

July 4th, 2003

To appear in Gazzaniga et al. (Ed)

"The cognitive neurosciences, 3rd edition"

## **Neural Mechanisms for Access to Consciousness**

Stanislas Dehaene <sup>1</sup> and Jean-Pierre Changeux <sup>2</sup>

<sup>1</sup> Unité INSERM 562 "Cognitive Neuroimaging", IFR 49, Service Hospitalier  
Frédéric Joliot, CEA/DSV, 4 place du Général Leclerc, 91401 Orsay cedex, France.

<sup>2</sup> CNRS URA 2182, "Récepteurs et Cognition", Institut Pasteur, 25 Rue du Dr  
Roux, 75015 Paris, France

## ***Introduction: the challenge of a science of consciousness***

Understanding consciousness has become the ultimate intellectual challenge of this new millennium. Even if philosophers now accept the notion that it is a “real, natural, biological phenomenon literally located in the brain” (Revonsuo, 2001), a view in harmony with the neuroscientist conception that “consciousness is entirely caused by neurobiological processes and realized in brain structures” (Changeux, 1983; Crick, 1994; Edelman, 1989), the real issue becomes: how to elaborate a science of consciousness?

This challenging problem raises two questions. A first one is how to empirically define experimental paradigms in order to delineate a relevant and ultimately causal relationship between subjective phenomena and objective measurements of neural activity. Cognitive psychologists have now defined a variety of minimal experimental protocols which allow a fair comparison between conscious and non-conscious processing of information (see Baars, 1989). Moreover, brain imaging and EEG methods in humans and monkeys as well as electrophysiological recordings at the single cell level in awake monkeys, provide access to reliable neural correlates of conscious versus nonconscious perceptual processes (see Lamme, 2003). In this chapter, we will restrict the discussion to experimental research on the subliminal processing of visual stimuli under masking and attentional blink paradigms.

A second question is more conceptual. Given the broad diversity of methods required to evaluate neural activity and the extreme, often unresolved, complexity of the neuronal architectures involved, it seems risky to draw conclusions simply on the basis of intuitive reasoning. In our opinion, in the present state of affair, a theoretical framework

appears necessary for an in-depth understanding of conscious phenomena. It will, for instance, consist in formal models, expressed in terms of neuronal networks, that link together the molecular, neuronal, behavioral and subjective data in a coherent, non contradictory though minimal form (Changeux & Dehaene, 1989). Such “bridging laws” implemented as formal automata should simultaneously account for the available data and produce experimentally testable predictions at all of those levels. Being minimal, they are not anticipated to give an exhaustive description of reality but, even if wrong, may give rise to novel theories and, as such, contribute to the progress of knowledge.

The context of consciousness is so broad and diverse and the issues often so muddled (see Chalmers, this volume), that we shall deliberately limit ourselves, in this review, to only one aspect of consciousness, the notion of *conscious access*. This is the observation that a piece of information, once conscious, becomes broadly available for multiple processes including action planning, voluntary redirection of attention, memory, evaluation, and verbal or non-verbal report. Like others (Weiskrantz, 1997), we emphasize *reportability* as a key property of conscious representations. This discussion will aim at characterizing the crucial differences between those aspects of neural activity that can be reported by a subject, and those that cannot. According to some philosophers, this constitutes an “easy problem” and is irrelevant to the more central issues of phenomenality and self-awareness (e.g. Block, 1995). Our view, however, is that conscious access is one of the few empirically tractable problems presently accessible to an authentic scientific investigation. We further hope that an understanding of the neural processes that lead to overt report will eventually result in a theory of covert acts of self-

report, and thus may ultimately contribute to an explanation of the nature of our private phenomenal world.

In the present chapter, we will examine conscious access in the framework of an integrative theory based on the hypothesis of a “conscious *neuronal* workspace” (Dehaene & Changeux, 2000; Dehaene, Kerszberg, & Changeux, 1998; Dehaene & Naccache, 2001; Dehaene, Sergent, & Changeux, 2003). The model emphasizes the role of distributed neurons with long-distance connections, particularly dense in prefrontal, cingulate, and parietal regions, which are capable of interconnecting multiple specialized processors and can broadcast signals at the brain scale in a spontaneous and sudden manner. Those neurons form what is referred to here as a conscious “global neuronal workspace” (see Baars, 1989), that breaks the modularity of the nervous system and allows the broadcasting of information to multiple neural targets. This broadcasting creates a global availability that, according to our hypothesis, is experienced as consciousness and results in reportability. The discussion will include the direct comparison of presently available experimental data with the theory and stress novel predictions concerning the neural correlates of access consciousness.

