



Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework

Stanislas Dehaene*, Lionel Naccache

Unité INSERM 334, Service Hospitalier Frédéric Joliot, CEA/DRM/DSV, 4, Place du Général Leclerc, 91401 Orsay Cedex, France

Received 8 February 2000; accepted 27 September 2000

Abstract

This introductory chapter attempts to clarify the philosophical, empirical, and theoretical bases on which a cognitive neuroscience approach to consciousness can be founded. We isolate three major empirical observations that any theory of consciousness should incorporate, namely (1) a considerable amount of processing is possible without consciousness, (2) attention is a prerequisite of consciousness, and (3) consciousness is required for some specific cognitive tasks, including those that require durable information maintenance, novel combinations of operations, or the spontaneous generation of intentional behavior. We then propose a theoretical framework that synthesizes those facts: the hypothesis of a global neuronal workspace. This framework postulates that, at any given time, many modular cerebral networks are active in parallel and process information in an unconscious manner. An information becomes conscious, however, if the neural population that represents it is mobilized by top-down attentional amplification into a brain-scale state of coherent activity that involves many neurons distributed throughout the brain. The long-distance connectivity of these ‘workspace neurons’ can, when they are active for a minimal duration, make the information available to a variety of processes including perceptual categorization, long-term memorization, evaluation, and intentional action. We postulate that this global availability of information through the workspace is what we subjectively experience as a conscious state. A complete theory of consciousness should explain why some cognitive and cerebral representations can be permanently or temporarily inaccessible to consciousness, what is the range of possible conscious contents, how they map onto specific cerebral circuits, and whether a generic neuronal mechanism underlies all of them. We confront the workspace model with those issues and identify novel experimental predictions. Neurophysiological, anatomical, and brain-imaging data strongly argue for a major role of prefrontal cortex,

* Corresponding author. Tel.: +33-1-69-86-78-73; fax: +33-1-69-86-78-16.
E-mail address: dehaene@shfj.cea.fr (S. Dehaene).

anterior cingulate, and the areas that connect to them, in creating the postulated brain-scale workspace. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Consciousness; Awareness; Attention; Priming

1. Introduction

The goal of this volume is to provide readers with a perspective on the latest contributions of cognitive psychology, neuropsychology, and brain imaging to our understanding of consciousness. For a long time, the word ‘consciousness’ was used only reluctantly by most psychologists and neuroscientists. This reluctance is now largely overturned, and consciousness has become an exciting and quickly moving field of research. Thanks largely to advances in neuropsychology and brain imaging, but also to a new reading of the psychological and neuropsychological research of the last decades in domains such as attention, working memory, novelty detection, or the body schema, a new comprehension of the neural underpinnings of consciousness is emerging. In parallel, a variety of models, pitched at various levels in neural and/or cognitive science, are now available for some of its key elements.

Within this fresh perspective, firmly grounded in empirical research, the problem of consciousness no longer seems intractable. Yet no convincing synthesis of the recent literature is available to date. Nor do we know yet whether the elements of a solution that we currently have will suffice to solve the problem, or whether key ingredients are still missing. By grouping some of the most innovative approaches together in a single volume, this special issue aims at providing the readers with a new opportunity to see for themselves whether a synthesis is now possible.

In this introduction, we set the grounds for subsequent papers by first clarifying what we think should be the aim of a cognitive neuroscience approach to consciousness. We isolate three major findings that are explored in greater detail in several chapters of this volume. Finally, we propose a synthesis that integrates them into what we view as a promising theoretical framework: the hypothesis of a global neuronal workspace. With this framework in mind, we look back at some of the remaining empirical and conceptual difficulties of consciousness research, and examine whether a clarification is in sight.

2. Nature of the problem and range of possible solutions

Let us begin by clarifying the nature of the problem that a cognitive neuroscience of consciousness should address. In our opinion, this problem, though empirically challenging, is conceptually simple. Human subjects routinely refer to a variety of conscious states. In various daily life and psychophysical testing situations, they use phrases such as ‘I was not conscious of X’, ‘I suddenly realized that Y’, or ‘I knew

that Z, therefore I decided to do X'. In other words, they use a vocabulary of psychological attitudes such as believing, pretending, knowing, etc., that all involve to various extents the concept of 'being conscious'. In any given situation, such conscious phenomenological reports can be very consistent both within and across subjects. The task of cognitive neuroscience is to identify which mental representations and, ultimately, which brain states are associated with such reports. Within a materialistic framework, each instance of mental activity is also a physical brain state.¹ The cognitive neuroscience of consciousness aims at determining whether there is a systematic form of information processing and a reproducible class of neuronal activation patterns that systematically distinguish mental states that subjects label as 'conscious' from other states.²

From this perspective, the problem of the cognitive neuroscience of consciousness does not seem to pose any greater conceptual difficulty than identifying the cognitive and cerebral architectures for, say, motor action (identifying what categories of neural and/or information-processing states are systematically associated with moving a limb). What is specific to consciousness, however, is that the object of our study is an introspective phenomenon, not an objectively measurable response. Thus, the scientific body of consciousness calls for a specific attitude which departs from the 'objectivist' or 'behaviorist' perspective often adopted in behavioral and neural experimentation. In order to cross-correlate subjective reports of consciousness with neuronal or information-processing states, the first crucial step is *to take seriously introspective phenomenological reports*. Subjective reports are the key phenomena that a cognitive neuroscience of consciousness purport to study. As such, they constitute primary data that need to be measured and recorded along with other psychophysiological observations (Dennett, 1992; Weiskrantz, 1997; see also Merikle, Smilek, & Eastwood, this volume).

The idea that introspective reports must be considered as serious data in search of a model does not imply that introspection is a privileged mode of access to the inner workings of the mind. Introspection can be wrong, as is clearly demonstrated, for instance, in split-brain subjects whose left-hemispheric verbal 'interpreter' invents a plausible but clearly false explanation for the behavior caused by their right hemisphere (Gazzaniga, LeDoux, & Wilson, 1977). We need to find a scientific explanation for subjective reports, but we must not assume that they always constitute accurate descriptions of reality. This distinction is clearest in the case of hallucinations. If someone claims to have visual hallucinations of floating faces, or 'out-of-body' experiences, for instance, it would be wrong to take these reports as unequi-

¹ We use the word 'state' in the present context to mean any configuration of neural activity, whether stable (a fixed point) or dynamic (a trajectory in neural space). It is an open question as to whether neural states require stability over a minimal duration to become conscious, although the workspace model would predict that some degree of stable amplification over a period of at least about 100 ms is required.

² One should also bear in mind the possibility that what naive subjects call 'consciousness' will ultimately be parceled into distinct theoretical constructs, each with its own neural substrate, just like the naive concept of 'warmth' was ultimately split into two distinct physical parameters, temperature and heat.

vocal evidence for parapsychology, but it would be equally wrong to dismiss them as unverifiable subjective phenomena. The correct approach is to try to explain how such conscious states can arise, for instance by appealing to an inappropriate activation of face processing or vestibular neural circuits, as can indeed be observed by brain-imaging methods during hallucinations (Fytche et al., 1998; Silbersweig et al., 1995).

The emphasis on subjective reports as data does not mean that the resulting body of knowledge will be inherently subjective and therefore non-scientific. As noted by Searle (1998), a body of knowledge is scientific ('epistemically objective') inasmuch as it can be verified independently of the attitudes or preferences of the experimenters, but there is nothing in this definition that prevents a genuinely scientific approach of domains that are inherently subjective because they exist only in the experience of the subject ('ontologically subjective' phenomena). "The requirement that science be objective does not prevent us from getting an epistemically objective science of a domain that is ontologically subjective." (Searle, 1998, p. 1937).

One major hurdle in realizing this program, however, is that "we are still in the grip of a residual dualism" (Searle, 1998, p. 1939). Many scientists and philosophers still adhere to an essentialist view of consciousness, according to which conscious states are ineffable experiences of a distinct nature that may never be amenable to a physical explanation. Such a view, which amounts to a Cartesian dualism of substance, has led some to search for the bases of consciousness in a different form of physics (Penrose, 1990). Others make the radical claim that two human brains can be identical, atom for atom, and yet one can be conscious while the other is a mere 'zombie' without consciousness (Chalmers, 1996).

Contrary to those extreme statements, contributors to the present volume share the belief that the tools of cognitive psychology and neuroscience may suffice to analyze consciousness. This need not imply a return to an extreme form of direct psycho-neural reductionism. Rather, research on the cognitive neuroscience of consciousness should clearly take into account the many levels of organization at which the nervous system can be studied, from molecules to synapses, neurons, local circuits, large scale networks, and the hierarchy of mental representations that they support (Changeux & Dehaene, 1989). In our opinion, it would be inappropriate, and a form of 'category error', to attempt to reduce consciousness to a low level of neural organization, such as the firing of neurons in thalamocortical circuits or the properties of NMDA receptors, without specifying in functional terms the consequences of this neural organization at the cognitive level. While characterization of such neural bases will clearly be indispensable to our understanding of consciousness, it cannot suffice. A full theory will require many more 'bridging laws' to explain how these neural events organize into larger-scale active circuits, how those circuits themselves support specific representations and forms of information processing, and how these processes are ultimately associated with conscious reports. Hence, this entire volume privileges cognitive neuroscientific approaches to consciousness that seem capable of addressing both the cognitive architecture of mental representations and their neural implementation.

3. Three fundamental empirical findings on consciousness

In this section, we begin by providing a short review of empirical observations that we consider as particularly relevant to the cognitive neuroscience of consciousness. We focus on three findings: the depth of unconscious processing; the attention-dependence of conscious perception; and the necessity of consciousness for some integrative mental operations.

3.1. *Cognitive processing is possible without consciousness*

Our first general observation is that a considerable amount of processing can occur without consciousness. Such unconscious processing is open to scientific investigation using behavioral, neuropsychological and brain-imaging methods. By increasing the range of cognitive processes that do not require consciousness, studies of unconscious processing contribute to narrowing down the cognitive bases of consciousness. The current evidence indicates that many perceptual, motor, semantic, emotional and context-dependent processes can occur unconsciously.

A first line of evidence comes from studies of brain-lesioned patients. Pöppel, Held, and Frost (1973) demonstrated that four patients with a partial blindness due to a lesion in visual cortical areas (hemianopsic scotoma) remained able to detect visual stimuli presented in their blind field. Although the patients claimed that they could not see the stimuli, indicating a lack of phenomenal consciousness, they nevertheless performed above chance when directing a visual saccade to them. This ‘blindsight’ phenomenon was subsequently replicated and extended in numerous studies (Weiskrantz, 1997). Importantly, some patients performed at the same level as control subjects, for instance in motor pointing tasks. Thus, unconscious processing is not limited to situations in which information is degraded or partially available. Rather, an entire stream of processing may unfold outside of consciousness.

Dissociations between accurate performance and lack of consciousness were subsequently identified in many categories of neuropsychological impairments such as visual agnosia, prosopagnosia, achromatopsia, callosal disconnection, aphasia, alexia, amnesia, and hemineglect (for reviews, see Köhler & Moscovitch, 1997; Schacter, Buckner, & Koutstaal, 1998; see also Driver & Vuilleumier, this volume). The current evidence suggests that, in many of these cases, unconscious processing is possible at a perceptual, but also a semantic level. For instance, Renault, Signoret, Debrulle, Breton, and Bolgert (1989) recorded event-related potentials to familiar and unknown faces in a prosopagnosic patient. Although the patient denied any recognition of the familiar faces, an electrical waveform indexing perceptual processing, the P300, was significantly shorter and more intense for the familiar faces. Similar results were obtained by recording the electrodermal response, an index of vegetative processing of emotional stimuli, in prosopagnosic patients (Bauer, 1984; Tranel & Damasio, 1985). Even clearer evidence for semantic-level processing comes from studies of picture–word priming in neglect patients (McGlinchey-Berroth, Milberg, Verfaellie, Alexander, & Kilduff, 1993). When two images are

presented simultaneously in the left and right visual fields, neglect patients deny seeing the one on the left, and indeed cannot report it beyond chance level. Nevertheless, when having to perform a lexical decision task on a subsequent foveal word, which can be related or unrelated to the previous image, they show the same amount of semantic priming from both hemifields, indicating that even the unreportable left-side image was processed to a semantic level.

