17 Two Neural Correlates of Consciousness

From Consciousness, Function and Representation, Collected Papers, Volume 1, MIT Press, 2007

This is a longer version of Block, N. (2005). "Two Neural Correlates of Consciousness." Trends in Cognitive Sciences 9(2): 46-52.

I have previously proposed a conceptual distinction between phenomenal consciousness and access consciousness (Block 1990, 1992, 1995). Phenomenally conscious content is what differs between experiences as of red and green, whereas access-conscious content is content, information about which is ''broadcast'' in the ''global workspace.'' Some have accepted the distinction but held that phenomenal consciousness and access consciousness coincide in the real world (Chalmers 1996, 1997; but see Block 1997). Others have accepted something in the vicinity of the conceptual distinction but argued that only access consciousness can be studied experimentally (Dehaene and Changeux 2004). Others (Dennett 1995) have disparaged the conceptual distinction itself. This article argues that the framework of phenomenal consciousness and access consciousness helps to make sense of recent results in cognitive neuroscience; we see a glimmer of an empirical case for thinking that they correspond to different NCCs.

Phenomenal NCC

Christof Koch (2004, 16) defines ''the'' NCC as ''the minimal set of neuronal events and mechanisms jointly sufficient for a specific conscious percept.'' However, since there is more than one concept of consciousness, this definition allows that a given percept may have more than one NCC. In my proposed framework, the Phenomenal NCC is the minimal neural basis of the phenomenal *content* of an experience, that which differs between the experience as of red and the experience as of green. I will start with an example: the neural basis of visual experiences as of motion is likely to be activation of a certain sort in area MT/V5.[1] (Philosophers often use the terminology ''as of motion'' instead of ''of motion'' since the experience can and does occur without motion.) The evidence includes

- Activation of MT/V5 occurs during motion perception (Heeger et al. 1999).
- Microstimulation to monkey MT/V5 while the monkey viewed moving dots influenced the monkey's motion judgments, depending on the directionality of the cortical column stimulated (Britten et al. 1992).

- Bilateral (both sides of the brain) damage to a region that is likely to include MT/V5 in humans causes akinetopsia—the inability to perceive—and to have visual experiences as of motion. Akinetopsic subjects see motion as a series of stills. (See Zihl, von Cramon, and Mai 1983; Rees, Kreiman, and Koch 2002.)
- The motion after effect—a moving afterimage—occurs when subjects adapt to a moving pattern and then look at a stationary pattern. These moving afterimages also activate MT/V5 (Huk, Ress, and Heeger 2001).
- Transcranial magnetic stimulation (TMS[2]) applied to MT/V5 disrupts these moving afterimages (Théoret et al. 2002).
- MT/V5 is activated even when subjects view "implied motion" in still photographs—for example, of a discus thrower in midthrow (Kourtzi and Kanwisher 2000).
- TMS applied to the visual cortex in the right circumstances causes phosphenes[3]—brief flashes of light and color (Kammer 1999). When TMS is applied to MT/V5, it causes subjects to experience moving phosphenes (Cowey and Walsh 2000).

Mere activation over a certain threshold in MT/V5 might not be enough for the experience as of motion: the activation probably has to be part of a feedback loop, what Lamme (Lamme and Roelfsema 2000; Lamme 2004) calls recurrent processing. Pascual-Leone and Walsh (2002) applied TMS to both MT/V5 and V1 in human subjects with the pulses placed so that the stationary phosphenes determined by the pulses to V1 and the moving phosphenes from pulses to MT/V5 overlapped in visual space. When the pulse to V1 was applied 5 to 45 msec later than to MT/V5, all subjects said that their phosphenes were mostly stationary instead of moving. (See Pascual-Leone and Walsh 2002 for references to single-cell recording in monkeys that comport with these results.) The delays are consonant with the time for feedback between MT/V5 and V1, which suggests that experiencing moving phosphenes depends not only on activation of MT/V5 but also on a recurrent feedback loop in which signals go back to V1 and then forward to MT/V5 (Pascual-Leone and Walsh 2002).

So recurrent activity in and around MT/V5, in the context of other brain areas functioning normally—exactly which brain areas are required is unknown at present—is a good bet for being the physical basis of visual experience as of motion. (But see box 17.1 as well as Zeki and ffytche 1998 and Sincich et al. 2004 for some data that complicate the conclusion.) Corresponding conclusions can be drawn for other types of contents of experience. For example, recurrent activation of the fusiform face area on the ventral (bottom) surface of the temporal lobe (again in context) may determine experience as of a face (Kanwisher 2001). The overall conclusion is that there are different Phenomenal NCCs for different phenomenal contents. (See Zeki 2001 on microconsciousness; also see Pins and ffytche 2003.)

Of course no one would take activation of MT/V5 + recurrent loops to V1 all by itself in a bottle as sufficient for experience of motion. (See box 17.2.) A useful distinction

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 344)

Box 17.1
Blindsight and MT/V5

The picture presented in the text is complicated by attention to studies involving blindsight patient GY, who has experiences of motion that may be visual but does not have the corresponding part of V1. GY does well in forced-choice guesses about stationary stimuli to his blind field that he says he does not see. But he says he is aware of some moving stimuli (Weiskrantz 1997). Functional magnetic resonance imaging (fMRI) shows that GY's area MT/V5 is activated when he is aware of moving stimuli presented to his blind field (Weiskrantz 1997). However, he does not experience moving phosphenes when TMS is applied to MT/V5 in the left hemisphere of his brain, where he is missing the corresponding V1 (Zeki and ffytche 1998). Recent neuroanatomy has shown that there is a pathway between the eyes and MT/V5 that bypasses V1 (directly from the LGN—the neural way station between the eyes and the cortex) (Sincich et al. 2004). GY has spoken to investigators about his experience. In 1994, GY said that his experience of motion in the blind field was "a 'feeling' of something happening in his blind field" (Zeki and ffytche 1998, 29). In 1996, he said his experience was that of "a black shadow moving on a black background" (p. 30). The shadow description comports with Riddoch's 1917 paper, which included studies of five patients who had gunshot wounds affecting V1 in World War I. (Zeki and ffytche (1998)—commendably and rarely in neuroscience—quote some of these patients.) The conclusion I would draw from reading what these subjects and GY say is that their experiences are very abstract, involving pure motion without any other experiential features such as color, light, shape, or contour. (Some philosophers I have mentioned this to wrongly think this description is incoherent!) It is not certain that these motion experiences should be described as visual. One suggestion is that activation of MT/V5 requires feedback loops to lower areas for experiences as of color, light, shape, and contour and for moving color, light, and so on, but not for pure motion. However, it may be that recurrent processes are necessary for all conscious experience, since there may be recurrent processes feeding back to MT/V5 from higher areas.

