The Influence of Selective Attention on Consciousness☆

Piers DL Howe, University of Melbourne, Melbourne, VIC, Australia
Karla Evans, The University of York, Heslington, United Kingdom
Riccardo Pedersini, Brigham and Women's Hospital, Cambridge, MA, United States
Michael A Cohen, Massachusetts Institute of Technology, Cambridge, MA, United States
Todd S Horowitz, National Cancer Institute, Rockville, MA, United States
Jeremy M Wolfe, Brigham and Women's Hospital, Cambridge, MA, United States

Glossary

Balint's syndrome A neurological disorder caused by bilateral damage to the parietal cortex that results in the inability to perceive more than one object at a time.
Binding problem The difficulty of mentally conjoining features that belong to the same object.
Blindsight A neurological disorder, typically caused by damage to the primary visual cortex, that results in the patient being unaware of any stimuli located in a particular part of the visual field while still being able to detect them.
Feature An attribute of an object (e.g., its color).
Illusory conjunction An illusory combination of features from different objects.
Receptive field The region of the retinal image to which a particular neuron responds.
Pop-out In visual search, the situation where the speed or accuracy of target detection is independent of the number of distracters.

Introduction

We are aware of the world around us, but not in a uniform fashion. We selectively attend to some stimuli and are consequently less aware of others. You are probably sitting, at present. If you selectively attend to the pressure of your posterior on the seat, you will become more aware of that sensation than you were a moment before. This article is concerned with the extent to which selective attention influences conscious awareness. There are three possibilities. It could be that selective attention has no effect on our conscious awareness. As the opening example makes clear, this is not a promising hypothesis and we will ignore it. At the alternative extreme, it has been proposed that conscious awareness is fully determined by selective attention—that we are conscious only of the current contents of attention. A more moderate position is that attention modulates awareness, but we have some awareness of unattended stimuli. This middle position reflects our view.

☆*Change History*: October 2015. PDL Howe, K Evans, R Pedersini, MA Cohen, TS Horowitz, and JM Wolfe updated the text, modified the title, inserted an Abstract and Keywords.

At the outset, justifying this position is made difficult by the many uses of terms like "consciousness" and "attention" in common speech and technical writing. In this article, we will restrict ourselves to the conscious awareness of visual stimuli, though the same questions arise in the other senses, between sensory domains and perhaps even when monitoring one's own thoughts.

We will often use the term "object"—another term with a problematic definition. For example, consider an image of a face. One could consider the entire face to constitute a single object. Alternatively, one could consider the eyes, the nose, the mouth, and so on to be objects. Indeed, each of these objects could in turn be decomposed as you attend to, say, a pupil or a nostril. As there is no general agreement on what constitutes an object, we refrain from attempting to precisely define it. Instead, we ask for the reader's indulgence and use the term as a layman would—imprecisely.

Colloquial speech tends to incorrectly treat "awareness" (like "attention") as a single entity. In fact, we can profitably distinguish between the type of awareness that accompanies attention and the type of awareness that seems to occur in the absence of attention. This is an old idea. In 1780, Étienne Bonnot de Condillac asked his readers to imagine arriving at a chateau late at night. The next morning, you wake in a completely darkened room. Then the curtains are thrown open for just a moment on the scene out the window with its farms, hills, forest, etc. Condillac argued that you would initially see something, perhaps just patches of color, throughout the scene, but you would be unable to identify what you were seeing until you had directed attention to different parts of the scene. Condillac's patches of color are what we are calling the awareness in the absence of attention. We contrast this level of awareness with that obtained from attending to one of those colored patches and consequently realizing that it represents, say, a meadow in the summertime.

We will argue that in the absence of attention we can, at most, be aware of object attributes but not how they are related. For example, if a object is composed of a red vertical bar and blue horizontal bar, then, in the absence of attention, we might be aware that there was a vertical bar and a horizontal bar and that there was red and blue. However, we would not know which bar was which color. To be able to relate (or "bind") a bar's color to a bar's orientation requires that the bars be attended. To understand why this might be the case we need to consider the "binding problem."

Feature Integration Theory, Object Recognition, and Awareness

When we attend to an object, we usually feel that we are aware of multiple features of that object. For example, we might be aware of a round, red, revolving disk. That type of awareness requires that we "bind" the roundness, redness, and motion to the same object. The neurons that analyze different attributes of an object are often located in different regions of the brain. Consequently, binding features together to form a coherent representation poses a problem. In a world filled with many objects, often in close proximity, how do we know that the red computed in this part of the brain goes with the motion analyzed in this other part? This issue is known as the binding problem. One proposed solution is that it is the act of attending to an object that allows different features of the same object to be conjoined and features from other objects to be excluded. Indeed, this may be the main function of selective attention.

In principle, our brains could have been constructed in such a way that we would not suffer from the binding problem. For example, the optic tectum of the common toad (*Bufo bufo*) contains a class of "fly-detector" neurons that signal the location of small, moving black dots. The method effectively avoids the binding problem because the toad can detect the fly directly without having to first measure the fly's individual attributes such as its motion, color and size. Unfortunately, this method allows for the detection of only a small number of different types of objects, as it needs a dedicated group of neurons for each type of object that it is to detect.

