June 2003

A Neurofunctional Theory of Consciousness

Jesse J. Prinz
jesse@subcortex.com
University of North Carolina, Chapel Hill

[Penultimate draft of chapter in A. Brook and K. Akins (Eds.) *Philosophy and neuroscience*. Cambridge: Cambridge University Press (forthcoming).]

Introduction: The Problems of Consciousness

Reading the philosophical literature on consciousness, one might get the idea that there is just one problem in consciousness studies, the hard problem. That would be a mistake. There are other problems; some are more tractable, but none are easy, and all interesting. The literature on the hard problem gives the impression that we have made little progress. Consciousness is just an excuse to work and re-work familiar positions on the mind-body problem. But progress is being made elsewhere. Researchers are moving towards increasingly specific accounts of the neural basis of conscious experience. These efforts will leave some questions unanswered, but they are no less significant for that.

To move beyond the hard problem, I would like to consider some real problems facing consciousness researchers. First, there is a What Problem. This is the problem of figuring out what we are conscious of. What are the contents of conscious experience? For those looking at the brain, it is closely tied to a Where Problem. Where in the brain does consciousness arise? Locating consciousness may not be enough. We need to address a How Problem. How do certain states come to be conscious? This can be construed as a version of the hard problem, by asking it with right intonation: How could certain physical states possibly be experienced? But there is another reading that is also worth investigating. The How Problem can be interpreted as asking, What are the psychological or neuronal mechanisms or processes that distinguish conscious states from unconscious states? The mechanisms and processes can be pinpointed by pursuing a closely related to a When Problem: Under what conditions do the states that are potentially conscious become conscious? Once we know what we are conscious of and when we are conscious, we can begin to address a Why Problem: Why do we have conscious states? And finally, there is a Who Problem: Who is conscious? Non-human animals? Human infants? Machines?

I will discuss these problems in turn. Some already have answers, others have answers on the way, and at least one may never be answered in full. What I offer here is a progress report on a theory of consciousness that I have presented elsewhere (Prinz, 2000; 2001). I hope to show that progress is indeed being made.

1. What Are We Conscious of?

I think the best answer to the What question was given by Ray Jackendoff in his 1987 book, *Consciousness and the Computational Mind*. Jackendoff began with the observation that perceptual systems are organized hierarchically, with different subsystems representing features at varying degrees of abstraction. There is a movement from very fine-grained local features, with minimal global integration, to very abstract categorical representations that are especially useful for capturing invariance across perceptual vantage points. Between these extremes, the disjointed local and the abstract categorical, there are postulated to be intermediate-level subsystems. These are vantage point specific, rather than invariant, and they also have global organization—the parts are bound together coherently. Jackendoff based this picture on perceptual psychology from the 1970s and '80s. He was particularly inspired by Marr's theory of vision, and by models of categorical speech perception in phonology. These models may not have held up in detail, but the overall approach remains current. Perception is still widely believed to be hierarchical. The major difference between the state of play now, and back in 1987 is that we have a much better understanding of the perceptual hierarchies, and that understanding has been fueled by advances in neuroscience. Jackendoff had virtually nothing to say about the brain.

To assess Jackendoff's conjecture we can start with the neuroscience of vision. Is there evidence that an intermediate level of processing in the visual brain is privileged with respect to consciousness? I think the answer is yes. First consider low-level vision, which is associated with activity in primary visual cortex (V1). Destruction of V1 ordinarily eliminates visual consciousness, but it may do so by preventing higher areas from receiving visual signals. People with V1 damage sometimes experience visual hallucinations, and some also experience residual visual experience when presented with rapidly moving, high contrast stimuli can cause residual visual experience (Seguin, 1886; Sahraie et al. 1997). Crick and Koch (1995 have argued that V1 cannot be the seat of consciousness, because it encodes the wrong information. Some research suggests that V1 does not show context effects that are typical of conscious color vision, and V1 may not be responsive to certain imagined contours that we experience in optical illusions. Rees et al. (2002) cite evidence that V1 activity drops during eye blinks, despite the fact that we are not consciously aware of visual interruptions while blinking. Some of these data are controversial. For example, Seghier et al. (2000) and other groups have found evidence for the processing of illusory contours in V1. It is not yet clear what role that processing plays, however. Higher visual areas are much more consistently responsive (Mendola et al., 1999). Moreover, Ramsden et al. (2001) found that the activity in V1 during the perception of illusory contours was the inverse of activity during perception of real contours. In higher visual areas, activity for real and illusory contours is alike. Also receptive fields in V1 may be too small to respond to illusory contours (or color context effects) over large gaps. Higher visual areas may be needed to perceive these. It is premature to rule conclusively on whether V1 is a locus of visual consciousness, but current evidence provides little reason to think that it is.

