Supplementary Note 2: Neural Correlates of Attention

Jesse Prinz

I have argued that perceptual states become conscious when and only when we attend, and that attention is the process by which perceptual states become available to, but not encoded, in working memory. This supplements the theory of the contents of consciousness that I defended elsewhere: the intermediate-level hypothesis. Intermediate-level representations are rendered conscious when attention makes them available. With this AIR Theory in hand, we are now in a position to address the question of neural correlates. To find the NCC, we need find the neural correlate of attention, and to find the neural correlate of attention, we need to find out what processes allow information to flow from perceptual pathways in the brain to pathways associated with working memory.

The search for NCCs is regarded as important because theories of consciousness stated at a psychological level of analysis often seem too liberal about the conditions under which consciousness can arise. The AIR theory is no exception. It would be relatively easy to program a computer to represent sensory (e.g., video) inputs in a hierarchical way, and have some of the intermediate-level states in the hierarchy be selectively available to the computer's working memory (e.g., RAM access). Intuitively, it seems outrageous to attribute consciousness to such a computer program. Folk intuitions resist the idea that consciousness can arise so easily in inorganic matter (Huebner et al., ***). Some philosophers are inclined to take such intuitions vary seriously as evidence against theories of consciousness stated at the psychological level. Bock (1978) uses such intuitions as the centerpiece of his argument against functionalism, saying that psychologically specified theories of consciousness could be implemented by the large collections of human being spread across geographical space, such as the population of China or the economy of Boliva. He takes it as obvious that such a group would not, as a collective, instantiate conscious experiences. But why should we assume that our intuitions are reliable about this? It would beg the question to simply assume that consciousness requires a particular kind of realizer (see also Huebner, for a critique of Block and a defense of group consciousness, ***).

The fact that intuitions are uninformative does not mean we should give up the quest for neural correlates. First of all, neural correlates *might* be necessary for consciousness. By definition, they correlate with consciousness, and are thus, like psychological processes, good candidates for the physical bases of experience. Second, we do know that changes in cellular processes in the brain can affect the quality of experience even if, at some gross level, function remains constant (for more discussion of this, see Prinz, ***; ***). If so, the AIR theory needs to be supplemented with a theory about neural correlates. One might put the point by saying that attention is not *merely* availability to working memory, but availability that comes about in a particular way. Neuroscience can help us discover the specific physical processes that occur when we attend, and those processes may be necessary for experience. (Elsewhere, I argue that both levels are necessary (***).)

The brain can be analyzed at many different levels, from whole systems down to cells and their components. As a starting place in our search for the neural correlates of AIRs, one may look for the large-scale brain structures associated with attention and working memory. Attention is usually associated with inferior portions of posterior parietal cortex (as noted in the discussion of neglect). If I am right about the nature of attention, then we might think about inferior parietal cortex, not as the locus of attention, but as a common control structure that plays a role in many attentional tasks. Attention itself is located in the perceptual pathways. It is not a structure, nor even a mechanism, per se, but rather a way of processing perceptual representations. Attention can be influenced by structures throughout the brain, but parietal cortex is often involved. Working memory is associated with lateral frontal cortex. It is believed that different sense modalities, and different kinds of executive processes recruit deferent working memory mechanisms (***; ***) but these are generally located in lateral frontal areas. So, at the level of gross anatomy, one might suppose that consciousness will characteristically involve a circuit that includes perceptual pathways, such as the visual and auditory systems in temporal cortex, parietal attention control centers, and lateral frontal structures associated with working memory.

This circuitry is borne out by neuroimaging studies. When people are conscious, all three components of the circuit tend to be active. This is the case, for example, in studies of conscious change detections (Rees, ***) and also in the study discussed earlier of visual experience in blindsight (Sahraie et al., 1997). When conscious experience arises in the portion of the visual field that is usually blind, activation can be seen in temporal, parietal, and lateral frontal cortex.

Such systems-level observations do not tell us anything about the neural processes underwriting conscious experience, so we need to delve deeper. One thing that has been established repeatedly is that attention correlates with increased in neuronal activity as measured by BOLD signals in fMRI. The locus of increase is in perceptual areas corresponding to the feature dimension being attended. For example, Corbetta and Shulman (2002) recorded increases in MT, the visual area associated with motion when subjects we asked to focus on the direction of a moving stimulus. Of course, it would be a mistake to simply identify attention with an increase in metabolic responses. Brain areas fluctuate in activation all the time, and many of the increases have little to do with attention. We need to know what is causing these increases at a cellular level, and, for that, BOLD signals are of limited value.