At this stage, it may be worth stressing a few basic distinctions. A first one separates the notions of *state* of consciousness from that of *content* of consciousness. The English language distinguishes an intransitive meaning of consciousness (e.g. “the patient was still conscious”) and a transitive meaning (e.g. “I was conscious *of* motion). The former refers to the state of consciousness, usually considered as a continuous variable (coma, sleep, drowsiness, awake state...). The latter refers to the temporary selection of a well-delimited content as the focus of conscious attention. The global neuronal

workspace is essentially a theory of conscious content. It specifies the neural conditions under which a given representation is made potentially available to a broad variety of neural processes, thus giving rise to a subjective feeling of conscious access. A prerequisite, however, is that the neuronal workspace within which this global broadcasting occurs is available in an appropriate state of awakeness or readiness. We speculate that the graded states of consciousness correspond to different levels of spontaneous thalamo-cortical homeostatic regulation contributing to a “conscious milieu” which includes long-distance cortical neurons under the influence of ascending neuromodulatory inputs from, for instance, cholinergic, noradrenergic, and/or dopaminergic neurons from the basal forebrain and brainstem. In the last few years, considerable progress has been made in identifying the electrophysiological correlates of such global state changes, which are consistent with a graded modulation of thalamic and fronto-cingulate networks (Llinas, Ribary, Contreras, & Pedroarena, 1998; Paus, 2000). A detailed review of those findings is beyond the scope of the present chapter, which will be limited to how a specific content gains access to consciousness (see however Schiff, this volume).

### ***The neuronal workspace hypothesis***

The concept of a “global neuronal workspace” (Dehaene, Kerszberg & Changeux, 1998; Dehaene & Naccache, 2001) is historically rooted in a long neuropsychological tradition, dating back to Hughlings Jackson and perpetuated among others by Baddeley, Shallice, Mesulam or Posner, which emphasizes the hierarchical organization of the brain and separates lower automatized systems from increasingly higher and more autonomous supervisory executive systems. It also builds upon Fodor’s distinction between the

vertical “modular faculties” and a distinct “isotropic central and horizontal system” capable of sharing information across modules. Finally, it relates to Baars’s cognitive theory of consciousness, which distinguishes a vast array of unconscious specialized processors running in parallel, and a single limited-capacity serial “workspace” that allowed them to exchange information (Baars, 1989).

Baars, however, did not specify how the psychological construct of a conscious workspace could be implemented in terms of neuronal networks. By contrast, our views arose progressively from the design of computational neural network models that aimed at specifying the contribution of prefrontal cortex to increasingly higher cognitive tasks (Dehaene & Changeux, 1989, 1991; Dehaene & Changeux, 1997; Dehaene, Kerszberg et al., 1998). We successively considered how a network could retain an active memory across a long delay (Dehaene & Changeux, 1989), how it could encode abstract rules that could be selected from external or internal rewards (Dehaene & Changeux, 1991), and finally how networks based on those principles could pass complex planning tasks such as the Tower of London test or the Stroop test (Dehaene & Changeux, 1997; Dehaene, Kerszberg et al., 1998). The neuronal workspace model is the last development of the neuronal architectures that we proposed to address those specific problems.

## **Two computational spaces in the brain**

The neuronal workspace hypothesis distinguishes two computational spaces in the brain, each characterized by a distinct pattern of connectivity (figure 1).

**1. The network of processors.** Subcortical networks and most of the cortex can be viewed as a collection of specialized processors, each attuned to the processing of a particular type of information. Processors vary widely in complexity, from the

elementary line segment detectors in area V1 or the motion processors in area MT, to the “visual word form” processor in the human fusiform gyrus, or the “mirror-neuron” system in area F5. In spite of this diversity, processors share characteristics of specialization, automaticity, and fast feedforward processing. Their function is made possible by a limited number of local or medium-range connections that bring to each processor the “encapsulated” inputs necessary to its function.