Now consider high level vision, which is associated with activity in areas of inferotemporal cortex. Logothetis and colleagues have been arguing that these areas are crucial for visual consciousness. When we look at two distinct images simultaneously, one in each eye, we experience only one or the other. Logothetis and colleagues presented such stimuli to monkeys and trained them to indicate which of two stimuli they were seeing. They measured cellular response in various visual areas as the monkeys made their reports. It turned out the 90% of the measured cells in high-level visual areas correlated with indicated percept, compared to 38% of the measured cells in intermediate level areas (Sheinberg and Logothetis, 1997; Leopold and Logothetis, 1996). Logothetis (1998) provisionally concludes that activity in inferotemporal cortex (IT) "may indeed be the neural correlate of conscious perception."

There are three difficulties with this conclusion. First, the fact that fewer than half of the intermediate cells correspond to the perceived object does not show that conscious perception is not located at the intermediate level. The cells in that 38% may be responding in a distinctive way, and that distinctive way of responding may mark the difference between conscious and unconscious vision. In fact, the intermediate cells that correspond to the perceived object are almost certainly responding distinctively. After all, they seem to be the only cells that are allowing information to propagate forward to higher areas. The 90% in IT reflects a selection process at the prior stage of processing. Second, there is independent reason to think IT is not the locus of visual awareness. Cells in IT encode the wrong information. They tend to abstract away from size and specific orientation (Vuilleumier et al. 2002). Indeed many IT cells respond the same way to a shape oriented to the left and the same shape oriented to the right (Baylis and

Driver, 2001). Third, damage in high level visual areas tends to cause deficits in recognition of images, but not visual experience as such. Damage to intermediate level visual areas impairs both. This is the difference between associative and apperceptive agnosia (Farah, 1990).

In contrast to cells in IT and in V1, cells in intermediate-level visual areas do correspond to the contents of experience. They represent illusory contours and color constancy, and they represent objects from a specific orientation, as they appear in consciousness. Damage to intermediate areas (V2, V3, V4, and V5) eliminates consciousness of the stimulus features processed in these areas (Zeki, 1993), and these areas are active during visual experiences and visual hallucinations (ffytche, 1998). This confirms Jackendoff's conjecture, that the intermediate level is the level of consciousness.

Jackendoff speculates in his book that the intermediate level is the locus of consciousness in all perceptual modalities, not just vision. Though hardly confirmed, that generalization is consistent with current date from neuroscience. All sense modalities seem to be hierarchical. In audition, there is a hierarchy extending from the superior temportal plane into the superior temporal gyrus. Researchers have divided this pathway up into core, belt, and parabelt regions, which correspond to high, intermeidiate, and low levels of processing (Kaas and Hackett, 2000). In touch, there is a hierarchy that extends from Brodmann areas 3a, 3b, 1, and 2 into areas 5 and 7 (Kaas, 1993; Friedman et al. 1986). There is a taste hierarchy as well, including the insula and moving into orbitofrontal cortex (Rolls, 1998). Even smell, the ancient sense, has hierarchical organization, which extends from piriform cortex into orbitofrontal and prefrontal areas (Savic et al., 2000). I have argued elsewhere that a sensory hierarchy can also be identified for the perception of internal bodily states, which forms the basis of both interceptive experience and emotional experience, a special case of interoception (Prinz, 2004).