Human visual systems (indeed mammalian visual systems, in general) have a flexible ability to represent arbitrary combinations of attributes like color, size, orientation, and so forth. In order to understand the relationship of attributes, these visual systems have had to solve the binding problem.

Our understanding of the structure and function of the visual system (Fig. 1) is obtained from multiple sources including neuroimaging techniques like functional magnetic resonance imaging (fMRI) in human observers and more invasive neuroanatomical and neurophysiological methods performed mainly in animal models such as the cat and the monkey.

Visual information flows from the retina to the lateral geniculate nucleus (LGN) of the thalamus. The LGN in turn relays that information to the primary visual cortex (V1) located on the rear surface of the brain, mostly inside the calcarine fissure. From here the pathway divides, with the dorsal and ventral streams being particularly important subdivisions. The dorsal stream includes visual areas V1 and V2, the middle temporal area (MT) and the medial superior temporal area (MST). The ventral stream also includes areas V1 and V2 and then proceeds to area V4 and to areas in the inferior temporal cortex (IT). As one progresses along either stream, the neural activity become less purely stimulus-driven, more readily modulated by changes in attentional state and increasingly likely to mirror the reported conscious percept.

The dorsal stream is often referred to as the "where" pathway as it is particularly sensitive to spatial information. For example, area MT is especially sensitive to the motion of an object. To a large extent, the impression of the object's motion is closely related to the activity in this area. Electrical microstimulation of neurons in MT in macaque monkeys influences their judgment of motion. Damage to MT can cause akinetopsia, an inability to perceive motion. Suffers of this condition can report that an object was in one position and is now in another. However, they have no conscious perception of the movement of the object. Conversely, for those that do not suffer from akinetopsia, it is possible to have a conscious percept of motion even when the visual stimulus does not move. For example, if one stares fixedly at a coherent moving pattern, such as a waterfall, and then fixates on a stationary object, the stationary object will appear to move (a motion aftereffect) and MT

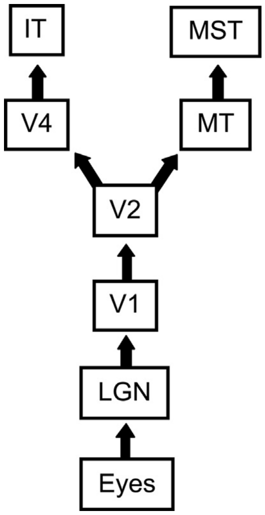

Figure 1 A schematic representation of the ventral (left) and dorsal (right) visual pathways. Also known as the “what” and “where” pathways. Abbreviations are explained in the text.

will be activated. If transcranial magnetic stimulation (TMS) is used to prevent MT from activating, then the motion aftereffect is not seen.

The ventral stream is often referred to as a “what” pathway, as it is particularly sensitive to the identity of the object. For instance, activity in IT closely reflects the subject’s impression of the object’s shape. This was elegantly demonstrated by David Sheinberg and Nikos Logothetis in a series of neurophysiological recordings in macaque IT. First, they would isolate a neuron and find an image to which it responded strongly and one that did not excite it. Then they would present one of these images to one of the monkey’s eyes while simultaneously presenting the other image to the other eye. The monkey had been trained to pull a lever to indicate which image it saw. As with humans, the percept reported by the monkey alternated between the two images, even though the images themselves were constant, a phenomenon known as “binocular rivalry.” They found that almost every IT cell responded only when the monkey reported seeing the image that had previously been shown to excite that cell. Crucially, these cells did not respond when the monkey reported seeing the cell’s non-preferred image even though the preferred stimulus was still present on the retina of the other eye. The activity of these cells reflected conscious perception, as opposed to the unchanging retinal image.

It should be stressed that the activity in MT and IT does not cause the perceptual awareness of motion and shape respectively. Indeed, when monkeys are rendered unconscious by anesthetic, MT and IT continue to be active. Instead, it seems that when a monkey is conscious of an object, much of its awareness of motion and shape is reflected by the activity in these areas. For present purposes, the important observation is that when you see a moving object, its shape and motion are typically bound into a single coherent percept. In this physiological framework, the binding problem is the problem of understanding how motion information from MT, shape information from IT and various other bits of information from other visual areas come to be unified in a bound percept.

Feature Integration Theory

Anne Treisman’s Feature Integration Theory (FIT), first proposed in 1980, holds that attention is critical to the formation of bound representations of objects and, by extension, it proposes that attention is critical to our conscious experience of those bound representations. In FIT, following the understanding of visual neurophysiology given above, the visual system first decomposes the visual scene into its composite features, arrayed in a set of “feature maps.” The pre-attentive description of a scene or object comprises a list of such features. The term “pre-attentive” has been controversial but it can be operationally defined here as the representation of a stimulus before selective attention is directed to that stimulus.