**2. The global neuronal workspace.** We postulate the existence of a distinct set of cortical “workspace” neurons characterized by their ability to send and receive projections to many distant areas through long-range excitatory axons. These neurons therefore no longer obey a principle of local, encapsulated connectivity, but rather break the modularity of the cortex by allowing many different processors to exchange information in a global and flexible manner. Information which is encoded in workspace neurons can be quickly made available to many brain systems, in particular the motor and speech-production processors for overt behavioral report. We hypothesize that the entry of inputs into this global workspace constitutes the neural basis of access to consciousness.

### **Top-down amplification and dynamic mobilization**

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

To enter consciousness, it is not sufficient for a processor to be activated; this activity must also be amplified and maintained over a sufficient duration for it to become accessible to multiple other processes. Without such « dynamic mobilization », a process may still contribute to cognitive performance, but only nonconsciously.

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

### **Criteria for conscious access**

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

1. **Active firing.** The object must be represented as a firing pattern of neurons.

There is, of course, considerable information which is already stored in the nervous system, in a latent form, for instance, in synaptic connections and weights, neurotransmitter release efficiencies, receptor densities, etc. The



model predicts that such information does not become conscious. It can only be read-out very indirectly through its contribution to neural firing.

2. **Long-distance connectivity.** The active neurons must possess a sufficient number of reciprocal anatomical connections to distributed workspace neurons, particularly in prefrontal, parietal and cingulate cortices. This criterion implies that the activity of many neurons, for instance in subcortical and brainstem nuclei, is excluded from conscious mobilization (e.g. circuits for respiration or emotion). In many cases, we only become aware of those circuits through their indirect effects on other representations, e.g. in somatic cortical areas.
3. **Dynamic mobilization.** At any given moment, workspace neurons can only sustain a single global representation, the rest of workspace neurons being inhibited. This implies that, out of the multiple active cortical representations that could become conscious, only one will receive the appropriate top-down amplification and be mobilized into consciousness. The other representations are temporarily nonconscious. It would only take a small reorientation of top-down signals to access them, but, according to our views, until this is achieved they do not participate in consciousness.

## **Consciousness and attention**

Within the framework of the workspace model, the relations between attention and consciousness can be clarified as follows. Top-down attentional selection and amplification is necessary for the access of a representation to consciousness. Thus, subjects cannot become aware of any sensory stimulus unless either (1) they are already

attending to the relevant cortical processors, or (2) the stimulus itself attracts top-down attention. The latter is possible because some specialized cortical systems, such as the frontal eye fields, can operate quickly and subliminally (Thompson & Schall, 1999), with the result of reorienting attention and thus forcing a change in the contents of the workspace.

Thus, two modes of conscious access to sensory information should be distinguished: a top-down mode in which workspace neurons become spontaneously activated and selectively amplify a possibly very small sensory signal; and a bottom-up mode in which initially unattended sensory signals carry sufficient strength as to cause a reorienting of top-down amplification towards them.

According to this view, then attention and consciousness cannot be equated for two reasons. First, the orientation of visuo-spatial attention results from the operation of cortical processors that may operate under voluntary conscious control, but also nonconsciously, as when attention is attracted by peripheral stimuli. Second, when attention is present, it may not always be sufficient for a stimulus to gain access to consciousness. In this situation, we expect attention to modulate the depth of subliminal processing, while still failing to make the stimuli conscious (Naccache, Blandin, & Dehaene, 2002, see below). Those distinctions are outlined in figure 2.

### **Workspace modulation and selection by reward**

Workspace neurons are assumed to be the targets of two different types of neuromodulatory inputs. First, workspace neurons display a constantly fluctuating spontaneous activity, whose *intensity* is modulated by ascending activating systems, for instance from cholinergic, noradrenergic and serotonergic nuclei in the brain stem,

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

### **Brain anatomy of the neuronal workspace**

The neuronal workspace hypothesis posits that, as a whole, the workspace neurons are reciprocally connected via long-distance axons to many if not all of the cortical processors, thus permitting locally available information to be brought into consciousness. Nevertheless, these neurons may be more densely accumulated in some areas than in others. Anatomically, long-range cortico-cortical tangential connections, including callosal connections, originate mostly from the pyramidal cells of layers II and III, which give or receive the so-called « association » efferents and afferents. Those layers are thicker in von Economo’s type 2 (dorsolateral prefrontal) and type 3 (inferior parietal) cortical structures. In the monkey, those areas entertain a very strong interconnection amongst themselves as well as with the anterior and posterior cingulate,

the association cortex of the superior temporal sulcus, and the parahippocampal region, thalamus and striatum (Goldman-Rakic, 1988). The high concentration of neurons with long-distance axons in those areas may explain why they frequently appear co-activated in neuroimaging studies of conscious effortful processing.

