4

Imaging Conscious and Subliminal Word Processing

Stanislas Dehaene

You have just begun to read this chapter. Your eyes are currently scanning this sentence at a rate of three or four saccades per second. Almost unfailingly, in a fraction of a second, your visual system manages to transform those black marks on paper into recognized letters, then words, then meanings. Yet you are unaware of the complexity of those mental operations: your consciousness is entirely focused on their end result, the meaning of the text. The long chain of unconscious computations that precedes conscious access is, by definition, inaccessible to introspection and, in the early days of scientific psychology, was thought to be inaccessible to experimentation. If anything, however, the conscious step now seems even more mysterious—it is hard to even imagine what magical wand could turn mechanical brain computations into consciously accessible states of mind.

In the past decades, we have learned a lot about the architecture of the word recognition system. The methods that have permitted a dissection of word recognition processes are attributable, to a large extent, to the work of Michael Posner. In the first half of his career, he and others such as Shepard and Sternberg made seminal contributions to modern mental chronometry (Posner, 1978). Later, Posner went on to explore the cerebral underpinnings of the mental representations that mental chronometry had inferred, using (and helping develop) the new tool of cognitive neuroimaging. He quickly saw the revolutionary potential of this new method: "The microscope and telescope opened vast domains of unexpected scientific discovery. Now that new imaging methods can visualize the brain systems used for normal and pathological thought, a similar opportunity may be available for human cognition" (Posner, 1993; p. 673).

Parts of this article have also been used in chapters by Dehaene and Changeux (2004) and Dehaene (2003). I thank Jean-Pierre Changeux, Lionel Naccache, and Claire Sergent for their essential contributions, Nancy Kanwisher and Ulrich Mayr for their useful feedback, and the Institut National de la Santé et de las Recherche Médicale, Commissariat à l'Energie Médicale, and McDonnell Foundation for their financial support.

Reading the first cognitive neuroimaging paper (Petersen, Fox, Posner, Mintun, & Raichle, 1988) was, for me, a highly memorable moment. Two key advantages of brain imaging were immediately apparent. First, imaging provided a new method of functional dissection of the hierarchy of cognitive operations (Posner, Petersen, Fox, & Raichle, 1988). Second, even the highest stages of processing that often eluded chronometric analysis, such as the complex semantic manipulations inherent in the verb generation task, were analyzable and associated with reproducible activation sites. In my opinion, this possibility of imaging the cerebral underpinnings of any mental activity, without requiring the research participant to perform an overt response or even an actual task, was a powerful factor in the resurgence of interest for higher cognition and consciousness, since it permitted the visualization of the substrates of covert attentive states (Pardo, Fox, & Raichle, 1991; Posner & Petersen, 1990) and revealed the striking contribution of previously neglected areas such as the anterior cingulate (Pardo, Pardo, Janer, & Raichle, 1990).

Both word recognition and consciousness have become huge fields of research that cannot possibly be reviewed in a short chapter. My purpose is more limited: to revisit the organization of the visual word recognition system, focusing on a single issue that my colleagues and I have addressed in several experiments: what patterns of activity characterize nonconscious and conscious words? Using pattern masking, words can be presented on the retina for several tens of milliseconds, and yet fail to be consciously seen by the participant. This raises two obvious questions. First, to what extent are such subliminal words processed in the brain? Second, what cerebral events distinguish subliminal words from words that are consciously reportable? Focusing on those well-defined questions narrows down the problem of consciousness to a more tractable level.

This chapter is divided in three parts. First, I review the empirical data on the neuroimaging of masked words. I then describe a putative integrative theory based on the hypothesis of a "conscious neuronal workspace" (Dehaene & Changeux, 2000; Dehaene, Kerszberg, & Changeux, 1998; Dehaene & Naccache, 2001; Dehaene, Sergent, & Changeux, 2003). This model emphasizes the role of distributed neurons with long-distance connections, particularly dense in prefrontal, cingulate, and parietal regions. These neurons are capable of interconnecting multiple specialized processors and can broadcast signals at the brain scale in a spontaneous and sudden manner, thus creating a global availability that, according to our hypothesis, is experienced as consciousness (see also Baars, 1989). In the third part, I examine recent experiments that tested some predictions of the workspace model.

Neuroimaging of Subliminal Word Processing

How can one experimentally distinguish the substrates of conscious access from the considerable amount of neural activity that occurs nonconsciously? In laboratory tests, our approach has consisted in exploring paradigms in which symbolic information is deliberately presented under subliminal conditions. By studying to what extent such information is processed, and what brain areas it contacts, one can progressively draw a negative picture, as it were, of which aspects of brain activity do not suffice to give rise to consciousness. Subliminal priming has also proven an important tool in decoding the stages of representation of words in various brain areas.

General Logic of Subliminal Priming Studies

In a typical subliminal priming experiment, each trial consists in the consecutive presentation, at the same screen location, of a random configuration of letters or geometrical shapes (premask), a first word (the prime), another random configuration of symbols or shapes (postmask), and a second word (the target). The prime word is presented briefly, typically 10 to 50 ms. The target is presented for a much longer duration, typically 500 ms. Finally, the stimulus onset asynchrony between the prime and target is short (typically 60 to 120 ms). Under those conditions, research participants report seeing only the masks and the target word, but not the prime.

This subjective invisibility can be confirmed by asking participants to perform an explicit task on the primes. In recent experiments, my colleagues and I have used a forced-choice identification task in which participants have to select, among two alternative words, the one that matches the prime. With a 29 ms prime presentation duration, performance typically does not differ from chance (Dehaene et al., 2001). With slightly longer primes, slightly better performance may be observed, but it typically does not correlate with the amount of subliminal priming found (Greenwald, Draine, & Abrams, 1996; Vorberg, Mattler, Heinecke, Schmidt, & Schwarzbach, 2003). Other tasks such as prime presence–absence judgment (Dehaene, Naccache, et al., 1998) or prime categorization (Naccache & Dehaene, 2001b) yield similar results.

In spite of their subjective invisibility, the prime words impact on the processing of target words. This can be demonstrated by varying the relation between the prime and target. The simplest design compares trials in which the same word is presented twice, as both prime and target, with trials in which different words are presented as prime and target. In this condition of repetition priming, response times to the target are consistently shorter on repeated than on nonrepeated trials (e.g., Forster & Davis, 1984). By systematically varying the physical, phonological, or even semantic proximity between the prime and target, it is then possible to probe which levels of word processing underlie this facilitation effect.

To identify the brain systems activated by masked words, functional imaging can be combined with masked priming (Dehaene, Jobert, et al., in press; Dehaene, Naccache, et al., 1998; Dehaene, Naccache, et al., 2001; Naccache & Dehaene, 2001a). Functional magnetic resonance imaging (fMRI) is currently too slow to separate the cerebral activity induced by the prime and by the target. Thus, one necessarily measures the total activity induced by the prime-target pair (relative, say, to a control situation in which only the masks are presented). In spite of this limitation, one can still acquire knowledge of the processing of the prime by varying the type of relation between the prime and target. When the prime and the target are the same word, there is a measurable reduction in brain activation compared to a situation in which the prime and the target are different words. Measuring where this subliminal repetition suppression effect occurs provides an indirect image of the brain areas that have been traversed by the hidden prime word. This can be supplemented by recordings of event-related potentials, which have an appropriate temporal resolution to follow the dynamics of prime- and target-induced activations.

Subliminal Priming in the Visual Word Form System

Dehaene and colleagues (2001) first examined the cerebral bases of subliminal word repetition priming. FMRI data were collected in a fast event-related paradigm while research participants performed a bimanual semantic classification task on visual words. Unbeknownst to them, each target word was preceded by a subliminal prime. My colleagues and I used a 2 × 2 design in which the prime and target could be the same word or different words and could appear in the same or different case, thus defining four types of events. The amount of activation in each of those events was identified relative to a fifth event in which only the masks were presented and no response was required (see Figure 4.1).

