Attention without awareness: A Brief Review.

R.W. Kentridge

Department of Psychology

University of Durham

Durham DH1 3LE

UK

robert.kentridge@durham.ac.uk

Introduction

In this chapter I review some of what we know about the relationship between conscious awareness and attention. I will focus on visual attention and visual awareness and, in particular, on the question of whether we can attend to something without becoming consciously aware of it.

The relationship between attention and conscious awareness is more complex than many theorists of the topic have expected. One reason why it has been usual to expect the relationship to be simple is that, as William James (1890) remarked, “Everyone knows what attention is” (p. 381), and the same dictum has often also been applied to consciousness (by Sir Roger Penrose and Gerald Edelman, amongst others). Because we feel that we understand both of these phenomena intimately, through our own experience of them, it has often been supposed that we should be able to discover the relationship between them simply by introspection.

One of the theories of the attention/consciousness relationship that has been arrived at largely on the basis of an introspective approach to theorizing is the theory according to which attention is the gateway to consciousness. William James suggests a theory along these lines in *The Principles of Psychology,* when he writes that 'my experience is what I agree to attend to' (p. 380), and he endorses something similar, although in a slightly weaker form, in his 1892 chapter on "The Stream of Consciousness" when he claims that "what is called our 'experience' is almost entirely determined by our habits of attention."

Comment [CRM1]: Maybe there's no risk of confusion here, but are the terms 'experience', 'awareness' and 'consciousness' being used synonymously here? If so it might be worth saying so.

A somewhat different introspection-based account of the attention/consciousness relation is given by Wilhelm Wundt in his 1912 "Introduction to Psychology". Chapter One of that book -- entitled "Consciousness & Attention" -- explains attention, not as the gateway to consciousness, but as a process that selects a subset of what is already conscious: "We call that psychical process, which is operative in the clear perception of a narrow region of the content of consciousness, attention" (p. 16).

The view that attention and awareness are closely related, and are perhaps facets of one and the same process, has persisted since the time when James and Wundt were writing. Many neuroscientists (myself included) would now argue that, contrary to the traditional views, attention and awareness dissociate in a number of ways, and so they cannot have a common basis. But the question of how exactly we should account

Deleted: that

for the attention/consciousness relation remains in contention, as many of the chapters in this volume show. As Smithies points out (this volume) purely philosophical approaches this issue, based only on our introspective grasp of the two phenomena, are prone to degenerate into merely verbal disagreements between advocates of alternative stipulative definitions. We should therefore prefer to approach the relationship between attention and consciousness empirically. But if we are to do so then, even if we know what both of these phenomena are, we shall still need to produce operational definitions of them. Getting these definitions is not straightforward.

Complexities of Attention and Awareness.

Since attention comes in many varieties and has many components the production of an operational definition for it is a complex task. The core of attention, as cognitive psychologists understand it, is the use of information to facilitate the execution of a task to which many stimuli might potentially provide the solution. The information that is used need not tell the subject anything about what the solution to the task is. The use of that information is facilitative, not because it primes the solution, but because it excludes some irrelevant stimuli from consideration.

Suppose, as an example, that we have a task in which subjects must scan each of a series of pictures, all of which depict many shapes, in order to determine which of the pictures include a triangle. If the different shapes are drawn in different colours then being told in advance that any triangles in the next picture will be drawn in red could help the subjects to solve this task without providing any information about whether

Comment [CRM2]: This seems to limit the definition so that it applies only to *voluntary* attention.

Comment [B3]: I'm not sure – you can still make use of information without being aware of it surely?

Comment [CRM4]: This seems to limit the definition so that it applies only to *perceptual* attention.

Comment [B5]: Probably true – perceptual attention is what I work on.

Comment [s6]: Is this to say that attention is the selection of information to facilitate the execution of a task?

or not a specific picture actually contains a triangle. The use of this information to help in the performance of the task would be a paradigmatic instance of attention.

Experiments of this sort have several components that it is important to distinguish from one another. In the example just given the triangle that is being looked for, should there be one, is called the *attentional target*. The other shapes are the *distracters* or *foils*. The shape of these stimuli is their *task relevant dimension*. The information that triangles in the next picture will be red is the *basis of attentional selection.* The manner in which this information is conveyed is the *attentional cue*. This cue might take several forms. In our example, for instance, it might take the form of a spot of colour, presented to subjects just before each of the pictures. The feature of the cue in virtue of which it carries facilitative information we may call its critical property.

Distinct psychological and neural processes may be involved in the treatment of each of these components. The question 'What is the attention/consciousness relation?' is therefore complex, since each of these processes may relate to conscious awareness in a different way. These complexities are multiplied when we consider the difficulties associated with attempts to produce an operational definition of consciousness.

One familiar proposal for the operational definition of consciousness is that a conscious state be defined as a state that is available for explicit verbal report. There are plenty of problems with this definition (see e.g. Block, 2007; Goldman, 1997). But so long as we restrict our attention to normal humans, without motor or language deficits, and with intact connections between their cerebral hemispheres, this

Comment [s7]: If we can't define consciousness, then aren't we going to be changing the subject as soon as we introduce an operational definition?

definition is good enough to provide an operationalization of consciousness that is adequate for most cognitive neuroscientific purposes. There are, nonetheless, several things that might be meant when we say that a subject *lacks* conscious awareness of an item. We may mean that the subject is unaware of its existence, unable to recognize some of its properties (such as where it stands on the task relevant dimension), or unaware of its meaning (such as whether it is a cue, and what it is cueing).

Comment [B8]: That is why I use the 2AFC test for discrimination of target existence as a 'backstop' test. Sorry, should have made the point earlier.

Comment [CRM9]: Oh no they don't. Block's cases are supposed to show that conscious states may not be available to verbal report, even in normal subjects.

There are, then, various things that might be meant when we ask whether attention can operate without awareness. We may be asking 'Can we process a cue without being aware of that cue's existence (or without being aware of its status as a cue, or of its critical properties)?' or we may be asking 'Can we show cue-facilitated processing of a target without being aware of that target's existence (or of the properties by which is selected, or of where it stands on the task relevant dimension)?'.