Similar priming studies indicate that a considerable amount of unconscious processing also occurs in normal subjects. Even a very brief visual stimulus can be perceived consciously when presented in isolation. However, the same brief stimulus can fail to reach consciousness when it is surrounded in time by other stimuli that serve as masks. This lack of consciousness can be assessed objectively using signal detection theory (for discussion, see Holender, 1986; Merikle, 1992; see also Merikle et al.).³ Crucially, the masked stimulus can still have a measurable influence on the processing of subsequent stimuli, a phenomenon known as masked priming. There are now multiple demonstrations of perceptual, semantic, and motor processing of masked stimuli. For instance, in various tasks, processing of a conscious target stimulus can be facilitated by the prior masked presentation of the same stimulus (repetition priming; e.g. Bar & Biederman, 1999). Furthermore, masked priming also occurs when the relation between prime and target is a purely semantic one, such as between two related words (Dehaene, Naccache et al., 1998; Klinger & Greenwald, 1995; Marcel, 1983; see also Merikle et al.). We studied semantic priming with numerical stimuli (Dehaene, Naccache et al., 1998; Koechlin, Naccache, Block, & Dehaene, 1999). When subjects had to decide whether target numbers were larger or smaller than five, the prior presentation of another masked number accelerated the response in direct proportion to its amount of similarity with the target, as measured by numerical distance (Koechlin et al., 1999). Furthermore, the same number-comparison experiment also provided evidence that processing of the prime occurs even beyond this semantic stage to reach motor preparation systems (Dehaene, Naccache et al., 1998). When the instruction specified that targets larger than five should be responded to with the right hand, for instance, primes that were larger than five facilitated a right-hand response, and measures of brain activation demonstrated a significant covert activation of motor cortex prior to the main overt response (see also Eimer & Schlaghecken, 1998; Neumann & Klotz, 1994). Thus, an entire stream of perceptual, semantic and motor processes, specified by giving arbitrary verbal instructions to a normal subject, can occur outside of consciousness.

The number priming experiment also illustrates that it is now feasible to visualize

³ Unfortunately, signal detection theory provides an imperfect criterion for consciousness. If subjects exhibit a d' measure that does not differ significantly from zero in a forced-choice stimulus detection or discrimination task, one may conclude that no information about the stimulus was available for conscious processing. Conversely, however, a non-zero d' measure need not imply consciousness, but may result from both conscious and unconscious influences. Experimental paradigms that partially go beyond this limitation have been proposed (e.g. Jacoby, 1991; Klinger & Greenwald, 1995). We concur with Merikle et al. (this volume), however, in thinking that subjective reports remain the crucial measure when assessing the degree of consciousness (see also Weiskrantz, 1997).

directly the brain areas involved in unconscious processing, without having to rely exclusively on indirect priming measures (Dehaene, Naccache et al., 1998; Morris, Öhman, & Dolan, 1998; Sahraie et al., 1997; Whalen et al., 1998; see also Driver & Vuilleumier and Kanwisher, this volume). In the Whalen et al. (1998) experiment, for instance, subjects were passively looking at emotionally neutral faces throughout. Yet the brief, unconscious presentation of masked faces bearing an emotional expression of fear, relative to neutral masked faces, yielded an increased activation of the amygdala, a brain structure known to be involved in emotional processing. We expect such brain-imaging studies to play an important role in mapping the cerebral networks implicated in unconscious processing, and therefore isolating the neural substrates of consciousness.

3.2. Attention is a prerequisite of consciousness

Experiments with masked primes indicate that some minimal duration and clarity of stimulus presentation are necessary for it to become conscious. However, are these conditions also sufficient? Do all stimuli with sufficient intensity and duration automatically gain access to consciousness? Evidence from brain-lesioned patients as well as normal subjects provides a negative answer. Conditions of stimulation, by themselves, do not suffice to determine whether a given stimulus is or is not perceived consciously. Rather, conscious perception seems to result from an interaction of these stimulation factors with the attentional state of the observer. The radical claim was even made that “there seems to be no conscious perception without attention” (Mack & Rock, 1998, p. ix).

Brain-lesioned patients suffering from hemineglect provide a striking illustration of the role of attentional factors in consciousness (Driver & Mattingley, 1998; see also Driver & Vuilleumier, 2001, this issue). Hemineglect frequently results from lesions of the right parietal region, which is thought to be involved in the orientation of attention towards locations and objects. Neglect patients fail to attend to stimuli located in contralesional space, regardless of their modality of their presentation. The focus of attention seems permanently biased toward the right half of space, and patients behave as if the left half had become unavailable to consciousness. This is seen most clearly in the extinction phenomenon: when two visual stimuli are presented side by side left and right of fixation, the patients report only seeing the stimulus on the right, and appear completely unconscious of the identity or even the presence of a stimulus on the left. Nevertheless, the very same left-hemifield stimulus, when presented in isolation at the same retinal location, is perceived normally. Furthermore, even during extinction, priming measures indicate a considerable amount of covert processing of the neglected stimulus at both perceptual and semantic levels (e.g. McGlinchey-Berroth et al., 1993). Hence, although the cortical machinery for bottom-up processing of left-lateralized stimuli seems to be largely intact and activated during extinction, this is clearly not sufficient to produce a conscious experience; a concomitant attentional signal seems compulsory.

In normal subjects, the role of attention in conscious perception has been the subject of considerable research (see Merikle et al., 1995). While there remains

controversy concerning the depth of processing of unattended stimuli, there is no doubt that attention serves as a filter prior to conscious perception (see Driver & Vuilleumier, 2001, this issue). Visual search experiments indicate that, given an array of items, the orienting of attention plays a critical role in determining whether a given item gains access to consciousness (Sperling, 1960; Treisman & Gelade, 1980). Objects that do not fall in an attended region of the visual field cannot be consciously reported. Furthermore, there are systematic parallels between the fate of unattended stimuli and the processing of masked primes. Merikle and Joordens (1997) describe three phenomena (Stroop priming, false recognition, and exclusion failure) in which qualitatively similar patterns of performance are observed in divided attention and in masked priming experiments. They conclude that “perception with and without awareness, and perception with and without attention, are equivalent ways of describing the same underlying process distinction” (p. 219).

Mack and Rock (1998) have investigated a phenomenon called inattentional blindness that clearly illustrates this point. They asked normal subjects to engage in a demanding visual discrimination task at a specific location in their visual field. Then on a single trial, another visual stimulus appeared at a different location. This stimulus clearly had sufficient contrast and duration (typically 200 ms) to be perceptible in isolation, yet the use of a single critical trial and of a distracting task ensured that it was completely unattended and unexpected. Under these conditions a large percentage of subjects failed to report the critical stimulus and continued to deny its presence when explicitly questioned about it. In some experimental conditions, even a large black circle presented for 700 ms in the fovea failed to be consciously perceived! Yet priming measures again indicated that the unseen stimulus was processed covertly. For instance, a word extinguished by inattentional blindness yielded strong priming in a subsequent stem completion task. Such evidence, together with similar observations that supra-threshold visual stimuli fail to be reported during the ‘attentional blink’ (Luck, Vogel, & Shapiro, 1996; Raymond, Shapiro, & Arnell, 1992; Vogel, Luck, & Shapiro, 1998), and that large changes in a complex visual display fail to be noticed unless they are attended (‘change blindness’; e.g. O’Regan, Rensink, & Clark, 1999), support the hypothesis that attention is a necessary prerequisite for conscious perception.⁴

3.3. *Consciousness is required for specific mental operations*

Given that a considerable amount of mental processing seems to occur uncon-

⁴ The notion that attention is required for conscious perception seems to raise a potential paradox: if we can only perceive what we attend to, how do we ever become aware of unexpected information? In visual search experiments, for instance, a vertical line ‘pops out’ of the display and is immediately detected regardless of display size. How is this possible if that location did not receive prior attention? Much of this paradox dissolves, however, once it is recognized that some stimuli can automatically and unconsciously capture attention (Yantis & Jonides, 1984, 1996). Although we can consciously orient our attention, for instance to search through a display, orienting of attention is also determined by unconscious bottom-up mechanisms that have been attuned by evolution to quickly orient us to salient new features of our environment. Pop-out experiments can be reinterpreted as revealing a fast attraction of attention to salient features.

sciously, one is led to ask what are the computational benefits associated with consciousness. Are there any specific mental operations that are feasible only when one is conscious of performing them? Are there sharp limits on the style and amount of unconscious computation? This issue is obviously crucial if one is to understand the computational nature and the evolutionary advantages associated with consciousness. Yet little empirical research to date bears on this topic. In this section, which is clearly more speculative than previous ones, we tentatively identify at least three classes of computations that seem to require consciousness: durable and explicit information maintenance, novel combinations of operations, and intentional behavior (see also Jack & Shallice, this volume for a similar attempt to identify ‘Type-C’ processes specifically associated with consciousness).

3.3.1. *Durable and explicit information maintenance*

The classical experiment by Sperling (1960) on iconic memory demonstrates that, in the absence of conscious amplification, the visual representation of an array of letters quickly decays to an undetectable level. After a few seconds or less, only the letters that have been consciously attended remain accessible. We suggest that, in many cases, the ability to maintain representations in an active state for a durable period of time in the absence of stimulation seems to require consciousness. By ‘in an active state’, we mean that the information is encoded in the firing patterns of active populations of neurons and is therefore immediately available to influence the systems they connect with. Although sensory and motor information can be temporarily maintained by passive domain-specific buffers such as Sperling’s iconic store, with a half-life varying from a few hundreds of milliseconds to a few seconds (auditory information being possibly held for a longer duration than visual information), exponential decay seems to be the rule whenever information is not attended (e.g. Cohen & Dehaene, 1998; Tiitinen, May, Reinikainen, & Naatanen, 1994).

Priming studies nicely illustrate the short-lived nature of unconscious representations. In successful masked priming experiments, the stimulus onset asynchrony (SOA) between prime and target is typically quite short, in the order of 50–150 ms. Experiments that have systematically varied this parameter indicate that the amount of priming drops sharply to a non-significant value within a few hundreds of milliseconds (Greenwald, 1996). Thus, the influence of an unconscious prime decays very quickly, suggesting that its mental representation vanishes dramatically as time passes.⁵ This interpretation is supported by single-unit recordings in the monkey infero-temporal (IT) cortex during masked and unmasked presentations of faces (Rolls & Tovee, 1994; Rolls, Tovee, & Panzeri, 1999). A very short and masked visual presentation yields a short-lasting burst of firing (~50 ms) in face-

⁵ Some incidental learning and mere exposure experiments have reported unconscious priming effects at a long duration (e.g. Bar & Biederman, 1999; Bornstein & D’Agostino, 1992; Elliott & Dolan, 1998). We interpret these findings as suggesting that even unseen, short-lived stimuli may leave long-lasting latent traces, for instance in the form of alterations in synaptic weights in the processing network. This is not incompatible, however, with our postulate that no active, explicit representation of a prime can remain beyond a few hundred milliseconds in the absence of conscious attentional amplification.

selective cells. However, an unmasked face presented for the same short duration yields a long burst whose duration (up to 350 ms) far exceeds the stimulation period. Physiological and behavioral studies in both humans and monkeys suggest that this ability to maintain information on-line independently of the stimulus presence depends on a working memory system associated with dorsolateral prefrontal regions (Fuster, 1989; Goldman-Rakic, 1987). By this argument, then, the working memory system made available by prefrontal circuitry must be tightly related to the durable maintenance of information in consciousness (e.g. Fuster, 1989; Kosslyn & Koenig, 1992; Posner, 1994; see below).

Another remarkable illustration of the effect of time delays on the ability to maintain active and accurate unconscious representation is provided by studies of the impact of visual illusions on reaching behavior (Aglioti, DeSouza, & Goodale, 1995; Daprati & Gentilucci, 1997; Gentilucci, Chieffi, Deprati, Saetti, & Toni, 1996; Hu, Eagleson, & Goodale, 1999). In the Müller-Lyer and Tichener illusions, although two objects have the same objective length, one of them is perceived as looking shorter due to the influence of contextual cues. Nevertheless, when subjects make a fast reaching movement toward the objects, their finger grip size is essentially unaffected by the illusion and is therefore close to objective size. Hence, the motor system is informed of an objective size parameter which is not available to consciousness, providing yet another instance of unconscious visuo-motor processing. Crucially, however, when one introduces a short delay between the offset of stimuli and the onset of the motor response, grip size becomes less and less accurate and is now influenced by the subjective illusion (Gentilucci et al., 1996). In this situation, subjects have to bridge the gap between stimulus and response by maintaining an internal representation of target size. The fact that they now misreach indicates that the accurate but unconscious information cannot be maintained across a time delay. Again, active information survives a temporal gap only if it is conscious.

3.3.2. *Novel combinations of operations*

The ability to combine several mental operations to perform a novel or unusual task is a second type of computation that seems to require consciousness. Conflict situations, in which a routine behavior must be inhibited and superseded by a non-automatized strategy, nicely illustrate this point. Merikle, Joordens, and Stolz (1995) studied subjects' ability to control inhibition in a Stroop-like task as a function of the conscious perceptibility of the conflicting information. Subjects had to classify a colored target string as green or red. Each target was preceded by a prime which could be the word GREEN or RED. In this situation, the classical Stroop effect was obtained: responses were faster when the word and color were congruent than when they were incongruent. However, when the prime–target relations were manipulated by presenting 75% of incongruent trials, subjects could strategically take advantage of the predictability of the target from the prime, and became faster on incongruent trials than on congruent trials, thus inverting the Stroop effect. Crucially, this strategic inversion only occurred when the prime was consciously perceptible. No strategic effect was observed when the word prime was masked (Merikle et al.,

1995) or fell outside the focus of attention (Merikle & Joordens, 1997). In this situation, only the classical, automatic Stroop effect prevailed. Thus, the ability to inhibit an automatic stream of processes and to deploy a novel strategy depended crucially on the conscious availability of information.