The behavioral results showed that response times were faster on repeated trials, whether or not the words shared the same case. In searching for the cerebral bases of this effect, two distinct types of brain regions were identified. The right extrastriate occipital cortex showed repetition suppression only for physically identical primes and targets, suggesting a role for right visual areas in coding the precise visual features of the letters (Marsolek, Kosslyn, & Squire, 1992). The left fusiform gyrus, however, showed repetition suppression whenever the same word was repeated, whether in the same case or not. Thus, this region appears to encode the word string in a case-independent fashion.

This region may provide the cerebral substrate of the visual word form system (Posner & Carr, 1992; Warrington & Shallice, 1980), a structural representation of visual words as an ordered sequence of abstract letter identities or multiletter graphemes, invariant for size, font, and case (Cohen et al., 2002; Dehaene, Le Clec'H, Poline, Le Bihan, & Cohen, 2002; Molko et al., 2002). It is roughly symmetrical to the right-hemispheric fusiform face area, and may play for visual word recognition the same role that similar or neighboring regions of the fusiform and lingual gyri play for other visual objects such as faces, objects, or places. In adults, it has become partially attuned to a specific script, as shown by its greater response to real words than to consonant strings of similar arbitrary shape (Cohen et al., 2002). Indeed, in children this area activates in direct proportion to the child's reading skills, and its response is absent in dyslexic readers who have not developed expertise in word recognition (Paulesu et al., 2001; Shaywitz et al., 2002).

To further specify the exact nature of the word representation attained by subliminal primes in the visual word form area. I recently performed two more repetition priming experiments (Dehaene, Jobert, et al., 2004). The first examined whether visual features alone could explain priming in this region.

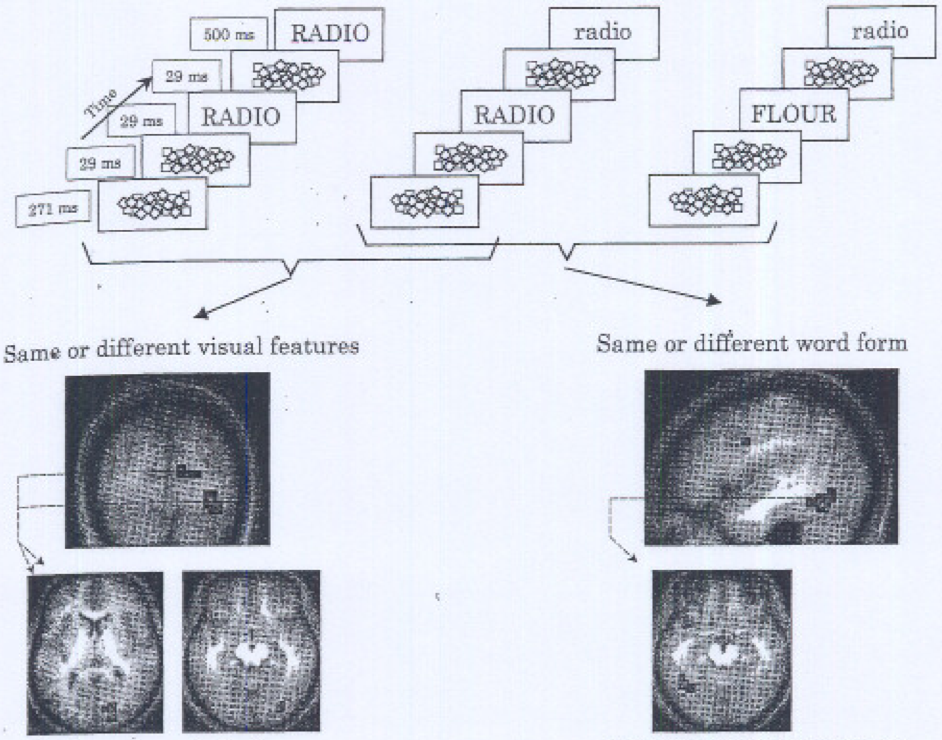

Feature-based priming
in occipital extrastriate cortex

Case-independent priming
in left fusiform gyrus

Figure 4.1. Imaging of subliminal priming in the visual word-recognition system. Repetition suppression is observed in occipital extrastriate cortex, particularly in the right hemisphere, when a word is physically repeated; and in the left fusiform gyrus (visual word form area) when a word is repeated with or without case change. Data from Dehaene et al. (2001).

To this end, the visual similarity of upper and lowercase letters was manipulated. Half of the prime-target pairs were words made of letters that are highly similar in upper and lower case (e.g., Oo, Pp). The other half used only highly dissimilar letters (e.g., Aa, Gg) for which the uppercase–lowercase association is essentially arbitrary. The same strip of left fusiform cortex as in Experiment 1 showed replicable subliminal repetition suppression that was present even when the letters were visually dissimilar (the same effect was also observed in response times, replicating earlier results, e.g., Bowers, Vigliocco, & Haan, 1998). This confirmed that this region is not solely concerned with visual shapes but encodes letter strings using a culturally acquired abstract letter code.

A second experiment examined whether single letters or larger units such as graphemes or whole words are encoded in the visual word form area. To repeat letters without repeating words, anagrams were used. For instance, by priming the French target word "REFLET" with the prime "trefle," almost all of the middle letters (r, e, f, l, e) could be repeated. By moving the prime relative to the target (e.g., "trefle" followed by "REFLET"), it was even possible to repeat those letters at the same retinal location without repeating the same word. By comparing this to a word-repeated trial, with or without a shift in letter position, the nature and position invariance of the neural codes underlying priming could be tested. Would priming depend on letter repetition, word repetition, or both? The results revealed an interesting dissociation between posterior and anterior areas. The posterior portion of the visual word form area ($y = -68$ in the Talairach coordinate system) showed repetition suppression only when the same letters were repeated at the same retinal location, regardless of their case. This region thus holds a case-invariant but position-selective letter code. At a more anterior position ($y = -56$), location-independent priming was found for both repeated words and anagrams compared to a control, nonrepeated condition. Thus, this region encodes a case- and position-invariant representation of visual units that are smaller that the whole word. Finally, in a still more anterior fusiform region ($y = -48$), priming became greater for repeated words than for anagrams, thus revealing a case- and position-independent whole-word code, or at least a code sensitive to the larger graphemic units that distinguish a word from its anagram. Behavioral response times were sensitive only to this whole-word code.

Two conclusions may be drawn from those studies at the visual word form level. First, behavioral priming effects provide only a coarse indication of the various levels of representations traversed by a subliminal prime. Functional imaging reveals a much richer variety of priming effects, ranging from feature-based to letter or whole-word based priming. Furthermore, this hierarchical organization is remarkably reproducible across subjects, even for a culturally acquired activity such as reading. I am reminded once again of the words of Posner: "It is a popularly held belief in psychology that the cognitive functions of the brain are widely distributed among different brain areas. Nevertheless, imaging studies reveal a startling degree of region-specific activity. When thought is analyzed in terms of component mental operations, a beautiful localization emerges" (Posner, 1993, p. 673).

A second conclusion is that subliminal primes can be processed quite far along the ventral visual identification pathway. Although several studies have identified a tight correlation between ventral fusiform activity and the contents of visual consciousness (e.g., Bar et al., 2001; Grill-Spector, Kushnir, Hendler, & Malach, 2000), fusiform activity is not sufficient for conscious reportability.

Subliminal Semantic Access: The Case of Numbers

Can subliminal words be processed beyond the visual system? The issue of semantic access from masked words remains controversial in psychology. In laboratory tests, we obtained positive evidence for subliminal semantic access by using a small set of high-frequency words with simple semantics: number words.

