report. They fail to consciously see objects presented on the left or, in some cases, the left half of objects presented centrally. Neglect can be demonstrated in many ways. When asked to copy line drawings, features on the left are ignored; when asked to form mental images of familiar places, landmarks on the left are forgotten; when asked to bisect lines, the midpoint is placed too far to the right; when asked to find target letters in a group of distractors, targets on the left are missed. Thus, they fail to demonstrate consciousness on both implicit and explicit measures. In the absence of positive reason to think these individuals are conscious, I think this widespread constellation of behaviors should be taken as evidence that they are not.

What about inattentional blindness? Chalmers says this might better be regarded as inattentional amnesia, but I see little reason to explain the results that way. Mack and Rock (1998) consider this possibility and devise a clever experiment to rule it out. If unattended stimuli were consciously seen, then they should able to integrate with other conscious stimuli presented consecutively and proximately to produce illusory motion. But they exert no such influence. Chalmers replies that motion perception involves binding, and binding is associated with attention. Thus, the absence of illusory motion is explained by the absence of attention, not the absence of consciousness. But this reply only serves to support the AIR theory. Conscious experiences are characteristically bound, and if attention is required for binding, that is another powerful reason for thinking that attention coincides with consciousness.

Moreover, there is good reason to reject the inattentional amnesia story. The main positive argument for taking the story seriously is given by Wolfe (1999). He demonstrates that visual memory doesn't last very long, and he points out that Mack and Rock's stimuli are quickly presented. But, the argument doesn't hold water. First of all, some inattentional blindness studies present stimili for long intervals (recall Most et al. (2001), whose stimuli were present for five seconds). Second of all, subjects in Wolfe's studies fail to retain shape information and object identities, but they are quite sure they have seen *something*, which is strikingly different from what we find in inattentional blindness studies. If the memory limitations measure by Wolfe were behind the Mack and Rock results, their subjects should be able to report that they saw a stimulus but forgot what it was.

The point can be extended to other cases. Generally speaking, visual memory limitations affect subjects reports of *what* they see, not *whether* they saw anything at all. Consider the profound difference between the cases developed by Sperling and the cases developed by Mack and Rock. In the Sperling studies, subjects are sure they have seen a letter array, even if they can't make out the letters. In the Mack and Rock studies, subjects have no idea there has been a surprise stimulus.

They would be at chance if they had to guess whether something was presented. Sperling's subjects would readily report that they have seen something and forgotten it. Not so with Mack and Rock's subjects. Even on interrogation, they deny having see anything. This difference cries out for an explanation. If both cases were just forms of amnesia, they should seem similar to subjects, but they do not. We can readily explain the contrast if we suppose that the Sperling's subjects are forgetting what they experienced whereas Mack and Rock's subjects are not experiencing anything at all.

This is an argument to the best explanation. It accounts for why Sperling cases and inattentional blindness cases seem so different (if you are in doubt, try the stimuli on yourself). The standard interpretation of inattentional blindness also hangs together elegantly with the standard interpretation of neglect, where Chalmers appeals to different failures (report and memory), the standard view says that both are direct results of absent attention. The standard story also explains why unattended stimuli cannot be picked up using any normal measure of awareness (self report, forced choice, drawing, spontaneous action, and so forth). Indeed, the only reason we even know these stimuli are perceived is that they induce priming effects. But priming is a hallmark of unconscious perception.

Of course, it is always *possible* that people with neglect or subjects in inattention studies are having conscious experiences that they cannot demonstrate using subjective or objective measures. I think this possibility is a bit like skeptical hypotheses. We cannot refute it, but we should not take it very seriously. We should not postulate conscious experience in cases where we have no strong evidence that there is experience. The standard interpretation of neglect and inattentional blindness studies are preferable to the alternatives that Chalmers would have us consider.

4. Conclusion

It has been known for a long time that perception can occur without consciousness. In this chapter, I've tried to identify that factor that makes a difference. I suggested that consciousness arises when and only when intermediate-level representations are modulated by attention—what I call the AIR theory. Attention is the difference maker. This conclusion is based on empirical research, but has also come under recent empirical attack. I deflected what I take to be the most powerful objections. Perhaps some further finding will force a revision of the view, but, for now, it looks like we have found the key to consciousness. Attend and you will see.