Whatever our answers to these questions there will be several further questions that a satisfactory account of the attention/consciousness relation would need to address: Does the cue-facilitated processing of a target act is such a way that the processing of any of the properties of the target is facilitated, or does it only facilitate discrimination of the task-relevant stimulus dimension? Does the presence of absence of consciousness affect all cue-types in the same ways?

There are similar complications of, instead of asking whether attention can operate without awareness, we ask whether awareness can be found in the absence of attention. In the latter case, however, these complications are not so severe (and in what follows I shall address this latter question relatively briefly).

Even having distinguished these various notions of ‘attention’ and ‘awareness’ we are not done with the conceptual preliminaries to investigating the relationship between attention and awareness empirically since there are different bases for selection, and different relationships between cues and their meanings. These are known to have different psychological properties and different underlying neural mechanisms. Again, then, they may relate to conscious awareness differently.

Varieties of Cue and Varieties of Attention

Attention may be paid to an item on account of that item’s spatial location or on account of its having some non-spatial feature. In spatial attention we may be cued that targets are more likely to appear at one location than at others, and so we attend to the items in that location. In feature-based attention (as in the example discussed above) we are cued that targets will have a feature, such as a particular colour, and so we attend to all the items that have that colour. Once their location, or other features, has determined *which* objects are to be attended it may then be that attention is allocated to the *objects themselves*, rather than their locations or features. In object-based attention facilitation occurs throughout the irregular boundaries of an object rather than in the regular circular ‘spotlight’ of spatial attention. We need, then, to consider different ways in which attention is allocated, and different ways in which objects may be partitioned from their backgrounds.

Deleted: are

Finally we need to consider how cues give rise to selection. The simplest relationship between cues and selection occurs when the cue itself exemplifies the property that is to serve as the basis for selection. In spatial selection a cue might appear in the location of the upcoming potential target. In feature-based selection the cue itself might have the feature that the to-be-attended items are defined by (e.g. a red cue might indicate that the red stimuli are the ones should be selected for facilitated processing). In more complex cases the relationship of cues to bases of selection might be arbitrary: words (e.g. 'RED' or 'GREEN'), symbols (e.g. left- or right-pointing arrows) or even information delivered in a different modality (e.g. high or low tones) might be used to arbitrarily specify the basis for selection. No specific relationship can be assumed between the cue and the basis of selection. In some experiments the word 'RED' might indicate that targets are likely to be green.

It is often assumed that cues which directly specify the basis for selection are processed automatically, and so that we cannot avoid attending on the basis of such cues, even when they are, in fact, uninformative. If, for example, there are flashes of light which appear either on the left and right hand sides of a display, and if those flashes are completely uninformative about the location of the forthcoming target, then processing of targets appearing on the same side of the display as the preceding cue is nonetheless facilitated relative to targets appearing on the opposite side from the cue. Cues which are processed automatically are often referred to as *exogenous* cues. They are contrasted with *endogenous* cues, which require interpretation rather than directly specifying the basis of selection. Because the use of endogenous cues is not obligatory, they are often said to be processed voluntarily.

Comment [B10]: If the advantage is based on attention rather than priming then the advantage is gained by virtue of reductions in the processed of unattended foils.

Comment [ww11]: To follow on Chris' question. What is meant by informational advantage? The feature of the cue captures attention and this has subsequent behavioural effects, decreasing RT or increasing accuracy.

Comment [ww12]: Suggested reorganization. Announce direct specification and exogenous/automatic cues together; then speak of arbitrary relationship, endogenous, interpretation.

Once again, however, we should be wary of collapsing distinctions that may prove important. It is not the case that all and only arbitrary cues are processed automatically. The direction of another person's gaze, for example, acts as a very effective attentional cue – we typically attend to locations at which other people are looking. Whatever signals we use to establish the direction of another's gaze they are certainly something other than a direct specification of that spatial location. Nevertheless, it appears that gaze cueing is automatic despite the fact that it does not directly specify the basis of selection. It may even be that arbitrary cue-selection relationships which are very well learned become automatic – if left-arrows indicate that targets are likely to subsequently likely to appear on the right side of the display reliably enough and for long enough then eventually we base our selection on them even if they subsequently become unpredictive. It therefore appears that direct and arbitrary cue-selection relationships might be orthogonal to the automatic or voluntary use of cues.

Deleted: eventually

To summarise this introductory survey of the scope of what appeared to be a simple question, there are an enormous number of relationships that could exist between stages and types of attention and the awareness of the existence, properties and meaning of the stimuli playing a part in the attentional process.

Types of Evidence

We have seen that there are a number of questions that will need to be answered by a complete account of the relationship between attention and conscious awareness. There are two types of evidence that might help us to answer these questions. First, and most obviously, there is behavioural evidence. We might attempt to construct tasks in which we measure whether a person is able to use an attentional cue, and we might also measure or manipulate what they are aware of. For example we might try to answer the question of whether a cue can influence target processing when one is unaware of the cue's existence by designing an experiment in which a spatial cue is masked, so that subjects cannot discriminate whether or not a cue was present during an experimental trial. If we can also show that this undetectable cue nevertheless influences processing of a subsequently presented target this would indicate that a cue can influence target processing when one is unaware of the cue's existence (at least for the kind of spatial cue used in the experiment).

Deleted:

In addition to straightforward behaviour evidence we might also consider neuropsychological approaches in which patients whose ability to attend or ability to be aware are compromised by damage to their brains. If, for example, patients who cannot orient their attention into a specific region of space are nevertheless aware of stimuli in that region of space we might conclude that it is possible to be aware of a target without attending to it. Of course, there may be room for debate over whether the extent to which the patients cannot attend is total. Dissociations in behavioural studies, whether or not they use neurological patients, are often going to be controversial because they require us to be confident that either some aspect of attention or of awareness is totally absent rather than merely attenuated.

One additional line of evidence that might disambiguate the status of a behavioural dissociation is a parallel neurophysiological dissociation. An ambiguous behavioural dissociation between attention and awareness might be resolved if, in the same experimental paradigm, states of attention and awareness have distinct neural correlates. One can also apply neurophysiological evidence to understanding the extent to which different aspects of attention are independent. For example examining whether attention modulates neural sensitivity to properties used in attentional selection, properties used in target discrimination or both.