These hierarchies differ in organizational detail, of course, but there is a general pattern. Highly specific, local features are registered first, then combinations of those features are registered, and, finally relatively abstract or categorical stimulus properties are registered. It may be a long time before we have good evidence pinpointing consciousness in these hierarchies, but the functional profiles invite speculation. The smallest units of perception, whether they are isolated tones or tiny features of texture, are too discrete to map onto units of experience. More complex features, especially those that abstract away from the idiosyncratic features of a particular perceptual episode, seem more abstract than the contents of experience. Only in the middle, where tones merge together into the components of a melody and minute edges merge into textural patterns, do we seem to have conscious experience. Existing findings are consistent with these speculaions. For example, in auditory processing, there is a distinction between disorders of recognition, where sounds can be heard intact, and disorders of perception, where sounds are disrupted or impoverished. This corresponds to the distinction between associative and apperceptive agnosia in vision, which supports the intermediate-level conjecture (Vignolo, 1982). Time will tell if the conjecture holds up. If it does, Jackendoff's answer to the What question will be upheld.

2. How Do We Become Conscious?

Jackendoff may have been right about the What Problem, but he does not offer a solution to the How Problem. He does not specify the conditions under which intermediate-level perceptual representations become conscious. The implied answer is that consciousness arises whenever intermediate-level representations are active. Consciousness is the tokening of intermediate level representations. This would be a seductively simple theory, but it cannot be right. There is good reason to think that activity can occur at the intermediate-level without conscious experience. Consider subliminal perception. One can extract information about the meaning, form, or identity of an object even if it is presented under conditions that prevent conscious awareness (e.g., Bar and Biederman, 1998; Dell'Acqua and Grainger, 1999). To extract such information, one must process the stimulus all the way through the perceptual hierarchy. Thus, intermediate-level activation is not sufficient for consciousness.

A more dramatic demonstration of this point comes from studies of patients with unilateral neglect. In subliminal perception studies, stimuli are presented very briefly. Such studies can be interpreted as demonstrating memory effects, rather than unconscious perception. Subjects may consciously perceive the stimuli and then forget them. This interpretation is unavailable for cases of neglect. In these cases, right inferior parietal injuries prevent patients from attending to the left side of their visual fields or the left side of visually presented objects, or the left side of their own bodies. They report no conscious experience of things on the left. This effect has no temporal cap. For example, patients can stare at an object for a long time and still be oblivious to its left side. Nevertheless, there is evidence that patients with neglect are registering information about the objects that they fail to experience (Marshall and Halligan, 1988; Bisiach, 1992; Driver, 1996). This suggests that they are processing neglected objects through (and perhaps beyond) the intermediate level. This conclusion has been supported by neuroimaging studies, which demonstrate intermediate-level activity in the visual systems of unconsciously perceived objects in patients with right parietal injuries (Rees et al, 2000; Vuilleumier et al. 2001).

Neglected provides two important lessons for consciousness. First, apparent cases of unconscious perception are not always readily interpretable as memory effects. There is every reason to think that patients with neglect are capable of perceiving, to some degree, with out conscious experience. This shows that mere activity in a perceptual system is not sufficient for consciousness. The second lesson comes from a moment's reflection on the cause of the deficit in neglect. Conscious experience of the left side is lost as a result of injuries in centers of the parietal cortex that have been independently associated with attention. Neglect is an attention disorder. This suggests that consciousness requires attention. This hypothesis gains experimental support from behavioral studies with normal subjects. Evidence suggests that people will fail to consciously perceive a centrally presented object if their attention is occupied by another task. The phenomenon was labeled inattentional blindness by Mack and Rock (1998). In their studies, many subjects fail to notice a centrally presented object that is briefly displayed, and unexpected. Inattentional blindness can also be sustained for a relatively long time. Chabris and Simons (1999) had subjects watch a video in which two teams were tossing a basket ball. Subjects were asked to count how many times the ball was passed by a particular team—an attention-demanding task. During the game, a person in a gorilla suit strolls across the center of the screen. The gorilla is highly salient to passive viewers, but 66% of the subjects who were counting passes failed to notice the gorilla. It is possible that these subjects consciously perceived the gorilla and simply didn't categorize it; but it is equally possible that they didn't perceive it consciously at all. The latter interpretation would be most consistent with the neglect findings. All these findings provide strong support for a link between attention and consciousness (see also Luck et al. 1996 on the attentional blink).