We tentatively suggest, as a generalization, that the strategic operations which are associated with planning a novel strategy, evaluating it, controlling its execution, and correcting possible errors cannot be accomplished unconsciously. It is noteworthy that such processes are always associated with a subjective feeling of ‘mental effort’, which is absent during automatized or unconscious processing and may therefore serve as a selective marker of conscious processing (Dehaene, Kerszberg, & Changeux, 1998).⁶

3.3.3. *Intentional behavior*

A third type of mental activity that may be specifically associated with consciousness is the spontaneous generation of intentional behavior. Consider the case of blindsight patients. Some of these patients, even though they claim to be blind, show such an excellent performance in pointing to objects that some have suggested them as a paradigmatic example of the philosopher’s ‘zombie’ (a hypothetical human being who would behave normally, but lacks consciousness). As noted by Dennett (1992) and Weiskrantz (1997), however, this interpretation fails to take into account a fundamental difference with normal subjects: blindsight patients never spontaneously initiate any visually-guided behavior in their impaired field. Good performance can be elicited only by forcing them to respond to stimulation.

All patients with preserved implicit processing seem to have a similar impairment in using the preserved information to generate intentional behavior. The experimental paradigms that reveal above-chance performance in these patients systematically rely on automatizable tasks (stimulus–response associations or procedural learning) with forced-choice instructions. This is also true for normal subjects in subliminal processing tasks. As noted above, masked priming experiments reveal the impossibility for subjects to strategically use the unconscious information demonstrated by priming effects. Given the large amount of information that has been demonstrated to be available with consciousness, this limitation on subliminal processing is not trivial. Intentionally driven behaviors may constitute an important class of processes accessible only to conscious information.

Introspective speech acts, in which the subject uses language to describe his/her mental life, constitute a particular category of intentional behaviors that relate to conscious processing. Consciousness is systematically associated with the potential ability for the subject to report on his/her mental state. This property of *reportability* is so exclusive to conscious information that it is commonly used as an empirical

⁶ An important qualification is that even tasks that involve complex series of operations and that initially require conscious effort may become progressively automatized after some practice (e.g. driving a car). At this point, such tasks may proceed effortlessly and without conscious control. Indeed, many demonstrations of unconscious priming involve acquired strategies that required a long training period, such as word reading.

criterion to assess the conscious or unconscious status of an information or a mental state (Gazzaniga et al., 1977; Weiskrantz, 1997).

4. A theoretical framework for consciousness

Once those three basic empirical properties of conscious processing have been identified, can a theoretical framework be proposed for them? Current accounts of consciousness are founded on extraordinarily diverse and seemingly incommensurate principles, ranging from cellular properties such as thalamocortical rhythms to purely cognitive constructions such as the concept of a ‘central executive’. Instead of attempting a synthesis of those diverse proposals, we isolate in this section three theoretical postulates that are largely shared, even if they are not always explicitly recognized. We then try to show how these postulates, taken together, converge onto a coherent framework for consciousness: the hypothesis of a global neuronal workspace.

4.1. *The modularity of mind*

A first widely shared hypothesis is that automatic or unconscious cognitive processing rests on multiple dedicated processors or ‘modules’ (Baars, 1989; Fodor, 1983; Shallice, 1988). There are both functional and neurobiological definitions of modularity. In cognitive psychology, modules have been characterized by their information encapsulation, domain specificity, and automatic processing. In neuroscience, specialized neural circuits that process only specific types of inputs have been identified at various spatial scales, from orientation-selective cortical columns to face-selective areas. The breakdown of brain circuits into functionally specialized subsystems can be evidenced by various methods including brain imaging, neuropsychological dissociation, and cell recording.

We shall not discuss here the debated issue of whether each postulated psychological module can be identified with a specific neural circuit. We note, however, that the properties of automaticity and information encapsulation postulated in psychology are partially reflected in modular brain circuits. Specialized neural responses, such as face-selective cells, can be recorded in both awake and anesthetized animals, thus reflecting an automatic computation that can proceed without attention. Increasingly refined analyses of anatomical connectivity reveal a channeling of information to specific targeted circuits and areas, thus supporting a form of information encapsulation (Felleman & Van Essen, 1991; Young et al., 1995).

As a tentative theoretical generalization, we propose that *a given process, involving several mental operations, can proceed unconsciously only if a set of adequately interconnected modular systems is available to perform each of the required operations*. For instance, a masked fearful face may cause unconscious emotional priming because there are dedicated neural systems in the superior colliculus, pulvinar, and right amygdala associated with the attribution of emotional valence to faces (Morris, Öhman, & Dolan, 1999). Our hypothesis implies that multiple unconscious operations can proceed in parallel, as long as they do not

simultaneously appeal to the same modular systems in contradictory ways. Note that unconscious processing may not be limited to low-level or computationally simple operations. High-level processes may operate unconsciously, as long as they are associated with functional neural pathways either established by evolution, laid down during development, or automatized by learning. Hence, there is no systematic relation between the objective complexity of a computation and the possibility of its proceeding unconsciously. For instance, face processing, word reading, and postural control all require complex computations, yet there is considerable evidence that they can proceed without attention based on specialized neural subsystems. Conversely, computationally trivial but non-automatized operations, such as solving $21 - 8$, require conscious effort.

4.2. *The apparent non-modularity of the conscious mind*

It was recognized early on that several mental activities cannot be explained easily by the modularity hypothesis (Fodor, 1983). During decision making or discourse production, subjects bring to mind information conveyed by many different sources in a seemingly non-modular fashion. Furthermore, during the performance of effortful tasks, they can temporarily inhibit the automatic activation of some processors and enter into a strategic or ‘controlled’ mode of processing (Posner, 1994; Schneider & Shiffrin, 1977; Shallice, 1988). Many cognitive theories share the hypothesis that controlled processing requires a distinct functional architecture which goes beyond modularity and can establish flexible links amongst existing processors. It has been called the central executive (Baddeley, 1986), the supervisory attentional system (Shallice, 1988), the anterior attention system (Posner, 1994; Posner & Dehaene, 1994), the global workspace (Baars, 1989; Dehaene, Kerszberg, & Changeux, 1998) or the dynamic core (Tononi & Edelman, 1998).

Here we synthesize those ideas by postulating that, besides specialized processors, the architecture of the human brain also comprises *a distributed neural system or ‘workspace’ with long-distance connectivity that can potentially interconnect multiple specialized brain areas in a coordinated, though variable manner* (Dehaene, Kerszberg, & Changeux, 1998). Through the workspace, modular systems that do not directly exchange information in an automatic mode can nevertheless gain access to each other’s content. The global workspace thus provides a common ‘communication protocol’ through which a particularly large potential for the combination of multiple input, output, and internal systems becomes available (Baars, 1989).

If the workspace hypothesis is correct, it becomes an empirical issue to determine which modular systems make their contents globally available to others through the workspace. Computations performed by modules that are not interconnected through the workspace would never be able to participate in a conscious content, regardless of the amount of introspective effort (examples may include the brainstem systems for blood pressure control, or the superior colliculus circuitry for gaze control). The vast amounts of information that we can consciously process suggests

that at least five main categories of neural systems must participate in the workspace: perceptual circuits that inform about the present state of the environment; motor circuits that allow the preparation and controlled execution of actions; long-term memory circuits that can reinstate past workspace states; evaluation circuits that attribute them a valence in relation to previous experience; and attentional or top-down circuits that selectively gate the focus of interest. The global interconnection of those five systems can explain the subjective unitary nature of consciousness and the feeling that conscious information can be manipulated mentally in a largely unconstrained fashion. In particular, connections to the motor and language systems allow any workspace content to be described verbally or non-verbally ('reportability'; Weiskrantz, 1997).

4.3. *Attentional amplification and dynamic mobilization*

A third widely shared theoretical postulate concerns the role of attention in gating access to consciousness. As reviewed earlier, empirical data indicate that considerable processing is possible without attention, but that attention is required for information to enter consciousness (Mack & Rock, 1998). This is compatible with Michael Posner's hypothesis of an attentional amplification (Posner, 1994; Posner & Dehaene, 1994), according to which the orienting of attention causes increased cerebral activation in attended areas and a transient increase in their efficiency.

Dehaene, Kerszberg, and Changeux (1998) have integrated this notion within the workspace model by postulating that *top-down attentional amplification is the mechanism by which modular processes can be temporarily mobilized and made available to the global workspace, and therefore to consciousness*. According to this theory, the same cerebral processes may, at different times, contribute to the content of consciousness or not. To enter consciousness, it is not sufficient for a process to have on-going activity; 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 unconsciously.

A consequence of this hypothesis is the absence of a sharp anatomical delineation of the workspace system. In time, the contours of the workspace fluctuate as different brain circuits are temporarily mobilized, then demobilized. It would therefore be incorrect to identify the workspace, and therefore consciousness, with a fixed set of brain areas. Rather, many brain areas contain workspace neurons with the appropriate long-distance and widespread connectivity, and at any given time only a fraction of these neurons constitute the mobilized workspace. As discussed below, workspace neurons seem to be particularly dense in prefrontal cortices (PFCs) and anterior cingulate (AC), thus conferring those areas a dominant role. However, we see no need to postulate that any single brain area is systematically activated in all conscious states, regardless of their content. It is the style of activation (dynamic long-distance mobilization), rather than its cerebral localization, which characterizes consciousness. This hypothesis therefore departs radically from the notion of a single central 'Cartesian theater' in which conscious information is displayed

(Dennett, 1992). In particular, information that is already available within a modular process does not need to be re-represented elsewhere for a ‘conscious audience’: dynamic mobilization makes it directly available in its original format to all other workspace processes.

The term ‘mobilization’ may be misinterpreted as implying the existence of an internal homunculus who decides to successively amplify and then suppress the relevant processes at will. Our view, however, considers this mobilization as a collective dynamic phenomenon that does not require any supervision, but rather results from the spontaneous generation of stochastic activity patterns in workspace neurons and their selection according to their adequacy to the current context (Dehaene, Kerszberg, & Changeux, 1998). Stochastic fluctuations in workspace neurons would result, at the collective neuronal assembly level, in the spontaneous activation, in a sudden, coherent, exclusive and ‘auto-catalytic’ (self-amplifying) manner, of a subset of workspace neurons, the rest being inhibited. This active workspace state is not completely random, but is heavily constrained and selected by the activation of surrounding processors that encode the behavioral context, goals, and rewards of the organism. In the resulting dynamics, transient self-sustained workspace states follow one another in a constant stream, without requiring any external supervision. Explicit, though still elementary, computer simulations of such ‘neuronal Darwinism’ are available, illustrating its computational feasibility (Changeux & Dehaene, 1989; Dehaene & Changeux, 1997; Dehaene, Kerszberg, & Changeux, 1998; Friston, Tononi, Reeke, Sporns, & Edelman, 1994).

5. Empirical consequences, reinterpretations, and predictions

The remainder of this paper is devoted to an exploration of the empirical consequences of this theoretical framework. We first examine the predicted structural and dynamical conditions under which information may become conscious. We then consider the consequences of our views for the exploration of the neural substrates of consciousness and its clinical or experimental disruption.

5.1. *Structural constraints on the contents of consciousness*

An important scientific goal regarding consciousness is to explain why some representations that are encoded in the nervous system are permanently impervious to consciousness. In the present framework, the conscious availability of information is postulated to be determined by two structural criteria which are ultimately grounded in brain anatomy. First, the information must be represented in an active manner in the firing of one or several neuronal assemblies. Second, bidirectional connections must exist between these assemblies and the set of workspace neurons, so that a sustained amplification loop can be established. Cerebral representations that violate either criteria are predicted to be permanently inaccessible to consciousness.

The first criterion – active representation – excludes from the contents of consciousness the enormous wealth of information which is present in the nervous

system only in latent form, for instance in the patterns of anatomical connections or in strengthened memory traces. As an example, consider the wiring of the auditory system. We can consciously attend to the spatial location of a sound, but we are oblivious to the cues that our nervous system uses to compute it. One such cue is the small time difference between the time of arrival of sound in the two ears (interaural delay). Interaural delay is coded in a very straightforward manner by the neural connection lengths of the medial superior olive (Smith, Joris, & Yin, 1993). Yet such connectivity information, by hypothesis, cannot reach consciousness.