here is that between a *core* and a *total* NCC (Shoemaker 1981; Chalmers 2002). The total NCC of a conscious state is—all by itself—sufficient for the state. The core NCC is the part of the total NCC that distinguishes one conscious content from another—the rest of the total NCC being considered as the background conditions, which supply the rest of the sufficient condition. (One interesting issue is whether there might be somewhat different background conditions for different experiential contents, or whether the background conditions—at least in a single sensory modality—are always the same.[4]) In these terms, then, the core Phenomenal NCC for the neural basis of the experience as of motion as opposed to the experience as of red or as of a face, is likely to be recurrent activation of MT/V5. See figure 17.1. (See box 17.3 for some doubts about the concept of an NCC.)

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 345)

Box 17.2
Area MT/V5 in a Bottle?

The total Phenomenal NCC for the experience as of motion is a sufficient condition all by itself for the experience. What might that turn out to be? I suggest approaching it by asking what we could remove from a normal brain and still have that experience. My suggestion is that we might be able to remove—at least—areas responsible for access to experiential contents and still have the heart of the same experiential contents. (In my approach, areas responsible for access to experiential contents probably also are responsible for conceptualization of those contents. So experiential contents without access might be nonconceptual, or may only involve purely sensory concepts.) Nakamura and Mishkin (1980, 1986) removed frontal, parietal, and superior temporal areas in one hemisphere of monkeys, leaving what is usually considered the visual system intact. They also disconnected visual inputs to the undamaged hemisphere. This preparation is sometimes said to cause blindness (Rees, Kreiman, and Koch 2002), but Nakamura and Mishkin are careful to say that this is shorthand for behavioral unresponsiveness to visual stimuli (at least temporarily), and should not be taken to show complete lack of visual sensation. One intriguing result is that when the limbic (emotional) system in the damaged hemisphere is intact, the monkeys showed eye and head movements as if engaged in visual exploration. This contrasts with monkeys in which V1 is ablated who stare fixedly.

Access NCC

We can distinguish between phenomenal contents of experience and access-conscious contents, information about which information is made available to the brain's "consumer" systems: systems of memory, perceptual categorization, reasoning, planning, evaluation of alternatives, decision making, voluntary direction of attention, and more generally, rational control of action. Wide availability motivates the idea that there is some mechanism via which producing systems can communicate with all the consuming systems at once, a "global workspace" (Baars 1997), and that information concerning conscious representations is "broadcast" in this global workspace. According to the global-workspace metaphor, the sensory systems are the "producers" of representations, and the aforementioned systems are the "consumers." The neural basis of information being sent to this global workspace is the "Access NCC."[5]

Rees, Kreiman, and Koch (2002) note that in studies of the neural correlates of bistable perception, in which there are spontaneous fluctuations in conscious contents, reports of conscious contents correlate with activation in frontal and parietal areas. Dehaene and Changeux (2004) suggest that a significant piece of the neural machinery of what they call "access to consciousness" (roughly equivalent to my access consciousness) is to be found in "workspace neurons" that have long-range excitatory

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 346)

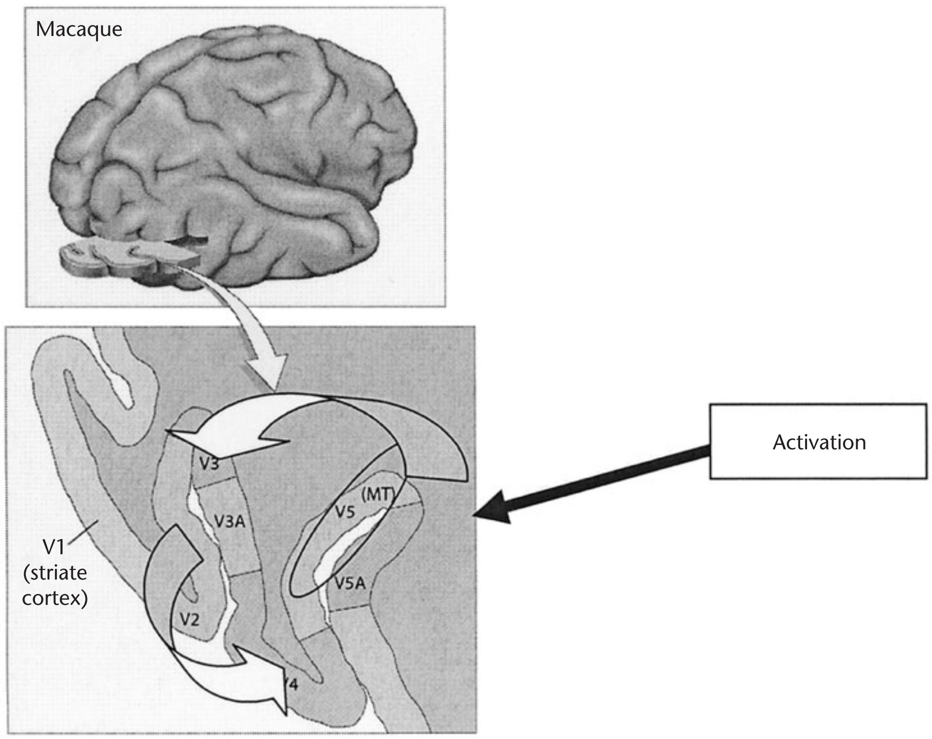

Figure 17.1
The core Phenomenal NCC for the visual experiential content as of motion: MT/V5 activation with recurrent loops to and from lower areas. The arrows are supposed to indicate recurrent loops. Adapted from S. Zeki, *A Vision of the Brain* (Oxford: Blackwell, 1993), 97, as modified by M. Gazzaniga, R. Ivry, and G. Mangun, *Cognitive Neuroscience*, 2nd ed. (New York: Norton, 2002). Arrows indicating recurrent loops were added.

axons allowing, for example, visual areas in the back of the head to communicate with frontal and parietal areas toward the front of the head. Thus it is a good guess that the Access NCC, the neural basis of access, is activation of these frontal and parietal areas by occipital (classic ''visual'') areas in the back of the head. (See figure 17.2.)