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

Our model leads to the prediction that long-distance connections have been the target of a recent evolutionary pressure in the course of hominization and are particularly developed in our species. In that respect, it can be noted that the relative anatomical expansion of cortical areas rich in long-axon neurons, such as the prefrontal cortex, may have contributed to important changes in the functional properties of the workspace (see Changeux, 2003). It is also noteworthy that a particular type of spindle cell, which establishes long-distance projections, is found in the anterior cingulate cortex of humans and great apes, but not other primates (Allman, Hakeem, Erwin, Nimchinsky, & Hof, 2001). Detailed anatomical studies of transcortical connectivity in the human brain have also revealed the presence of distant transcortical projections, that for instance link directly the right fusiform gyrus to multiple areas of the left-hemisphere including Broca's and Wernicke's areas (Di Virgilio & Clarke, 1997). It is anticipated that those key components of the verbal reportability system are connected to many cortical areas,

given the variety of percepts and concepts that we can name or understand through language.

### ***Comparison with data on subliminal processing of masked words***

How can one experimentally test the model's distinction between the substrates of conscious access and the considerable amount of neural activity which can occur spontaneously and nonconsciously within the specialized processors? Our approach has consisted in exploring paradigms in which sensory information is deliberately presented under subliminal conditions. By studying to what extent such information is processed, and what brain areas it contacts, one can progressively draw a negative picture, as it were, of which aspects of brain activity do not suffice to give rise to consciousness. Conversely, one may then ask which particular processes are associated with the crossing of the threshold for consciousness. Here, we briefly review those empirical findings and subsequently examine how they fit with the global neuronal workspace theory.

### **Neuroimaging of subliminal word processing**

In a classical psychological paradigm, the pattern masking of visual words, a word is flashed on a computer screen for a duration of less than 50 ms. If the word is presented alone, it typically remains readable with some effort. However, when the same word is preceded and followed by random geometrical shapes or letter at the same retinal location, it may become totally inaccessible to consciousness. In spite of this invisibility, behavioral priming experiments have repeatedly indicated that masked words are processed nonconsciously at orthographic, phonological, and possibly semantic levels

(see e.g. Forster & Davis, 1984; Greenwald, Draine, & Abrams, 1996; Neely & Kahan, 2001).

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

Using this method, several cortical stages of word processing have been shown to be activated by subliminal words (figure 3):

**Early visual activity.** Extrastriate visual areas reduce their activation in the subliminal repetition priming paradigm, but only if the prime and target words are repeated in the same case and font (Dehaene et al., 2001). Those areas are therefore thought to extract small features of the letter shapes.

**Visual word recognition.** Subliminal repetition priming has a major effect on the activation of a subarea of the left fusiform gyrus which has been termed the visual word form area (Cohen et al., 2000). Contrary to occipital extrastriate cortex, the fusiform visual word form area reduces its activation even when a word is repeated twice in a different case (e.g. prime=RADIO, target=radio). This suggests that a case-independent representation of letter strings can be accessed nonconsciously (Dehaene et al., 2001). By using words made of letters whose upper and lower case shapes are arbitrarily related (e.g. A/a, G/g), we have demonstrated that this representation comprises culturally-specific information laid down in the course of learning to read (Dehaene et al., in press).

**Subliminal binding.** The binding of letters in a coherent word is required for reading, because different words can be made up of the same letters. Can such binding occur nonconsciously? We addressed this question by preceding a target word by either a repetition of itself, or by an anagram made of the same letters in a different order (Dehaene et al., in press). In response times, repetition priming occurred for words but not for anagrams, indicating that precise information about the configuration of letters was extracted unconsciously. Furthermore, fMRI separated a posterior fusiform region sensitive only to the component letters, from a more anterior region that began to be sensitive to letter combinations. Thus, the evidence suggests that the binding of letters into larger units is independent of consciousness, and an organized structural representation of a word can be constructed in the absence of consciousness.