More generally, the ‘active representation’ criterion may explain the observation that we can never be conscious of the inner workings of our cerebral processes, but only of their outputs. A classical example is syntax: we can become conscious that a sentence is not grammatical, but we have no introspection on the inner workings of the syntactical apparatus that underlies this judgment, and which is presumably encoded in connection weights within temporal and frontal language areas.

There is, however, one interesting exception to this limit on introspection. Subjects’ verbal reports do provide a reliable source of information on a restricted class of processes that are slow, serial, and controlled, such as those involved in solving complex arithmetic problems or the Tower of Hanoi task (Ericsson & Simon, 1993). We propose that what distinguishes such processes is that they are not encoded in hardwired connectivity, but rather are generated dynamically through the serial organization of active representations of current goals, intentions, decisions, intermediate results, or errors. The firing of many prefrontal neurons in the monkey encodes information about the animal’s current goals, behavioral plans, errors and successes (e.g. Fuster, 1989). According to the workspace model, such active representations of on-going performance can become available for conscious amplification and communication to other workspace components, explaining, for instance, that we can consciously report the strategic steps that we adopted. The model implies that this is the only situation where we can have reliable conscious access to our mental algorithms. Even then, such access is predicted to be limited. Indeed, when multiplying 32 by 47, we are conscious of our goals, subgoals, main steps (multiplying 2 by 7, then 3 by 7, etc.), and possible errors, but we have no introspection as to how we solve each individual problem.

Our second criterion – bidirectional connectivity with the workspace – implies that some representations, even though they are encoded by an active neuronal assembly, may permanently evade consciousness. This may occur if the connectivity needed to establish a reverberating loop with workspace units is absent or damaged. Consider, for example, the minimal contrast between patients with visual neglect, patients with a retinal scotoma, and normal subjects who all have a blind spot in their retina. Superficially, these conditions have much in common. In all of them, subjects fail to consciously perceive visual stimuli presented at a certain location. Yet the objective processing abilities and subjective reports associated with those visual impairments are strikingly different (see Table 1). Patients with visual neglect typically cannot see stimuli in the neglected part of space, but are not conscious that they are lacking this information. Patients with a retinal scotoma also cannot see in a specific region of their visual field, but they are conscious of their

Table 1

Three classical perceptual conditions in which conscious vision is affected, and their proposed theoretical interpretation. A plus sign indicates an available ability, while a minus sign indicates an absent or deteriorated ability

Condition	Symptoms			Theoretical interpretation	
	Consciousness of visual stimulus	Consciousness of visual impairment	Capacity for unconscious processing	Conscious amplification	Modular processing
Visual neglect	–	–	+	–	+
Retinal scotoma	–	+	–	+	–
Blind spot in normal subjects	–	–	–	–	–

blindness in this region. Finally, all of us have a blind region devoid of photoreceptors in the middle of our retinas, the blind spot, yet we are not conscious of having a hole in our vision.

Table 1 shows how the hypothesis of a conscious mobilization of visual processes through top-down amplification can explain these phenomena. In the case of parietal neglect patients, it has been shown that considerable information about the neglected stimulus is still being actively processed (cf. *supra*). The lesion, however, is thought to affect parietal circuits involved in spatial attention. According to our framework, this may have the effect of disrupting a crucial component in the top-down amplification of visual information, and therefore preventing this information from being mobilized into the workspace. This predicts that recordings of activity evoked by a neglected stimulus in the intact occipito-temporal visual pathway would reveal a significant short-lived activation (sufficient to underlie residual unconscious processing), but without attentional amplification and with an absence of cross-correlation with other distant areas (correlating with the subjects' inability to bring this information to consciousness). Very recently, data compatible with those predictions have been described (Rees et al., 2000; see also Driver & Vuilleumier, 2001, this issue).

The case of retinal scotomas is essentially symmetrical:⁷ subjects lack peripheral visual input, but they have an intact network of cortical areas supporting the attentional amplification of visual information into consciousness. Thus, the information

⁷ We do not attempt here a full theoretical treatment of visual scotomas of cortical origin, which are more complex than those of retinal origin. In both types of pathologies, patients are conscious of their visual impairment, presumably because their intact attentional system allows them to detect the absence of visual inputs. However, patients with cortical scotomas sometimes exhibit residual unconscious processing abilities in their blind field (blindsight). One possibility is that those residual abilities rely on subcortical circuits such as the superior colliculus (Sahraie et al., 1997) which, for lack of workspace neurons, would remain permanently inaccessible to consciousness. Alternatively, they may be supported by cortical activity (e.g. in area V5; Zeki & Ffytche, 1998), but of a weakened and transient nature insufficient to establish a sustained closed loop with the workspace and therefore to enter consciousness.

that visual inputs are no longer available in their scotoma can be made available to the workspace and, from there, contact long-term memory and motor intention circuits. Retinal patients can therefore recognize that they are impaired relative to an earlier time period, and they can report it verbally or non-verbally. According to this account, becoming conscious that one is blind occurs when a discrepancy is detected, within the preserved cortical visual representations, between the activity elicited when attending to long-term memory circuits, and the absence of activity elicited when attempting to attend to the outside world. (Note that this would predict that a person blinded from birth would not experience blindness in the same form; being unable to elicit memories of prior seeing, he/she would have a more ‘intellectual’ or verbal understanding of blindness, perhaps similar to the one that sighted people have.)

Finally, normal subjects’ lack of consciousness of their blind spot can be explained by the lack of both perceptual *and* attentional resources for this part of the visual field (Dennett, 1992). Our visual cortex receives no retinal information from the blind spot, but then, neither can we orient attention to this absence of information. We therefore remain permanently unaware of it. It is noteworthy that the presence of the blind spot can only be demonstrated indirectly: closing one eye, we attend to a small object and move it on the retina until we suddenly see it disappear as it passes over the blind spot. In this situation, object-oriented attention is used to detect an anomalous object disappearance. Spatial attention, however, is permanently unable to let us perceive the blind spot as a hole in our visual field.⁸

As should be clear from the above discussion, whether or not a given category of information is accessible to consciousness cannot be decided a priori, but must be submitted to an empirical investigation. Indeed, recent research has begun to reveal brain circuits that seem to be permanently inaccessible to consciousness. For instance, psychophysical experiments indicate that some information about visual gratings, though extracted by V1 neurons, cannot be consciously perceived (He, Cavanagh, & Intriligator, 1996). Likewise, the dorsal occipito-parietal route involved in guiding hand and eye movements makes accurate use of information about object size and shape, of which subjects can be completely unaware (Aglioti et al., 1995; Daprati & Gentilucci, 1997; Gentilucci et al., 1996; Hu et al., 1999). Although such experiments bear on the contents, rather than on the mechanisms, of consciousness, they may provide crucial tests of theories of consciousness (Crick & Koch, 1995). If the present hypothesis is correct, they should always reveal that the unconscious information is either not explicitly encoded in neural firing, or is encoded by neural populations that lack bidirectional connectivity with the workspace.

5.2. *Dynamical constraints on consciousness*

The previous section dealt with permanently inaccessible information. However, information can also be temporarily inaccessible to consciousness for purely dynamical reasons, as seen in masking paradigms. A masked stimulus presented for a very short duration, even if subject to considerable scrutiny, cannot be brought to

consciousness; a slightly longer duration of presentation, however, is sufficient to render the prime easily visible.

According to workspace theory, conscious access requires the temporary dynamical mobilization of an active processor into a self-sustained loop of activation: active workspace neurons send top-down amplification signals that boost the currently active processor neurons, whose bottom-up signals in turn help maintain workspace activity. Establishment of this closed loop requires a minimal duration, thus imposing a temporal ‘granularity’ to the successive neural states that form the stream of consciousness. This dynamical constraint suggests the existence of two thresholds in human information processing, one that corresponds to the minimal stimulus duration needed to cause any differentiated neural activity at all, and another, the ‘consciousness threshold’, which corresponds to the significantly longer duration needed for such a neural representation to be mobilized in the workspace through a self-sustained long-distance loop. Stimuli that fall in between those two thresholds cause transient changes in neuronal firing and can propagate through multiple circuits (subliminal processing), but cannot take part in a conscious state.

Fig. 1 shows an elementary neural network that illustrates this idea (though it is obviously too simplistic to represent more than a mere aid to intuition). A cascading series of feed-forward networks initially receives a short burst of firing, similar to the phasic response that can be recorded from IT neurons in response to a masked face (Rolls & Tovee, 1994). As seen in Fig. 1, this burst can then propagate through a large number of successive stages, which might be associated with a transformation of the information into semantic, mnemonic or motor codes. At all these levels, however, only a short-lived burst of activity is seen. Although a closed loop involving a distant workspace network is present in the circuit, it takes a longer or stronger stimulus to reliably activate it and to place the network in a long-lasting self-sustained state. Short stimuli cause only a weak, transient and variable activation of the workspace. It is tempting to view such transient firing as a potential neuronal basis for the complex phenomenology of masking paradigms. At intermediate prime durations (40–50 ms), subjects never characterize their perception of the masked prime as ‘conscious’, but many of them report that they occasionally experience a glimpse of the stimulus that seems to immediately recede from their grasp, thus leaving them unable to describe what they saw.

In a more realistic situation, the chain of subliminal processing need not be as rigid as suggested in Fig. 1. In humans at least, verbal instructions can induce a rapid reorganization of existing processors into a novel chain through top-down amplification and selection (Fig. 2). Could such a dynamic chain also be traversed by unconscious information? The model leads us to answer positively as long as the instruction or context stimulus used to guide top-down selection itself is conscious.

⁸ A more thorough discussion of unawareness of the blind spot would require mention of filling-in experiments (e.g. Ramachandran, 1992). The parts of objects or textures that fall in the blind spot seem to be partially reconstructed or ‘filled in’ further on in the visual system, and the results of this reconstruction process can then be consciously attended. None of those experiments, however, refute the basic fact that we are not conscious of having a hole in our retinas.

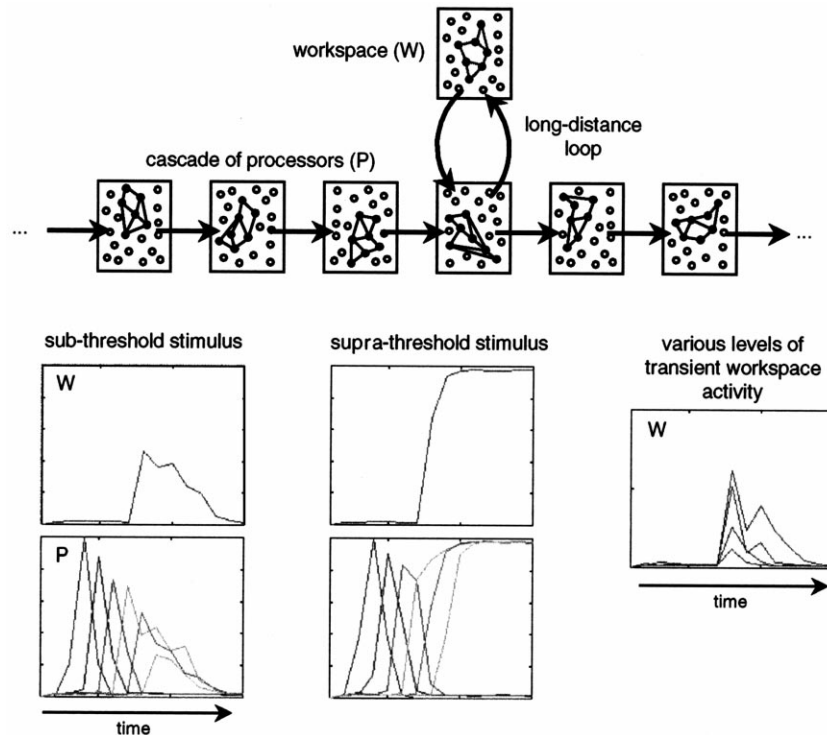


Fig. 1. A simple neural network exhibits dynamical activation patterns analogous to the processing of visual stimuli below and above the consciousness threshold in human subjects. In this minimal scheme, a series of processors are organized in a feed-forward cascade. One of them can also enter in a self-sustained reciprocal loop with a distant workspace network. For simplicity, the evolution of the activation of each assembly (processors and workspace) is modeled by a single McCulloch–Pitts equation, which assumes that activation grows as a non-linear sigmoidal function of the sum of inputs to the assembly, including a small self-connection term, a noise term, and a threshold. Across a wide range of parameters, the same basic findings are reproduced. A short, transient input burst propagates through the processors while causing only a minimal, transient workspace activation (bottom left panels, plots of activation as a function of time). This illustrates how a masked prime, which causes only a transient burst of activation in perceptual neurons of the ventral visual stream (Rolls & Tovee, 1994), can launch an entire stream of visual, semantic and motor processes (Dehaene, Naccache et al., 1998) while failing to establish the sustained coherent workspace activation necessary for consciousness. A slightly more prolonged input causes a sharp transition in activation, with the sudden establishment of a long-lasting activation of both workspace and processor units (middle panels). Thus, the system exhibits a perceptual threshold that stimuli must exceed in order to evoke sustained workspace activation. The right panel illustrates how the very same subthreshold stimulus, on different trials, can evoke transient workspace activity of variable intensity and duration. Similarly, subjects presented with masked primes report a variable phenomenology that ranges from total blindness to a transient feeling that the prime may be on the brink of reportability.