As Dehaene and his colleagues (2004) have emphasized, there is a winner-take-all competition among representations to be broadcast in the global workspace.[6] This point is crucial to the nature of the Access NCC and the difference between it and the Phenomenal NCC. One item of evidence for winner-take-all processes derives from the attentional blink paradigm, in which the subject is given a string of very brief visual stimuli, most of which are distractors. The subject is asked to report on one or two

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 347)

Box 17.3
NCC or NDC?

I have been talking about the ''neural correlates of consciousness.'' But the evidence of the sort just described argues for something both weaker and stronger than correlation:

- *Weaker*, because none of the evidence cited has anything to say about whether there is some other sort of physical constitution—an alternative biology, or even silicon chips—that is sufficient for the same experiences. The evidence supports a *one*-way connection, neural → experiential, not a *two*-way connection, neural ↔ experiential.
- *Stronger*, because it is evidence for determination, not just correlation. There is a correlation between the temperature in Brooklyn and Manhattan, but there is no necessity to it. The relation between recurrent MT activation and experience as of motion appears to be a necessary one: you cannot have (recurrent) activation of MT/V5 (together with certain unknown supporting areas) without visual experience as of motion.

Thus we should really be thinking about ''NDC'' for ''neural determiner of consciousness'' instead of NCC. (I will continue to use the acronym ''NCC'' since it is established terminology.)

''targets'' after the sequence of rapid visual stimuli. If there are two targets separated by an appropriate delay, the subject does not report seeing the second one, even though the second one would have been likely to be reported if the subject had not been given the first target. Dehaene, Sergent, and Changeux (2003) used a modified attentional blink paradigm, in which subjects were asked to indicate on a continuous scale the visibility of the second target. The second target was at its peak of invisibility when the targets were separated by 260 msec. The result of interest here is that the subjects almost never used the *intermediate* cursor positions (at the 260 msec delay)—that is, they rated the ''blinked'' stimulus as either totally unseen or as maximally seen almost all the time. Thus Phenomenal NCC activations compete for dominating the Access NCC. Importantly, it is not the case that the Phenomenal NCC representation that is highest in initial activation will dominate, because domination can be the result of ''biasing'' factors such as expectations or preferences (Lamme 2003, 2004).

Although the winning Phenomenal NCC will in general be amplified by the recurrent loop, a losing Phenomenal NCC may itself involve recurrent loops to lower areas that will be sufficient for an experiential or phenomenal content. For example, an activation of area MT/V5 might have recurrent interactions with V1, making it the neural basis of an experiential content, but nonetheless lose in the winner-take-all competition and so not be accessed (Lamme 2004). The general point is that the simplest and most explanatory theory may be one in which recurrent MT/V1 loops are sufficient for an experiential content despite not being accessible when they lose the winner-take-all

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 348)

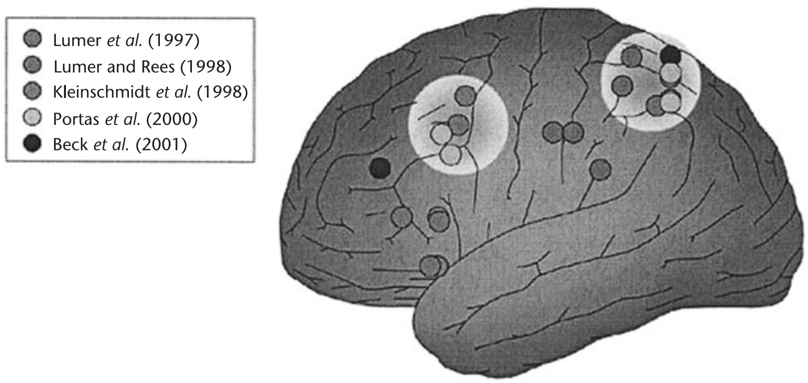

Figure 17.2
Suggestion for the core Access NCC for visual experiences, from G. Rees, G. Kreiman, and C. Koch, "Neural correlates of consciousness in humans," *Nature Reviews Neuroscience* 3, 4 (2002): 261–270. Activations cluster in superior parietal and dorsolateral prefrontal cortex as indicated by large light circles. These are frontal and parietal areas that fluctuate spontaneously in binocular rivalry and other bistable perception in a way that is time-locked to fluctuation in reported experience. The core Access NCC may be activation of these areas by neural firing in the occipital cortex in the back of the head. Do we count the Phenomenal NCC as part of the Access NCC—in which case this figure pictures the Access NCC minus the Phenomenal NCC? Or do we regard the Access NCC as not including the Phenomenal NCC, in which case this figure pictures the Access NCC? This is a terminological issue—assuming that phenomenal consciousness is the gateway to full-fledged access consciousness.

competition. Thus the winner-take-all process that is part of the nature of global broadcasting also strongly suggests that the Phenomenal NCC can be instantiated without the Access NCC, so global broadcasting does not encompass all of consciousness. This idea is further bolstered by evidence that there is brief parallel processing of many objects in the ventral visual stream[7] (up to inferotemporal cortex) before zooming in on one or two of them (Rousselet, Thorp, and Fabre-Thorp 2004).

But Is the Phenomenal NCC Really the Neural Basis of a Kind of Consciousness?

You may ask, "If the Phenomenal NCC can perhaps occur without the Access NCC, how do we know that the Phenomenal NCC is really the neural basis of anything conscious?" A quick answer is that, since the Phenomenal NCC determines the contents of experience, what it determines is ipso facto a kind of consciousness. The Phenomenal NCC for visual motion determines the experiential content of visual motion—as distinct from the experiential content of seeing something as a face. That content itself is a kind of phenomenology, a kind of consciousness. If there could be a phenomenal

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 349)

content without anything that could be called awareness of it, some might not want to apply the word ''consciousness'' to it. For this reason, Burge (1997) distinguishes between phenomenality—which he is uncomfortable about calling a kind of ''consciousness''—and phenomenal consciousness, which is phenomenality that is the subject of some kind of access.[8] If one accepts Burge's terminology, though, it is important to realize that phenomenality is the important and puzzling phenomenon that is the heart of the mind-body problem and what we do not understand how to explain in neurological terms. If we could solve the ''Hard Problem of consciousness'' (Chalmers 1996) for phenomenality in Burge's sense, there would be no ''Hard Problem'' left for phenomenal consciousness in Burge's sense.