In FIT, the approximate position of each feature is recorded on its pre-attentive feature map. For example, if the visual scene contains two red objects, the feature map corresponding to redness would be activated at two points roughly corresponding to the locations of the red objects. If each feature were associated with a precise region in space, this might solve the binding problem. Features that correspond to the same region in space could be automatically conjoined thus guaranteeing veridical perception. Unfortunately, the location of many features is measured in an imprecise fashion. For example, the smallest receptive fields in IT, the region whose activity correlates well with shape perception, have a spatial extent of a few degrees of visual angle. Within this region, the cell will respond to an object in an approximately translation invariant manner. Thus, a neuron in IT cannot signal the location of a particular shape with a precision of better than a few degrees, while the perception of coherent objects requires a much finer resolution.

Because of the poor resolution of these feature maps, if two objects are close together, then there is the potential that the features from one object may become conjoined with the features of the other object thus creating a percept of an object that did not in fact exist. For example, if the visual scene contains a red vertical bar and a blue horizontal bar then one might see a blue vertical bar and a red horizontal bar. Such inappropriate combinations of features are known as illusory conjunctions. FIT suggests that attention hinders the formation of illusory conjunctions.

Supporting this assertion is a series of classic experiments by Treisman and her colleagues showing that, if attention is occupied elsewhere, illusory conjunctions are, in fact, reported. In one version of the experiment, observers viewed a display of five characters aligned horizontally. The outer two characters were always digits and the inner three characters were always letters. While the digits were always black, the letters were colored. The observer's primary task was to name the digits. After doing that, the observer reported the letters and their associated colors. When the display was presented sufficiently rapidly, observers would often report seeing an incorrect conjunction of a color and a letter. For example, if the display contained a red X and a green T, they might report seeing a red T. Crucially, these illusory conjunctions occurred at a much higher rate than could be attributed to the observer simply misperceiving a given feature. Generally, the observer correctly perceived the features present in the display. It was the conjoining of features that proved to be problematic.

When asked to report how confident they were that they had actually seen an object, observers were just as confident when they reported seeing an illusory conjunction as they were when they correctly reported the features of an object. Indeed, although all observers were told that the digits would always be black (and in fact always were) about half the observers spontaneously reported that the digits sometimes appeared to be colored, sometimes even going as far as to argue with the experimenter about the issue! This raises an interesting problem in the study of attention and awareness. In tasks of this sort, one can only ask about what *was* seen, after the fact. If one asks about the current status of a visible object, the observer will attend to it in order to answer the question and will be unable to give an accurate report of the unattended state. Nevertheless, the phenomenology of illusory conjunctions does show that, within a fraction of a second of the disappearance of a display, observers can be quite convinced that they have seen something that was not, in fact, present. Subsequent studies have shown that illusory conjunctions can be perceived even when the subject attends to the objects, especially if the objects are perceptually grouped. Clearly, attention does not always succeed in solving the binding problem.

There is neuropsychological evidence, from studies of patients with Balint's syndrome, which supports the idea that attention can inhibit the formation of illusory conjunctions. This syndrome occurs when both the left and right parietal lobes are damaged. As these areas help govern the deployment of attention, such patients have great difficulty in directing their attention to a given object, resulting in the inability to perceive more than one object at a time. As would be expected, they are also prone to suffer from illusory conjunctions, experiencing them even when the image is displayed for several seconds.

Neurophysiological support also comes from work by Robert Desimone and colleagues. They performed a series of extracellular studies in area V4 of the macaque monkey that have shown that attention can help solve the binding problem. First they would find a stimulus that, when presented on its own, would elicit a strong response from the neuron in question (the preferred stimulus), and another that would elicit only a weak response (the non-preferred stimulus). They would then present both stimuli simultaneously so that they both were within the neuron's receptive field. In the absence of attention, the cell would simultaneously respond to both stimuli, with its response (spike rate) lying between that generated by each stimulus when presented on its own. In other words, the response reflected contributions from both stimuli, meaning that the cell could not distinguish between the two. However, when the monkey attended to one of the stimuli, the situation changed and the cell responded primarily to the attended stimuli. Specifically, when the monkey attended to the preferred stimulus, the cell would respond strongly, but when the non-preferred stimulus was attended, only a weak response was elicited. In this case, attention was able to solve the binding problem, at least at the neuronal level, by shrinking the receptive field of the cell to include just the selected item, thereby removing the influence of the unattended item.

This constriction of the receptive field does not explain how signals about one feature analyzed in one cortical area can be bound to signals about another feature from another area. Other mechanisms have been suggested to account for this aspect of binding. Several of these are based on the idea that neurons in different cortical areas that respond to the same object synchronize their activity, so that they create action potentials at the same time. Consequently, a third brain area could determine whether two neurons in two different parts of the brain are responding to different features of the same object by being sensitive to this synchrony. As attention is known to increase neural synchrony, theories based on synchrony are consistent with the notion that attention is needed to solve the binding problem.

An alternative theory put forward by Edward Vul and Anina Rich suggests that features are not directly bound together but rather are consciously perceived as being bound when they are perceived to share a location in space and time. Attending to a feature allows it to be more precisely localized, thereby increasing the chances that the correct features are perceived together and decreasing the chances of illusory conjunctions forming. Thus they propose that attention solves the binding problem only in an indirect fashion by increasing the precision by which features can be localized in space and time.