More information can be gleaned from cell recordings in humans or, more typically, in laboratory animals. One promising lead comes form a recent study by Mitchell et al. (2007) who trained macaque monkeys to visually track two cued shapes in a group of four shapes that moved around on a computer screen. During this task the recorded cells in intermediate-level visual area V4. The researchers were interested in whether different kinds of cells behaved differently during this attention-demanding task. Cells were divided into two distinct groups based on the waveforms they produced. There were interneurons with basket cell and chandelier cell morphology, on the one hand, and pyramidal cells, on the other. The former have short-lasting action potentials, and can therefore be called narrow-spiking neurons, and the latter have longer action potentials, and can be called broad-spiking. Both types of cells changed from baseline firing rates during the task, but interestingly, it was only the fast- or narrow-spiking neurons that consistently increased from baseline, in contrast to the slow- or broad-spiking cells, which both increased and decreased. This suggests that the net gains in activity during attention may owe more to narrow-spiking neurons than to broad.

One thing that makes this finding intriguing is that narrow-spiking interneurons are not excitatory, like pyramidal cells, but inhibitory. Interneurons also communicate across shorted distances. So the net increases in neuron activity during attention owes to increases in neurons that inhibit activity and don't send signals that travel very far. This is counter-intuitive. On the AIR theory, consciousness is supposed to involve signals being sent across long distances, from perception to working memory. Why would short-distance inhibition help with that? One answer has to do with an effect that narrow-spiking inhibitory neurons have on the cells to which they are connected. The inhibitory signals increase neural synchrony, and they do this in all types of cells (Sohal and Huguenard, 2005). In fact, the synchrony they bring about tends to be within the same frequency—the so-called gamma band. Within the gamma band, neuronal activity oscillates at a rate of 30–90Hz. This is significant because there has been speculation for a long time that gamma oscillations (often abbreviated as 40Hz) are associated with consciousness (von der Malsburg, ***; Singer, ***; Singer and Gray, ***; Singer, ***; Crick and Koch, ***; Hameroff, ***).

The association between gamma synchrony and consciousness has been controversial. Crick and Koch (1991) originally advanced a hypothesis that anticipates the story I am telling here, but they subsequently abandoned the idea, citing insufficient evidence (***). I think their change of heart was pre-mature. The cumulative evidence gamma synchrony is now quite impressive. For example, gamma synchrony is disrupted when people undergo anesthesia (John and Prichep, 2005); it is observed during dream sleep (Llinas and Ribary, 1993); it increases with deviant auditory signals (Edwards et al. 2005) and with bright high contrast visual displays (Biederlack et al., 2006; ***); and it correlates with stimulus selection in binocular rivalry (Engel et al., 1996; Fries et al. 2002).

That hypothesis fits nicely with the AIR theory, because gamma synchrony has been associated with attention. Fries et al. (2001) found that V4 cells fire in the gamma band when a monkey is attending to gradients presented in those cells' receptive fields. Nakatani et al. (2005) found that gamma oscillations correlated with perceived targets in attentional blink. Steinmetz et al. (2000) found that synchrony correlated with attention across sense modalities, in a vision to touch matching task. Schurger et al. (2006) found that gamma band activity occurs in blindsight under conditions where a stimulus is consciously seen, as compared to when it is invisible. Gamma synchrony is also associated with working memory, which, I argued, is the target of attention (Jokisch & Jenson, 2007).

Thus, it looks like there is good evidence for a link between gamma synchrony and attention, and a promising explanation for that connection stemming from the increased activity in inhibitory interneurons. But this just raises another question. What's so important about synchrony? The answer has to do with communication (see, e.g., Fries, 2005). Imagine you are on a balcony above a big crowd and everyone is engaged in conversation with the people in their immediate surround. The result will be a indecipherable cacophony of speech sounds. But now suppose that some group within the crowd starts saying the same thing at the same time. Their collective voice will rise above the multitudes, and be audible and discernable up above. Likewise, when neurons fire together, they are able to send signals that reach other areas of the nervous system. There is good evidence that synchrony increases output (Salinas and Sejnowski, 2001) and facilitates connectivity (see e.g. Vanni, 1999). Melloni et al. (2007) found evidence for this in a masking task. They measured electrical activity on the scalp while subject perceived a word either visibly (followed by a low luminance mask) or invisibly (followed by a high luminance mask). The found that in the visible condition, there was a much higher degree of synchrony across distantly connected brain areas.