References

Arnell, Karen M., Kassandra V. Killman, and David Fijavz. 2007. Blinded by emotion: Target misses follow attention capture by arousing distractors in RSVP. *Emotion* 7: 465-477.

Baars, Bernard. 1988. *A cognitive theory of consciousness*. Cambridge, MA:

Cambridge University Press.
Bar, Moshe. 2000. Conscious and non-conscious processing of visual object identity. In *Beyond dissociation: interaction between dissociated implicit and explicit processing*, ed. Yves Rosetti and Antti Revonsuo, pp. 153-174. Amsterdam: John Benjamins.
Block, Ned. 2007. Consciousness, accessibility, and the mesh between psychology and neuroscience. Behavioral and Brain Sciences 30: 481–548.
Chalmers, David. 2008. Is there consciousness outside attention?: Comments on Jesse Prinz. Address given at the annual SPAWN conference, University of Syracuse, August 2008.
Crick, Francis, and Christof Koch. 1990. Towards a neurobiological theory of consciousness. *Seminars in Neuroscience* 2: 263-275.
De Brigard, Felipe, and Jesse J. Prinz. Forthcoming. Attention and consciousness. *Wiley Interdisciplinary Reviews: Cognitive Science*.
Dehaene, Stanislas, and Lionel Nacchache. 2001. Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition* 79: 1-37.
De Gelder, Beatrice, Jeffrey S. Morris, and Raymond J. Dolan. 2005. Unconscious fear influences emotional awareness of faces and voices. *Proceedings of the National Academy of Sciences* 102: 18682-18687.
Enns, James T., and Vincent DiLollo. 2000. What's new in visual masking? *Trends in Cognitive Sciences* 4: 345- 352.
Fang, Fang. and Sheng He. (2005). Cortical responses to invisible objects in the human dorsal and ventral pathways. *Nature Neuroscience* 8: 1380-1385.
Fougnie, Daryl, and René Marois. 2007. Executive load in working memory induces inattentional blindness. *Psychonomic Bulletin & Review* 14: 142-147.
Goldberg, Ilan, Micha Harell, and Rafael Malach. (2006). When the brain loses its self: prefrontal inactivation during sensorimotor processing. *Neuron* 50: 329-339.
Hamm, Alfons O., Almut I.Weike, Harold T. Schupp, Thomas Treig, Alexander Dressel, and Christof Kessler. 2003. Affective blindsight: Intact fear conditioning to a visual cue in a cortically blind patient. *Brain* 126: 267-275.
Halsey, Rita, and Alphonse Chapanis. (1951) On the number of absolutely identifiable spectral hues. *Journal of the Optical Society of America* 41: 1057-1058.
Jackendoff, Ray (1987). *Consciousness and the computational mind*. Cambridge, MA: MIT Press.
Jiang, Yi, Patricia Costello, Fang Fang, Minor Huang, and Sheng He. 2006. A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Proceedings of The National Academy of Science*. 103: 17048-17052.
Kentridge, Robert W., Tanja C. W. Nijboer, and Charles A. Heywood. 2008. Attended but unseen: Visual attention is not sufficient for visual awareness. *Neuropsychologia* 46: 864-69.
Kirk, Robert. 1994. *Raw feeling: A philosophical account of the essence of consciousness*. New York, NY: Oxoford University Press.