Awareness without attention.

The first issue that I will use these various sources of evidence to address is the question of whether it is possible to be aware of things without attending to them. This is a relatively straightforward place to start because, unlike the reverse question -- Is it possible to attend without being aware? -- we do not need to distinguish between aspects of types of attention. We just need to be confident that it is absent.

Two paradigms have been used to address this question, the first is inattentional blindness. The second is the dual-task attentional operating characteristics (AOC) paradigm (Sperling & Dosher, 1986). Mack and Rock (1998) review inattention blindness in their book of that name. Their basic paradigm involves subjects undertaking an attentionally highly demanding task (typically making a judgement about which is the longer of two crossed lines of very similar lengths). Subjects undertake a number of such trials until, on the critical trial, an unexpected probe stimulus is presented along with the target stimulus (the crossed lines). If the probe stimulus is presented at a different location to the attended target stimulus Mack and Rock found that it often went unnoticed. As many as 30% of subjects denied having seen the probe when debriefed, despite the fact that it was presented in clear view on the plain background of the display. Perhaps surprisingly the percentage of subjects who are unaware of the probe is even greater when the target stimulus is presented at a fixed location in the peripheral visual field, and when the probe appears at fixation where subjects are directly looking at it. Mack and Rock's explanation for this effect is that people are unaware of items at unattended locations in the visual scene. Not all subjects can (or perhaps need to) attend to the target cross exclusively and so those subjects might see the probe stimulus despite the fact that it appears at an unattended location. Some subjects, however, concentrate their attention so much on the target cross that there is insufficient attention elsewhere for them to become aware of the probe. In the case where the target cross is presented in the periphery the line length task is even more demanding than it is when the cross is presented centrally. The additional task difficulty means that there is a greater necessity to focus attention on the cross and so the probe is seen less often despite the fact that the unattended location in this version of the experiment includes the direction of gaze.

These are impressive demonstrations, perhaps slightly tempered by the fact that they do not work in all subjects. They certainly suggest that attention facilitates awareness, as does the related phenomenon of change-blindness (Rensink, O'Regan and Clark, 1997). A single demonstration of a person being aware of a stimulus which they are not attending to would, however, demolish the stronger general claim that awareness is not possible without attention. The dual-task AOC paradigm appears to provide just such a demonstration.

One criticism of Mack and Rock's work is that their paradigm does not permit the extent to which attention is withdrawn from areas away from the target to be measured. The dual-task AOC paradigm is designed to produce just such a measure. In a dual-task AOC experiment subjects are presented with a primary task whose difficulty can be adjusted so as to vary the extent to which it requires focussed attention (it might, for example be a visual search task within a small region of visual space in which the number of distractors or the similarity between the target and distractors can be adjusted). Subjects are also required to complete a secondary probe task on some trials. The probe task involves making a discrimination of stimuli presented at a location remote from the location of the primary task stimuli. The discrimination might be something like distinguishing whether a pair of discs both coloured half red and half green are in the same orientation as one another. Unlike in Mack and Rock's design the subjects in these experiments know in advance that there will be a probe task although they do not know on which trials probes will appear. It is possible to measure performance on the probe task as a function of the difficulty of the primary task. The result is known as the AOC curve. At some level of primary task difficulty the performance of the probe task falls to chance. This occurs for all sorts of relatively simple visual discriminations. It is an assumption of this research that when this level of difficulty is reached attention must be so focussed on the primary task that no attention can be diverted to the probe task. In this way the extent to which attention is withdrawn from a part of visual scene can be measured and manipulated. Li et al (2002) made a counterintuitive finding in an experiment using this task: When the AOC curve implied that attention was completely focussed on the primary task, and simple probe tasks were impossible to perform at a rate better than chance, subjects nevertheless performed extremely well on much more complex but naturalistic probe tasks. They could, for example, tell whether a briefly presented probe photograph contained an animal. Li et al. conclude that some visual tasks can

Comment [s13]: Is it obvious that there's no degree of attention left over to be captured by the appearance of an animal photo?

Comment [B14]: No, see later…

be performed in the near absence of attention. One might be tempted to extend this conclusion and infer that subjects were aware of the probe images despite not attending to them, contrary to Mack and Rock's hypothesis. This extrapolation is dangerous, however. Li et al measured subjects' performance on the probe image categorisation task, they did not explicitly measure awareness. It is well known that these sort of categorisation tasks can be performed extremely rapidly – so rapidly that the signal passing from the eyes to the hand (for making a manual response) through the brain only just has time to pass down each axon and cross each synapse in the shortest neural circuit available once (Thorpe & Fabre-Thorpe, 2001). It therefore seems likely that this categorisation response can be made before a subject has time to become conscious of the stimulus to which they responded. Categorisation performance and visual awareness might well dissociate – if so, the Li et al result does not falsify Mack and Rock's hypothesis that attention is necessary for awareness.

Comment [s15]: As I remember, Koch and Tsuchiya in their review article hedge this claim: some visual tasks can be performed in the *near* absence of attention.

Comment [B16]: The exact strength of claim varies between papers and authors.

Attention without awareness.

The question of whether attention can influence the processing of a stimulus in the absence of awareness is more complex than the question we have just considered, because, as we have seen, it may need to receive different answers, depending on which varieties of attention, and which varieties of processing, we have in mind. To make progress we need to separate out the different senses of the question that were identified above, and to address them in turn. We can begin by asking: Can an attentional cue influence target processing when one is not consciously aware of the cue's meaning?

There is some reason to think that this question should be answered in the affirmative. It is well known that quite complex relationships between stimuli can often be learned implicitly (Reber, 1993). It is also true that exogenous spatial attention (where the cue directly specifies the likely target location) occurs independently of conscious voluntary control (Spence & Driver, 1994). It would not be surprising, then, if attention could be directed without conscious voluntary control by cues whose meaning had been learned implicitly. Tony Lambert has conducted a series of experiments showing just this (Lambert & Sumich, 1996; Lambert et al, 1999).