But this answer is too quick, since the doubt that motivates the question is a doubt that the Phenomenal NCC really does determine the contents of experience, and since the Phenomenal NCC was defined in terms of the contents of experience, the doubt challenges the evidence presented earlier for a Phenomenal NCC. The doubter may say that without access, there can be no true phenomenal contents but only *protocontents* that become contents when globally broadcast. But how does the doubter claim to know that? Some are motivated by a terminological point—that we should not call something ''phenomenal'' or ''conscious'' if it is not broadcast for access (Kanwisher 2001). However, the substantive empirical question is the following: If our evidence always concerns phenomenal contents that are actually accessed, how can the Phenomenal and Access NCC ever be empirically distinguished?

The answer is that it is not true that our evidence always concerns experiential contents that are accessed. There are a variety of paradigms in which we can use convergent evidence involving varying degrees of access to try to separate out the Phenomenal from Access NCC. One such paradigm is signal-detection theory.

Signal-Detection Theory (SDT) Approaches

Suppose a subject is shown a series of stimuli at around threshold level and asked to press one button if a target is visible and another if not. SDT models the subject's behavior in terms of two factors: the extent to which the subject sees the target and the criterion the subject implicitly sets for reporting seeing it. The criterion is famously influenceable by features of the experimental setup that affect subjects' expectations or motivation—such as by the proportion of ''catch trials'' (where no stimulus is presented) and by rewards for hits and penalties for false alarms. We know from standard SDT analyses that subjects' reports of whether there was a target or whether they saw it *do not just reflect the extent to which they did see it* (i.e., did have a visual phenomenal state), but also the subjects' threshold for reporting and even for believing that they did see it. Two experimental setups in which there are the *same* experiential contents may result in *different* beliefs and different reports.

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 350)

A dramatic example is a series of experiments concerning the "exclusion" paradigm (Debner and Jacoby 1994), in which subjects are instructed to complete a word stem with something other than the end of a masked word that has just been presented to them. If the word "reason" is presented "unconsciously" at 50 msec, subjects are more likely than baseline to *disobey* the exclusion instructions, completing "rea____" with "son," whereas if "reason" is presented "consciously" at 250 msec, subjects are more likely than baseline to choose some other ending (e.g., as in "reader"). This paradigm has impressed many because it appears to yield opposite results for unconscious and conscious stimuli. However, Visser and Merikle (1999) showed that changing the motivation of the subjects by using a reward structure can change the degree of exclusion. They started subjects with a $15 credit and docked them $1 for each error. Visser and Merikle interpret the result in terms of the effect of reward/punishment on increased attention, accepting the idea that the 50 msec/250 msec difference engenders an unconscious/conscious difference. But there is an alternative SDT interpretation suggested by Snodgrass (2002), in which the results in part *reflect a criterion shift rather than a difference in consciousness*. The idea is that punishment for errors of failing to exclude pushes the criterion level (the degree of phenomenal experience that the subject implicitly sets as a condition for action) for inhibiting the immediate response so low that weak conscious perception of "reason" blocks use of "son" even though the subjects are so lacking in confidence that they say and think they do not see the word. That is, their criterion level for inhibiting the immediate response is lower than their criterion level for believing that they saw a word, and the phenomenal level is in between the two criteria. A subject's state of mind when successfully excluding one of the 50 msec stimuli could be articulated—overarticulated, no doubt—as "I probably didn't see a word but if I did, it was 'reason,' so I'd better complete the stem with 'reader'" (Block 2001). And the SDT interpretation is confirmed by the effect on "inclusion" instructions. With "inclusion" instructions, the subjects see "reason" and then are given "rea____" but are told to complete the stem with *the word they saw* if possible. In this paradigm, SDT predicts no shift with change in reward or punishment, because there is no issue of a criterion level. There is no degree of experience that subjects implicitly set as a condition of acting: rather, they just use the first word that comes to mind regardless of level of confidence that it is the word they saw. And the result (Visser and Merikle 1999) is just that: the difference in reward/punishment structure makes no difference in the result under "inclusion" instructions. Thus there is a striking difference in the effect of reward on exclusion as compared with inclusion instructions.

There is, therefore, evidence in the "exclusion" case of experiential contents (e.g., as of seeing "reason") without the kind of access required for report, planning, decision making, evaluation of alternatives, memory, and voluntary direction of attention. Some of the 50 msec stimuli are weakly conscious although not broadcast in the global

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 351)