Features

While it is easy to say that the visual system decomposes a visual object into its constituent features, it is harder to be precise about what this statement might mean. In particular, there is imperfect agreement about the list of features that might be available to be bound. Various tests have been proposed, most of them based on the premise that features can be analyzed to some extent in the

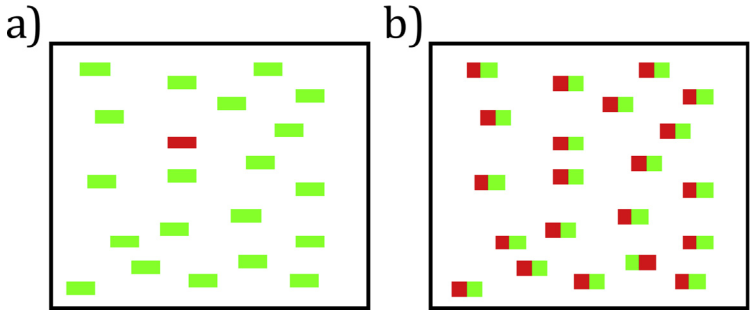

Figure 2 Two visual search experiments. (A) The target (the *red rectangle*) is easy to find. (B) The target (the *green-on-the-left-red-on-the-right rectangle*) is hard to find.

absence of attention. In particular it has been assumed that in the absence of attention one can detect and localize examples of a particular feature. For example, if an item is the only item in a display that has a particular feature, that item will tend to "pop-out" of the display, automatically summoning attention to it (as long as the other items are not too similar to that target item nor too different from each other). According to this logic, by measuring which items pop-out one can determine what the features are Fig. 2A shows a case where the target has a unique feature. It is the only red object in the scene. Consequently, it pops-out and can be located very quickly, independent of the number of other items. Conversely, in Fig. 2B, the target and distracters share the same features, so the target is not uniquely defined as the only object with a particular feature. The target is the rectangle that is green on the left but red on the right whereas the distracters are red on the left and green on the right. In this case, finding the target is a slow process.

Texture segmentation is another test. Consider two regions of a display, one with a putative feature and the other without. If a border between those regions can be effortlessly detected, indicating that the process can occur in the absence of attention, one could declare that there is a feature difference that permits the segmentation. Unfortunately, these and other methods for identifying features agree imperfectly. It is quite clear that some attributes, like color, motion, and orientation pass all tests. Other attributes (e.g., various aspects of form) are more problematic.

Feedforward Models of Object Recognition

Even if attention is needed for binding, it is not necessary—and probably incorrect—to hold that attention is always needed for object recognition. A class of feedforward models of object recognition show how some quite sophisticated object recognition could occur without explicitly invoking selective attention. As we will see later, this fits well with evidence that humans have some ability to detect objects in one part of the field even when their attention is occupied elsewhere. For example, Maximillian Riesenhuber and Tomaso Poggio developed a theory of object recognition based on an idealization of the hierarchy of the monkey visual system. As one progresses through the monkey visual system, cells become selective for increasingly complex visual stimuli. For example, in the LGN some cells have an on-center off-surround receptive field organization. The stimulus that most excites these cells is a spot of light on a black surround. On entering the primary visual cortex, we find cells whose optimum stimulus is more complex, perhaps a bar of a particular orientation and length. In V2, there are cells whose optimum stimulus is 2 bars in a particular configuration. In V4, some cells are most excited by a collection of bars joined together in a particular manner. Riesenhuber and Poggio were able to build a feedforward model that was able to explain how these selectivities were generated. Crucially, this was achieved without invoking any feedback mechanisms. Since attention must be mediated by feedback, they argued that this showed that at least some recognition can occur in the absence of attention. The original model was applied to shapes that resembled bent paper-clips. In subsequent work, they and others have developed models that can recognize objects like cars and faces. Does this mean that attention is unnecessary? Probably not since the models tend to fail when there are many objects in the display. These models do show that it is possible in theory to have some degree of recognition without attention. This, in turn, makes it plausible that one might have some awareness an object even if that object was not the target of selective attention.

Reverse Hierarchy Theory

Reverse Hierarchy Theory (RHT), proposed by Shaul Hochstein and Merav Ahissar, is an example of a model combining feedforward and feedback components. The feedforward component is similar to the Riesenhuber and Poggio model. It explains how the hierarchy of the visual system allows for visual scenes to be processed to some degree in a feedforward manner in the absence of attention. It is the feedback component that allows for more detailed perception to occur (hence reverse hierarchy). For example, it suggests that to appreciate fine differences in orientation, the brain would deploy feedback from a high level representation of the stimulus in order to pay attention to the detailed orientation information held in cells in the early visual cortex. One interesting aspect of this proposal is that it suggests that the high-level information (e.g., animal or face) might reach awareness before information about low-level features.

To summarize the argument to this point; precise binding of features to objects is a problem that the visual system seems to solve by the use of selective attention. It follows that attention is required for awareness of those bindings and for awareness of object identities that rely on those bindings. However, the example of Condillac's chateau indicates that there will be awareness of something in regions not yet visited by selective attention. Moreover, feedforward models show that, at least in theory, the unattended awareness of something need not be limited to raw local features. Some quite sophisticated analysis and awareness might be possible away from the current focus of attention. It is to that awareness without attention that we turn to in the next section.