Initially, evidence for subliminal semantic access came from the observation of motor priming (Dehaene, Naccache, et al., 1998). We engaged research

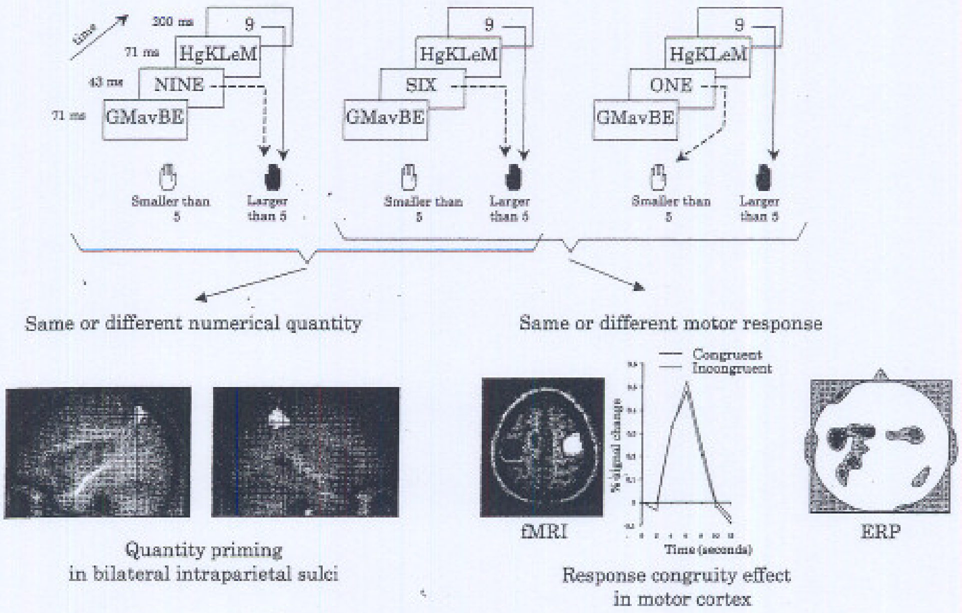

Quantity priming
in bilateral intraparietal sulci

Response congruity effect
in motor cortex

Figure 4.2. Imaging of subliminal semantic and motor priming during number comparison. Repetition suppression is observed in the left and right intraparietal sulci when the same numerical quantity is repeated (regardless of number notation). Motor priming is observed in the left and right motor cortices as a function of the response induced by the subliminal primes. Data from Dehaene, Naccache, et al. (1998); Naccache and Dehaene (2001a). ERP = event related potential; fMRI = functional magnetic resonance imaging.

participants in a number comparison task, where they classified target numbers as larger or smaller than five using their left or right hand. Unbeknowst to them, subliminal numerical primes were presented just before the target, and these primes could also be larger or small than five. Thus, the size of the primes could be congruent or incongruent with the subsequent target. This response congruity factor was found to interfere with participants' response times. Furthermore, it yielded a response conflict in motor cortex that was measurable in both event related potentials (ERPs) and fMRI (see Figure 4.2; Dehaene, Naccache, et al., 1998).

That participants could activate the motor cortex of the hand that would have been appropriate for responding to the prime seemed to imply that participants had unconsciously categorized the prime as larger or smaller than five. Thus, motor priming was initially taken as a clear, though indirect, proof of semantic processing of subliminal numbers (Dehaene, Naccache, et al., 1998). However, there was an alternative interpretation: the observed motor activation could be due to direct motor specification (Neumann & Klotz, 1994). Because a small number of stimuli—the digits 1, 4, 6, and 9 and the corresponding words—were repeatedly used as both primes and targets, participants could have learned to associate each visual stimulus with the corresponding response, thus bypassing semantic access. Indeed, such a shortcut was recently demonstrated by Abrams and Greenwald (2000). In an affective categorization task, new primes that were made of fragments of previous seen targets yielded subliminal motor priming solely based on visual fragments, not on whole-word meaning. For instance the prime word SMILE, created from the targets words SMUT and BILE, ended up paradoxically priming the negative rather than the positive response. Thus, although the task required semantic categorization, and although a priming effect was observed, the primes only received a shallow, nonsemantic analysis of their component letters and the associated responses (see also Damian, 2001).

New experiments and reanalyses have now demonstrated that the number priming results do not easily fall prey to a similar nonsemantic interpretation. First, other analyses (Naccache & Dehaene, 2001b) demonstrated that the motor priming effect was present in the first block of trials and that neither motor priming nor the classical semantic distance effect were affected by practice. Those findings are inconsistent with the idea that the task is increasingly being performed via a nonsemantic route. Second, we replicated our original behavioral experiments with novel numbers that are only presented as primes, never as targets (Naccache & Dehaene, 2001b). Because those numbers are never seen consciously and are never responded to, they cannot be associated with motor responses. Yet in two different experiments, those novel primes were found to cause significant motor priming, indicating that at least part of the motor priming effect arises from a genuinely semantic route. This positive effect of novel numerical primes has now been replicated and extended by others (e.g., Greenwald, Abrams, Naccache, & Dehaene, 2003; Kunde, Kiesel, & Hoffmann, 2003; Reynvoet, Caessens, & Brysbaert, 2002). Importantly, both Greenwald and colleagues (2003) and Kunde and colleagues (2003) have found that it depends partly on the instructions and stimulus content. Research participants may perform using direct stimulus-response associations if the instructions emphasize fast responding to a small number of stimuli (Kunde et al., 2003); or they may learn to respond only to the individual component digit of two-digit numerals, if those digits are entirely predictive of the motor response (Greenwald et al., 2003). If the stimuli and instructions force participants to perform a genuine number comparison, however, subliminal semantic priming that generalizes to novel primes is found.

Finally, we obtained more direct evidence for semantic-level repetition priming. Independently of motor priming, when participants were engaged in the number comparison task, their responses were also accelerated when the prime and the target represented the same quantity, possibly in different notations (e.g., prime NINE, target 9), relative to congruent but nonrepeated trials (e.g., prime SIX, target 9; Dehaene, Naccache, et al., 1998; Koechlin, Naccache, Block, & Dehaene, 1999). Furthermore, this form of priming, which we termed quantity priming, depended on the numerical distance between the prime and target, clearly suggesting semantic mediation (Koechlin et al., 1999; Naccache & Dehaene, 2001b; Reynvoet, Brysbaert, & Fias, 2002; Reynvoet, Caessens, et al., 2002). A neuroimaging version of this paradigm showed that this notation-independent quantity priming in a bilateral intraparietal region was thought to be involved in the semantic representation and manipulation of numerical quantities (Naccache & Dehaene, 2001a).

In summary, priming effects with subliminal numbers have been observed at both the semantic (intraparietal) and motor levels. This provides a clear indication that semantic-level processing of masked primes is possible. It should be noted that digits are some of the most frequent visual symbols and are semantically unambiguous. The ease and speed of visual-to-semantic transduction may explain why it seems easier to obtain semantic priming with numbers than with other types of words (Abrams & Greenwald, 2000; Damian, 2001).

Differences Between Unconscious and Conscious Word Processing

The masking paradigm also provides a way to address the second key issue associated with subliminal processing: What are the changes in brain activity that distinguish subliminal and conscious situations? A simple experiment consists in measuring the fMRI or ERP correlates of brain activity evoked by words in minimally different masked and unmasked situations (Dehaene et al., 2001). In this experiment, participants viewed a constant stream of visual shapes. Occasionally, a readable word would appear. Unbeknownst to the participants, however, there were other moments in which subliminal words were presented. The difference between the visible and invisible conditions was minimal. It merely involved shuffling the order of the mask and blank screen that preceded and followed the word presentation. In the masked condition, the order was blank-mask-word-mask-blank, and the 29-ms word was thus made invisible by the masks; in the visible condition, the order was mask-blank-word-blank-mask, and the word was unmasked. Two types of control trials were also included in which the same temporal context sequences were presented, but a blank appeared instead of the word. Contrasting word-present and word-absent trials subtracted away the effect of the masking context, even with temporally sensitive methods such as ERP recordings, and allowed us to image only the activations induced by a visible or invisible word. Note that, contrary to the above masking experiments, this method looks at the brain activation caused by masked words directly, rather than indirectly through their priming effects on subsequent words.