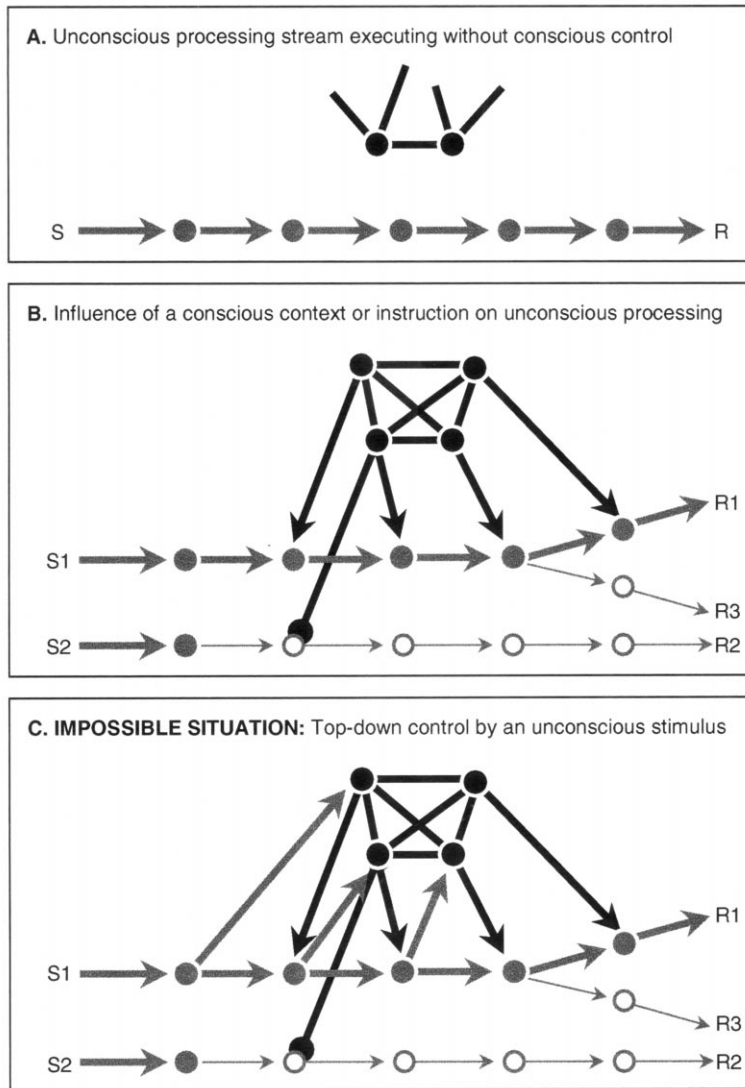


Fig. 2. Which tasks may or may not proceed unconsciously? In these schemas, the gray lines represent the propagation of neural activation associated with the unconscious processing of some information, and the black lines the activation elicited by the presently active conscious workspace neurons. The workspace model predicts that one or several automated stimulus–response chains can be executed unconsciously while the workspace is occupied elsewhere (A). Even tasks that require stimulus and processor selection may be executed unconsciously once the appropriate circuit has been set up by a conscious instruction or context (B). However, it should be impossible for an unconscious stimulus to modify processing on a trial-by-trial basis through top-down control (C). A stimulus that contacts the workspace for a duration sufficient to alter top-down control should always be globally reportable.

Thus, the workspace model makes the counter-intuitive prediction that even a complex task that calls for the setting up of a novel non-automatized pathway, once prepared consciously, can be applied unconsciously. Support for this prediction comes from the above-cited masked priming experiments, which indicate that a novel and arbitrary task instruction can be applied to masked primes (Dehaene, Naccache et al., 1998), even when the instruction changes on every trial (Neumann & Klotz, 1994). The same argument leads us to expect that neglect and blindsight patients should be able to selectively attend to stimuli in their blind field and even to react differentially to them as a function of experimenter instructions, all the while denying seeing them.

What should *not* be possible, however, is for an unconscious stimulus itself to control top-down circuit selection (Fig. 2C). Thus, tasks in which the instruction stimulus itself is masked, or is presented in the blind field of a neglect or blindsight patient, should not be applicable unconsciously. Although this prediction has not been tested explicitly, the finding that normal subjects cannot perform exclusion and strategic inversion tasks with masked primes (Merikle et al., 1995) fits with this idea, since those tasks require an active inhibition that varies on a trial-by-trial basis.

5.3. Neural substrates of the contents of consciousness

The new tools of cognitive neuroscience, particularly brain-imaging methods, now make it possible to explore empirically the neural substrates of consciousness. Rather than attempting a exhaustive review of this fast-growing literature, we use the workspace framework to help organize existing results and derive new predictions. The framework predicts that multiple processors encode the various possible contents of consciousness, but that all of them share a common mechanism of coherent brain-scale mobilization. Accordingly, brain-imaging studies of consciousness can be coarsely divided into two categories: studies of the various contents of consciousness, and studies of its shared mechanisms.

We first examine studies of the cerebral substrates of specific contents of consciousness. Brain-imaging studies provide striking illustrations of a one-to-one mapping between specific brain circuits and categories of conscious contents. For instance, Kanwisher (this volume) describes an area, the fusiform face area (FFA), that seems to be active whenever subjects report a conscious visual percept of a face. This includes non-trivial cases, such as rivalry and hallucinations. When subjects are presented with a constant stimulus consisting of a face in one eye and a house in the other, they report alternatively seeing a face, then a house, but not both (binocular rivalry). Likewise, the FFA does not maintain a constant level of activity, but oscillates between high and low levels of activation in tight synchrony with the subjective reports (Tong, Nakayama, Vaughan, & Kanwisher, 1998). Furthermore, when patients hallucinate faces, the FFA activates precisely when subjects report seeing the hallucination (Ffytche et al., 1998). Although the FFA may not be entirely *specific* for faces, as suggested by its involvement in visual expertise for cars or birds

(Gauthier, Skudlarski, Gore, & Anderson, 2000), its activation certainly correlates tightly with conscious face perception.⁹

Another classical example is motion perception. The activation of the human area V5 (or MT) correlates systematically with the conscious perception of motion (Watson et al., 1993), even in non-trivial cases such as visual illusions. For instance, V5 is active when subjects are presented with static paintings of ‘kinetic art’ that elicit a purely subjective impression of motion (Zeki, Watson, & Frackowiak, 1993). V5 is also active when subjects report experiencing a motion-aftereffect to a static stimulus, and the duration of the activation matches the duration of the illusion (Tootell et al., 1995).

Such experiments indicate a tight correlation between the activation of a specific neural circuit (say, V5) and the subjective report of a conscious content (motion). However, correlation does not imply causation. To establish that a given brain state represents the *causal substrate* of a specific conscious content, rather than a mere *correlate* of consciousness, causality must be established by demonstrating that alterations of this brain state systematically alter subjects’ consciousness. In the case of face or motion perception, this can be verified with several methods. Some patients happen to suffer from brain lesions encompassing the FFA or area V5. As predicted, they selectively lose the conscious visual perception of faces or motion (for review, see Young, 1992; Zeki, 1993). Transcranial magnetic stimulation can also be used to temporarily disrupt brain circuits. When applied to area V5, it prevents the conscious perception of motion (Beckers & Zeki, 1995; Walsh, Ellison, Battelli, & Cowey, 1998). Finally, implanted electrodes can be used, not only to disrupt consciousness, but even to change its contents. In the monkey, microstimulation of small populations of neurons in area V5 biases the perception of motion towards the neurons’ preferred direction (Salzman, Britten, & Newsome, 1990). While the interpretation of this particular study raises the difficult issue of animal consciousness, similar experiments can also be performed in conscious humans in whom electrodes are implanted for therapeutic purposes. It is known since Penfield that human brain stimulation can elicit a rich conscious phenomenology, including dream-like states. Stimulation can even induce highly specific conscious contents such as feelings of profound depression (Bejjani et al., 1999) or hilarity (Fried, Wilson, MacDonald, & Behnke, 1998). Such experiments, together with the others reported in this section, begin to provide evidence for a form of ‘type–type physicalism’, in which the major categories of contents of consciousness are causally related, in a systematic manner, to categories of physical brain states that can be reproducibly identified in each subject.

⁹ The workspace theory predicts that the same modular processors are involved in both conscious and unconscious processing. Hence, the systematic correlation between FFA activation and conscious face perception is predicted to break down under conditions of subliminal face perception or inattentional blindness for faces. In those situations, there should be a small but significant FFA activation (relative to non-face stimuli) without consciousness. Driver and Vuilleumier (this volume) report results compatible with this prediction (see Rees et al., 2000).

5.4. *Neural substrates of the mechanisms of consciousness: prefrontal cortex (PFC), anterior cingulate (AC), and the workspace hypothesis*

While the various contents of consciousness map onto numerous, widely distributed brain circuits, the workspace model predicts that all of these conscious states share a common mechanism. The mobilization of any information into consciousness should be characterized by the simultaneous, coherent activation of multiple distant areas to form a single, brain-scale workspace. Areas rich in workspace neurons should be seen as ‘active’ with brain-imaging methods whenever subjects perform a task which is feasible only in a conscious state, such as one requiring a novel combination of mental operations. Finally, conscious processing should be accompanied by a temporary top-down amplification of activity in neural circuits encoding the current content of consciousness. The cognitive neuroscience literature contains numerous illustrations of these principles, and many of them point to PFC and AC as playing a crucial role in the conscious workspace (Posner, 1994).

5.4.1. *Brain imaging of conscious effort*

Although this was not their goal, many early brain-imaging experiments studied complex, effortful tasks that presumably cannot be performed without conscious guidance. A common feature of those tasks is the presence of intense PFC and AC activation (Cohen et al., 1997; Pardo, Pardo, Janer, & Raichle, 1990; Paus, Koski, Caramanos, & Westbury, 1998). Importantly, PFC and AC activations do not seem needed for automatized tasks, but appear suddenly whenever an automatized task suddenly calls for conscious control. In the verb generation task, for instance, Raichle et al. (1994) demonstrated that PFC and AC activation is present during initial task performance, vanishes after the task has become automatized, but immediately recovers when novel items are presented. Furthermore, in a variety of tasks, AC activates immediately after errors and, more generally, whenever conflicts must be resolved (Carter et al., 1998; Dehaene, Posner, & Tucker, 1994). In the Wisconsin card sorting test, PFC activates suddenly when subjects have to invent a new behavioral rule (Konishi et al., 1998). Both PFC and AC possess the ability to remain active in the absence of external stimulation, such as during the delay period of a delayed-response task (Cohen et al., 1997), or during internally driven activities such as mental calculation (Chochon, Cohen, van de Moortele, & Dehaene, 1999; Rueckert et al., 1996). Finally, concomitant to PFC and AC activation, a selective attentional amplification is seen in relevant posterior areas during focused-attention tasks (Corbetta, Miezin, Dobmeyer, Smulman, & Petersen, 1991; Posner & Dehaene, 1994).

In those experiments, whose goal was not to study consciousness, conscious control is correlated with a variety of factors such as attention, difficulty, and effort. A stronger test of the neural substrates of consciousness requires contrasting two experimental conditions that differ minimally in all respects except for the subjects’ state of consciousness (Baars, 1989). In the last few years, many such paradigms have been developed, contrasting, for instance, levels of anesthesia (Fiset et al., 1999) or implicit versus explicit memory tasks (Rugg et al., 1998). Here we discuss only two examples (but see Frith, Perry, & Lumer, 1999, for review).

5.4.2. *Contrasting conscious and unconscious subjects*

An elegant approach consists of using exactly the same stimuli and tasks, but in separating a posteriori subjects who became conscious of an aspect of the experimental situation and subjects who did not. Using this method, McIntosh, Rajah, and Lobaugh (1999) showed an increased activation in left PFC (and, to a lesser degree, in bilateral occipital cortices and left thalamus) only in subjects who became aware of a systematic relation between auditory and visual stimuli. Importantly, this activation was accompanied by a major increase in the functional correlation of left PFC with other distant brain regions including the contralateral PFC, sensory association cortices, and cerebellum. This long-distance coherence pattern appeared precisely when subjects became conscious and started to use their conscious knowledge to guide behavior. Using a sequence learning task, Grafton, Hazeltine, and Ivry (1995) also identified a large-scale circuit with a strong focus in the right PFC, which was only observed in subjects who became aware of the presence of a repeated sequence in the stimuli.