**Semantic access.** The issue of semantic access from subliminal masked words remains controversial in psychology. Nevertheless, positive evidence for subliminal semantic access was obtained using a small set of high-frequency words with simple semantics:

number words. When subjects were engaged in a number comparison task, their responses were accelerated when the prime and the target represented the same quantity, possibly in different notations (e.g. prime NINE, target 9) (Dehaene, Naccache et al., 1998; Koechlin, Naccache, Block, & Dehaene, 1999). In this paradigm, fMRI showed “quantity priming” in a bilateral intraparietal region thought to be involved in the semantic representation and manipulation of numerical quantities (Naccache & Dehaene, 2001a).

**Motor activation.** Masked words and digits have a measurable influence down to the motor preparation level. In the above number comparison task, where subjects classified targets as larger or smaller than 5 using their left or right hand, the numerical primes could also be larger or small than 5, and thus they could induce a motor response congruent or incongruent with the subsequent target. This response congruity factor interfered with subjects’ response times, and yielded a response conflict in motor cortex which was measurable in both ERPs and fMRI (Dehaene, Naccache et al., 1998).

In brief, an entire series of processing stages, highly attuned to the processing of words at the perceptual, semantic and even the motor level, can be successively activated in a feedforward manner by a subliminal prime. Those findings agrees with a major tenet of the neuronal workspace theory, the existence of a broad set of distributed specialized processors that, most of the time, process information nonconsciously. Furthermore, examination of the effects of attention on subliminal processing reveals a remarkable dissociation between attention and conscious access. All of the above priming experiments allowed subjects to deploy attention to the target. We recently showed that when the prime-target pair occurs at an unpredictable moment, thus preventing the



deployment of temporal attention, then subliminal priming effects disappear (Naccache et al., 2002). Thus, the idea that subliminal priming reflects a purely passive process of spreading activation can be rejected. Rather, subliminal primes benefit from an attentional amplification, although this is not sufficient from them to enter awareness. Those data imply that subliminal primes are actively processed along consciously prepared routes.

Nevertheless, subliminal processing is not as flexible as conscious processing. As assumed by the neuronal workspace model, functions that depend on central executive control, such as inhibition or conflict detection, appear to require consciousness. When asked to inhibit the dominant strategy of naming the prime words, subjects are unable to comply with this instruction unless the words are consciously accessed (Debner & Jacoby, 1994). Similarly, we contrasted the motor conflict effects generated by subliminal and supraliminal numerical primes (Dehaene, Artiges, Naccache, Viard, & Martinot, submitted). The anterior cingulate showed a conflict effect only with supraliminal primes, not with subliminal primes. Furthermore, schizophrenic subjects with known anterior cingulate pathology showed normal subliminal priming, but abnormal supraliminal motor interference. These results indirectly suggest that executive control processes associated with prefrontal and cingulate cortices can only operate on consciously perceived stimuli.

### **Physiological correlates of nonconscious and conscious processing**

The masking paradigm also provides a more direct way of identifying the changes in brain activity that distinguish subliminal and conscious situations. A simple experiment consists in measuring the fMRI or ERP correlates of brain activity evoked by

masked or unmasked words, relative to word-absent trials, under conditions that are as comparable as possible in every other respect (Dehaene et al., 2001). This method looks at the brain activation caused by masked words directly, rather than indirectly through their priming effects on subsequent words. Using this design, we found that the masked words caused a small transient bottom-up activation which was increasingly smaller as one moved from extrastriate cortex to fusiform gyrus and precentral cortex. When the words were unmasked, activation greatly increased in the same areas, but also extended to a distributed network which included distant parietal, inferior prefrontal and midline precentral/cingulate cortices. Unmasking also enhanced the long-distance correlation between those sites and, in ERP recordings, was associated with an enhanced late positive complex (P300) that was absent or greatly reduced in the masked situation.

Electrophysiological studies of masking in awake monkeys converge adequately with the research in humans. Lamme et al. (2002) recorded from V1 neurons while monkeys reported the presence or absence of a visual stimulus with variable degrees of masking. They observed a first peak of neural firing shortly following visual stimulation (~60 ms after stimulus onset). This peak remained unchanged and equally selective for stimulus orientation whether the stimulus was reportable or not, indicating subliminal processing of orientation in V1. However, a second peak of moderate amplification of firing rate (~90-150 ms after stimulus onset) was seen only when the stimulus was reportable. According to Lamme et al. (2002), the first period may reflect bottom-up propagation of activation, while the second would reflect top-down signals arising from higher cortical areas. Analogous results were obtained from recordings in monkey infero-temporal cortex (Kovacs, Vogels, & Orban, 1995; Rolls, Tovee, & Panzeri, 1999) and

frontal eye field (Thompson & Schall, 1999): a first firing peak (~80 ms after stimulus onset) was left unchanged or only slightly reduced by masking, and the cells maintained their stimulus selectivity; but masking drastically interrupted a later phase of firing starting about 100 ms after the stimulus.