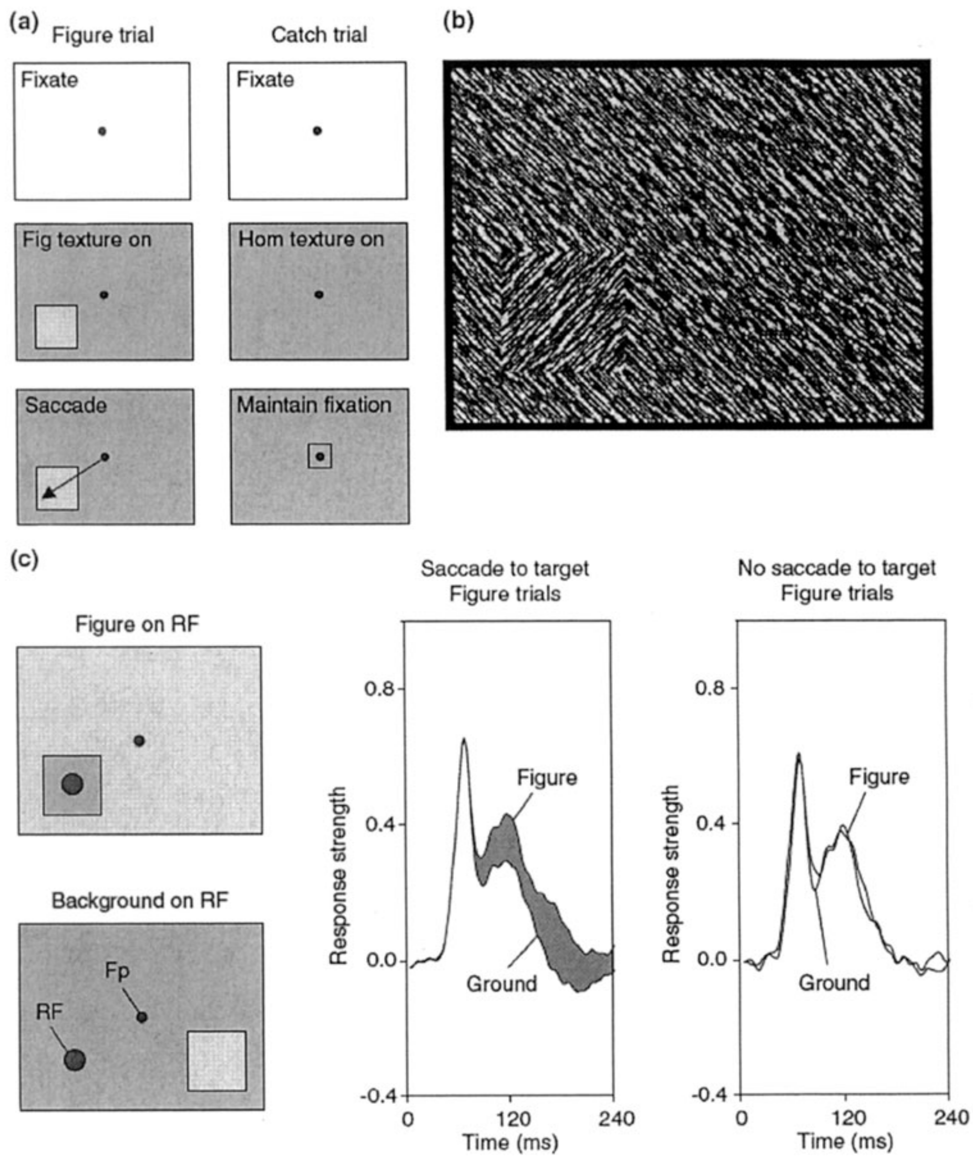

Figure 17.3

(a) Supèr, Spekreijse, and Lamme (2001) trained monkeys to saccade from a fixation point to a target (bottom left of (a)). Initially, a fixation point was presented (top). Then a target texture was presented (''Fig texture on,'' left) or there was a homogeneous pattern with no target (''Hom texture on,'' right). If there was no target, the monkey was rewarded for maintaining fixation for 500 msec (right panels). The target could be in one of three locations. (b) The targets were areas of an workspace. Thus SDT gives us reason to think that experiential content—based on the Phenomenal NCC—can be instantiated without the kind of access that is based in the Access NCC.

Neural SDT

In a landmark series of experiments, Supèr, Spekreijse, and Lamme (2001) recorded from V1 (which, you will recall, is the first classic "visual" area in the cortex) during a task in which monkeys were rewarded for saccading to a target if there was one or continuing to look at the fixation point if not. (A saccade is an eye movement whose function is to make a region of interest project to the densest part of the retina; in natural visual exploration, there are roughly two per second, although the movement itself takes only 30 msec.) Supèr and colleagues manipulated whether the locations in V1 corresponded to figure or ground. When the monkey detected (saccaded to) the target, there was an increased V1 response for figure as compared with ground. See figure 17.3, in which this increased figure response is referred to as "modulation."

Supèr and colleagues were able to manipulate the modulation by varying the saliency of the stimulus (i.e., the number of pixels in line segments in the target; see figure 17.3b) and the proportion of "catch trials" in which there was no target. For high-saliency stimuli and small numbers of catch trials, there was a near-perfect correlation

overall pattern in which the lines were orthogonal to the rest of the pattern. (c) Supèr et al. recorded from sites in V1 whose receptive fields (RF) included the three locations in which targets could occur. When the monkey saccaded from the fixation point (Fp) to the target, the neural response from the target counted as "figure" and the other two sites were counted as "ground." Figure responses were greater than ground responses after ~90 msec, as indicated in the shaded area (central panel). The shaded area indicates the degree of "modulation." When the targets were highly salient and the number of catch trials were few, modulation disappeared when the monkey did not detect the target (right panel). That is, when the monkey did not saccade to the target and the saliency was high and catch trials low, there was little difference between the activity in the part of V1 corresponding to the target and the two other locations, as indicated in the right-most panel of (c). (However, when the saliency of the target was low or catch trials high, there was a substantial difference.) Modulation also disappeared under anesthesia. Supèr et al. manipulated the saliency of the target by decreasing the size of the line segments used. The target shown in (b) is 16 pixels on a side, but they also used 8- and 4-pixel targets. For 16-pixel targets, modulation is present as shown in (c) when the target is detected and absent when the target is absent. But as the number of pixels is decreased, the difference between the case when the target is detected and not detected decreases, so long as the number of catch trials is held constant. When the pixel count is 4, there is no significant difference in modulation between detection and nondetection. Figures (courtesy of Victor Lamme) redrawn with permission from Nature Neuroscience.

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 353)

between modulation and saccades to the target, and in that sense modulation and access to the target corresponded well. But moving the saliency down or the percentage of catch trials up boosted the modulation when the animal did not saccade to the target to the 50 percent range. That is, with low saliency or a high number of catch trials, the monkey's criterion level for saccading was close enough to the visual "signal" that the modulation averaged the same whether the animal saccaded to the target or not. For example, this happened when the pixel count was reduced from 16 to 4, maintaining catch trials at 20 percent, and also when the pixel count was 16 and the catch trials were raised to 50 percent. If the pixel count was reduced to 4 but the catch-trial percentage was also reduced to zero, then the correlation between modulation and access was restored. These results show that the modulation does not reflect access to the target (since it was the same whether the target was or was not accessed). Nor does the modulation reflect the saccade, so it is on the sensory rather than motor side of the decision process. It also does not reflect attention, since the detected targets can be assumed to draw more attention. The modulation seems to reflect something intermediate between the stimulus and access. In a classic signal-detection analysis, Supèr, Spekreijse, and Lamme indeed showed that the modulation is an intermediate-level representation that can be disconnected from access either by raising the perceptual decision criterion or by decreasing saliency of the stimulus, lowering the visual "signal" to the range of the decision criterion.

The modulation shown by Supèr and colleagues disappears under anesthesia (Lamme, Zipser, and Spekreijse 1998) and is probably produced by recurrent processes (Lamme, Supèr, and Spekreijse 1998), unlike other V1 representations such as direction and orientation tuning. So there is some plausibility to taking it as an indication of if not directly part of a Phenomenal NCC for the experiential content of seeing the target. (See also Ress and Heeger 2003.)

Can the Phenomenal NCC Be Studied Empirically?