Using this design, my colleagues and I found that the masked words caused a small transient bottom-up activation in extrastriate cortex, fusiform gyrus, and precentral cortex. In particular, the presence of a subliminal activation in the left visual word form area was replicated. This activation was small, however, and decreased with distance from the primary visual cortex. When the words were unmasked, however, activation greatly increased in the same areas. Furthermore, it also extended to other distant parietal, inferior prefrontal and midline precentral/cingulate areas. Functional correlation studies revealed that unmasking also enhanced the long-distance correlation between the left fusiform gyrus and several of those anterior sites. Finally, ERP recordings in the same paradigm demonstrated that the P1 and N1 waves were present for the masked words, though attenuated and more focal. Conscious processing was associated with a drastic enhancement of those components as well as with the emergence of a N400 and late positive complex (P300), two waves that were absent or greatly reduced in the masked situation.

Interpretation of those findings should be cautious for several reasons. First, masking itself may have a rather peripheral effect, perhaps even at the level of the retina. Hence, some of the differences between visible and invisible words, such as the reduction in early visual activation, may be related to peculiarities of the masking paradigm rather than to consciousness itself. Indeed, data from the inattentional blink suggest that lack of consciousness may occur without any reduction in the P1, N1, and even the N400 components of event-related potentials (Vogel, Luck, & Shapiro, 1998). Second, research participants were engaged in a mental naming task that, of course, could only be done on the visible words. Hence, it is not known which of the observed differences between visible and invisible words related to the task rather than to consciousness per se (if the latter notion is even valid—by definition, consciousness always affords access to a greater variety of processes than does subliminal processing).

With those limitations in mind, this study nevertheless yielded several cerebral correlates of consciousness, all of which have also been observed in several other studies (see, e.g., Rees, Kreiman, & Koch, 2002). Those include amplification of perceptual activity, presence of prefrontal and cingulate activation, long-distance correlations with prefrontal cortex and cingulate, and presence of a P300 waveform.

The Neuronal Workspace Hypothesis

The theory of the global neuronal workspace (Dehaene, Kerszberg, et al., 1998; Dehaene & Naccache, 2001; Dehaene, Sergent, et al., 2003) was proposed to tentatively explain why these various characteristics tend to frequently cohere together as correlates of conscious reportability. This theory is historically rooted in a long neuropsychological tradition, dating back to Hughlings Jackson and perpetuated among others by Baddeley, Shallice, Mesulam, and Posner, which emphasizes the hierarchical organization of the brain and separates lower automatized systems from increasingly higher and more autonomous supervisory executive systems. It also builds on Fodor's distinction between the vertical modular faculties and a distinct isotropic central and horizontal system that is capable of sharing information across modules. Finally, it relates to Baars's cognitive theory of consciousness, which distinguishes a vast array of unconscious specialized processors running in parallel, and a single limited-capacity serial workspace that allows them to exchange information (Baars, 1989).

Baars, however, did not specify how the psychological construct of a conscious workspace could be implemented in terms of neuronal networks. By contrast, Changeux and I have had a long-standing collaboration on the development of computational neural network models that aimed at specifying the contribution of prefrontal cortex to increasingly higher cognitive tasks (Dehaene & Changeux, 1989, 1991, 1997; Dehaene, Kerszberg, et al., 1998).

We successively considered how a network could retain an active memory across a long delay (Dehaene & Changeux, 1989); how it could encode abstract rules that could be selected from external or internal rewards (Dehaene & Changeux, 1991); and finally how networks based on those principles could pass complex planning tasks such as the Tower of London test or the Stroop test (Dehaene & Changeux, 1997; Dehaene, Kerszberg, et al., 1998). The neuronal workspace model is the last development of the neuronal architectures that we proposed to address those specific problems.

Two Computational Spaces in the Brain

The neuronal workspace hypothesis distinguishes two computational spaces in the brain, each characterized by a distinct pattern of connectivity.

1. The network of processors. Subcortical networks and most of the cortex can be viewed as a collection of specialized processors, each attuned to the processing of a particular type of information. Processors vary widely in complexity, from the elementary line segment detectors in area V1 or the motion processors in area MT, to the visual word form processor in the human fusiform gyrus, or the mirror–neuron system in area F5. In spite of this diversity, processors share characteristics of specialization, automaticity, and fast feedforward processing. Their function is made possible by a limited number of local or medium-range connections that bring to each processor the encapsulated inputs necessary to its function.
2. The global neuronal workspace. We postulate the existence of a distinct set of cortical workspace neurons characterized by their ability to send and receive projections to many distant areas through long-range excitatory axons. These neurons therefore no longer obey a principle of local, encapsulated connectivity, but rather break the modularity of the cortex by allowing many different processors to exchange information in a global and flexible manner. Information that is encoded in workspace neurons can be quickly made available to many brain systems, in particular the motor and speech-production processors for overt behavioral report. We hypothesize that the entry of inputs into this global workspace constitutes the neural basis of access to consciousness (see Figure 4.3).

Top-Down Amplification and Dynamic Mobilization

Among the long-distance connections established by workspace neurons, top-down connections play an essential role in the temporary mobilization of a given content into consciousness. Top-down attentional amplification is the mechanism by which modular processors can be temporarily mobilized and made available to the global workspace, and therefore enter into consciousness. According to this view, the same brain processes may, at different times, contribute to the content of consciousness or not. To enter consciousness, it is not