In summary, the following findings fit with the postulates of the neuronal workspace model. First, the same cortical area, such as the left fusiform gyrus, can participate in both subliminal and conscious streams of processing. Second, conscious processing is associated with an amplification of late perceptual activity and its functional correlation with distant parietal, prefrontal, and cingulate sites. Third, consciousness allows the deployment of executive control processes; those seem to be amongst the few processes that cannot be deployed subliminally.

### **All-or-none dynamics of workspace activity: comparison of subjective report and objective measurements**

The neuronal workspace model states that the distinct anatomical connectivity of workspace neurons leads to qualitatively distinct patterns of activity. Because of their global recurrent connectivity, workspace neurons have the capacity of “igniting” suddenly in a self-amplifying manner as soon as a minimum subset of them is activated. At any given moment, the state of activity of workspace is therefore characterized by the intense activation of a subset of workspace neurons, the rest of workspace neurons being actively inhibited. This particular set of active workspace neurons may be viewed as a neuronal correlate of the content of consciousness. For instance, the conscious report of a word might be constituted by the simultaneous, coordinated activation of workspace

neurons in a specific processor (the fusiform visual word form area) and in distributed temporal, parietal, prefrontal and cingulate sites associated with speech production.

Because the entire workspace is globally interconnected, only one such workspace representation can be active at any given time. This property distinguishes it from peripheral processors in which, due to local patterns of connections, several representations with different formats may coexist. Furthermore, the reciprocal excitatory connectivity between workspace neurons imposes a self-amplification of activity which creates a dynamic threshold. Neural activity is either sufficient to trigger a reverberating loop of bottom-up and top-down activity which quickly amplifies it to a high and self-sustained level; or it remains below this threshold (i.e. “sub-liminal”) and only a briefly decaying bottom-up activity is seen. Thus, access to the workspace is *all-or-none* and *exclusive* of other representations.

Recently, those non-linear properties were explicitly demonstrated in a detailed simulation of realistic thalamo-cortical networks (Dehaene, Sergent & Changeux, 2003) (see figure 4). The simulations show that a brief thalamic stimulation T1 could lead to the ignition of a large set of distant cortical areas, which remain active by self-sustaining reverberatory loops for tens of milliseconds beyond the initial stimulus duration. This establishes a clear link between the content of working memory and of consciousness, and may explain why the maintenance of active information over a short delay is only feasible when the information is conscious.

Crucially, during this period of workspace occupancy by stimulus T1, another T2 could still be processed by peripheral thalamo-cortical processors, but often could not activate workspace neurons until the representation of T1 had vanished. This temporary

inability showed many parallels with the “attentional blink”, a well-known psychological paradigm in which subjects are temporarily unable to report stimuli while they are attending to another task (Chun & Potter, 1995; Vogel, Luck, & Shapiro, 1998).

An original property of the model is the prediction of a dynamic all-or-none bifurcation in neuronal activity. Across simulated trials, depending on random fluctuations in spontaneous activity prior to stimulus arrival, ascending activity could be sufficient to trigger self-amplifying recurrent activity, or it remained below threshold and only transient bottom-up activity was seen. Thus, for a fixed T1-T2 lag, simulated firing rates in higher areas and other indices of global activity (gamma-band power, long-distance cross-correlation) were found to be distributed bimodally across trials – either global and long-lasting, or local and very short-lived.

The neuronal workspace model therefore predicts that the apparent gradual drop in reportability observed during the attentional blink may be an artificial consequence of averaging across trials with full access awareness and others with no awareness. We tested this prediction experimentally using a modified attentional blink paradigm in which subjects reported to what extent they had seen a word (T2) within a rapid letter stream that contained another target letter string (T1). To obtain a continuous measure of subjective perception, we asked subjects to move a cursor on a continuous scale, from “not seen” on the left to “maximal visibility” on the right. The results indicated that subjective perception during the blink is indeed all-or-none. At the peak of the blink, which occurred ~260 ms after T1, the very same stimulus T2 was either fully perceived (cursor placed on “maximal visibility”; ~50% of trials), or totally unseen (cursor placed on “not seen”). Participants almost never used intermediate cursor positions, although

controls showed that they were able to in other psychophysical situations. This experiment substantiates the hypothesis that conscious states are associated with a fast all-or-none dynamic phase transition in a large-scale neuronal network. More generally, the concept of a sudden “ignition”, self-amplified by recurrent top-down/bottom-up interaction, may begin to explain the very notion of a threshold or “limen” of consciousness.