These findings present a solution to the how problem. Conscious seems to arise in intermediate-level perceptual subsystems when and only when activity in those systems is modulated by attention. When attention is allocated, perception becomes conscious. Attentional modulation of intermediate level representations is both necessary and sufficient for consciousness. This also makes sense of the Logothetis studies. When presented with conflicting simultaneous stimuli, we may be able to attend to only one. The shifts in attention lead to shifts in consciousness.

Some researchers object to the thesis that attention is necessary. Christof Koch (personal communication) gives the example of staring a solid expanse of color, spanning the visual field. Do we really need to attend to experience the hue? I think the answer is yes. Contrary intuitions stem from the fact that attention is often associated with effort and focus. In this case attention is spread evenly across the field, or perhaps it bounces around, as if scanning fruitlessly for points of interest. In either case, the allocation of attention might not require effort. It does, however, involve selection. Attention is generally in the business of picking out salient objects. This is a limiting case. There are no objects, so the whole field, or large portions are selected. I don't have demonstrative proof of this conjecture. My point is that Koch's alleged counterexample is not decisive. There is a plausible interpretation that involves attention. If all else were equal, there would be no way to decide between these interpretations. But all else is not equal. Evidence from the inattentional blindness literature suggests we are not conscious when attention is withdrawn. In the color case, we must be attending.

So far, I have said very little about what attention actually is. We need an answer to this question before we can make any serious progress. Fortunately, psychologists and neuroscientists have developed numerous models of attention. The general theme of these models is that attention is a selection process that allows information to be sent to outputs for further processing. It is well-established that visual attention modulates activity through the visual hierarchy, and changes can include increase in activity for cells whose receptive fields include attended objects, and decrease in activity in cells responding to unattended objects (Fries et al. 2001; Corchs and Deco, 2002). There is some evidence that the relevant changes involve changes in rates of neuronal oscillation (Niebur et al., 1993; Fries et al. 2001). Whatever the specific mechanism, changes in perceptual processing seem to have an impact on what can propagate forward for further processing. Attention thus affects a virtual change in connectivity between neuronal populations (Olshausen et al., 1994). Attention is a routing gate.

This raises another question. When perceptual information can propagate forward, where does it go? The most plausible answer is that attention allows perceptual information to access working memory stores. "Working memory" refers to our ability to retain information for brief periods of time, and to manipulate it in various ways. In the brain, a working memory area is understood as a neuronal population that responds to features of a perceived stimulus, but maintains activity when that stimulus is removed. Working memory areas have been identified in frontal and temporal cortex (e.g., Miller et al. 1996; Nakamura and Kubota, 1995). When we perceive, what we perceive is made available for temporary storage. Attention is gate keeper to working memory.

If all of this is right, the solution to the How Problem can be stated as follows: consciousness arises when intermediate-level perception representations are made available to working memory via attention. I call this the AIR theory of consciousness, for attended intermediate-level representations. The suggestion that working memory access is important to consciousness has many defenders (e.g., Crick and Koch, 1995; Baars, 1997; Sahraie et al., 1997). Rees (2001) has defended a view that is especially close to the AIR theory. He uses neuroimaging evidence to support the conjecture that consciousness involves a large-scale neural network that includes perception centers along with prefrontal working memory center and parietal attention centers. I take this convergence to be an encouraging sign that we are closing in on the How Problem

3. Why Are We Conscious?

The AIR theory is a functionalist account of what consciousness is. It says that consciousness can be understood in terms of an information-processing role. As such, it lends itself to a straightforward answer to the Why Problem. In asking why we are conscious, we are really wondering what consciousness is for or what purpose it serves. Does it do any work for us? That depends on what it is. If consciousness is the property of inner representations of a certain sort being made available in a certain way, then the answer depends on what these representations are and what their availability does for us.