Lambert showed that subjects' performance was modulated by symbolic cues whose type had a systematic relationship with the location at which subsequent targets were presented even when subjects were unaware of this relationship. In a simple target detection experiment, for example, the centrally presented letter 'W' might be followed more often by a target on the left than on the right whereas the letter 'S' might be followed more often by targets on the right than on the left. After sufficient exposure to the contingency subjects showed significant effects of cue validity even when, in a post-experiment questionnaire, they showed no knowledge of the contingency. This shows that cues can influence target processing even when subjects are unaware of their meaning.

Comment [ww17]: Question: is this sort of influence a top-down effect? The issue is that these experiments might also help us modify the way we think about attention and influences on attention.

Comment [B18]: The letters are being discriminated somewhere else in the brain from the likely locus of spatial facilitation. So the answer is probably yes.

We saw above that we should distinguish Lambert's question from the question of whether a cue can influence target processing when one is unable to discriminate the cue's critical properties.

Many experiments have used visual masking or brief presentation in order to impair subjects' awareness of cues in attentional tasks. Such procedures can have two effects. They might render subjects completely unaware of the cues. Alternatively, subjects may be aware that a cue was presented but be incapable of discriminating the nature of the cue. A recent example of the latter is an elegant experiment by Manon Mulckhuyse et al. (2007), showing that an uninformative peripheral cue on one side of a display whose presentation began a centi-second or so before a similar central item and a foil on the opposite side of the display speeded reaction time to a target that was subsequently presented on its side of the display (despite a lack of contingency – targets were equally likely to appear on the same of opposite side from the cue). The cue was, of course, in plain view. Subjects were aware of its presence. Nevertheless, subjects were largely incapable of discriminating on which side of the display the cue, that is the item with the slightly early onset, was presented. This indiscriminable peripheral cue also produced the slowing of reaction time known as inhibition of return when the time between cue and target presentation was long – a classic signature of exogenously controlled spatial attention, and a clear sign that these cues can influence target processing without the subject being able to discriminate the cues' critical properties.

Comment [ww19]: Could you say more what the experiment was? Task instructions etc.? Having a summary of the literature is one very helpful aspect of the discussion, so detail up front would aid reader….NB: you report task on the next page, but put in here.

Comment [ww20]: i.e. explicitly report. There is, of course, discrimination in the minimal sense of responding to the spatial location of the target.

The experiments of Lambert and of Mulckhuyse enable us to return affirmative answers to two of our questions about the attention/consciousness relation. They show that it is possible for a cue to have an attention-like influence on target processing even when one is unaware of the cue's significance, and even when one is unaware of the cue's critical properties. As we saw above, however there are several other questions to address. Let us turn now to the question of whether it is possible for a cue to influence target processing when one is unaware even of the cue's existence.

Mulckhuyse's experiment, described above, is something of an exception in that most of the experiments in which cues are masked, are briefly presented, or are presented at very low levels of contrast are experiments that aim to render those cues entirely invisible. Many such experiments have successfully shown cueing effects in response to invisible cues, both symbolic and peripheral. There is, however, considerable dispute about the mechanism through which invisible cues act and whether they are truly effective in both endogenous and exogenous attentional tasks.

Ulrich Ansorge and his colleagues have argued that the effectiveness of masked cues depends on there being some sort of match between the intentions that the cue itself would elicit in a subject and those elicited by the target that they cue (e.g. Ansorge et al. 2002). That is, the cues should, in themselves, elicit a response that is compatible with that required to indicate the presence of the target they are cueing. If this is so then the facilitative effects of wholly unseen cues has something of the feel of a priming effect, rather than a paradigmatically attentional one. Ansorge et al. (2002) also show that the processing of unseen cues differs from that of normal cues in its flexibility: If the contingency between visible cues and targets change then subjects are quick to adapt their behaviour, but subjects did not make a similar adaptation in response to changes in the contingency between unseen cues and targets.

Comment [ww21]: A bit unsure what is meant here. Could you give a concrete example from the paper? Also, the role of the cue begins to sound like a prime in the sense of disposing to action (here "elicit a response").

Ansorge's group has a wealth of data showing that response specification is an important factor in determining the effectiveness of unseen cues, and showing that the processing of such cues is less flexible than might otherwise be expected, but there is evidence that while both of these factors might influence cue effectiveness, neither are absolute. The Mulckhuyse et al (2007) experiment described previously is one in which the cues surely do *not* specify the response, because the task was one of simple detection – the response was the same whether the target dot appeared on the left or right side of the display. Some component of unseen cue processing therefore seems to be acting at a purely visual level. It also appears that the attentional system can adapt to changes in the valence of cues resulting from changes in cue-target contingencies. Lambert et al (1999) describes a task in which visible peripheral cues predict targets on their own side of the display (with 80% probability) whereas cues whose small size renders them difficult to see (some subjects could just see these small cues, they were invisible to other subjects) predict a reverse contingency – targets appear on the opposite side of the display to unseen cues with 80% probability. Subjects who were unaware of the small cues successfully learned to attend according to the reverse contingency. Subjects who were aware of the small cues failed to learn the reverse contingency – apparently they could not suppress the automatic orienting of attention to the sudden onset cue despite the fact that it carried information that the opposite location should be attended.

Comment [ww22]: The cue shouldn't specify a response, yes? Otherwise, it sounds like a prime. So, how to make the point here. Could you clarify.

Comment [ww23]: Not sure I understand this shift to "purely visual level" and how it follows from the previous point (your "therefore"). A bit more set up for the relevant contrast would be helpful.

Comment [ww24]: Break this into new paragraph.

Comment [ww25]: The "unseen" cues are the ones you just referred to as "difficult" to see, yes? So perhaps the first sentence should change "difficult" to something like "practically invisible".

Comment [ww26]: Another explanation: the other visible cues cued expectation that target appears on the same side, so the subject is focusing on this correlation when they can see the smaller cues.

I have found a similar effect in a patient with blindsight. This patient learned to attend according to a reverse peripheral contingency when unaware of the presence of cues (Kentridge et al, 1999). In some circumstances, then, changes in target contingency of unseen cues do induce adaptation. It is worth noting, however, that all of these studies (those of Mulckhuyse, of Lambert, and of Kentridge) involved detection tasks with a single response, and so do not constitute evidence that changes in unseen cue contingencies can lead subjects to switch from one response to another. Ansorge himself has demonstrated that automatic and rule-governed attentional responses to unseen cues follow different time courses and differ in their strength according to the type of response measured, amongst other things (Ansorge & Heumann, 2006). These variations may well account for the variety of results that I have described.