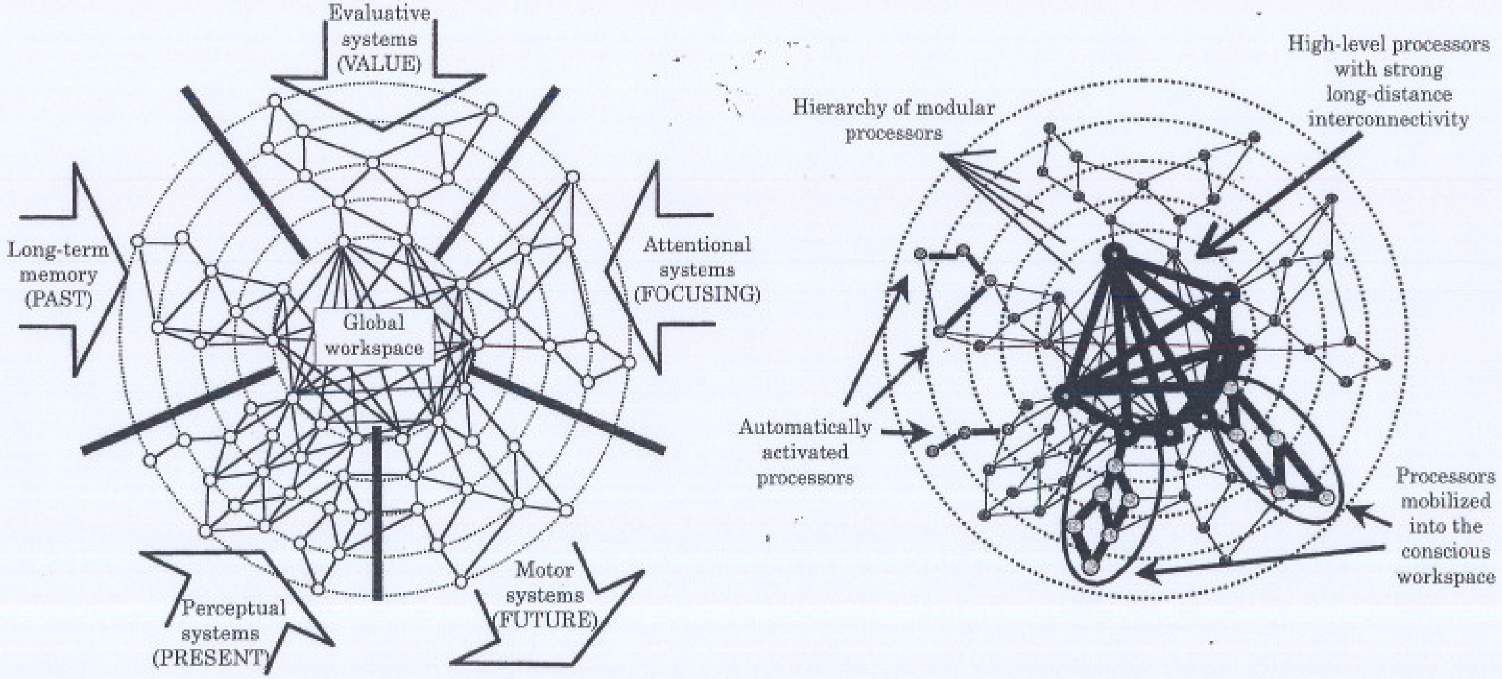

Figure 4.3. Schematic architecture of the workspace model (after Dehaene, Kerszberg, et al., 1998). In this diagram inspired by Mesulam, cerebral processors (circles) and their connections are represented in a hierarchy (primary areas on the outer circle, central associative areas in the center). Areas rich in workspace neurons (center) are tightly interconnected by a divergent set of connections. This long-range connectivity supports brain-scale states of activity (right diagram) in which distant processors are mobilized, amplified, and interconnected through the workspace. From "A Neuronal Model of a Global Workspace in Effortful Cognitive Tasks," by S. Dehaene, M. Kerszberg, and J. P. Changeux, 1998, *Proceedings of the National Academy of Sciences of the United States of America, 95*, p. 14530. Copyright 1998 by the Academy of Sciences of the United States of America. Adapted with permission.

sufficient for a processor to be activated; this activity must also be amplified and maintained over a sufficient duration for it to become accessible to multiple other processes. Without such "dynamic mobilization," a process may still contribute to cognitive performance, but only nonconsciously.

A consequence of this hypothesis is the absence of a sharp anatomical delineation of the workspace representations. In time, the contours of the workspace fluctuate as different brain circuits are temporarily mobilized, then abandoned by a given global representation. Workspace neurons are present in many areas, but at any given time only a particular set of these neurons contribute to the mobilized workspace content. They are part of what may be referred to, in a selectionist framework, as a generator of diversity (Changeux & Dehaene, 1989). As time elapses, the activity of workspace neurons is characterized by a series of discrete episodes of spontaneous metastable coherent activation separated by sharp transitions. This would fit with the introspective feeling of a stream of consciousness, compared by William James to a sequence of flights and perchings of a bird.

Criteria for Conscious Access

To be able to be mobilized in the conscious workspace, a mental object must meet three criteria:

1. Active firing. The object must be represented as a firing pattern of neurons. There is, of course, considerable information that is already stored in the nervous system, in a latent form, for instance, in synaptic connections and weights, neurotransmitter release efficiencies, receptor densities, and so on. The model predicts that such information does not become conscious. It can only be read-out indirectly through its contribution to neural firing.
2. Long-distance connectivity. The active neurons must possess a sufficient number of reciprocal anatomical connections to distributed workspace neurons, particularly in prefrontal, parietal, and cingulate cortices. This criterion implies that the activity of many neurons, for instance in subcortical and brainstem nuclei, is excluded from conscious mobilization (e.g., circuits for respiration or emotion). In many cases, we only become aware of those circuits through their indirect effects on other representations, for example, in somatic cortical areas.
3. Dynamic mobilization. At any given moment, workspace neurons can only sustain a single global representation, the rest of workspace neurons being inhibited. This implies that, out of the multiple active cortical representations that could become conscious, only one will receive the appropriate top-down amplification and be mobilized into consciousness. The other representations are temporarily nonconscious. It would only take a small reorientation of top-down signals to access them, but, according to our views, until this is achieved they do not participate in consciousness.

Workspace Modulation and Selection by Reward

Workspace neurons are assumed to be the targets of two different types of neuromodulatory inputs. First, workspace neurons display a constantly fluctuating spontaneous activity, whose intensity is modulated by ascending activating systems, for instance from cholinergic, noradrenergic, and serotoninergic nuclei in the brain stem, basal forebrain, and hypothalamus. Those systems therefore modify the state of consciousness through different levels of arousal. Second, the stability of workspace activity is modulated by ascending reward inputs arising from the limbic system (via connections to the anterior cingulate, orbitofrontal cortex, and the direct influence of ascending dopaminergic inputs). External or internal goals and rewards may thus stabilize or destabilize particular contents of the conscious workspace. Active representations that fit with the current goal of the organism are selected and maintained over a longer period. Conversely, active representations that lead to error are rejected. This mental selection process has been simulated in former models, which account for classical cognitive tasks such as the Wisconsin card sorting test (Dehaene & Changeux, 1991), the Tower of London (Dehaene & Changeux, 1997), and the Stroop task (Dehaene, Kerszberg, & Changeux, 1998).

Brain Anatomy of the Neuronal Workspace

The neuronal workspace hypothesis posits that, as a whole, the workspace neurons are reciprocally connected via long-distance axons to many if not all of the cortical processors, thus permitting locally available information to be brought into consciousness. Nevertheless, these neurons may be more densely accumulated in some areas than in others. Anatomically, long-range corticocortical tangential connections, including callosal connections, originate mostly from the pyramidal cells of layers II and III, which give or receive the so-called "association" efferents and afferents. Those layers are thicker in von Economo's type 2 (dorsolateral prefrontal) and type 3 (inferior parietal) cortical structures. In the monkey, those areas entertain a strong interconnection among themselves as well as with the anterior and posterior cingulate, the association cortex of the superior temporal sulcus, and the parahippocampal region, thalamus, and striatum (Goldman-Rakic, 1988). The high concentration of neurons with long-distance axons in those areas may explain why they frequently appear coactivated in neuroimaging studies of conscious effortful processing.

Although we emphasize cortico-cortical connectivity, it should be noted that cortico-thalamic columns are the processing units in the brain and in our recent simulation (Dehaene, Sergent, et al., 2003). Thus, long-distance connections between thalamic nuclei may also contribute to the establishment of a coherent brain-scale state (Llinas, Ribary, Contreras, & Pedroarena, 1998). Studies of split-brain patients should be particularly helpful in delineating the relative contribution of cortical and subcortical connections to workspace coherence.

Our model leads to the prediction that long-distance connections have been the target of a recent evolutionary pressure in the course of hominization and are particularly developed in the human species. In that respect, it can be noted that the relative anatomical expansion of cortical areas rich in long-axon neurons, such as the prefrontal cortex, may have contributed to important changes in the functional properties of the workspace (see Changeux, 2002). It is also noteworthy that a particular type of spindle cell, which establishes long-distance projections, is found in the anterior cingulate cortex of humans and great apes, but not other primates (Allman, Hakeem, Erwin, Nimchinsky, & Hof, 2001). Detailed anatomical studies of transcortical connectivity in the human brain have also revealed the presence of distant transcortical projections, that for instance link directly the right fusiform gyrus to multiple areas of the left-hemisphere including Broca's and Wernicke's areas (Di Virgilio & Clarke, 1997). It is anticipated that those key components of the verbal reportability system are connected to many cortical areas, given the variety of percepts and concepts that we can name or understand through language.