It should be evident that consciousness is extremely valuable if the AIR theory is right. Consciousness serves the crucial function broadcasting viewpoint specific information into working memory. Viewpoint specific representations are important for making certain kinds of decisions. If we encounter a predator, for example, it is useful to know whether it is facing us or facing in another direction. Consciousness provides us with this information by making it available to working memory. Working memory is not just a transient storage bin. It is the store from which decisions are made and actions are chosen. It is crucial for decisions that have not been rehearsed much in the past. Without working memory, we would be reduced to reflexive response. There is also evidence that working memory is a gateway to episodic memory (Buckner et al. 1999). Information is sent to working memory before getting encoded in a long-term store. If that is right, then, without consciousness, we would have no biographies, no memories of our past. In these ways, consciousness is what distinguishes us from paramecia and jellyfish.

It might be objected that the functions I have been describing could be achieved without consciousness. In another possible cognitive system, viewpoint specific representations could be broadcast to working memory without any experience or phenomenology. This is perfectly true, but the "could" renders the point irrelevant. Fly swatters serve the function of killing flies. That function could also be achieved by insect-eradicating lamps. It would be a howler to infer from this that fly swatters have no function. The same function can be carried out in different ways. In us, the function of broadcasting viewpoint specific information to working memory is achieved by mechanisms that give us conscious experiences. Those experiences have the function they serve.

The objector might press the point, however. If the function of broadcasting information can be realized without consciousness, then consciousness cannot be identified with that function. Rather, it must be identified with something that happens to realize that function in us. This objection is well taken, but it is not fatal to the AIR theory. It requires a clarification. Consciousness does not arise whenever intermediate-level representations are broadcast to working memory; it arises when that is done via the kinds of mechanisms found in us. Just as a fly swatter kills flies in a certain way, consciousness is broadcasting in a certain way. To figure out what way, we need to turn to the science of attention. We need to look at specific models of how attention routs information into working memory. Various models of attention have been worked out in physiological detail. Above, I mentioned excitation, inhibition, and changes in oscillation frequency as examples. The idea is that we will ultimately be able to give a true and complete neurocomputational model to go along with the AIR theory. The neurocomputation model will specify the mechanisms that matter for consciousness. Specified in this way, the AIR theory is different from other functionalist theories in the philosophy of mind. It is a neurofunctional theory: consciousness is identified with a functional role that is implemented by mechanisms specifiable in the language of computational neuroscience. Computational neuroscience is a functional theory; it describes what neurons and neural networks do. But it is also a neuroscientific theory, because it describes those functions at a level of architectural specificity that can be investigated using the tools of neuroscience. Such functions might be realizable by materials other than those found in the nervous system, but they cannot be realized by everything under the sun. The computations that our brains performed, with their specific temporal profiles, place structural and organizational constraints on the things that can perform them. In describing things at a fine level of neurocomputational detail, we should not forget, however, that all of these neuronal mechanisms are in the business of implementing functions that can be described at a psychological level of analysis. As stated above, the AIR theory is a characterization of that psychological level. The psychological level implemented by the neurocomputational level specifies what the neurocomputations are achieving, and, thus, the psychological level tells us what consciousness is for.

4. Who Is Conscious?

Now only the Who Problem remains. On the face of it, the problem is easy. Here's a formula for answering it. Find out what the mechanisms of consciousness are in us. Then see what other creatures have those mechanisms. On the story I have been telling, the mechanisms of consciousness are the attentional mechanisms that make intermediate-level perceptual representations available to working memory. Now we can ask: what creatures have those?

Higher mammals, especially primates, probably have neural mechanisms that are very much like ours. Much of what we know about the cellular neurophysiology of human attention and working memory comes from studies of monkeys. There also seems to be great continuity between primates and rats. Rats are used as animal models for attention and working memory abnormality in schizophrenia and attention deficit / hyperactivity disorders. The functional neuroanatomy of a rat is similar to our own in many respects. For example, rat prefrontal cortex plays a key role in rat working memory (e.g., Izaki et al. 2001).