There has also been disagreement over the effectiveness of unseen symbolic cues. This stems from a now classic study by Peter McCormick (1997) in which he showed that, although subjects' attention could be directed by unseen peripheral cues (exogenous orienting), when those cues had a reversed spatial contingency subjects persisted in *orienting* to the cue location rather than the location indicated by the reverse contingency. The experiments of Lambert et al. and Kentridge et al., described above, both contradict this finding. The key difference may be that McCormick's subjects had relatively few trials in which they were presented with unseen cues to learn the reverse contingency. Instead, McCormick presented his subjects with blocks of practice trials with visible cues in order to allow them to learn the reverse contingency. It may be the case that learning a contingency consciously does not transfer to the processing of unseen stimuli. Certainly in Kentridge et al (1999) the subject experienced over 400 trials before attentional control by the reverse contingency began to emerge. Lambert et al.'s (1999) result, in which a seen reverse contingency could not be learned whereas an unseen one could, provides further support for the hypothesis that different learning systems, or different representations, mediate the use of conscious and unconscious cue-target contingency information.

Comment [ww27]: Perhaps I have misunderstood the experiment but as described, why isn't the fact that the subject sees the smaller stimuli confounding since if he focuses on the feature of visibility then that feature is at best at chance correlated with position (i.e. if he sees all cases of small stimuli) or if not, then weighted towards being on the same side? So, the evidence doesn't on its own point towards different learning mechanisms (at least from the Lambert result alone).

Comment [B28]: The confusion probably arises because some subjects see the small cue whereas other do not – Lambert is comparing these groups of subjects – I've added a clarification.

Unawareness of the Target

Having considered the different ways in which cues of which we are variously unaware may influence the processing of targets, we turn now to consider cases in which it is the target of which the subject is unaware. To my mind these cases are the crucial ones for our attempt to understand the relationship between attention and consciousness.

Following the lead of Bernard Baars' (1988) Global Workspace Model, or following Jesse Prinz (this volume), one might think that attention serves to select parts of the world upon which specialised mental processes (e.g. working memory) operate and one might think that the application of these processes gives rise to consciousness. If so then one will think that the things that are attended should also be consciously experienced. One might, alternatively, follow a lead suggested by Victor Lamme's (2006) model, in which attention serves to amplify a number of sensory signals, but in which it is only those signals that are strong enough to activate feedback processes initiated in pre-frontal cortex which reach consciousness. In that case one will think that attention serves to amplify mental processes which are also applied less effectively to non-attended aspects of the world, and so one will think that it should be possible to attend to something without becoming aware of it.

Comment [CRM29]: Not necessarily. Consider Jesse Prinz's view, according to which attention is the process by which stimuli get transmitted to working memory, but there is no guarantee that the transmission will be successfully received.

In 1999, following chance remarks made by our subject in the course of running another experiment, my colleagues Charlie Heywood, Larry Weiskrantz and I conducted a study in which we tested whether a patient who was unaware of visual stimuli, yet still capable of making systematic and accurate responses to them could be influenced by cues in Posner's (1980) classic attentional paradigms. The patient suffered from the rare neurological condition 'blindsight': He is capable of accurately detecting visual stimuli and of making simple discriminations about their properties, despite reporting that he is subjectively blind to these stimuli. Like other blindsight patients he describes his responses in these experiments as 'guesses', and denies seeing the stimuli to which he responds. This condition arises when primary visual cortex, or its immediate afferents, are damaged while extrastriate cortex is spared. It is thought that the residual visual function demonstrated by blindsight patients is mediated through projections from subcortical structures receiving visual input (superior colliculus, pulvinar, LGN) which terminate directly in extrastriate cortical areas without passing though striate cortex.

When we were mapping the spatial extent of blindsight in the patient GY (Kentridge, Heywood & Weiskrantz, 1997) he spontaneously mentioned that he had realised, from our discussion, that we were testing his ability to detect stimuli presented high up in the visual field and so had decided to try to attend up there. In our 1997 experiment we had no way of knowing whether GY's intention to attend to his upper visual field had any effect on his performance but his remark struck as so odd, requiring, as it did, for him to attend to things he could not see, that we set out to investigate it. We adapted Posner's classic peripheral and central cueing paradigms to make them suitable for a blindsight patient. The target, was, of course, in GY's area of blindness (his cortical damage only affected one of his visual hemifields). Our modification of Posner's paradigm was minimal. The only major change being the addition of an auditory tone that sounded at the time a target stimulus may, or may not, have been presented. GY was instructed to respond as soon as possible after the

Comment [s30]: Interesting question what he takes himself to be doing, e.g. just *thinking* about the relevant locations in the blind field? It's quite hard to make sense of, e.g. can I try to attend to what's going on in the apartment next door or in the space behind my head? Any insights would be fascinating!

Comment [B31]:

tone sounded if he guessed that a stimulus in his blind field had accompanied the tone. We selected two target locations in his blind hemifield where we knew his detection abilities were comparable and demonstrated to him in his seeing visual field where these two locations were. GY's central vision is spared so we could use arrows pointing towards one or other of the target locations as centrally presented symbolic cues. The peripheral cues were pairs of bars bracketing one or other target location. Since these locations fall in GY's blind field he was not consciously aware of these cues. We also tested a reverse contingency, in which a peripheral cue in one location was followed on 80% of trials by a target in the other location. We obtained highly significant reaction time advantages for valid over invalid cues in all three conditions. As there was no apparent increase in false positive rate for valid compared to invalid trials we had demonstrated selective spatial attention in blindsight in just the same way as Posner had in normal observers. Our only other modification to Posner's design was to ask GY to make a second response on each trial indicating whether he had had any awareness of a target during that trial. He steadfastly denied any knowledge of targets and, indeed, suggested at one point that there were no targets and that we were running some control condition. In terms of Posner's operational measure of spatial attention we had therefore demonstrated that attention could selectively modulate the processing of a target without that target entering awareness. We subsequently repeated the experiment, this time using a discrimination (rather than detection) task where GY had to guess whether the target was a horizontal or vertical line. Again, we demonstrated a significant cueing effect with no trade-off against accuracy for the faster reactions to validly cues targets. Despite the simplicity of these designs and their similarity to Posner's classic paradigm for investigating spatial attention, there is some controversy over the interpretation of our results.