Doubts about whether phenomenal consciousness (and hence its neural basis, the Phenomenal NCC) can be studied empirically are common (see box 17.4), and are often based on the idea that ultimately, introspective reports—that is, reports about one's conscious experience—are the fundamental epistemological basis of theories of consciousness, the gold standard (Dehaene and Changeux 2004; Weiskrantz 1997; Papineau 2002, especially chap. 7). Reports are not supposed to be infallible, but any discounting of reports as reporting too much or too little will supposedly have to be based solely on *other* reports. Reports inevitably reflect the Access NCC, not just the Phenomenal NCC. When people tell you about their conscious states, you only hear about the ones that have won the winner-take-all competition. Hence we can only study "access to consciousness" (Dehaene and Changeux 2004)—that is, access to ex-

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 354)

Box 17.4
Questions for Future Research

1. In visual extinction due to right parietal damage, the subject reports not seeing a stimulus on the left when there is a competing stimulus on the right. Rees et al. (2002) showed that the fusiform face area (in the relevant hemisphere) of an extinction patient can be activated robustly when the patient says he does not see the face (because of a competing stimulus), though not quite as strongly as when the subject says he does see the face. One question is: is there *recurrent* activation of the relevant part of V1 in such a patient? A related question is: does the fusiform-face-area activation in such a patient show the enhanced figure modulation response described by Supèr et al.? If the answer to both turns out to be yes, that is evidence that recurrent fusiform-face activation is a genuine core Phenomenal NCC for face experience, even though the subject says he does not see a face.
2. If indeed recurrent activation of sensory areas creates the core Phenomenal NCCs, why? For example, why is recurrent activation of area MT/V5 (together with the unknown background activation) sufficient for visual experience of motion instead of some other experiential content or no content? That is a form of the infamous Hard Problem of consciousness (Chalmers 1996).

periential content, not experiential content itself. I do not agree with this methodological view for a number of reasons.

First, observed electrons can provide evidence about electrons that cannot in principle be observed—for example, electrons that are too distant in space and time (i.e., outside our light cone) to be observed. Why should we suppose matters are any different for consciousness?

Second, there is no gold standard of evidence, here or in any area of science. We should go for the simplest theory compatible with *all* the evidence. *No evidence is privileged.* In particular, it is not true that our theory of consciousness should be completely determined by the introspective reports of subjects. An analogy: it is trivial to program two computers to yield the same input-output function via different algorithms. No theory of what goes on in computers based wholly on the computers' "reports"—that is, input-output relations—stands a chance of success. Why should we suppose consciousness is any different? Just as two computationally different computers can have the same input-output function, two brains that are different in conscious structure might at least in principle have the same input-output function.

Third, any neuroscientific approach that bases everything on reports about a subject's own experience is in danger of focusing on the *neural basis of higher-order thought*—thought to the effect that I myself have an experience—rather than the neural basis of experiential content or even access to experiential content. To give an introspective report, the subject has to have a higher-order thought—so to insist

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 355)

on introspective reportability as the gold standard is to encourage leaving out cases in which subjects have experiences that are not adequately reflected in higher-order thoughts.[9]

Finally, even those who assimilate experiential content to its accessibility should not accept introspective reports as a gold standard. Animals have plenty of access to their experiences, but probably little in the way of higher-order thought about them of the sort that could be the basis of an introspective report. Cowey and Stoerig (1997) showed that monkeys that had been made blindsighted on one side and trained to make a visual discrimination in their sighted field, could make the discrimination in their blind field. However, when given the option, they preferred a third ''nothing'' response. This is evidence about the monkey's perceptual state that does not depend on any introspective reports.

But is the monkeys' button pushing just a *nonverbal introspective report*? Nonhuman primates that have learned symbolic systems for communication may not even make spontaneous reports about the world (Terrace 2004; Wallman 1992), so there is little ground for supposing that they are prone to reports about their own experience.[10] If a human were to push the ''nothing'' button, we might guess whether there is a thought underlying the response. We might consider two hypotheses: first, the *introspective* report, ''I am having no visual experience,'' and second, the *environmental* report, ''There is nothing on the screen.'' If the subject were a child of 3–4, the introspective report would be unlikely since children have a great deal of difficulty with states of mind about their own mental states (Esbensen, Taylor, and Stoess 1997; Gopnik and Graf 1988). Given that the environmental report would be preferable even for a child, we can hardly suppose the introspective report would be preferable in the case of a macaque! The take-home message is that you do not need reports *about the subject's experiences* to get good evidence about what the subject is experiencing: indications of what the subject takes to be in front of him or her will do just fine.

Where are we? I have proposed a distinction between a Phenomenal NCC and an Access NCC. The ''single NCC'' framework does not do as well in making sense of the empirical data, in particular signal-detection theory data as an account in which there are two NCCs, a Phenomenal NCC and an Access NCC. Of course both NCCs are to be firmly distinguished from perceptual representations that are not conscious in any sense (as in the right-most panel of figure 17.3c). More generally, rather than asking ''What is the direct evidence about the Phenomenal NCC independently of the Access NCC?'', we should instead ask ''What framework makes the most sense of the data?''

Notes

This is a longer version of a paper in *Trends in Cognitive Sciences* 9(2): 46–52, February 2005.

1. The first classical ''visual'' cortical area is V1; later classic ''visual'' areas include V2, V3, V4, and V5. V5 has two names because it was identified and named by two groups. I put ''visual'' in scare

quotes because there is some debate as to whether some of the classic "visual" areas are best thought of as multimodal and spatial. In the United States, the area I am talking about is usually called "MT+" because it includes structures that adjoin MT.

2. TMS delivers an electromagnetic jolt to brain areas when placed appropriately on the scalp. The effect is to disrupt organized signals but also to create a signal in a quiescent area. Thus TMS can both disrupt moving afterimages and create phosphenes. A comparison is to hitting a radio: the static caused might interrupt good reception going on but also cause a noise when there is no reception. (I am indebted here to Nancy Kanwisher and Vincent Walsh.)

3. To experience phosphenes for yourself, close your eyes and exert pressure on your eye from the side with your finger. Or if you prefer not to put your eyeball at risk, look at the following website for an artist's rendition: http://www.reflectingskin.net/phosphenes.html.