5.4.3. *Binocular rivalry*

In binocular rivalry, subjects are presented with two dissimilar images, one in each eye, and report seeing only one of them at a time. The dominant image, however, alternates with a period of a few seconds. This paradigm is ideal for the study of consciousness because the conscious content changes while the stimulus remains constant. One can therefore study the cerebral activity caused by the dominant image and, a few seconds later, contrast it with the activity when the very same image has become unconscious. Neuronal recordings in awake monkeys trained to report their perception of two rivaling stimuli indicate that early on in visual pathways (e.g. areas V1, V2, V4, and V5), many cells maintain a constant level of firing, indicating that they respond to the constant stimulus rather than to the variable percept (Leopold & Logothetis, 1999). As one moves up in the visual hierarchy, however, an increasing proportion of cells modulate their firing with the reported perceptual alternations. In IT, as many as 90% of the cells respond only to the perceptually dominant image. Brain imaging in humans indicates that IT activity evoked by the dominant image of a rivaling stimulus is indistinguishable from that evoked when the same image is presented alone in a non-rivaling situation (Tong et al., 1998). Importantly, however, IT activity is accompanied by a concomitant widespread increase in AC, prefrontal, and parietal activation (Lumer, Friston, & Rees, 1998). Thus, the point in time when a given image becomes dominant is characterized by a major brain-scale switch in many areas including AC and PFC. This is not merely a coincidental activation of several unrelated neural systems. Rather, posterior and anterior areas appear to transiently form a single large-scale coherent state, as revealed by increases in functional correlation in fMRI (Lumer & Rees, 1999), high-frequency coherences over distances greater than 10 cm in MEG (Srinivasan, Russell, Edelman, & Tononi, 1999), and transient synchronous neuronal firing in V1 neurons (Fries, Roelfsema, Engel, König, & Singer, 1997).

In humans, similar increases in coherence and phase synchrony in the EEG and MEG gamma band (30–80 Hz) have been evidenced in a variety of conscious

perception paradigms besides rivalry (Rodriguez et al., 1999; Tallon-Baudry & Bertrand, 1999). According to the workspace model, such a long-range coherence should be systematically observed whenever distant areas are mobilized into the conscious workspace. Conversely, however, it cannot be excluded that temporal coding by synchrony is also used by modular processors during non-conscious processing. Thus, increased high-frequency coherence and synchrony are predicted to be a necessary but not sufficient neural precondition for consciousness.

5.4.4. *Anatomy and neurophysiology of the conscious workspace*

Is the workspace hypothesis compatible with finer-grained brain anatomy and physiology? The requirements of the workspace model are simple. Neurons contributing to this workspace should be distributed in at least five categories of circuits (high-level perceptual, motor, long-term memory, evaluative and attentional networks). During conscious tasks, they should, for a minimal duration, enter into coherent self-sustained activation patterns in spite of their spatial separation. Therefore, they must be tightly interconnected through long axons. Again, all three criteria point to PFC, AC, and areas interconnected with them as playing a major role in the conscious workspace (Fuster, 1989; Posner, 1994; Shallice, 1988).

Consider first the requirement for long-distance connectivity. Dehaene, Kerszberg, and Changeux (1998) noted that long-range cortico-cortical tangential connections, including interhemispheric connections, mostly originate from the pyramidal cells of layers 2 and 3. This suggests that the extent to which an area contributes to the global workspace might be simply related to the fraction of its pyramidal neurons that belong to layers 2 and 3. Those layers, though present throughout the cortex, are particularly thick in dorsolateral prefrontal and inferior parietal cortical structures. A simple prediction, then, is that the activity of those layers may be tightly correlated with consciousness. This could be tested using auto-radiography and other future high-resolution functional imaging methods in primates and in humans, for instance in binocular rivalry tasks.

In monkeys, Goldman-Rakic (1988) and her collaborators have described a dense network of long-distance reciprocal connections linking dorsolateral PFC with premotor, superior temporal, inferior parietal, anterior and posterior cingulate cortices as well as deeper structures including the neostriatum, parahippocampal formation, and thalamus (Fig. 3). This connectivity pattern, which is probably also present in humans, provides a plausible substrate for fast communication amongst the five categories of processors that we postulated contribute primarily to the conscious workspace.¹⁰ Temporal and parietal circuits provide a variety of high-level perceptual categorizations of the outside world. Premotor, supplementary

¹⁰ For her studies of monkey anatomy, Goldman-Rakic (1988) proposes a strictly modular view of PFC, according to which multiple such circuits run in parallel, each of them specialized for a stimulus attribute. The non-modularity of the workspace, however, leads us to postulate that humans may differ from monkeys in showing greater cross-circuit convergence towards common prefrontal and cingulate projection sites as well as heavier reciprocal projections between various sectors of PFC. The degree of cross-circuit convergence in the monkey PFC might also be greater than was initially envisaged (e.g. Rao, Rainer & Miller, 1997)




Fig. 3. Neural substrates of the proposed conscious workspace. (A) Symbolic representation of the hierarchy of connections between brain processors (each symbolized by a circle) (after Dehaene, Kerszberg, & Changeux, 1998). Higher levels of this hierarchy are assumed to be widely interconnected by long-distance interconnections, thus forming a global neuronal workspace. An amplified state of workspace activity, bringing together several peripheral processors in a coherent brain-scale activation pattern (black circles), can coexist with the automatic activation of multiple local chains of processors outside the workspace (gray circles). (B) Possible anatomical substrate of the proposed workspace: long-distance network identified in the monkey and linking dorsolateral prefrontal, parietal, temporal, and AC areas with other subcortical targets (from Goldman-Rakic, 1988). (C) Neural dynamics of the workspace, as observed in a neural simulation of a connectivity pattern simplified from (A) (see Dehaene et al., 1998 for details). The matrix shows the activation level of various processor units (top lines) and workspace units (bottom lines) as a function of time. Increased workspace unit activity is seen whenever a novel effortful task is introduced and after errors. Although processor unit activity is continuously present, selective amplification can also be seen when workspace activity is present. (D) Examples of functional neuroimaging tasks activating the postulated workspace network: generation of a novel sequence of random numbers (top left, from Artiges et al., 2000), effortful arithmetic (Chochon et al., 1999), and error processing (left, fMRI data from Carter et al., 1998; right, ERP dipole from Dehaene et al., 1994).

motor and posterior parietal cortices, together with the basal ganglia (notably the caudate nucleus), the cerebellum, and the speech production circuits of the left inferior frontal lobe, allow for the intentional guidance of actions, including verbal

reports, from workspace contents. The hippocampal region provides an ability to store and retrieve information over the long term. Direct or indirect connections with orbitofrontal cortex, AC, hypothalamus, amygdala, striatum, and mesencephalic neuromodulatory nuclei may be involved in computing the value or relevance of current representations in relation to previous experience. Finally, parietal and cingulate areas contribute to the attentional gating and shifting of the focus of interest. Although each of these systems, in isolation, can probably be activated without consciousness, we postulate that their coherent activity, supported by their strong interconnectivity, coincides with the mobilization of a conscious content into the workspace.¹¹

Physiologically, the neural dynamics of those areas are compatible with the role of consciousness in the durable and explicit maintenance of information over time. The dynamics of prefrontal activity are characterized by periods of long-lasting self-sustained firing, particularly obvious when the animal is engaged in a delayed response task (Fuster, 1989). Sustained firing can also be observed in most cortical regions belonging to the above-mentioned circuit, and prefrontal cooling abolishes those distant sustained responses, suggesting that self-sustained states are a functional property of the integral circuit (Chelazzi, Duncan, Miller, & Desimone, 1998; Fuster, Bauer, & Jervey, 1985; Miller, Erickson, & Desimone, 1996). Lesion studies in monkeys and humans indicate that prefrontal lesions often have little effect on the performance of automatized tasks, but strongly impact on exactly the three types of tasks that were listed earlier as crucially dependent on consciousness: the durable maintenance of explicit information (frontal patients suffer from impairments in delayed response and other working memory tests); the elaboration of novel combinations of operations (frontal patients show perseveration and impaired performance in tasks that call for the invention of novel strategies, such as the Tower of London test; Shallice, 1982); and the spontaneous generation of intentional behavior (patients with frontal or cingulate lesions may perform unintended actions that are induced by the experimenter or the context, as seen in utilization and imitation behavior; Lhermitte, 1983; Shallice, Burgess, Schon, & Baxter, 1989).

Given that workspace neurons must be distributed in widespread brain areas, it should not be surprising that prefrontal lesions do not altogether suppress consciousness, but merely interfere, with variable severity, with some of its functions. However, workspace neurons must be specifically targeted by diffuse neuromodulator systems involved in arousal, the sleep/wake cycle, and reward (Dehaene,

¹¹ Some have postulated that the hippocampus plays a central role in consciousness (e.g. Clark & Squire, 1998). However, recent evidence suggests that the hippocampus also contributes to implicit learning (Chun & Phelps, 1999) and can be activated by subliminal stimuli such as novel faces (Elliott & Dolan, 1998). Thus, it seems more likely that the hippocampus and surrounding cortices support a modular system that at any given time may or may not be mobilized into the workspace. Gray (1994) proposes that one possible role for the hippocampus, in relation to consciousness, is to serve as a 'novelty detector' that automatically draws attention when the organism is confronted with an unpredicted situation.

Kerszberg, & Changeux, 1998). Impairments of those ascending systems may thus cause a global alteration of conscious workspace activity. For instance, upper brainstem lesions affecting reticular ascending systems frequently cause vigilance impairments or coma (see Parvizi & Damasio, this volume). Interestingly, brain imaging of patients in this vegetative state has revealed a partial preservation of cortical activation, for instance to speech stimuli (de Jong, Willemsen, & Paans, 1997; Menon et al., 1998), indicating that modular processors may still be partially functional at a subliminal level.

6. Final remarks

The present chapter was aimed at introducing the cognitive neuroscience of consciousness and proposing a few testable hypotheses about its cerebral substrates. While we think that a promising synthesis is now emerging, based on the concepts of global workspace, dynamic mobilization, attentional amplification, and frontal circuitry, some readers may feel that those ideas hardly scratch the surface. What about the so-called ‘hard problems’ posed by concepts such as voluntary action, free will, qualia, the sense of self, or the evolution of consciousness? Our personal view is that, in the present state of our methods, trying to address those problems head-on can actually impede rather than facilitate progress (but see Block, this volume, for a different view). We believe that many of these problems will be found to dissolve once a satisfactory framework for consciousness is achieved. In this conclusion, we examine how such a dissolution might proceed.

6.1. *Voluntary action and free will*

The hypothesis of an attentional control of behavior by supervisory circuits including AC and PFC, above and beyond other more automatized sensorimotor pathways, may ultimately provide a neural substrate for the concepts of voluntary action and free will (Posner, 1994). One may hypothesize that subjects label an action or a decision as ‘voluntary’ whenever its onset and realization are controlled by higher-level circuitry and are therefore easily modified or withheld, and as ‘automatic’ or ‘involuntary’ if it involves a more direct or hardwired command pathway (Passingham, 1993). One particular type of voluntary decision, mostly found in humans, involves the setting of a goal and the selection of a course of action through the serial examination of various alternatives and the internal evaluation of their possible outcomes. This conscious decision process, which has been partially simulated in neural network models (Dehaene & Changeux, 1991, 1997), may correspond to what subjects refer to as ‘exercising one’s free will’. Note that under this hypothesis free will characterizes a certain type of decision-making algorithm and is therefore a property that applies at the cognitive or systems level, not at the neural or implementation level. This approach may begin to address the old philosophical issue of free will and determinism. Under our interpretation, a physical system whose successive states unfold according to a deterministic rule can

still be described as having free will, if it is able to represent a goal and to estimate the outcomes of its actions before initiating them.¹²

6.2. *Qualia and phenomenal consciousness*

According to the workspace hypothesis, a large variety of perceptual areas can be mobilized into consciousness. At a microscopic scale, each area in turn contains a complex anatomical circuitry that can support a diversity of activity patterns. The repertoire of possible contents of consciousness is thus characterized by an enormous combinatorial diversity: each workspace state is ‘highly differentiated’ and of ‘high complexity’, in the terminology of Tononi and Edelman (1998). Thus, the flux of neuronal workspace states associated with a perceptual experience is vastly beyond accurate verbal description or long-term memory storage. Furthermore, although the major organization of this repertoire is shared by all members of the species, its details result from a developmental process of epigenesis and are therefore specific to each individual. Thus, the contents of perceptual awareness are complex, dynamic, multi-faceted neural states that cannot be memorized or transmitted to others in their entirety. These biological properties seem potentially capable of substantiating philosophers’ intuitions about the ‘qualia’ of conscious experience, although considerable neuroscientific research will be needed before they are thoroughly understood.