Comment [ww32]: Meaning you pointed out the mirror image locations in the region where he could still see.

Comment [B33]: We just let him look at the fixation spot and the two target locations in his good field – that is, he could fixate to the right of all of them.

One very straightforward objection to concluding from our studies that attention does not necessarily give rise to awareness of the attended targets is that our only subject was a patient with a rare neurological condition who, moreover, had been the subject of extensive testing over many years. Although logically we had demonstrated that attention does not always give rise to awareness it is something of a leap to extrapolate this finding to the general population. We therefore devised an experiment in which we could investigate whether attentional cues modulated the processing of stimuli which were rendered invisible to normal observers using meta-contrast masking (Kentridge, Nijboer & Heywood, 2008). Rather than asking our normal subjects to make speeded guesses about targets they could not see we used a priming paradigm and tested whether attention modulated the effect of unseen primes on subsequently presented visible targets. Again, we were able to show that attention, directed by a visible, central, symbolic cue, modulated the processing of unseen primes in such a way that a prime in an attended location had more of an effect on the response of the subject to a visible target than did a prime in an unattended location.

In our experiments with blindsight and in those with normal observers one might worry that our effects were not mediated by shifts of attention but rather by explicit orienting responses. Such a worry would be mistaken: We monitored subjects' eye movements in all experiments and could therefore show that no overt orienting responses were elicited by our attentional cues. It also cannot be the case that our cues were serving an alerting function, generally predisposing subjects to respond to stimuli in a cued location rather than selectively enhancing sensitivity when making a decision about stimuli presented at the attended location. A non-specific enhancement of a tendency to respond to stimuli in a cued location cannot explain the reversed contingency effect we observed in our 1999 blindsight experiments.

Mole (2008) argues that facilitating processing of a stimulus in a specific spatial location through selective spatial attention is not the same as attending to the stimulus itself. According to this position our subjects are unaware of objects, but they are not attending to those objects. They are just attending to the region of space in which those objects fall. There is, however, good evidence that spatial attention only serves to enhance goal-specific properties of stimuli falling in attended space (Remington & Folk, 2001). If, for example, a subject's task involves discriminating the colour of a target then spatial attention will affect sensitivity to the colour of objects in attended space but have no effect on processing their shape. It is hard to reconcile the notion that stimuli in attended space are not themselves being attended with the fact that only certain properties of objects in attended space are subject to enhanced processing. Those selectively enhanced properties are properties of objects, not properties of space. The most compelling evidence for this position is from experiments in which spatial attention affects the perception of attended objects rather than simply speeding responses to them. It is well known that spatial attention can enhance constrast sensitivity, acuity and even the perceived hue of attended stimuli (Carrasco, Ling and Reed, 2004) and, indeed, that these perceptual changes can be accompanied by changes in neural responses (Liu, Pestilli & Carrasco, 2005). Carrasco, Williams and Yeshurun (2002) have even shown that spatial attention enhanced acuity to masked stimuli.

Comment [ww34]: Couldn't you also respond to Chris by appeal to the invalid trials in your original experiments. In those cases, the cue presumably directs attention to a specific region of space but the object appears elsewhere. In those cases, the detection of stimuli is still above chance. It would seem that the target itself is drawing the subject's attention despite the effects of the cue which pulls the subject's spatial attention away form the target. To explain why the subject can answer correctly in the invalid cue case, the best explanation is that the object is drawing attention.

The results I have discussed so far apply only to spatial attention and, in normal observers at least, only to central symbolic cueing. In the time between our experiments with blindsight and our masking experiment in normal observers a number of other groups investigated the relationship between other forms of attention and target awareness. Kanai et al (2006) used the continuous flash suppression paradigm to render normal subjects unaware of stimuli presented to one of their eyes. They investigated feature-based attention in which attending to an item that has a specific feature modulates processing of other items sharing the same feature. They were able to show that attending to a seen item possessing a specific feature selectively enhanced processing of unseen items sharing the same feature when compared to unseen items possessing a different feature. Kanai et al (2006) also attempted to apply their methods to evaluating the effects of spatial attention in normal observers. They, unlike Kentridge et al (2008), were unsuccessful. There are, however, important procedural differences between these studies which are likely to explain the contradictory results. First, the masking procedures are different and probably operate through different neural mechanisms. Second, the effects of attentional modulation are measured differently. Kanai et al (2006) measured changes in the magnitude of the tilt after-effect with and without attention. The tilt after-effect depends heavily on the selective adaptation of neural responses in striate cortex (Jin et al, 2005) and may even be subcortically mediated (Ye, Li, Yang & Zhou, 2009). Kentridge et al (2008) used a colour priming measure. Again, the neural underpinnings of colour priming and orientation adaptation are likely to be very different.

Sumner et al (2006) conducted a masked priming experiment in which peripheral, rather than symbolic cues directed the attention of normal observers to an unseen prime. Critically, Sumner et al primed a discrimination task in which there was a strong sensori-motor link between the target and the response: Subjects had to press a button matching the direction in which a target arrow pointed. In such tasks, although a clearly seen prime produces a normal priming effect, one often observes a negative priming effect with weak primes. Sumner et al demonstrated that attention modulated this negative priming effect with unseen masked primes. Importantly, they were also able to demonstrate that this effect could not be solely attributed to an effect of attention on perceptual salience. Attention must have modulated the efficacy of the target in eliciting its effect upon action – it must have modulated a sensori-motor process, not only a perceptual one.