Putting this all together, attention seems to involve an increase in inhibitory interneuron activity, which increases synchrony, which amplifies output, and allows for long-distance communication, including communication from perceptual pathways to working memory centers.

As stated, this account remains incomplete in three important ways. First, it leaves us with two candidates for the neural correlate of attention. Attention is associated both with the increased interneuron activity and with the resultant neural synchrony. Both correlate with consciousness, and either could be the coveted NCC. Are both necessary? Is one enough? Ideally one could find dissociations between the two and show that one of them correlates more perfectly with attention and consciousness. This has not been done.

Second, both of these neural correlates co-occur with more specific changes, including changes at lower levels of analysis. Recent research suggests that different kinds of interneurons make different contributions to gamma synchrony. Measuring in the hippocampus, Tukker et al. (2007) found that bistratified cells were more highly associated with synchrony in pyramidal cells, than Parvalbumin-expressing interneurons, which had a modest influence. Cholecystokinin-expressing interneurons seemed to play a thresholding function on pyramidal cell response, and this may contribute to gamma initiation. There are computational models in which synchrony begins with sparse pyramidal excitation, followed by interneuron inhibition, which then leads to a pattern of re-entry between these two classes of neurons and that results in the measure oscillations (Mann et al. 2005). Future work is needed to test models of this kind, and arrive at a precise characterization of how synchrony is initiated and sustained.

Synchrony itself has neural correlates at a lower level of analysis. It has traditionally been understood in terms of neural spikes, but Hameroff (2007) has recently argued that that synchrony is better understood in terms of events at the gap junctions in “dendritic webs” (the connections between dendrites). Moreover, he speculates that quantum events play a role in these. This proposal has met with tremendous controversy (see, e.g., ***; ***). At this point, the evidence for quantum contributions to consciousness meager. Hameroff conjectures that quantum events contribute to long-distance synchonization, but there is no direct evidence for this, and extant computational models can account for long-distance synchronization without appeal to quantum activity that has yet to be correlated with conscious experience (***). It is also troubling that the appeal to quantum events seems to have been motivated not be direct empirical observation, but because quantum processes are indeterministic, and that fact leads defenders of quantum consciousness to think that their approach can be used to defend the existence of free will. I think such a defense of free will is unpromising on multiple grounds. As Hume taught, chance is not the same as freedom. And, more to the point, there is no actual evidence for free will in any robust sense; the belief is just an article of faith, and that means quantum consciousness is grounded on a flagrant wishful thinking, rather than being driven by empirical findings. This is not to say that the quantum approach has been ruled out. Hameroff (***) has been adroit in addressing critics. It's just that the approach has not been ruled in, and it is currently too speculative to endorse with any confidence. The point of these remarks is that there is, as yet, no settled view about the precise means by which interneurons create synchrony, nor the precise cellular dynamics underlying synchrony.

Third, there is an even more serious limitation of current research. Both interneuron inhibition and gamma synchrony are widespread in the brain (see, e.g., ***). They are not restricted to attentional processing, and much less to intermediate-level perceptual areas. If consciousness is to be identified with a physical process, then these cannot be sufficient correlates, because they arise in brain areas that are not associated with consciousness. There may be specific sources of synchrony that are unique to conscious experience. For example, Grossberg (2008) has been developing a complex model that involves thalamocortical connections and specific laminar dynamics (on the latter, see also Crick, ***). Once we start talking about interlaminar microcircuits, we get may see a neural signature of consciousness emerge. The science is not at that point. Alternatively, it may turn out that there is no neural signature at this fine level of analysis. Just as we should not necessarily expect conscious to involve a unique class of neurotransmitters, we might not expect consciousness to involve a unique class of neurocomputations. Perhaps consciousness depends on what such computations achieve at a more macro level. I explore this theme elsewhere (***).

A tentative an incomplete conclusion then, is that the NCCs probably include synchronic oscillations or the interneuron inhibition processes that bring these about. More specific correlates underlying these processes are also correlates of consciousness, but their exact identity is still be actively explored. Which, if any, of these neural process is necessary or sufficient for conscious experience is a topic I take up elsewhere.