4. The distinction between core and total NCC as I defined it depends on the assumption that at least some core NCCs share background conditions. Suppose the background condition for the experience as of red and the experience as of green are the same and are the same as other visual experiences, but not the same as the background condition for taste experiences—for instance, the experience as of saltiness. Then the core NCC for visual experiences will have to be defined as the part of the total visual NCC that distinguishes one visual content from another.

5. The "made-available" terminology is supposed to capture both the occurrent nature of the experience (when something is *made* available, something happens) and the dispositional aspect (availability). There are many somewhat different ways of making access consciousness precise in this picture. One might think of the crucial feature as representations being *sent*, or else *received*, or else *translated* from the system of representation of the producing systems to the system of representation of the consuming systems.

6. The idea is not that the auditory signals from a voice compete with the visual signals from the person's mouth moving, but rather that a "coalition" that involves neural processing of both of those signals competes with other coalitions.

7. Milner and Goodale (1995) distinguish between a conscious visual pathway from the classic visual areas in the back of the head feeding into the temporal lobe on the side of the head (ventral stream) and an unconscious "dorsal" action-oriented stream starting in the back of the head and feeding to the top of the head.

8. More specifically, Burge argues that there is a kind of primitive of-ness of a phenomenally conscious state that is not reducible to higher-order thought (and not reducible to any other cognitive notion). In Block 1995, I argue that "phenomenal consciousness" in my sense of the term can be either transitive (take an object of which the subject is conscious) or intransitive. My intransitive phenomenal consciousness corresponds to Burge's phenomenality, and my transitive phenomenal consciousness corresponds to Burge's phenomenal consciousness.

9. Armstrong, Carruthers, Lycan, and Rosenthal have argued for seeing consciousness in terms of higher-order thought. In some versions of this view—for example, Rosenthal's—experiential content can exist without higher-order thought. Anyone who takes such a view should agree with me

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 357)

that a methodology focused exclusively on introspective report alone will be in danger of finding the neural basis of higher-order thought rather than the neural basis of experiential content. The difference between the Rosenthal-type view and mine has to do in part with the issue of whether the term "consciousness" refers to a higher-order state or to a first-order state. (My view is that the term "consciousness" is ambiguous and in one sense refers to a higher-order state and in another sense a first-order state.) But that difference about how the term is used is itself dependent (I believe) on a difference of opinion on whether the "Hard Problem" applies to experiential content. For someone who does not believe in the Hard Problem for experiential content, a higher-order thought about such contentful states may seem a more worthy bearer of the term "consciousness."

10. There have been many claims of reports by nonhuman primates—for example, by Savage-Rumbaugh—but it is controversial whether those claims are based on trained-up responses given in the expectation of reward.

References

Baars, B. J. 1997. *In the Theater of Consciousness: The Workspace of the Mind*. New York: Oxford University Press.

Beck, D., Rees, G., Frith, C. D., and Lavie, N. 2001. Neural correlates of change detection and change blindness. *Nature Neuroscience* 4: 645–650.

Block, N. 1990. Consciousness and accessibility. *Behavioral and Brain Sciences* 13: 596–598.

Block, N. 1992. Begging the question against phenomenal consciousness. *Behavioral and Brain Sciences* 15: 205–206. (Reprinted in N. Block, O. Flanagan, and G. Güzeldere, eds., *The Nature of Consciousness: Philosophical Debates*. Cambridge, MA: MIT Press, 1997.) 175–179.

Block, N. 1995. On a confusion about a function of consciousness. *Behavioral and Brain Sciences* 18(2): 227–247. (Reprinted in N. Block, O. Flanagan, and G. Güzeldere, eds., *The Nature of Consciousness: Philosophical Debates*. Cambridge, MA: MIT Press, 1997.)

Block, N. 1997. Biology versus computation in the study of consciousness. *Behavioral and Brain Sciences* 20(1): 159–165. http://www.nyu.edu/gsas/dept/philo/faculty/block/papers/Reply1997.pdf.

Block, N. 2001. Paradox and cross purposes in recent findings about consciousness. *Cognition* 79(1–2): 197–219.

Britten, K., Shadlen, M., Newsome, W., and Movshon, A. 1992. The analysis of visual motion: A comparison of neuronal and psychophysical performance. *Journal of Neuroscience* 12: 4745–4765.

Burge, Tyler. 1997. Two kinds of consciousness. In N. Block, O. Flanagan, and G. Güzeldere, eds., *The Nature of Consciousness: Philosophical Debates*. Cambridge, MA: MIT Press.

Chalmers, D. 1996. *The Conscious Mind*. Oxford: Oxford University Press.

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 358)

Chalmers, D. 1997. Availability: The cognitive basis of experience. *Behavioral and Brain Sciences* 20(1): 148–149. (Reprinted in N. Block, O. Flanagan, and G. Güzeldere, eds., *The Nature of Consciousness: Philosophical Debates*. Cambridge, MA: MIT Press, 1997.)

Chalmers, D. 2002. What is a neural correlate of consciousness? In T. Metzinger, ed., *Neural Correlates of Consciousness: Empirical and Conceptual Questions*. Cambridge, MA: MIT Press.

Cowey, A., and Stoerig, P. 1997, July. Visual detection in monkeys with blindsight. *Neuropsychologia* 35(7): 929–939.

Cowey, A., and Walsh, V. 2000. Magnetically induced phosphenes in sighted, blind and blindsighted subjects. *NeuroReport* 11: 3269.

Debner, J. A., and Jacoby, L. L. 1994. Unconscious perception: Attention, awareness, and control. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20: 304–317.

Dehaene, S., and Changeux, J.-P. 2004. Neural mechanisms for access to consciousness. In M. Gazzaniga, ed., *The Cognitive Neurosciences*, vol. 3. Cambridge, MA: MIT Press.

Dehaene, S., Sergent, C., and Changeux, J.-P. 2003. A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences* 100(14): 8520–8525.

Dennett, D. 1995. The path not taken. *Behavioral and Brain Sciences* 18(2): 1995: 252–253. (Reprinted in N. Block, O. Flanagan, and G. Güzeldere, eds., *The Nature of Consciousness: Philosophical Debates*. Cambridge, MA: MIT Press, 1997.)

Esbensen, B. M., Taylor, M., and Stoess, C. J. 1997. Children's behavioral understanding of knowledge acquisition. *Cognitive Development* 12: 53–84.

Gopnik, A., and Graf, P. 1988. Knowing how you know: Children's understanding of the sources of their knowledge. *Child Development* 59: 1366–1371.