It remains to be seen whether other forms of attentional selection, such as object-based attention, or functions that may depend on attention, such as feature binding, occur in the absence of awareness. There is one piece of evidence that suggests that binding may occur in the absence of attention. Wojciulik & Kanwisher (1998) tested a patient with the neurological condition simultanagnosia in a 'Stroop' task (in which the meaning of a word interferes with naming the colour of ink in which it is written). This effect must depend upon the shape and the colour of the stimulus being bound together. This binding together of the different features of an object on the basis of common spatial location has long been thought to depend upon attention (Triesman & Gelade, 1980). Wojciulik & Kanwisher found that their patient was incapable of making explicit judgements about binding – when shown the words 'brown' and 'green' written in green and brown ink he could not say which word was written in which colour ink. When asked to name the ink colour of the upper work of the pair he did, however, show the Stroop effect, suggesting that unconscious binding is possible.

In summary, despite many possible objections, behavioural evidence suggests that attention and awareness are mediated by distinct psychological processes both in terms of cue and target processing.

Neurophysiological Evidence.

The existence of dissociations between attention and awareness has particular consequences for the search for the neural correlates of consciousness. If attention and awareness did not dissociate then one could substitute relatively straightforward and objective manipulations and measures of attention in place of ambiguous and subjective measures and manipulations of consciousness. If the behavioural evidence I have summarised above is correct, however, it should be possible (although perhaps not easy) to demonstrate the existence of separate neural correlates of attention and of consciousness.

Many neural models of consciousness do draw a distinction between attention and consciousness. Lamme (2006), for example, distinguished between (1) effects of attention on the initial response to a stimulus as activity proceeds forward from primary visual cortex to more anterior regions, and (2) a putative consciousness-eliciting signal, feeding back from frontal cortex to occipital sensory areas. Finding direct evidence to test these models is difficult but recent evidence distinguishing neural responses to masked and unmasked signals (e.g. Gaillard et al., in press) does reveal that different patterns of neural response are elicited by the stimuli which affect behaviour depending on whether or not they are seen consciously. In the present context, however, the critical question is whether the difference in neural response between attended and unattended stimuli and the difference between seen and unseen stimuli are distinguishable. That is, is there a dissociation between the neural correlates of attention and the neural correlates of consciousness?

Two experiments conducted in Catherine Tallon-Baudry's laboratory suggest just such a neural dissociation between attention and awareness. In one series of experiments (Schurger et al., 2008) she tested a blindsight subject using a procedure (similar to that of Kentridge et al., 2004) in which a central cue directed the patient's attention to the location at which a line whose orientation must be discriminated was most likely to appear. The key difference was that the targets were much more salient than those used by Kentridge et al. Instead of being low contrast lines, ramped gradually on and off, Schurger et al. used high contrast abruptly onsetting checked bars. Although their subject GY usually denies seeing stimuli in his blind visual field he is sometimes aware that visual events have taken place in his blind field, particularly if those events involve high-contrast rapidly changing stimuli. It was therefore possible to categorise targets in these experiments in two different ways. First some targets were validly cued, and hence attended, whereas others were misleadingly cued, and hence unattended. Second, given the high contrast abrupt nature of the target, GY reported some awareness of a blind field event for some targets but not others. It was therefore possible to determine whether the same differences in neural events occurred when comparing the two attention conditions and the two awareness conditions. Schurger et al measured GY's neural activity on every trial using magnetoencephalography (MEG). They discovered that the MEG signal that correlated with awareness had an earlier onset, and a higher temporal frequency of neural activity, than that which correlated with attention, the latter signal having an onset about 150ms later, and having a much lower dominant temporal frequency.

In just the same way as one might object to drawing general conclusions from behavioural tests on GY one might also object that the dissociation between neural correlates of attention and awareness that Schurger et al. find is some peculiar characteristic of GY. More recently, however, Wyatt and Tallon-Baudry (2008) repeated this experiment using normal observers and target stimuli that were titrated in contrast so as to fall at the threshold of awareness: Subjects were aware of some stimuli and failed to see others. Again, the neural correlate of awareness was earlier and had a higher dominant frequency than the correlate of attention. In this experiment Wyatt and Tallon-Baudry could also measure a distinct response to the cue, which they interpret as an orienting signal, and which had a much lower frequency and quite distinct topography when compared to either the attentional or awareness modulations of the target response.

Conclusion.

The conclusions one draws from this wealth of data still depend upon how one defines the terms ‘attention’ and ‘awareness’. If we take the meaning of ‘attention’ to be that used by experimental psychologists, and if we measure awareness in terms of subjects’ willingness to report having seen stimuli, or their ability to succeed in unbiased forced choice detection tasks of those stimuli, then it seems unlikely that attention and awareness are two views of the same coin. It is clear that attention can modulate the processing of stimuli which remain unseen. This behavioural dissociation is echoed in a dissociation between the neural correlates of attention and awareness. It is also clear that awareness of attentional cues, or of their meaning, is not necessary for those cues to guide attention. None of this is to say that attention and awareness are not related. It certainly seems to be true that attention enhances aspects of awareness (see, e.g., Carrasco, Ling and Reed, 2004). But, although attention and awareness may be related, they are not the same.

Comment [s35]: Does that mean your claims are not in conflict with those who claim that attention is sufficient for consciousness, but who understand attention not in terms of the operational definition used in psychology, but as something like James' focalization of consciousness or Wundt's narrow region of the content of consciousness?

References

Ansorge, U. & Heumann, M. (2006) Shifts of visuospatial attention to invisible (metacontrast-masked) singletons: Clues from reaction times and event-related potential. Advances in Cognitive Psychology, 2, 61-76.

Ansorge, U., Heumann, M. & Scharlau, I. (2002) Influences of visibility, intentions, and probability in a peripheral cuing task. Consciousness & Cognition, 11, 528–545.

Baars, B.J. (1988), A Cognitive Theory of Consciousness. Cambridge, MA: Cambridge University Press.

Block, N. (2007) Consciousness, accessibility, and the mesh between psychology and neuroscience. Behavioral & Brain Sciences, 30, 481-548.

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

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

Drew, T., McCollough, A., Horowitz, T., & Vogel, E. (2009). Attentional enhancement during multiple-object tracking. Psychonomic Bulletin & Review, 16, 411 - 417.