Creatures that are more distantly related resemble us in functional respects. For example, researchers have found working memory capacities in octopuses (Mather, 1991), pigeons (Diekamp et al. 2002), bees (Menzel, 1984), and slugs (Yamada et al. 1992). But such creatures have neuronal mechanisms that are quite different from ours, and this raises a question. Is there any reason to think they are conscious? This question is far more difficult than the questions I have so far considered, and I see no obvious way to answer it. In other words, I see no obvious way to know how similar a mechanism must be to those found in our own brains in order to support conscious experiences. The problem can be put in terms of levels of analysis. Let's suppose we find that octopuses have viewpoint-specific representations, and attentionally-gated working memory. At the psychological level, they are like us. But suppose they differ at the neurocomputational level. Or, to be more realistic, suppose there are numerous levels of abstraction at which we can provide neurocomputational models, and octopuses differ on some, but not all of these. At which level must an octopus resemble us to be a candidate for consiousness? Elsewhere, I raise the same question for machines (Prinz, 2003). How close to the human brain must a computer be before we can say that it is probably conscious? I think there is no way to answer this question. To do so would require identifying the level of analysis that was crucial for consciousness in ourselves. That would require that we be able to look at the contribution of different levels individually. We have no way to do that, and if we did it wouldn't help. Suppose we could keep the psychological level of analysis constant while changing the neural mechanisms in a human volunteer. We might do that by replacing brain cells with microchips that work differently at a low level of analysis. The problem is that the human volunteer would report being conscious no matter what. Preserving function at the psychological level would guarantee it (see Prinz, 2003, for more discussion). So we cannot tell if the low-level alteration makes a difference vis-à-vis consciousness. I call this level-headed mysterianism.

There may also be another mystery that faces anyone trying to answer the Who Problem. Suppose we find a creature that exhibits signs of intelligence but is radically unlike us functionally. Suppose it has nothing that we can identify as a working memory store (all memories are always equally accessible and long-lasting) and no attention capacity (all perceived stimuli get processed equally). Suppose further that its perceptual systems are not hierarchical, and it perceives through senses unlike our own. Now we can ask whether this creature has any conscious experiences. We can reasonably speculate that it does not have conscious experiences like ours. But we cannot, it seems, conclusively speculate that it has no conscious life at all. The creature may have very complex sensory systems, and it may respond to the environment intelligibly. We would certainly want to call it a perceiver. The difficulty is in deciding whether its perceptions have phenomenal character. We cannot demonstratively rule out the possibility that experience is realized in radically different ways. This might lead one to an agnostic conclusion, which one might call radical realization mysterianism. I cannot think of a reason to take such a position especially seriously. We know of no such creatures, and if we met one, its radically different psychology would be grounds for holding off ascriptions of consciousness. Moreover, I am inclined to interpret the discovery of a material basis for consciousness in us as the discovery of an identity. Consciouness just is the operation of a neurofunctional mechanism found in us. To hold otherwise, is to regard consciousness as an emergent property of such a mechanism—a property that could emerge in other ways. I suspect that this kind of emergentism isn't true, but I can think of no decisive way to rule it out. The evidence for the material basis of consciousness in us is not decisive evidence against consciousness arising in different ways. And of course, this is the road to the hard problem.

If we are saddled with either of these two kinds of mysterianism (I am much more confident about the level-headed variety), then the theory of consciousness will remain forever incomplete. We will be able to discover conditions that are sufficient for consciousness, such as those found in our brains, but we will never have a perfectly settled list of the conditions necessary for consciousness. We will not know what level of physiology is essential, and perhaps we won't even know what psychological processes are essential. That is a humbling discovery.

But, lest we get too humble, we should take heart in the enormous progress that has been made. Discovering the material basis of consciousness in us is no mean feat. We are not quite there yet. A full account would include a correct neurocomputational theory of attention. That goal is certainly within reach. When it has been achieved, I think we will be in a position to say that we have cracked problems of profound importance. Finding the neurocomputational basis of consciousness will be like finding the material basis of genetic inheritance. It will be equally worthy of an invitation to Stockholm. There may be residual mysteries pertaining to realizations of consciousness in other creatures, but that should not distract us from the success, which may already be at hand.

Acknowledgements
I want to thank the audience of the McDonell Philosophy & Neuroscience Conference and, especially, Andy Brook for helpful comments.