To put this argument in a slightly different form, the workspace model leads to a distinction between three levels of accessibility. Some information encoded in the nervous system is permanently inaccessible (set I_1). Other information is in contact with the workspace and could be consciously amplified if it was attended to (set I_2). However, at any given time, only a subset of the latter is mobilized into the workspace (set I_3). We wonder whether these distinctions may suffice to capture the intuitions behind Ned Block’s (Block, 1995; see also Block, this volume) definitions of phenomenal (P) and access (A) consciousness. What Block sees as a difference in essence could merely be a qualitative difference due to the discrepancy between the size of the potentially accessible information (I_2) and the paucity of information that can actually be reported at any given time (I_3). Think, for instance, of Sperling’s experiment in which a large visual array of letters seems to be fully visible, yet only a very small subset can be reported. The former may give rise to the intuition of a rich phenomenological world – Block’s P-consciousness – while the latter corresponds to what can be selected, amplified, and passed on to other processes (A-consciousness). Both, however, would be facets of the same underlying phenomenon.

6.3. *Sense of self and reflexive consciousness*

Among the brain’s modular processors, some do not extract and process signals

¹² This argument goes back to Spinoza: “men are mistaken in thinking themselves free; their opinion is made up of consciousness of their own actions, and ignorance of the causes by which they are conditioned. Their idea of freedom, therefore, is simply their ignorance of any cause for their actions.” (Ethics, II, 35).

from the environment, but rather from the subject's own body and brain. Each brain thus contains multiple representations of itself and its body at several levels (Damasio, 1999). The physical location of our body is encoded in continuously updated somatic, kinesthetic, and motor maps. Its biochemical homeostasis is represented in various subcortical and cortical circuits controlling our drives and emotions. We also represent ourselves as a person with an identity (presumably involving face and person-processing circuits of the inferior and anterior temporal lobes) and an autobiography encoded in episodic memory. Finally, at a higher cognitive level, the action perception, verbal reasoning, and 'theory of mind' modules that we apply to interpret and predict other people's actions may also help us make sense of our own behavior (Fletcher et al., 1995; Gallese, Fadiga, Fogassi, & Rizzolatti, 1996; Weiskrantz, 1997). All of those systems are modular, and their selective impairment may cause a wide range of neuropsychological deficits involving misperception of oneself and others (e.g. delusions of control, Capgras and Fregoli syndrome, autism). We envisage that the bringing together of these modules into the conscious workspace may suffice to account for the subjective sense of self. Once mobilized into the conscious workspace, the activity of those 'self-coding' circuits would be available for inspection by many other processes, thus providing a putative basis for reflexive or higher-order consciousness.

6.4. The evolution of consciousness

Any theory of consciousness must address its emergence in the course of phylogenesis. The present view associates consciousness with a unified neural workspace through which many processes can communicate. The evolutionary advantages that this system confers to the organism may be related to the increased independence that it affords. The more an organism can rely on mental simulation and internal evaluation to select a course of action, instead of acting out in the open world, the lower are the risks and the expenditure of energy. By allowing more sources of knowledge to bear on this internal decision process, the neural workspace may represent an additional step in a general trend towards an increasing internalization of representations in the course of evolution, whose main advantage is the freeing of the organism from its immediate environment.

This evolutionary argument implies that 'having consciousness' is not an all-or-none property. The biological substrates of consciousness in human adults are probably also present, but probably in partial form in other species (or in young children or brain-lesioned patients). It is therefore a partially arbitrary question as to whether we want to extend the use of the term 'consciousness' to them. For instance, several mammals, and possibly even young human children, exhibit greater brain modularity than human adults (Cheng & Gallistel, 1986; Hermer & Spelke, 1994). Yet they also show intentional behavior, partially reportable mental states, some working memory ability – but perhaps no theory of mind. Do they have consciousness, then? Our hope is that once a detailed cognitive and neural theory of the various aspects of consciousness is available, the vacuity of this question will become obvious.

References

- Aglioti, S., DeSouza, J. F., & Goodale, M. A. (1995). Size-contrast illusions deceive the eye but not the hand. *Current Biology*, 5 (6), 679–685.
- Artiges, E., Salame, P., Recasens, C., Poline, J. B., Attar-Levy, D., De La Raillere, A., Pailere-Martinot, M. L., Danion, J. M., & Martinot, J. L. (2000). Working memory control in patients with schizophrenia: a PET study during a random number generation task. *American Journal of Psychiatry*, 157 (9), 1517–1519.
- Baars, B. J. (1989). *A cognitive theory of consciousness*. Cambridge: Cambridge University Press.
- Baddeley, A. D. (1986). *Working memory*. Oxford: Clarendon Press.
- Bar, M., & Biederman, I. (1999). Localizing the cortical region mediating visual awareness of object identity. *Proceedings of the National Academy of Sciences USA*, 96 (4), 1790–1793.
- Bauer, R. M. (1984). Autonomic recognition of names and faces in prosopagnosia: a neuropsychological application of the Guilty Knowledge Test. *Neuropsychologia*, 22 (4), 457–469.
- Beckers, G., & Zeki, S. (1995). The consequences of inactivating areas V1 and V5 on visual motion perception. *Brain*, 118 (Pt. 1), 49–60.
- Bejjani, B. P., Damier, P., Arnulf, I., Thivard, L., Bonnet, A. M., Dormont, D., Cornu, P., Pidoux, B., Samson, Y., & Agid, Y. (1999). Transient acute depression induced by high-frequency deep-brain stimulation. *New England Journal of Medicine*, 340 (19), 1476–1480.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18 (2), 227–287.
- Bornstein, R. F., & D'Agostino, P. R. (1992). Stimulus recognition and the mere exposure effect. *Journal of Personality and Social Psychology*, 63 (4), 545–552.
- Carter, C. S., Braver, T. S., Barch, D., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, 280, 747–749.
- Chalmers, D. (1996). *The conscious mind*. New York: Oxford University Press.
- Changeux, J. P., & Dehaene, S. (1989). Neuronal models of cognitive functions. *Cognition*, 33, 63–109.
- Chelazzi, L., Duncan, J., Miller, E. K., & Desimone, R. (1998). Responses of neurons in inferior temporal cortex during memory-guided visual search. *Journal of Neurophysiology*, 80 (6), 2918–2940.
- Cheng, K., & Gallistel, C. R. (1986). A purely geometric module in the rat's spatial representation. *Cognition*, 23, 149–178.
- Chochon, F., Cohen, L., van de Moortele, P. F., & Dehaene, S. (1999). Differential contributions of the left and right inferior parietal lobules to number processing. *Journal of Cognitive Neuroscience*, 11, 617–630.
- Chun, M. M., & Phelps, E. A. (1999). Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. *Nature Neuroscience*, 2 (9), 844–847.
- Clark, R. E., & Squire, L. R. (1998). Classical conditioning and brain systems: the role of awareness. *Science*, 280 (5360), 77–81.
- Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., & Smith, E. E. (1997). Temporal dynamics of brain activation during a working memory task. *Nature*, 386, 604–608.
- Cohen, L., & Dehaene, S. (1998). Competition between past and present. Assessing and explaining verbal perseverations. *Brain*, 121, 1641–1659.
- Corbetta, M., Miezin, F. M., Dobmeyer, S., Smulman, G. L., & Petersen, S. E. (1991). Selective and divided attention during visual discriminations of shape color and speed: functional anatomy by positron emission tomography. *Journal of Neuroscience*, 11, 2383–2402.
- Crick, F., & Koch, C. (1995). Are we aware of neural activity in primary visual cortex? *Nature*, 375, 121–123.
- Damasio, A. (1999). *The feeling of what happens*. New York: Harcourt Brace.
- Daprati, E., & Gentilucci, M. (1997). Grasping an illusion. *Neuropsychologia*, 35 (12), 1577–1582.
- 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 USA*, 94, 13293–13298.

- 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 USA*, *95*, 14529–14534.
- Dehaene, S., Naccache, L., Le Clec'H, G., Koechlin, E., Mueller, M., Dehaene-Lambertz, G., van de Moortele, P. F., & Le Bihan, D. (1998). Imaging unconscious semantic priming. *Nature*, *395*, 597–600.
- Dehaene, S., Posner, M. I., & Tucker, D. M. (1994). Localization of a neural system for error detection and compensation. *Psychological Science*, *5*, 303–305.
- de Jong, B. M., Willemsen, A. T., & Paans, A. M. (1997). Regional cerebral blood flow changes related to affective speech presentation in persistent vegetative state. *Clinical Neurology and Neurosurgery*, *99* (3), 213–216.
- Dennett, D. C. (1992). *Consciousness explained*. London: Penguin.
- Driver, J., & Mattingley, J. B. (1998). Parietal neglect and visual awareness. *Nature Neuroscience*, *1* (1), 17–22.
- Driver, J., & Vuilleumier, P. (2001 this issue). Perceptual awareness and its loss in unilateral neglect and extinction. *Cognition*, *79*, 39–88.
- Eimer, M., & Schlaghecken, F. (1998). Effects of masked stimuli on motor activation: behavioral and electrophysiological evidence. *Journal of Experimental Psychology: Human Perception and Performance*, *24* (6), 1737–1747.
- Elliott, R., & Dolan, R. J. (1998). Neural response during preference and memory judgments for subliminally presented stimuli: a functional neuroimaging study. *Journal of Neuroscience*, *18* (12), 4697–4704.
- Ericsson, K. A., & Simon, H. A. (1993). *Protocol analysis: verbal reports as data*. Cambridge, MA: MIT Press.
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, *1* (1), 1–47.
- Ffytche, D. H., Howard, R. J., Brammer, M. J., David, A., Woodruff, P., & Williams, S. (1998). The anatomy of conscious vision: an fMRI study of visual hallucinations. *Nature Neuroscience*, *1* (8), 738–742.
- Fiset, P., Paus, T., Daloze, T., Plourde, G., Meuret, P., Bonhomme, V., Hajj-Ali, N., Backman, S. B., & Evans, A. C. (1999). Brain mechanisms of propofol-induced loss of consciousness in humans: a positron emission tomographic study. *Journal of Neuroscience*, *19* (13), 5506–5513.
- Fletcher, P. C., Happé, F., Frith, U., Baker, S. C., Donlan, R. J., Frackowiak, R. S. J., & Frith, C. D. (1995). Other minds in the brain: a functional imaging study of theory of mind in story comprehension. *Cognition*, *57*, 109–128.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Fried, I., Wilson, C. L., MacDonald, K. A., & Behnke, E. J. (1998). Electric current stimulates laughter. *Nature*, *391* (6668), 650.
- Fries, P., Roelfsema, P. R., Engel, A. K., Konig, P., & Singer, W. (1997). Synchronization of oscillatory responses in visual cortex correlates with perception in interocular rivalry. *Proceedings of the National Academy of Sciences USA*, *94* (23), 12699–12704.
- Friston, K. J., Tononi, G., Reeke, G. N., Sporns, O., & Edelman, G. M. (1994). Value-dependent selection in the brain: simulation in a synthetic neural model. *Neuroscience*, *59*, 229–243.
- Frith, C., Perry, R., & Lumer, E. (1999). The neural correlates of conscious experience: an experimental framework. *Trends in Cognitive Science*, *3*, 105–114.
- Fuster, J. M. (1989). *The prefrontal cortex*. New York: Raven Press.
- Fuster, J. M., Bauer, R. H., & Jervey, J. P. (1985). Functional interactions between inferotemporal and prefrontal cortex in a cognitive task. *Brain Research*, *330* (2), 299–307.
- Gallese, V., Fadiga, L., Fogassi, L., & Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain*, *119* (Pt. 2), 593–609.
- Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, *3* (2), 191–197.
- Gazzaniga, M. S., LeDoux, J. E., & Wilson, D. H. (1977). Language, praxis, and the right hemisphere: clues to some mechanisms of consciousness. *Neurology*, *27* (12), 1144–1147.