The Relationship Between Attention and Awareness

In this section, we wish to distinguish between the hypothesis that some awareness occurs outside of the current focus of attention and the hypothesis that we are aware only of the current contents of attention. Recall the phenomenon, described above, of illusory conjunctions in which observers correctly report the colors and letters in a display but fail to correctly report which letter goes with which color. It could be that the experience of unbound colors and letters represents awareness without attention. However, there is a contrary point of view. Perhaps the imperfect awareness of the letters arises from imperfect attention to the letters. There is no guarantee that naming the digits in an illusory conjunction experiment withdraws all attention from the letters. Perhaps if all attention had been really withdrawn from the letters then the observers would not have been able to report any features of the letters at all. That would be the prediction if awareness cannot occur in the complete absence of attention.

When Awareness Requires Attention

Although we will argue against this extreme viewpoint, we will outline a set of phenomena that have been used to argue for this strong link between attention and awareness: inattentional blindness, change blindness, and the attentional blink.

Inattentional blindness was first described by Arien Mack and Irv Rock. They had observers performing an attentionally demanding perceptual task (e.g., which of two lines is longer?). On one critical trial, the briefly presented display contained an unexpected item. Observers were frequently unable to report that it had been presented. Any awareness of that item left no trace that could be reported after it was gone. Perhaps, attention to the primary task, prevented irrelevant items from ever rising to conscious awareness. As an experimental tool, one problem with this task is that it produces only one trial per observer. Once you ask about the unexpected item on one trial, other unexpected items on other trials tend to be successfully reported.

Change blindness is a more resilient phenomenon. While the phenomenon was initially discovered in the late 1950s and early 1960s, a major renaissance on the topic emerged in the mid-1990s. Dan Simons and Dan Levin, as well as Ronald Rensink and Kevin O'Regan and their colleagues presented observers with complex natural scenes (e.g., a photo of an airplane on the tarmac) and measured the ability to detect fairly large changes to these scenes (e.g., the plane's engine disappearing and reappearing). Critically, the visual transient generated by the change was masked by an eye movement, a brief blank interval, or some other visual transient. Observers thus had to actually detect the change in the image, rather than just the transient caused by the change. Change detection under these circumstances turns out to be very difficult. Observers could fail to notice changes even though they spent many seconds examining both versions of the display. If an object was attended during the transition between two frames, the change could be noted. Otherwise observers were unable to report it. One interpretation of these data would be that the observer was only truly aware of the currently attended object, while the apparent awareness of the rest of the display was, in some sense, an illusion.

The attentional blink, originally described by Donald Broadbent, and later characterized by Jane Raymond, Kim Shapiro, and Karen Arnell, is a quite different phenomenon that might point to a similar conclusion. In a typical attentional blink experiment, observers monitor a stream of images, letters for example, appearing at fixation at a rate of one every 100 ms. The observers are looking for particular targets, say "E" and "X." At this rate of presentation, an observer can easily report a single target letter appearing anywhere in the stream. However, if there are two targets, the second one is much more likely to be missed if it appears 200–500 ms after the first. This is not simply a matter of perceptual masking from the first target to the second. Given the same stream of letters (e.g., " J W E B P X L"), the target "X" will be easily detected if the observer does not have to report the "E," but will likely be missed if she does. Something about the attention to the first target causes an "attentional blink" that makes it harder to report the second. Interestingly, "blinked" items can be shown to have effects on the observer. A "blinked" word can produce semantic priming effects indicating that it has been read. Perhaps the type of attention that is tied up by the first target in an attentional blink display is the type of attention that permits awareness of an object.

Phenomena such as inattentional blindness, change blindness, and the attentional blink provide some evidence that, in the absence (or, at least, near-absence) of attention, the observer may be unable to recognize or even to see an object. In its strongest form, this argument proposes a tight linkage between selective attention and visual awareness. The fact that in some situations an observer is unable to report an unattended object does not prove that she is never able to do so. In the following, we describe phenomena that indicate that the proposed strong linkage of attention and awareness is too strong. The linkage of attention, binding, and awareness is challenged by studies that seem to show that there can be some degree of recognition and awareness of objects in the near-absence of selective attention.

When Awareness Does Not Require Attention

In a series of experiments done by Fei-Fei Li and her collaborators, observers performed an attentionally demanding visual search task in one location while concurrently monitoring another portion of the visual field for some class of targets (e.g., animal or vehicle). The goal was to tie up selective attention with the search task and to determine what, if anything, could be detected elsewhere at the same time. Some tasks (e.g., determining if red a square was to the right of a conjoining green square or vice versa) are profoundly disrupted when attention is thus engaged. Interestingly, however, detection of the presence of animals or vehicles in a briefly presented scene is no worse when selective attention is occupied than when it is not.

Note that detecting that a scene contained an animal is not the same as determining exactly what that animal was. Observers in Fei-Fei Li's experiments were not necessarily sure what type of animal they had detected or where it was in the display. Moreover, when Karla Evans and Anne Treisman asked observers to find animals in a rapid sequence of scenes, they found that observers were significantly impaired if the stream also contained humans. In the near-absence of attention, some image statistics seem to permit awareness of the presence of an animal (human or other) or a vehicle but it would be going too far to argue that these data make the proposed feature-binding role of attention unnecessary. Awareness of this specific animal or vehicle, which would presumably require feature-binding, appears to require attention.