Heeger, D., Boynton, G., Demb, J., Seideman, E., and Newsome, W. 1999. Motion opponency in visual cortex. *Journal of Neuroscience* 19: 7162–7174.

Huk, A., Ress, D., and Heeger, D. 2001. Neuronal basis of hte motion aftereffect reconsidered. *Neuron* 32: 161–172.

Kammer, T. 1999. Phosphenes and transient scotomas induced by magnetic stimulation of the occipital lobe: Their topographic relationship. *Neuropsychologia* 37: 191–198.

Kanwisher, N. 2001. Neural events and perceptual awareness. *Cognition* 79(1–2): 89–113.

Kleinschmidt, A., Buchel, C., Zeki, S., and Frackowiak, R. S. J. 1998. Human brain activity during spontaneously reversing perception of ambiguous figures. *Proc. R. Soc. Lond. B* 265: 2427–2433.

Koch, C. 2004. *The Quest for Consciousness: A Neuroscientific Approach*. Granwood Village, Co. Roberts and Company.

(AutoPDF V7 9/1/07 10:37) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 359)

Kourtzi, Z., and Kanwisher, N. 2000. Activation in human MT/MST by static images with implied motion. *Journal of Cognitive Neuroscience* 12(1): 48–55.

Lamme, V. 2003. Why visual attention and awareness are different. *Trends in Cognitive Science* 7: 12–18.

Lamme, V. 2004. Separate neural definitions of visual consciousness and visual attention: A case for phenomenal awareness. *Neural Networks* 17: 861–872.

Lamme, V., and Roelfsema, P. 2000. The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neuroscience* 23(11): 571–579.

Lamme, V. A. F., Supèr, H., and Spekreijse, H. 1998. Feedforward, horizontal, and feedback processing in the visual cortex. *Current Opinion in Neurobiology* 8: 529–535.

Lamme, V. A. F., Zipser, K., and Spekreijse, H. 1998. Figure-ground activity in primary visual cortex is suppressed by anaesthesia. *Proceedings of the National Academy of Sciences of the USA* 95: 3263–3268.

Lumer, E. D., Friston, K. J., and Rees, G. 1998. Neural correlates of perceptual rivalry in the human brain. *Science* 280: 1930–1934.

Lumer, E. D., and Rees, G. E. 1999. Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proc. Natl Acad. Sci. USA* 96: 1669–1673.

Milner, A. D., and Goodale, M. A. 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.

Nakamura, R., and Mishkin, M. 1980. Blindness in monkeys following non-visual cortical lesions. *Brain Research* 188: 572–577.

Nakamura, R., and Mishkin, M. 1986. Chronic "blindness" following lesions of nonvisual cortex in the monkey. *Experimental Brain Research* 63: 173–184.

Papineau, D. 2002. *Thinking about Consciousness*. Oxford: Oxford University Press.

Pascual-Leone, A., and Walsh, V. 2002, April. Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science* 292: 510–512.

Pins, D., and ffytche, D. 2003. The neural correlates of conscious vision. *Cerebral Cortex* 13: 461–474.

Portas, C. M., Strange, B. A., Friston, K. J., Dolan, R. J., and Frith, C. D. 2000. How does the brain sustain a visual percept? *Proc. R. Soc. Lond. B* 267: 845–850.

Rees, G., Kreiman, G., and Koch, C. 2002, April. Neural correlates of consciousness in humans. *Nature Reviews Neuroscience* 3(4): 261–270.

Rees, G., Wojciulik, E., Clarke, K., Husain, M., Frith, C. D., and Driver, J. 2002. Neural correlates of conscious and unconscious vision in parietal extinction. *Neurocase* 8: 387–393.

(AutoPDF V7 9/1/07 10:38) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 360)

Ress, D., and Heeger, D. 2003. Neuronal correlates of perception in early visual cortex. *Nature Neuroscience* 6: 4414–4420.

Rousselet, G., Thorpe, S., and Fabre-Thorpe, M. 2004. How parallel is visual processing in the ventral pathway? *Trends in Cognitive Sciences* 8(8): 363–370.

Shoemaker, S. 1981. Some varieties of functionalism. *Philosophical Topics* 12: 93–119. (Reprinted in S. Shoemaker, *Identity, Cause, and Mind*. Cambridge: Cambridge University Press, 1984.)

Sincich, L., Park, K. F., Wohlgemuth, M. J., and Horton, J. C. 2004. Bypassing V1: A direct geniculate input to area MT. *Nature Neuroscience* 7(10): 1123–1128.

Snodgrass, M. 2002. Disambiguating conscious and unconscious influences: Do exclusion paradigms demonstrate unconscious perception? *American Journal of Psychology* 115: 545–580.

Supèr, H., Spekreijse, H., and Lamme, V. 2001. Two distinct modes of sensory processing observed in monkey primary visual cortex (V1). *Nature Neuroscience* 4(3): 304–310.

Terrace, H. 2004. Metacognition and the Evolution of Language. In H. Terrace and J. Metcalfe, eds., *The Missing Link in Cognition: Origins of Self-Knowing Consciousness*. New York: Oxford University Press.

Théoret, H., Kobayashi, M., Ganis, G., Di Capua, P., and Pascual-Leone, A. 2002. Repetitive transcranial magnetic stimulation of human area MT/V5 disrupts perception and storage of the motion aftereffect. *Neuropsychologia* 40(13): 2280–2287.

Visser, T., and Merikle, P. 1999. Conscious and unconscious processes: The effects of motivation. *Consciousness and Cognition* 8: 94–113.

Wallman, J. 1992. *Aping Language*. Cambridge: Cambridge University Press.

Weiskrantz, L. 1997. *Consciousness Lost and Found*. Oxford: Oxford University Press.

Zeki, S. 2001. Localization and globalization in conscious vision. *Annual Reviews of Neuroscience* 24: 57–86.

Zeki, S., and ffytche, D. H. 1998. The Riddoch Syndrome: Insights into the neurobiology of conscious vision. *Brain* 121: 25–45.

Zihl, J., von Cramon, D., and Mai, N. 1983. Selective disturbance of movement vision after bilateral brain damage. *Brain* 106: 313–340.

(AutoPDF V7 9/1/07 10:38) MIT (Stone 7×9") StoneSerif&Sans J-1567 Block AC1: WSL 29/12/2006 pp. 343–362 1567_17 (p. 361)