- Gentilucci, M., Chieffi, S., Deprati, E., Saetti, M. C., & Toni, I. (1996). Visual illusion and action. *Neuropsychologia*, *34* (5), 369–376.
- Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational knowledge. In F. Plum, & V. Mountcastle (Eds.), *Handbook of physiology* (Vol. 5, pp. 373–417). Bethesda, MD: American Physiological Society.
- Goldman-Rakic, P. S. (1988). Topography of cognition: parallel distributed networks in primate association cortex. *Annual Review of Neuroscience*, *11*, 137–156.
- Grafton, S. T., Hazeltine, E., & Ivry, R. (1995). Functional mapping of sequence learning in normal humans. *Journal of Cognitive Neuroscience*, *7*, 497–510.
- Gray, J. A. (1994). The contents of consciousness: a neuropsychological conjecture. *Behavioral and Brain Sciences*, *18*, 659–722.
- Greenwald, A. G. (1996). Three cognitive markers of unconscious semantic activation. *Science*, *273* (5282), 1699–1702.
- He, S., Cavanagh, P., & Intriligator, J. (1996). Attentional resolution and the locus of visual awareness. *Nature*, *383* (6598), 334–337.
- Hermer, L., & Spelke, E. S. (1994). A geometric process for spatial reorientation in young children. *Nature*, *370*, 57–59.
- Holender, D. (1986). Semantic activation without conscious identification in dichotic listening, parafoveal vision and visual masking: a survey and appraisal. *Behavioral and Brain Sciences*, *9*, 1–23.
- Hu, Y., Eagleson, R., & Goodale, M. A. (1999). The effects of delay on the kinematics of grasping. *Experimental Brain Research*, *126* (1), 109–116.
- Jacoby, L. L. (1991). A process dissociation framework: separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*, 513–541.
- Klinger, M. R., & Greenwald, A. G. (1995). Unconscious priming of association judgments. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *21* (3), 569–581.
- 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.
- Köhler, S., & Moscovitch, M. (1997). Unconscious visual processing in neuropsychological syndromes: a survey of the literature and evaluation of models of consciousness. In M. D. Rugg (Ed.), *Cognitive neuroscience* (pp. 305–373). Hove: Psychology Press.
- Konishi, S., Nakajima, K., Uchida, I., Kameyama, M., Nakahara, K., Sekihara, K., & Miyashita, Y. (1998). Transient activation of inferior prefrontal cortex during cognitive set shifting. *Nature Neuroscience*, *1*, 80–84.
- Kosslyn, S. M., & Koenig, O. (1992). *Wet mind: the new cognitive neuroscience*. New York: Macmillan.
- Leopold, D. A., & Logothetis, N. K. (1999). Multistable phenomena: changing views in perception. *Trends in Cognitive Science*, *3*, 254–264.
- Lhermitte, F. (1983). “Utilization behaviour” and its relation to lesions of the frontal lobe. *Brain*, *106*, 237–255.
- Luck, S. J., Vogel, E. K., & Shapiro, K. L. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature*, *383* (6601), 616–618.
- Lumer, E. D., Friston, K. J., & Rees, G. (1998). Neural correlates of perceptual rivalry in the human brain. *Science*, *280*, 1930–1934.
- Lumer, E. D., & Rees, G. (1999). Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proceedings of the National Academy of Sciences USA*, *96* (4), 1669–1673.
- Mack, A., & Rock, I. (1998). *Inattention blindness*. Cambridge, MA: MIT Press.
- Marcel, A. J. (1983). Conscious and unconscious perception: experiments on visual masking and word recognition. *Cognitive Psychology*, *15*, 197–237.
- McGlinchey-Berroth, R., Milberg, W. P., Verfaellie, M., Alexander, M., & Kilduff, P. (1993). Semantic priming in the neglected field: evidence from a lexical decision task. *Cognitive Neuropsychology*, *10*, 79–108.
- McIntosh, A. R., Rajah, M. N., & Lobaugh, N. J. (1999). Interactions of prefrontal cortex in relation to awareness in sensory learning. *Science*, *284* (5419), 1531–1533.

- Menon, D. K., Owen, A. M., Williams, E. J., Minhas, P. S., Allen, C. M., Boniface, S. J., & Pickard, J. D. (1998). Cortical processing in persistent vegetative state. *Lancet*, 352 (9123), 200.
- Merikle, P. M. (1992). Perception without awareness: critical issues. *American Psychologist*, 47, 792–796.
- Merikle, P. M., & Joordens, S. (1997). Parallels between perception without attention and perception without awareness. *Consciousness and Cognition*, 6 (2–3), 219–236.
- Merikle, P. M., Joordens, S., & Stolz, J. A. (1995). Measuring the relative magnitude of unconscious influences. *Consciousness and Cognition*, 4, 422–439.
- Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience*, 16 (16), 5154–5167.
- Morris, J. S., Öhman, A., & Dolan, R. J. (1998). Conscious and unconscious emotional learning in the human amygdala. *Nature*, 393, 467–470.
- Morris, J. S., Öhman, A., & Dolan, R. J. (1999). A subcortical pathway to the right amygdala mediating “unseen” fear. *Proceedings of the National Academy of Sciences USA*, 96 (4), 1680–1685.
- 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.), *Conscious and non-conscious information processing* (Vol. XV, pp. 123–150). Cambridge, MA: MIT Press.
- O’Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change-blindness as a result of ‘mudsplashes’. *Nature*, 398 (6722), 34.
- 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 USA*, 87, 256–259.
- Passingham, R. (1993). *The frontal lobes and voluntary action* (Vol. 21). New York: Oxford University Press.
- Paus, T., Koski, L., Caramanos, Z., & Westbury, C. (1998). Regional differences in the effects of task difficulty and motor output on blood flow response in the human anterior cingulate cortex: a review of 107 PET activation studies. *NeuroReport*, 9, R37–R47.
- Penrose, R. (1990). *The emperor’s new mind. Concerning computers, minds, and the laws of physics*. London: Vintage Books.
- Pöppel, E., Held, R., & Frost, D. (1973). Residual visual function after brain wounds involving the central visual pathways in man. *Nature*, 243 (405), 295–296.
- Posner, M. I. (1994). Attention: the mechanisms of consciousness. *Proceedings of the National Academy of Sciences USA*, 91, 7398–7403.
- Posner, M. I., & Dehaene, S. (1994). Attentional networks. *Trends in Neuroscience*, 17, 75–79.
- Raichle, M. E., Fiez, J. A., Videen, T. O., MacLeod, A. K., Pardo, J. V., Fox, P. T., & Petersen, S. E. (1994). Practice-related changes in human brain functional anatomy during non-motor learning. *Cerebral Cortex*, 4, 8–26.
- Ramachandran, V. S. (1992). Filling in the blind spot. *Nature*, 356 (6365), 115.
- Rao, S. C., Rainer, G., & Miller, E. K. (1997). Intergration of what and where in the primate prefrontal cortex. *Science*, 276 (5313), 821–824.
- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: an attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, 18 (3), 849–860.
- Rees, G., Wojciulik, E., Clarke, K., Husain, M., Frith, C., & Driver, J. (2000). Unconscious activation of visual cortex in the damaged right hemisphere of a parietal patient with extinction. *Brain*, 123 (Pt. 8), 1624–1633.
- Renault, B., Signoret, J. L., Debruille, B., Breton, F., & Bolgert, F. (1989). Brain potentials reveal covert facial recognition in prosopagnosia. *Neuropsychologia*, 27 (7), 905–912.
- Rodriguez, E., George, N., Lachaux, J. P., Martinerie, J., Renault, B., & Varela, F. J. (1999). Perception’s shadow: long-distance synchronization of human brain activity. *Nature*, 397 (6718), 430–433.
- Rolls, E. T., & Tovee, M. J. (1994). Processing speed in the cerebral cortex and the neurophysiology of visual masking. *Proceedings of the Royal Society of London, Series B, Biological Sciences*, 257 (1348), 9–15.

- Rolls, E. T., Tovee, M. J., & Panzeri, S. (1999). The neurophysiology of backward visual masking: information analysis. *Journal of Cognitive Neuroscience*, *11* (3), 300–311.
- Rueckert, L., Lange, N., Partiot, A., Appollonio, I., Litvar, I., Le Bihan, D., & Grafman, J. (1996). Visualizing cortical activation during mental calculation with functional MRI. *NeuroImage*, *3*, 97–103.
- Rugg, M. D., Mark, R. E., Walla, P., Schloerscheidt, A. M., Birch, C. S., & Allan, K. (1998). Dissociation of the neural correlates of implicit and explicit memory. *Nature*, *392* (6676), 595–598.
- Sahraie, A., Weiskrantz, L., Barbur, J. L., Simmons, A., Williams, S. C. R., & Brammer, M. J. (1997). Pattern of neuronal activity associated with conscious and unconscious processing of visual signals. *Proceedings of the National Academy of Sciences USA*, *94*, 9406–9411.
- Salzman, C. D., Britten, K. H., & Newsome, W. T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. *Nature*, *346* (6280), 174–177.
- Schacter, D. L., Buckner, R. L., & Koutstaal, W. (1998). Memory, consciousness and neuroimaging. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *353* (1377), 1861–1878.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing. 1. Detection, search, and attention. *Psychological Review*, *84*, 1–66.
- Searle, J. R. (1998). How to study consciousness scientifically. *Philosophical Transactions of the Royal Society of London, Series B*, *353*, 1935–1942.
- Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London, Series B*, *298*, 199–209.
- Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge: Cambridge University Press.
- Shallice, T., Burgess, P. W., Schon, F., & Baxter, D. M. (1989). The origins of utilization behaviour. *Brain*, *112*, 1587–1598.
- Silbersweig, D. A., Stern, E., Frith, C. D., Cahill, C., Holmes, A., Grootoink, S., Seaward, J., McKenna, P., Chua, S. E., Schnoor, L., Jones, T., & Frackowiak, R. S. J. (1995). A functional neuroanatomy of hallucinations in schizophrenia. *Nature*, *378*, 176–179.
- Smith, P. H., Joris, P. X., & Yin, T. C. (1993). Projections of physiologically characterized spherical bushy cell axons from the cochlear nucleus of the cat: evidence for delay lines to the medial superior olive. *Journal of Comparative Neurology*, *331* (2), 245–260.
- Sperling, G. (1960). The information available in brief visual presentation. *Psychological Monographs*, *74*, 1–29.
- Srinivasan, R., Russell, D. P., Edelman, G. M., & Tononi, G. (1999). Increased synchronization of neuromagnetic responses during conscious perception. *Journal of Neuroscience*, *19* (13), 5435–5448.
- Tallon-Baudry, C., & Bertrand, O. (1999). Oscillatory gamma activity in humans and its role in object representation. *Trends in Cognitive Science*, *3*, 151–162.
- Tiitinen, H., May, P., Reinikainen, K., & Naatanen, R. (1994). Attentive novelty detection in humans is governed by pre-attentive sensory memory. *Nature*, *372* (6501), 90–92.
- Tong, F., Nakayama, K., Vaughan, J. T., & Kanwisher, N. (1998). Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron*, *21* (4), 753–759.
- Tononi, G., & Edelman, G. M. (1998). Consciousness and complexity. *Science*, *282* (5395), 1846–1851.
- Tootell, R. B. H., Reppas, J. B., Dale, A. M., Look, R. B., Sereno, M. I., Malach, R., Brady, T. J., & Rosen, B. R. (1995). Visual motion aftereffect in human cortical area MT revealed by functional magnetic resonance imaging. *Nature*, *375*, 139–141.
- Tranel, D., & Damasio, A. R. (1985). Knowledge without awareness: an autonomic index of facial recognition by prosopagnosics. *Science*, *228* (4706), 1453–1454.
- Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*, 97–136.
- 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* (6), 1656–1674.
- Walsh, V., Ellison, A., Battelli, L., & Cowey, A. (1998). Task-specific impairments and enhancements induced by magnetic stimulation of human visual area V5. *Proceedings of the Royal Society of London, Series B, Biological Sciences*, *265* (1395), 537–543.

- Watson, J. D. G., Myers, R., Frackowiak, R. S. J., Hajnal, J. V., Woods, R. P., Mazziotta, J. C., Shipp, S., & Zeki, S. (1993). Area V5 of the human brain: evidence from a combined study using positron emission tomography and magnetic resonance imaging. *Cerebral Cortex*, 3, 79–94.
- Weiskrantz, L. (1997). *Consciousness lost and found: a neuropsychological exploration*. New York: Oxford University Press.
- Whalen, P. J., Rauch, S. L., Etcoff, N. L., McInerney, S. C., Lee, M. B., & Jenike, M. A. (1998). Masked presentations of emotional facial expressions modulate amygdala activity without explicit knowledge. *Journal of Neuroscience*, 18, 411–418.
- Yantis, S., & Jonides, J. (1984). Abrupt visual onsets and selective attention: evidence from visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 10 (5), 601–621.
- Yantis, S., & Jonides, J. (1996). Attentional capture by abrupt onsets: new perceptual objects or visual masking? *Journal of Experimental Psychology: Human Perception and Performance*, 22 (6), 1505–1513.
- Young, A. W. (1992). Face recognition impairments. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, 335 (1273), 47–54.
- Young, M. P., Scannell, J. W., O’Neill, M. A., Hilgetag, C. C., Burns, G., & Blakemore, C. (1995). Non-metric multidimensional scaling in the analysis of neuroanatomical connection data and the organization of the primate cortical visual system. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, 348 (1325), 281–308.
- Zeki, S. (1993). *A vision of the brain*. London: Blackwell.
- Zeki, S., & Ffytche, D. H. (1998). The Riddoch syndrome: insights into the neurobiology of conscious vision. *Brain*, 121 (Pt. 1), 25–45.
- Zeki, S., Watson, J. D., & Frackowiak, R. S. (1993). Going beyond the information given: the relation of illusory visual motion to brain activity. *Proceedings of the Royal Society of London, Series B, Biological Sciences*, 252 (1335), 215–222.