A number of other phenomena also challenge the attention-binding-awareness linkage. For instance, Mary Potter and others have shown that high-order representations (i.e., gist) can be accessed very rapidly from natural scenes presented at rates of 10 or more per second, far too little time for selective attention to be directed to more than a small handful of the items in that scene. "Gist" in this case refers to a broad categorical label for the scene: beach, kitchen, etc. Does this challenge the relationship of attention to binding to object recognition? It certainly would if recognizing the gist of a scene involved promiscuous binding of multiple objects without attention. Alternatively, these results might be demonstrating that performance of this scene-categorizing task does not require binding. If information in the unbound feature statistics could support performance of the task, then it would not be necessary to assume binding. Aude Oliva and Antonio Torralba have shown that this can be done, in principle, for scenes. They have created filters whose output can be used to put a label on a scene (indoor, outdoor urban, beach, etc.) based on the raw image statistics of the scene; without the need to parse the image into objects, regions, and so forth. This "unbound" analysis is not adequate to identify a specific scene, for example Crane's Beach in Ipswich, MA, but it could provide the gist of a scene in the near absence of attention.

As with the detection of animals with selective attention occupied elsewhere, findings of this sort suggest that awareness of a visual stimulus is not a unitary, all-or-none experience. Awareness of an unattended scene may be different than awareness of a well-attended scene, but the scene is seen in both cases. This distinction may be reflected in neurophysiological findings showing that categorization and perception are mediated by different cortical areas. Specifically, those neurons that can categorize a target are not necessarily the same neurons that can, say, locate it in the image. Neurons in the visual cortex are able to signal a target's location but are generally insensitive to whether an object is a target or not (i.e., they cannot categorize it). However, David Freedman and colleagues have elegantly demonstrated neurons in prefrontal cortex can be highly sensitive to categorical status. They used a photographic morphing technique to create a picture of an animal that represented a combination of a cat and a dog. By varying the morphing parameters, they could vary the similarity between this computer-generated animal and prototypical cat and dog images. They presented a series of these computer-generated pictures to monkeys that had previously been trained to indicate whether each picture more closely resembled a cat or a dog. They found cells in the lateral prefrontal cortex that encoded the monkey's categorization. Such cells responded similarly to images that belonged to the same category, even when the images appeared very different. Conversely, the cells responded very differently to images that appeared very similar, but which belonged to different categories. These cells therefore responded to the categorization of the images, as opposed to the visual image itself.

We have seen that a strict linkage of attention, binding, object recognition, and awareness leads us to a theory that does not have room for the full range of phenomena. Still, it seems likely that attention is required for the recognition of specific objects and that, as Condillac argued, this act of attention changes the state of our visual awareness. Not being privy to more recent developments, Condillac does not tell us what he thinks we would see if we were performing an attentionally demanding task at fixation when the curtains were thrown wide, revealing the scene outside the chateau for the first time. However, it seems likely that his answer would have been much the same. You would have some impression of the outside world, but you would not understand what you were seeing until you had attended to the scene.

As he became familiar with the current literature, Condillac might agree that some information about the gist of the scene might be available in that first moment but his description of the initial state of awareness, modified by subsequent attention would remain essentially unchanged.

When Attention Does Not Imply Awareness

It is hard to gain any introspective access to this with real scenes because we are too good at analyzing them. However, the very unreal scene of Fig. 3 may serve.

Note that when you first look at this scene, you are aware of the patches of color and orientation across the entire stimulus. More than that, you are aware of some structure in the scene. There are plusses everywhere with a scattering of items with more than four line terminations. There are blue and yellow objects in the upper left and red-green elsewhere. However, if you are asked to detect red vertical components, you will need to direct your attention to a succession of specific items over a period of time and, having

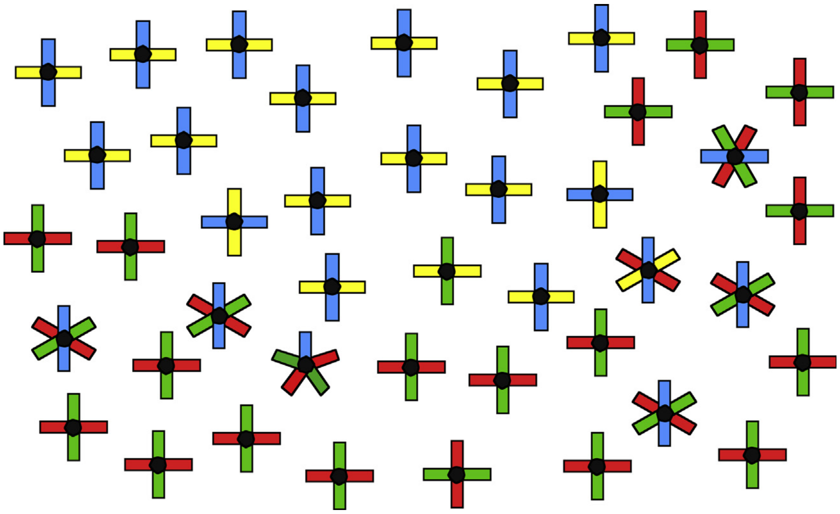

Figure 3 A display that gives some idea of one's level of visual awareness. Please see the text for details.

deployed your attention, your awareness of the stimulus will change. Now you will find that the red and green plusses are not all the same. There is a region of red verticals in the upper right and an isolated example at the bottom center.