Koch, Christof and Jochen Braun (1996). Towards a neuronal correlate of visual awareness. *Current Opinion in Neurobiology* 6: 158–164.
Koch, Chistif, and Naotsugu Tsuchiya. 2007. Attention and consciousness: Two different processes. *Trends in Cognitive Science* 11: 16-22.
Kouider, Sid, Stanislas Dehaene, Antoinette Jobert, and Denis Le Bihan. In press. Cerebral bases of subliminal and supraliminal priming during reading. *Cerebral Cortex*.
Lamme, Victor A.F. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences* 7: 12-18.
Landman, Rogier, Henk Spekreijse, and Victor A. F. Lamme. 2003. Large capacity storage of integrated objects before change blindness. *Vision Research* 43: 149–164.
Li, Fei F., Rufin VanRullen, Christof Koch, and Pietro Perona. 2002. Rapid natural scene categorization in the near absence of attention. *Proceedings of the National Academy of Sciences* 99: 8378 – 8383.
Lycan, William G. 2001. A Simple Argument for a Higher-Order Representation Theory of Consciousness. *Analysis* 61: 3-4.
Macdonald, James, and Nilli Lavie. 2008. Load induced blindness. *Journal of Experimental Psychology: Human Perception and performance.* 34: 1078-1091.
Mack, Arien, and Irvin Rock. (1998). *Inattentional blindness*. Cambridge, MA: MIT Press.
Mack, Arien, Zissis Pappas, Michael Silverman, and Robin Gay. 2002. What we see: Inattention and the capture of meaning. *Consciousness and Cognition* 11: 488-506.
Marr, David. 1982. *Vision*. San Francisco, CA: Freeman.
Marshall, John C., and Peter W. Halligan. 1988. Blindsight and insight in visiospatial neglect. *Nature* 336: 766 – 767.
Mitchell, June F., Gene R. Stoner, and John H. Reynolds. 2004. Object-based attention determines dominance in binocular rivalry. *Nature* 429: 410–413.
Most, Steven B., Brian J. Scholl, Erin Clifford, and Daniel J. Simons. 2005. What you see is what you set: Sustained inattentional blindness and the capture of awareness. *Psychological Review* 112: 217-242.
Posner, Michel I. 1980. Orienting of attention. *Quarterly Journal of Experimental Psychology* 32: 3-25.
Prinz, Jesse J. 2000. A neorofunctional theory of visual consciousness. *Consciousness and Cognition* 9: 243-59.
Prinz, Jesse J. 2005. A neurofunctional theory of consciousness. In *Cognition and the brain: Philosophy and neuroscience movement,* ed. Andy Brook and Kathleen Akins, 381-396. Cambridge: Cambridge University Press.
Prinz, Jesse J. 2007. Mental pointing: Phenomenal knowledge without concepts. *Journal of Consciousness Studies,* 14, 184-21.
Prinz, Jesse J. Forthcoming. *The conscious brain*. New York, NY: Oxford University Press.

Raymond, Jane E, Kimron L. Shapiro, Karen M. Arnell. 1992. Temporary suppression of visual processing in an RSVP task: an attentional blink? *Journal of experimental psychology. Human perception and performance* 18: 849–60.
Rees, Geraint, Gabriel Kreiman, and Christof. 2002. Neural correlates of consciousness in humans. *Nature Reviews Neuroscience* 3: 261-270.
Rock, Irvin, and Daniel Gutman. 1981. The effect of inattention on form perception. *Journal of Experimental Psychology: Human Perception & Performance* 7: 275-285.
Rosenthal, Daven M. (1997). A theory of consciousness. In *The nature of consciousness: Philosophical debates*, ed. Ned Block, Owen Flanagan, and Guven Güzeldere, 729-753. Cambridge, MA: MIT Press.
Salinas, Emilio, and Terrence J. Sejnowski. 2001. Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience* 2: 539-550.
Simons, Daniel J., and Daniel T. Levin. 1997. Change blindness. *Trends in Cognitive Science* 1: 261-267.
Sohal, Vicaas S., and John R. Huguenard. 2005. Inhibitory coupling specifically generates emergent gamma oscillations in diverse cell types. *Proceedings of The National Academy of Science* 102: 18638-18643.
Sperling, George. 1960. The information available in brief visual presentations. *Psychological Monographs* 74. 1-29.
Tye, Michael. 1995. *Ten problems of Consciousness*. Cambridge, MA: MIT Press.
Weiskrantz, Lawrence. 1986. *Blindsight: A case study and implications*. Oxford: Oxford University Press.
Winkielman, Piotr, Kent C. Berridge, and Julia L. Wilbarger. 2005. Unconscious affective reactions to masked happy versus angry faces influence consumption behavior and judgments of value. *Personality and Social Psychology Bulletin* 1: 121-135.
Wolfe, Jeremy M. 1999. Inattentional amnesia. In *Fleeting Memories*, ed. Veronika Coltheart, 71-94. Cambridge, MA: MIT Press.
Zeki, Semir. 1993. *A vision of the brain*. Oxford: Blackwell.