References

Baars, B.J. (1997). *In the theater of consciousness.* New York: Oxford University Press.
Bar, M. and Biederman, I., (1998). Subliminal visual priming. *Psychological Science, 9,* 464–469.
Baylis, G. C. and Driver, J. (2001). Shape-coding in IT cells generalizes over contrast and mirror reversal, but not figure-ground reversal. *Nature Neuroscience, 4,* 937-942.
Bisiach, E. (1992). Understanding consciousness: Clues from unilateral neglect and related disorders. In A. D. Milner and M. D. Rugg (eds.) *The Neuropsychology of Consciousness* (113-139). London: Academic Press.
Buckner, R.L., Kelley, W.M., and Petersen, S.E. (1999) Frontal cortex contributes to human memory formation. *Nature Neuroscience, 2,* 311-314.
Corchs S and Deco G. (2002). Large-scale neural model for visual attention: integration of experimental single-cell and fMRI data. *Cerebral Cortex, 12,* 339–48.
Crick, F. and Koch, C. (1995). Are we aware of activity in primary visual cortex? *Nature, 37,* 121-123.
Dell'Acqua, R. and Grainger, J. (1999). Unconscious semantic priming from pictures. *Cognition, 73,* B1-B15.
Diekamp, B., Kalt, T., and Güntürkün, O. (2002). Working memory neurons in pigeons. *Journal of Neuroscience, 22,* RC210.
Driver, J. (1996). What can visual neglect and extinction reveal about the extent of "preattentive" processing. In A.F. Kramer, M.G.H. Coles, and G.D. Logan (eds.) *Convergent Operations in the Study of Visual Selective Attention.* Washington, DC: APA Press.
Ffytche D.H., Howard R.J., Brammer M.J., David A., Woodruff P.W., and Williams S. (1998). The anatomy of conscious vision: an fMRI study of visual hallucinations. *Nature Neuroscience, 1,* 738-742.
Friedman D.P., Murray E.A., O'Neil B., Mishkin M. (1986) Cortical connections of the somatosensory fields of the lateral sulcus of Macaques: evidence of a corticolimbic pathway for touch. *Journal of Comparative Neurology, 252,* 323-347.
Fries, P., Reynolds, J. H., Rorie, A. E., and Desimone, R. (2001). Modulation of Oscillatory Neuronal Synchronization by Selective Visual Attention. *Science, 291,* 1560-1563.
Izaki, Y., Maruki, K., Hori, K. and Nomura, M. (2001) Effects of rat medial prefrontal cortex temporal inactivation on a delayed alternation task. *Neuroscience Letters, 315,* 129-132.
Kaas, J. H. and Hackett, T. A. (2000). Subdivisions of auditory cortex and processing streams in primates. *Proceedings of the National Academy of Science, 97,* 11793-11799.
Kaas, J.H. (1993). The functional organization of somatosensory cortex in primates. *Annals of Anatomy, 175,* 509–518.
Leopold D.A. and Logothetis N.K. (1996). Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature, 379,* 549-553.
Logothetis, N. (1998). Object vision and visual awareness. *Current Opinion in Neurobiology, 8,* 536–544.
Luck, S.J., Vogel, E.K., and Shapiro, K.L. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature, 383,* 616–618.
Mack, A., and Rock, I. (1998). Inattentional Blindness. Cambridge, MA: MIT Press.
Marshall, J.C. & Halligan, P.W. (1988). Blindsight and insight in visuospatial neglect.