You have awareness with and without selective attention. What about attention without awareness? Can you select and bind an object without being aware of it? Returning to the figure, imagine you are asked to locate the five-pointed item. It is entirely possible that you had already attended to that object during a search for red vertical without becoming aware of it. Thus, it seems possible that you can attend to an object without ever becoming aware of it. The point is tricky since it could be argued that you were aware of the five-pointed item when you putatively attended to it but you forgot about it prior to being queried about it. It is often difficult to distinguish between having been unaware and being amnesic.

Studies of blindsight patients provide converging evidence for the hypothesis that attention is not always sufficient for visual awareness. In blindsight, damage to the primary visual cortex results in a condition where patients report being unaware of part of the visual field (the unaware area). Interestingly, when asked to guess at what stimuli are located in the unaware area, patients can perform at above chance levels. Furthermore, when spatially cued to the unaware area, they exhibit speeded discrimination of targets that subsequently appear in that area, demonstrating that they can attend to stimuli that they cannot see. Evidently, attention does not necessarily result in awareness.

Post-attentive Awareness

Consider the red-vertical plus at bottom center in the figure. We can posit that you had some representation of it before it was selected. That representation can be called "pre-attentive." You also had a representation of the plus while it was attended. Call that an attended representation. What is the representation of that plus when you then move your attention to the blue-vertical star up and to the right? This can be called the "post-attentive" representation. A series of experiments show that changes to the plus (e.g., changing it to red-horizontal, green-vertical) once attention has been shifted elsewhere will go unnoticed until attention is directed back to the item (providing that the transients produced by such changes are masked by, say, a blink, a saccade or a visual transient). Observers show no more awareness of the current binding of features in a post-attentive object than of a pre-attentive object. At the same time, you—the observer—are aware that there *was* a red vertical plus at that location so, in that sense, your post-attentive awareness of that particular plus is different than your pre-attentive awareness.

Of course, you would also be aware of the plus, in that sense, if the lights went out and you relied on memory. The already difficult topic of visual awareness becomes more difficult once we admit a role for memory. Your awareness of a familiar face is different from your awareness of a new face. That difference is tied to memory and it is hard to know whether one should consider this to be part of the definition of visual awareness. Whatever one concludes about this question, the post-attentive vision research suggests that selective attention affects post-attentive awareness of an object through memory and not through some sort of persistent binding that continues once selective attention is disengaged from an object.

Awareness of Awareness

The account outlined above, suggests that we think we are aware of more than we actually are. We greatly overestimate our own awareness. Our naïve impression of our visual awareness is that we are aware of a large visual scene at a high resolution. Yet this is not so. At any moment, the only part of the visual scene we can see in high resolution is the small area around the current point of fixation. A particularly striking demonstration of this (beloved by people who sell eye trackers!) is to use an eye tracker to

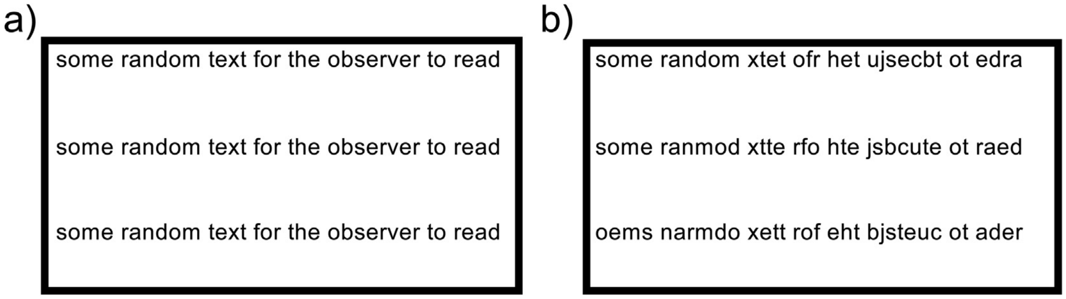

Figure 4 The observer starts to read the text shown in (A). During an eye-movement, all the text that is not near the observer's point of fixation (assumed to be at the top-left corner) becomes jumbled (B). The observer does not notice the change and so cannot differentiate between (A) and (B), thereby demonstrating how limited visual awareness really is.

monitor the observer's point of fixation. The observer's task is to read some text presented on a computer monitor, similar to that shown in Fig. 4A.

At some point, the eye tracker salesperson pushes a button and, during the observer's next saccade, the letters in the display became jumbled except for those in words near the point of fixation (Fig. 4B), which, in this figure, is assumed to be in the top left corner. Every time the observer saccades to a different point, the letters in the words near that point became unjumbled, while the letters in all other words either become or remain jumbled. Provided all changes to the display occur during a saccade, the observer is unaware of the scrambling. She reports that she is simply reading normal text. Similarly, if the image away from fixation is appropriately blurred, an observer would be unaware of this degradation and will have the impression of looking at the usual, apparently well-focused scene.