### **Spontaneous fluctuations and precursors of conscious access**

What causes the same target T2 to be occasionally perceived or not perceived? In our simulation, thalamic and cortical neurons are permanently subject to spontaneous oscillations, even prior to stimulus presentation. When intrinsic fluctuations are in phase with stimulus presentation, bottom-up activation is enhanced. This coincidence has a cascading effect on subsequent areas and eventually affects the probability of the entire network falling into a global active state. Thus, the resonance of incoming stimuli with spontaneous brain activity is essential for perception. Simulations show that access to consciousness can be partially predicted by random fluctuations in the size of the initial T2-induced bottom-up peak in early areas, or even by the amount of depolarization of T2 neurons 100-200 ms *prior* to stimulus presentation.

Those observations may shed some light on a recent controversy concerning the earliest correlates of consciousness during visual perception. Many experiments indicate that conscious perception, when contrasted to a non-conscious control, is associated with a sudden increase in parieto-fronto-cingulate activity, in agreement with the predictions of the neuronal workspace model (e.g. Beck, Rees, Frith, & Lavie, 2001; Dehaene et al., 2001; Lumer & Rees, 1999). However, some studies have also found early differences

within occipital cortex. For instance, Pins and ffytche (2003) sorted out trials as a function of whether the same near-threshold grating was or was not reported by human subjects. Using ERPs, they observed an early difference difference, about 100 ms after stimulus onset, which was traced back by fMRI-seeded dipole analysis to area V1 or to the surrounding occipital pole. They also found a later (~260 ms) difference in parietal and frontal regions. Pins and ffytche claimed that the early effect was a “primary correlate of conscious perception”, and that the later differences were secondary and reflected contingent processes of attention and report made possible by the first step.

Our model leads to the exactly opposite claim. In our simulation, although conscious access is *defined* by coherent long-lasting activity in higher-cortical areas, an early difference between seen and unseen trials is present even before this state is attained. This clearly illustrates that this early difference is an indirect consequence of selective averaging over a fluctuating baseline. Likewise, we suggest that the early difference in V1 activation seen by Pins and ffytche (2003) is only a (modest) predictor of subsequent target reportability, but does not, in and of itself, constitute the neuronal basis of a conscious state.

According to our view, some neural changes may show statistically significant correlations with the presence or absence of consciousness, without however participating in a *state* of consciousness. This distinction becomes clearer when such correlations are found even before the stimulus is presented. For instance, Super et al. (2003) found that changes in the firing rate of V1 neurons 100 ms prior to stimulus presentation partially predicted whether a stimulus would or would not be reported by a macaque monkey. How could the monkey possibly have a conscious state of a target

prior to its presentation? The only reasonable interpretation, adopted by Super et al. (2003), is that the pre-stimulus differences do not directly indicate access to consciousness, but merely a state of attentive readiness that makes it more likely for a subsequent stimulus to become conscious. We suggest that a similar interpretation explains the finding of Pins and ffytche (2003).

In summary, it is not sufficient to merely observe neural *correlates* of consciousness. Ultimately, what is needed is a direct, causal and contemporaneous link between neuronal activity and conscious perception. Early ventral visual activity cannot be conscious because it still occurs in an unchanged form during extinction of visual stimuli in patients with parietal neglect (Vuilleumier et al., 2002). Like us, Vuilleumier et al. emphasize that sensory activity alone is not sufficient, and that functional interactions between parietal, frontal, and sensory areas appear necessary for access to consciousness.

## **Conclusion**

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



Given this broad convergence, it appears even more critical to keep the theoretical differences among various models clear, to the benefit of decisive experimental tests. First, the present theory departs from other approaches that view the “dynamic core” of consciousness as eminently variable and not directly related to a special subset of neurons (Edelman, 1989; Tononi & Edelman, 1998). On the opposite, Herbert Jasper has insisted for years that “an anatomically and electrophysiologically separate neuronal system is involved in brain mechanisms of consciousness” (Jasper, 1998). Consistent with Jasper’s views, we emphasize that most cerebral processes are nonconscious, and that the neural mechanisms of access to consciousness involve a specific subset of neurons that can be delineated by minimal contrasts between subliminal and supraliminal stimuli.