Nature, 336, 766-67.
Mather, J.A. (1991) Navigation by spatial memory and use of visual landmarks in octopuses. *Journal of Comparative Physiology, A, 168,* 491-497.
Mendola, J.D., Dale, A.M., Fischl, B., Liu, A.K., Tootell, R.B.H. (1999). The representation of illusory and real contours in human cortical visual areas revealed by functional magnetic resonance imaging. *Journal of Neuroscience, 19,* 8560-8572.
Menzel, R., (1984). Short-term memory in bees. In Alkon, D.L., Farley, J. (eds.) *Primary Neural Substrates of Learning and Behavior Change* (259–274). Cambridge: Cambridge University Press.
Miller, E. K., Erickson, C. A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience, 16,* 5154-5167.
Nakamura, K., and Kubota, K. (1995). Mnemonic firing of neurons in the monkey temporal pole during a visual recognition task. *Journal of Neurophysiology, 74,* 162-178.
Niebur, E., Koch, C., and Rosin, C. (1993) An oscillation-based model for the neuronal basis of attention. *Vision Research, 18,* 2789–2802.
Olshausen, B. A., Anderson, C. H., and van Essen, D. C. (1994). .A neurobiological model of visual attention and invariant pattern recognition based task. *Journal of Neuroscience, 14,* 6171–6186.
Prinz, J. J. (2000). A neurofunctional theory of visual consciousness. *Consciousness and Cognition, 9,* 243-259.
Prinz, J. J. (2001), 'Functionalism, Dualism and the Neural Correlates of Consciousness', in W. Bechtel, P. Mandik, J. Mundale, and R. Stufflebeam, eds., *Philosophy and the neurosciences: A reader* (Oxford: Blackwell).
Prinz, J. J. (2003). Level-headed mysterianism and artificial experience. *Journal of Consciousness Studies, 10,* 111-132.
Prinz, J. J. (2004). *Gut reactions: A perceptual theory of emotion.* New York, NY: Oxford University Press.
Ramsden, B.M., Chou, P.H., and Roe, A.W. Real and illusory contour processing in area V1 of the primate: a cortical balancing act. *Cerebral Cortex, 11,* 648–65, 2001.
Rees G., Kreiman G., and Koch C. (2002). Neural correlates of consciousness in humans. *Nature Reviews Neuroscience, 3,* 261-270.
Rees G., Wojciulik E., Clarke K., Husain M., Frith C. and Driver J. (2000). Unconscious activation of visual cortex in the damaged right hemisphere of a parietal patient with extinction. *Brain, 123,* 1624-1633.
Rees, G. (2001). Neuroimaging of visual awareness in patients and normal subjects. *Current Opinion in Neurobiology, 11,* 150–156.
Rolls E.T. (1998). The orbitofrontal cortex. In Roberts A.C., Robbins T.W., AND Weiskrantz L. (eds.) *The prefrontal cortex* (67–86). Oxford: Oxford University Press.
Sahraie, A., Weiskrantz, L., Barbur, J. L., Simmons, A., Williams, S. C. R., & Brammer, M. J. (1997). Pattern of neuronal activity associated with conscious and unconscious processing of visual signals. *Proceedings of the National Academy of Science, 94,* 9406–9411.
Savic, I., Gulyas, B., Larsson, M., and Roland, P. (2000). Olfactory functions are mediated by parallel and hierarchical processing. *Neuron, 26,* 735–745.
Seghier, M., Dojat, M., Delon-Martin, C., Rubin, C., Warnking, J., Segebarth C., and Bullier, J. (2000). Moving Illusory Contours Activate Primary Visual Cortex: an fMRI Study. *Cerebral Cortex, 10,* 663-670.
Seguin, E. G. (1886). A contribution to the pathology of hemianopsis of central origin (cortex-hemianopsia). *Journal of Nervous and Mental Diseases, 13,* 1-38.

Sheinberg D.L. and Logothetis N.K. (1997). The role of temporal cortical areas in perceptual organization. *Proceedings of the National Academy of Science, 94,* 3408-3413.
Vuilleumier P, Sagiv N, Hazeltine E, Poldrack, R. A. Swick, D, Rafal, R .D., and D. E. Gabrieli, J. (2001). Neural fate of seen and unseen faces in visuospatial neglect: a combined event-related functional MRI and event-related potential study. *Proceedings of the National Academy of Science, 98,* 3495–500
Vuilleumier, P., Henson, R. N., Driver, J., and Dolan, R. J. (2002). Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. *Nature Neuroscience, 5,* 491-499.
Wilson, F. A. W., S. P. O'Scalaidhe, and Goldman-Rakic, P. S. (1993). Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science, 260,* 1955-1958.
Zeki, S. (1993). *A vision of the brain*. Oxford: Blackwell.