Gaillard R., Dehaene S., Adam C., Clémenceau S., Hasboun D., Baulac, M., Cohen, L. & Naccache, L. (2009) Converging Intracranial Markers of Conscious Access. PLoS Biol 7: e1000061. doi:10.1371/journal.pbio.1000061

Giersh, A & Caparos, S. (2005) Focused attention is not enough to activate discontinuities in lines, but scrutiny is. Consciousness & Cognition, 14, 613-32.

Giersh, A. & Fahle, M. (2002) Modulations of the processing of line discontinuities under selective attention conditions? Perception & Psychophysics, 64, 67-88.

Goldman, A.I. (1997) Consciousness, Folk Psychology, and Cognitive Science. in N. Block, O. J. Flanagan & G. Güzeldere (Eds.) The Nature of Consciousness. Boston MA: Bradford Books.

James, W. (1890) Principles of Psychology. London: Macmillan.

James, W. (1892) Psychology. Cleveland & New York, World.

Jin, D.Z., Dragoi, V., Sur, M. & Seung, J.S. (2005) Tilt aftereffect and adaptation-induced changes in orientation tuning in visual cortex. J. Neurophysiol. 94, 4038-4050.

Kanai, R., Tsuchiya, N. & Verstraten, F.A. (2006) The scope and limits of top-down attention in unconscious visual processing. Curr. Biol., 16, 2332–2336.

Kentridge, R.W., Heywood, C.A. & Weiskrantz, L. (1997) Residual vision in multiple retinal locations: Implications for blindsight. Journal of Cognitive Neuroscience, 9, 191-202.

Kentridge, R.W., Heywood, C.A. & Weiskrantz, L. (1999) Attention without awareness in blindsight. Proceedings of the Royal Society (London) Series B, 266, 1805-1811.

Kentridge, R.W., Heywood, C.A. & Weiskrantz, L. (2004) Spatial attention speeds discrimination without awareness in blindsight. Neuropsychologia, 42, 831-835.

Kentridge, R.W., Nijboer, T.C.W. & Heywood CA (2008) Attended but unseen: Visual attention is not sufficient for visual awareness. Neuropsychologia, 46, 864-869.

Lambert, A., Naikar, N., McLachlan, K., & Aitken, V. (1999). A new component of visual orienting:Implicit effects of peripheral information and subthreshold cues on covert attention. Journal of Experimental Psychology: Human Perception and Performance, 25, 321–340.

Lambert, A., & Sumich, A.L. (1996). Spatial orienting controlled without awareness: A semantically based implicit learning effect. Quarterly Journal of Experimental Psychology, 49A, 490–518.

Lamme, V.A.F. (2006) Towards a true neural stance on consciousness. Trends in Cognitive Sciences, 10, 494-501.

Li, F-F., VanRullen, R., Koch, C. & Perona, P. (2002). Natural scene categorization in the near absence of attention. Proc Natl Acad Sci USA 99, 9596-9601.

Liu, T., Pestilli, F., & Carrasco, M. (2005) Transient attention enhances perceptual performance and fMRI response in human visual cortex. Neuron, 45, 469-477.

Mack, A. & Rock, I. (2000) Inattention Blindness. Boston MA: MIT Press.

McCormick, P.A. (1997) Orienting attention without awareness. Journal of Experimental Psychology: Human Perception and Performance, 23, 168–180.

Mole, C. (2008) Attention and Consciousness. Journal of Consciousness Studies, 15, 86-104.

Mulckhuyse, M., Talsma, D. & Theeuwes, J. (2007) Grabbing attention without knowing: Automatic capture of attention by subliminal spatial cues. Visual Cognition, 15, 779-788.

O'Regan, J.K. & Noë, A. (2001) A sensorimotor account of vision and visual consciousness. Behavioral & Brain Sciences, 24, 939-1031.

Posner, M.I. (1980). Orienting of attention. Quarterly Journal of Experimental Psychology, 32, 3–25.

Posner, M.I. & Raichle, M.E. (1994) Images of mind. New York NY, W.H. Freeman.

Pylyshyn, Z., Haladjian, H., King, C., & Reilly, J. (2008). Selective nontarget inhibition in multiple object tracking (MOT). Visual Cognition, 16, 1011 - 1021.

Reber, A.S. (1993). Implicit learning and tacit knowledge: An essay on the cognitive unconscious. Oxford: Oxford University Press.

Remington, R.W. & Folk, C.L. (2001) A dissociation between attention and selection. Psychological Science, 12, 511–15.

Rensink, R., O'Regan, J. K. & Clark, J. J. (1997) To See or Not to See: The need for attention to perceive changes in scenes. Psychological Science 8(5), 368-373

Schurger, A., Cowey, A., Cohen, J.D., Treisman, A. & Tallon-Baudry, C. (2008) Distinct and independent correlates of attention and awareness in a hemianopic patient. Neuropsychologia, 46, 2189–2197.

Spence, C., & Driver, J. (1994). Covert spatial orienting in audition: Exogenous and endogenous mechanisms. Journal of Experimental Psychology: Human Perception & Performance, 20, 555–574.

Sperling,G. & Dosher, B.A. (1986) Strategy and optimization in human information processing. in K. Boff, L. Kaufman & J. Thomas (Eds) Handbook of perception and performance. Wiley: New York.

Sumner, P., Tsai, P., Yu, K & Nachev, P. (2006) Attentional modulation of sensorimotor processes in the absence of perceptual awareness. Proceedings of the National Academy of Science USA, 103, 10520–25.

Thorpe SJ, Fabre-Thorpe M. 2001. Neuroscience. Seeking categories in the brain. Science 291: 260-3.

Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. Cognitive Psychology, 12, 97–136.

Wojciulik, E. & Kanwisher, N. (1998) Implicit but not Explicit Feature Binding in a Balint's patient. Visual Cognition, 5, 157-181.

Wundt, W. (1912) An Introduction to Psychology. London: George Allen & Unwin.

Wyatt, V. & Tallon-Baudry, C. (2008) Neural Dissociation between Visual Awareness and Spatial Attention. Journal of Neuroscience, 28, 2667–2679.

Ye, X., Li, G., Yang, Y. & Zhou, Y. (2009) The effect of orientation adaptation on responses of lateral geniculate nucleus neurons with high orientation bias in cats. Neuroscience, 164, 760-769.