Some Predictions for Word and Number Processing

The framework of the neuronal workspace model applies readily to the above data on subliminal word processing. One must simply assume that the fusiform visual word form system and, in the case of numbers, the parietal quantity representation and comparison systems are automatized processors that can function in a purely bottom-up or feedforward fashion. The activation of motor cortex by subliminal primes also indicates that even a temporary chain of processors, prepared according to task instructions, can be traversed by subliminal information. This is compatible with the workspace model inasmuch as it is assumed that the prime-induced activation is sufficient for feedforward propagation through a task-prepared circuit but is too brief to be stabilized by descending amplification from workspace neurons. Finally, the differences observed between conscious and unconscious word processing (Dehaene et al., 2001)—amplification of processor activity, correlated activation of distant prefrontal, parietal, and cingulate sites, emission of a broad P300—conform closely to the expectations of the workspace model.

More recently, my colleagues have been inspired by the workspace model to perform additional experiments that tested critical predictions of the model. Three of them are described below: the relation between attention and consciousness; the role of the anterior cingulate; and the all-or-none character of conscious access.

Attention Without Consciousness

Within the framework of the workspace model, top-down attentional selection and amplification is necessary for the access of a representation to consciousness. However, attention and consciousness cannot be equated. Even when attention is present, it may not always be sufficient for a stimulus to gain access to consciousness. Thus, the model predicts that attention may modulate the depth of subliminal processing, while still failing to make the stimuli conscious.

Naccache, Blandin, and I recently tested this prediction using our number priming paradigm. We reasoned that all of the priming experiments described above allowed research participants to deploy attention to the target. Would the prime effect be modulated or even disappear if participants were unable to attend at the time of prime presentation? In three different experiments, we recently showed that when the prime-target pair occurs at an unpredictable moment, thus preventing the deployment of temporal attention, then subliminal priming effects disappear (Naccache, Blandin, & Dehaene, 2002). This indicates that the idea that subliminal priming reflects a purely passive process of spreading activation can be rejected. Rather, subliminal primes benefit from an attentional amplification, although this may not be sufficient for them to enter into the conscious workspace.

Consciousness and the Anterior Cingulate

Following previous proposals (e.g., Posner, 1994), the neuronal workspace model proposes that functions that depend on central executive control, such as inhibition or conflict detection, require consciousness. Recently, my colleagues and I applied this prediction to the anterior cingulate activation. Although the anterior cingulate activates in many conflict tasks (Botvinick, Braver, Barch, Carter, & Cohen, 2001), these tasks typically involve conscious conflicts. We predicted that a subliminal form of conflict should not lead to anterior cingulate activation.

To test this idea, we capitalized on our previous research showing that the number priming paradigm could generate a motor conflict without consciousness (Dehaene, Naccache, et al., 1998). We, therefore, contrasted the motor conflict effects generated by subliminal and supraliminal numerical primes (Dehaene, Artiges, et al., 2003). Both effects generated behavioral response interference, although the effect was smaller for subliminal primes (14 ms) than for supraliminal primes (34 ms). As predicted, however, the anterior cingulate showed a conflict effect only with supraliminal primes, not with subliminal primes. Furthermore, patients with schizophrenia with known anterior cingulate and prefrontal cortex impairments showed normal subliminal priming, but abnormal supraliminal motor interference. Thus, the subliminal motor conflict generated by masked numbers is resolved spontaneously, without requiring the deployment of executive attention resources. Our results support the early insight that anterior cingulate is tightly associated with conscious processing (e.g., Posner, 1994) and suggest that the executive control processes associated with prefrontal and cingulate cortices can only operate on consciously perceived stimuli. At the very least, they indicate that there is more to anterior cingulate activation than a mere mechanical activation whenever discrepant representations are present in motor cortex (Botvinick et al., 2001). There is a nonlinear relation between the amount of response conflict and the amount of cingulate activation, which implies that a threshold level of conflict is necessary for the anterior cingulate to activate. The neuronal workspace theory also predicts that this threshold should coincide with the threshold for conscious perception of the primes.

All-or-None Dynamics of Consciousness

In the neuronal workspace model, the distinct anatomical connectivity of workspace neurons leads to qualitatively distinct patterns of activity. Because of their global recurrent connectivity, workspace neurons have the capacity of igniting suddenly in a self-amplifying manner as soon as a minimum subset of them is activated. At any given moment, the state of activity of workspace is therefore characterized by the intense activation of a subset of workspace neurons, the rest of workspace neurons being actively inhibited (see Figure 4.4). This particular set of active workspace neurons may be viewed as a neuronal correlate of the content of consciousness. For instance, the conscious report of a visual word might be constituted by the simultaneous, coordinated activation of workspace neurons in the fusiform visual word form area and in distributed temporal, parietal, prefrontal, and cingulate sites associated with speech production.

For stimuli close to threshold, this predicts that conscious access should be all-or-none. Neural activity is either sufficient to trigger a reverberating loop of bottom-up and top-down activity that quickly attains a self-sustained level; or it remains below this threshold and only a briefly decaying bottom-up activation is seen. This nonlinear response was explicitly demonstrated in a detailed simulation of realistic thalamo-cortical networks (Dehaene, Sergent, et al., 2003). The simulations showed that a brief thalamic stimulation can lead to the ignition of a large set of distant cortical areas, which remain active through self-sustaining reverberatory loops for tens of milliseconds beyond the initial stimulus duration. Crucially, during this period of workspace occupancy by a first stimulus T1, another T2 can still be processed by peripheral thalamo-cortical processors but often cannot activate workspace neurons until the representation of T1 had vanished. This temporary inability showed many parallels with the attentional blink, a well-known psychological paradigm in which research participants are temporarily unable to report stimuli while they are attending to another task (Chun & Potter, 1995; Vogel et al., 1998).

Further simulations lead to the prediction that the attentional blink is a dynamic all-or-none phenomenon. In my simulation, depending on random fluctuations in spontaneous activity prior to stimulus arrival, ascending activity can be sufficient to trigger self-amplifying recurrent activity, or it can remain below threshold and only transient bottom-up activity was seen. Thus, for a fixed T1-T2 lag, simulated firing rates in higher areas and other indices of global activity (gamma-band power, long-distance cross-correlation) are distributed bimodally across trials—either global and long-lasting, or local and short-lived.

The theory, therefore, predicts that the apparent gradual drop in reportability observed during the attentional blink may be an artificial consequence of averaging across trials with full access awareness and others with no awareness. Sergent and I tested this prediction experimentally using a modified attentional blink paradigm in which research participants reported to what extent they had seen a word (T2) within a rapid letter stream that contained another target letter string (T1; Dehaene, Sergent, et al., 2003). To obtain a continuous measure of subjective perception, we asked participants to move a