The observer in these illustrations is aware of something, but it turns out not to be a "true" assessment of the contents of the current visual representation. Beyond a simple contribution of memory to awareness, this indicates a contribution of theory to awareness. In the reading example, the observer is not aware of scrambling occurring away from fixation. Wherever she fixates on the page, the letters form themselves into English words. It is a reasonable theory that the page consists of readable English words and our observer's awareness incorporates that theory. Her visual awareness is affected by what she thinks she knows. Returning to the role of attention in awareness, we see that selective attention alters not only the awareness of the attended object but, potentially, the awareness of other unattended objects, potentially divorcing that awareness from the actual perceptual facts.

Awareness of Attention

If one were inclined to propose a tight linkage of attention and awareness, one might propose not only that observers are only aware of the objects of attention, but that observers are aware of all of the deployments and, thus, all the objects of attention. However, in visual search experiments, it is estimated that observers can attend to 20–50 items per second. This rapid selection seems to occur without a clear awareness of which items in a display have or have not been selected. At least observers do not use any such awareness to guide their search. For example, Todd Horowitz and Jeremy Wolfe conducted a visual search experiment in which all the items in the display were randomly relocated every 111 ms. This made it impossible for observers to keep track of which items they had already attended to. Surprisingly, search was no less efficient in this case than in the control condition where the items were static. This showed that, at least in some cases, visual search had no memory. Observers acted as if they were unaware of what they have attended to.

Conclusions

In this article we have considered the relationship of conscious awareness and attention from the perspective of vision. Following Condillac, we found it helpful to differentiate between the awareness that results from attending to an object or group of features and that which occurs in the absence of attention. Condillac is more famous for his statue than for his chateau. He asked his readers to imagine the mental life of a statue with no senses. In his honor, we can imagine a statue with senses but without attention or, perhaps better, with attention disabled. In the absence of attention, the evidence indicates that our statue would retain some visual awareness, but would be unable to form any percepts that would require the binding of two of more features. This level of awareness might allow our statue to classify scenes (beach, mountains, etc.) or declare that they did or did not contain an animal, but would not allow it to determine specifics such which animal occurred in a given scene. If we now endow this statue with selective attention, it can then solve the binding problem, and have a more complete awareness. Specific objects can be selected, perceived and identified. If we now allow the statue to have a memory, the statue will know that a given object was at a particular location and, in the absence of contradicting information, the statue is likely to assume that the object continues in that location. If we add a theory-building capability, the statue can generalize from the fact that all selectively attended samples are seen in sharp focus to the assumption that all objects really are in sharp focus and then use this assumption to modify visual awareness accordingly. The statue now has an approximation of human visual awareness.

Further Reading

Evans, K., Treisman, A., 2005. Perception of objects in natural scenes: is it really attention free? J. Exp. Psychol. Hum. Percept. Perform. 31, 1476–1492.
Freedman, D.J., Riesenhuber, M., Poggio, T., Miller, E.K., 2001. Categorical representation of visual stimuli in the primate prefrontal cortex. Science 291, 312–316.
Freedman, D.J., Riesenhuber, M., Poggio, T., Miller, E.K., 2003. A comparison of primate prefrontal and inferior temporal cortices during visual categorization. J. Neurosci. 23, 5235–5246.
Hochstein, S., Ahissar, M., 2002. View from the top: hierarchies and reverse hierarchies in the visual system. Neuron 36, 791–804.
Horowitz, T.S., Wolfe, J.M., 1998. Visual search has no memory. Nature 394, 575–577.
Kentridge, R.W., Heywood, C.A., Weiskrantz, L., 2004. Spatial attention speeds discrimination without awareness in blindsight. Neuropsychologia 42, 831–835.
Li, F.F., VanRullen, R., Koch, C., Perona, P., 2002. Rapid natural scene categorization in the near absence of attention. Proc. Natl. Acad. Sci. U.S.A. 99, 9596–9601.
Moran, J., Desimone, R., 1985. Selective attention gates visual processing in the extrastriate cortex. Science 229, 782–784.
Potter, M.C., Faulconer, B.A., 1975. Time to understand pictures and words. Nature 253, 437–438.
Riesenhuber, M., Poggio, T., 1999. Are cortical models really bound by the 'binding problem'? Neuron 24, 87–93.
Sheinberg, D.L., Logothetis, N.K., 1997. The role of temporal cortical areas in perceptual organization. Proc. Natl. Acad. Sci. U.S.A. 94, 3408–3413.
Treisman, A.M., Gelade, G., 1980. A feature-integration theory of attention. Cogn. Psychol. 12, 97–136.
Treisman, A., Schmidt, H., 1982. Illusory conjunctions in the perception of objects. Cogn. Psychol. 14, 107–141.
Vul, E., Rich, A.N., 2010. Independent sampling of features enables conscious perception of bound objects. Psychol. Sci. 21 (8), 1168–1175.