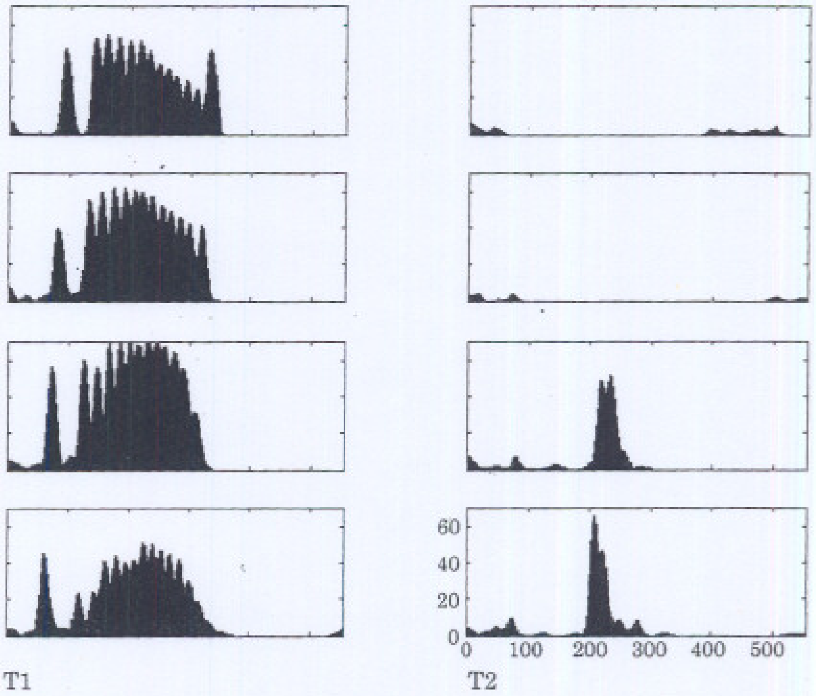

Figure 4.4. Instantaneous firing rate of simulated neurons in four hierarchically organized areas as observed in a recent neuronal implementation of the workspace model (Dehaene, Sergent, et al., 2003). A first stimulus (T1, left), although presented for only 40 ms, yields a short bottom-up activation followed by a long-lasting reverberating global state of activity. The latter is hypothesized to make T2 available for report by multiple satellite systems (not shown). The presentation of a second, competing stimulus (T2, right) during this period of global activity yields only a short-lived, bottom-up propagation of activity; thus, T2 is not reportable. This sequence of activity may provide a preliminary characterization of the neuronal events responsible for the attentional blink, in which processing of a first target T1 yields a temporary lack of consciousness of a subsequent stimulus T2. From "A Neuronal Model Linking Subjective Reports and Objective Physiological Data During Conscious Perception," by S. Dehaene, C. Sergent, and J. P. Changeux, 2003, *Proceedings of the National Academy of Sciences of the United States of America, 100*, p. 8523. Copyright 2003 by the Academy of Sciences of the United States of America. Adapted with permission.

cursor on a continuous scale, from "not seen" on the left to "maximal visibility" on the right. The results indicated that subjective perception during the blink is indeed all-or-none. At the peak of the blink, which occurred ~260 ms after T1, the very same stimulus T2 was either fully perceived (cursor placed on maximal visibility; ~50% of trials), or totally unseen (cursor placed on "not seen"). Participants almost never used intermediate cursor positions, although controls showed that they were able to in other psychophysical situations. This experiment substantiates the hypothesis that conscious states are associated with a fast all-or-none dynamic phase transition in a large-scale neuronal network. More generally, the concept of a sudden ignition, self-amplified by recurrent top-down/bottom-up interaction, may begin to explain the very notion of a threshold or "limen" of consciousness.

Conclusion

It is encouraging that there is increasing empirical and theoretical agreement about the essential ingredients for a theory of consciousness. The proposed neuronal workspace theory, indeed, can be seen as a physiological implementation of the classical concept of a central executive, supervisory attentional, or self-regulation system (e.g., Norman & Shallice, 1980; Posner & Rothbart, 1998) that accesses and modulates lower-level processors. At the neuronal network level, a key role is given to connections with the prefrontal cortex and anterior cingulate, in agreement with Bianchi (1922), Crick and Koch (1995), and Posner and Rothbart (1998). Finally, the concept of reverberatory, recurrent or reentrant projections in perceptual awareness has been abundantly mentioned in the past (Changeux, 1983; Di Lollo, Enns, & Rensink, 2000; Edelman, 1993; Lamme & Roelfsema, 2000).

Not surprisingly, all of these properties have been part and parcel of Posner's views on consciousness for a long period of time. The emphasis on the anterior cingulate as a crucial node in the conscious executive attention system, in particular, is a remarkable discovery that we owe to recent advances in brain imaging, as this region was largely neglected in neuropsychological research prior to the cognitive neuroimaging experiments of Posner, Petersen, Raichle, and colleagues. More broadly, many of the themes in this chapter, such as the idea that there are identifiable architectures of consciousness, that most cerebral processes are nonconscious, and that consciousness is associated with attentional amplification, can be traced back to Michael Posner.

In the précis of their book *Images of Mind,* Posner and Raichle (1995) described some of the criteria that a putative complete theory of cognitive neuroscience should fulfill:

> A theory of cognitive neuroscience would tell us how the brain works. It would be integrated in two senses. It would give an account at all the levels described in our general framework, from a specification of the cognitive systems to an understanding of the cellular mechanisms that support them. It would also describe how these cognitive systems achieve the subjective conscious experience that we call mind.

They ended by modestly stating: "We do not think that such an integrated theory is at hand, at least not in our hands." Clearly, such a theory is not in anyone's hands as yet. However, if we all somehow feel that this remote goal has become a bit closer in the past decades, we owe it in large part to the revolutionary conceptual and methodological tools devised by Posner.

References

Abrams, R. L., & Greenwald, A. G. (2000). Parts outweigh the whole (word) in unconscious analysis of meaning. *Psychological Science, 11,* 118–124.

Allman, J. M., Hakeem, A., Erwin, J. M., Nimchinsky, E., & Hof, P. (2001). The anterior cingulate cortex. The evolution of an interface between emotion and cognition. *Annals of the New York Academy of Sciences, 935,* 107–117.

Baars, B. J. (1989). *A cognitive theory of consciousness.* Cambridge, England: Cambridge University Press.

Bar, M., Tootell, R. B. H., Schacter, D. L., Greve, D. N., Fischl, B., Mendola, J. D., et al. (2001). Cortical mechanisms specific to explicit visual object recognition. *Neuron, 29,* 529–535.

Bianchi, L. (1922). *The mechanism of the brain and the functions of the frontal lobes.* New York: W. Wood.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review, 108,* 624–652.

Bowers, J. S., Vigliocco, G., & Haan, R. (1998). Orthographic, phonological, and articulatory contributions to masked letter and word priming. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 1705–1719.

Changeux, J. P. (1983). *Neuronal man.* Paris: Fayard.

Changeux, J. P. (2002). *The physiology of truth: Neuroscience and human knowledge.* Paris: Odile Jacob.

Changeux, J. P., & Dehaene, S. (1989). Neuronal models of cognitive functions. *Cognition, 33,* 63–109.

Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 109–127.

Cohen, L., Lehericy, S., Chochon, F., Lemer, C., Rivaud, S., & Dehaene, S. (2002). Language-specific tuning of visual cortex? Functional properties of the visual word form area. *Brain, 125*(Pt. 5), 1054–1069.

Crick, F., & Koch, C. (1995). Are we aware of neural activity in primary visual cortex? *Nature, 375,* 121–123.

Damian, M. F. (2001). Congruity effects evoked by subliminally presented primes: Automaticity rather than semantic processing. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 154–165.

Dehaene, S. (2003). The neural bases of subliminal priming. In N. Kanwisher & J. Duncan (Eds.), *Attention and performance: Functional neuroimaging of visual cognition* (Vol. 20, p. 20). New York: Oxford University Press.

Dehaene, S., Artiges, E., Naccache, L., Martelli, C., Viard, A., Schürhoff, F., et al. (2003). Conscious and subliminal conflicts in normal and schizophrenic subjects: The role of the anterior cingulate. *Proceedings of the National Academy of Sciences of the United States of America, 100,* 13722–13727.

Dehaene, S., & Changeux, J. P. (1989). A simple model of prefrontal cortex function in delayed-response tasks. *Journal of Cognitive Neuroscience, 1,* 244–261.

Dehaene, S., & Changeux, J. P. (1991). The Wisconsin Card Sorting Test: Theoretical analysis and modelling in a neuronal network. *Cerebral Cortex, 1,* 62–79.

Dehaene, S., & Changeux, J. P. (1997). A hierarchical neuronal network for planning behavior. *Proceedings of the National Academy of Sciences of the United States of America, 94,* 13293–13298.

Dehaene, S., & Changeux, J. P. (2000). Reward-dependent learning in neuronal networks for planning and decision making. *Progress in Brain Research, 126,* 217–229.

Dehaene, S., & Changeux, J. P. (2004). Neural mechanisms for access to consciousness. In M. Gazzaniga (Ed.), *The cognitive neurosciences* (3rd ed.). New York: Norton.

Dehaene, S., Jobert, A., Naccache, L., Ciuciu, P., Poline, J. B., Le Bihan, D., et al. (in press). Letter binding and invariant recognition of masked words: Behavioral and neuroimaging evidence. *Psychological Science, 15,* 307–313.

Dehaene, S., Kerszberg, M., & Changeux, J. P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences of the United States of America,* 95, 14529–14534.

Dehaene, S., Le Clec'H, G., Poline, J. B., Le Bihan, D., & Cohen, L. (2002). The visual word form area: A prelexical representation of visual words in the fusiform gyrus. *NeuroReport, 13,* 321–325.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition, 79,* 1–37.

Dehaene, S., Naccache, L., Cohen, L., Le Bihan, D., Mangin, J. F., Poline, J. B., et al. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience, 4,* 752–758.

Dehaene, S., Naccache, L., Le Clec'H, G., Koechlin, E., Mueller, M., Dehaene-Lambertz, G., et al. (1998). Imaging unconscious semantic priming. *Nature, 395,* 597–600.

Dehaene, S., Sergent, C., & 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 of the United States of America, 100,* 8520–8525.

Di Lollo, V., Enns, J. T., & Rensink, R. A. (2000). Competition for consciousness among visual events: The psychophysics of reentrant visual processes. *Journal of Experimental Psychology: General, 129,* 481–507.

Di Virgilio, G., & Clarke, S. (1997). Direct interhemispheric visual input to human speech areas. *Human Brain Mapping, 5,* 347–354.

Edelman, G. M. (1993). Neural Darwinism: Selection and reentrant signaling in higher brain function. *Neuron, 10,* 115–125.

Forster, K. I., & Davis, C. (1984). Repetition priming and frequency attenuation in lexical access. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 680–698.

Goldman-Rakic, P. S. (1988). Topography of cognition: Parallel distributed networks in primate association cortex. *Annual Review of Neuroscience, 11,* 137–156.

Greenwald, A. G., Abrams, R. L., Naccache, L., & Dehaene, S. (2003). Long-term semantic memory versus contextual memory in unconscious number processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 235–247.

Greenwald, A. G., Draine, S. C., & Abrams, R. L. (1996). Three cognitive markers of unconscious semantic activation. *Science, 273,* 1699–1702.

Grill-Spector, K., Kushnir, T., Hendler, T., & Malach, R. (2000). The dynamics of object-selective activation correlate with recognition performance in humans. *Nature Neuroscience, 3,* 837–843.

Koechlin, E., Naccache, L., Block, E., & Dehaene, S. (1999). Primed numbers: Exploring the modularity of numerical representations with masked and unmasked semantic priming. *Journal of Experimental Psychology: Human Perception and Performance, 25,* 1882–1905.

Kunde, W., Kiesel, A., & Hoffmann, J. (2003). Conscious control over the content of unconscious cognition. *Cognition, 88,* 223–242.

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

Llinas, R., Ribary, U., Contreras, D., & Pedroarena, C. (1998). The neuronal basis for consciousness. *Philosophical Transactions: Biological Sciences, 353,* 1841–1849.

Marsolek, C. J., Kosslyn, S. M., & Squire, L. R. (1992). Form-specific visual priming in the right cerebral hemisphere. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 492–508.

Molko, N., Cohen, L., Mangin, J. F., Chochon, F., Lehericy, S., Le Bihan, D., et al. (2002). Visualizing the neural bases of a disconnection syndrome with diffusion tensor imaging. *Journal of Cognitive Neuroscience, 14,* 629–636.

Naccache, L., Blandin, E., & Dehaene, S. (2002). Unconscious masked priming depends on temporal attention. *Psychological Science, 13,* 416–424.

Naccache, L., & Dehaene, S. (2001a). The priming method: Imaging unconscious repetition priming reveals an abstract representation of number in the parietal lobes. *Cerebral Cortex, 11,* 966–974.

Naccache, L., & Dehaene, S. (2001b). Unconscious semantic priming extends to novel unseen stimuli. *Cognition, 80,* 215–229.

Neumann, O., & Klotz, W. (1994). Motor responses to non-reportable, masked stimuli: Where is the limit of direct motor specification. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV: Conscious and non-conscious information processing* (pp. 123–150). Cambridge, MA: MIT Press.

Norman, D. A., & Shallice, T. (1980). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation* (Vol. 4, pp. 1–18). New York: Plenum Press.

Pardo, J. V., Fox, P. T., & Raichle, M. E. (1991). Localization of a human system for sustained attention by positron emission tomography. *Nature, 349,* 61–64.

Pardo, J. V., Pardo, P. J., Janer, K. W., & Raichle, M. E. (1990). The anterior cingulate cortex mediates processing selection in the Stroop attentional conflict paradigm. *Proceedings of the National Academy of Sciences of the United States of America, 87,* 256–259.

Paulesu, E., Demonet, J. F., Fazio, F., McCrory, E., Chanoine, V., Brunswick, N., et al. (2001). Dyslexia: cultural diversity and biological unity. *Science, 291,* 2165–2167.

Petersen, S. E., Fox, P. T., Posner, M. I., Mintun, M., & Raichle, M. E. (1988). Positron emission tomographic studies of the cortical anatomy of single-word processing. *Nature, 331,* 585–589.

Posner, M. I. (1978). *Chronometric explorations of the mind.* Hillsdale, NJ: Erlbaum.

Posner, M. I. (1993). Seeing the mind. *Science, 262,* 673–674.

Posner, M. I. (1994). Attention: The mechanisms of consciousness. *Proceedings of the National Academy of Sciences of the United States of America, 91,* 7398–7403.

Posner, M. I., & Carr, T. H. (1992). Lexical access and the brain: Anatomical constraints on cognitive models of word recognition. *American Journal of Psychology, 105,* 1–26.

Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience, 13,* 25–42.

Posner, M. I., Petersen, S. E., Fox, P. T., & Raichle, M. E. (1988). Localization of cognitive operations in the human brain. *Science, 240,* 1627–1631.

Posner, M. I., & Raichle, M. E. (1995). Precis of images of mind. *Behavioral and Brain Sciences, 18,* 327–383.

Posner, M. I., & Rothbart, M. K. (1998). Attention, self-regulation and consciousness. *Philosophical Transactions: Biological Sciences, 353,* 1915–1927.

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

Reynvoet, B., Brysbaert, M., & Fias, W. (2002). Semantic priming in number naming. *Quarterly Journal of Experimental Psychology A, 55,* 1127–1139.

Reynvoet, B., Caessens, B., & Brysbaert, M. (2002). Automatic stimulus–response associations may be semantically mediated. *Psychonomic Bulletin and Review, 9,* 107–112.

Shaywitz, B. A., Shaywitz, S. E., Pugh, K. R., Mencl, W. E., Fulbright, R. K., Skudlarski, P., et al. (2002). Disruption of posterior brain systems for reading in children with developmental dyslexia. *Biological Psychiatry, 52,* 101–110.

Vogel, E. K., Luck, S. J., & Shapiro, K. L. (1998). Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 1656–1674.

Vorberg, D., Mattler, U., Heinecke, A., Schmidt, T., & Schwarzbach, J. (2003). Different time courses for visual perception and action priming. *Proceedings of the National Academy of Sciences of the United States of America, 100,* 6275–6280.

Warrington, E. K., & Shallice, T. (1980). Word-form dyslexia. *Brain, 103,* 99–112.