Second, our theory is not compatible with the statements that prefrontal regions function as an “unconscious homunculus” (Crick & Koch, 2003), that area V1 would not be mobilized by conscious processing (Crick & Koch, 1995), and that strong recurrent loops are avoided in the cortex (Crick & Koch, 1998). On the contrary, we think that prefrontal regions densely contribute to workspace neurons and that their activity usually betrays conscious processing. Moreover, we emphasise that almost all cortical processors (including area V1) can be mobilized into the workspace through joint bottom-up and top-down excitatory links. During fine-grained visual imagery, for instance, our theory predicts that there should be top-down mobilization of V1 neurons and long-distance correlations with prefrontal cortex neurons, and that interference with V1 activity, for instance with transcranial magnetic stimulation, should disrupt the conscious mental image.

Finally, our theory contrasts with the view that consciousness precedes attention (Lamme, 2003). We do not think that a form of “phenomenal consciousness” can be attributed to pre-attentive contents that cannot be reported (Block, 2001), and that inattentional blindness phenomena such as the attentional blink reflect “brief consciousness followed by amnesia” (Lamme, 2003). Rather, we propose that masked or blinked stimuli never lead to a strong activation of workspace neurons in the first place and thus are simply nonconscious.

The neuronal workspace model proposes that the neural basis of conscious access is a sudden self-amplifying bifurcation leading to a global brain-scale pattern of activity. We shall now underline several critical predictions of this view. First, when a stimulus is presented above threshold, we predict that following an initial period of subliminal perceptual processing lasting ~100-200 ms, there is a sudden non-linear transition towards a temporarily metastable state of globally increased brain activity lasting ~200-300 ms. This sudden increase should be particularly evident in prefrontal, cingulate, and parietal cortices. It should be accompanied by a synchronous amplification of posterior perceptual activation and by thalamo-cortical brain-scale synchrony in the gamma range (20-100 Hz). Finer scale electrophysiological or optical recordings should demonstrate that this state is selective to a subset of neurons coding for the stimulus, and is accompanied in higher areas by broad inhibition of neurons coding for other stimuli. Even higher resolution experiments may reveal a layer-specific distribution of active workspace neurons, with intense top-down activity originating from a subpopulation of cells with long axons in supra- and infra-granular layers of prefrontal, cingulate and parietal cortices. The model also predicts that, when presented with a stimulus at

threshold, workspace neurons respond in an all-or-none manner, either highly activated or totally inactivated. This all-or-none character should be detectable macroscopically as a bimodal distribution of parameters such as the P300 component of event-related potentials. Finally, pharmacological interventions or lesions that affect top-down connections or inhibitory interactions in higher cortical areas should alter the dynamics of workspace ignition, and therefore should modify subjective perception in a predictable manner, while possibly leaving subliminal bottom-up processing unaffected (e.g. Granon, Faure, & Changeux, 2003). Through predictions such as these, the road is now paved for a neuroscientific approach to consciousness.

## **Figure Legends**

**Figure 1.** Schematic depiction of the workspace model (redrawn from Dehaene, Kerszberg & Changeux, 1998). Cortical processors are shown in a state of activity in which a stimulus T1 has gained access to the conscious workspace, while another stimulus T2 is only processed nonconsciously up to a limited level.

**Figure 2.** Types of interactions between an automatized stream of processors linking stimulus S and response R, and the workspace system. **A,B,C:** three types of nonconscious processing. In **A**, a processing chain is not connected to workspace neurons and therefore remains permanently inaccessible to consciousness. In **B**, a processing chain is connected reciprocally to workspace neurons, but is temporarily not mobilized by top-down amplification. In **C**, the processing chain is attended and subliminal processing is amplified, yet activation is too brief to establish a bidirectional reverberating loop. **D, E, F:** conscious processing. In **D**, the loop is closed, allowing stimulus information to be held on-line and broadcasted to multiple systems R', R'', etc. **E** and **F** illustrate the two orders in which this conscious state can be achieved: a top-down mode (**E**) in which workspace neurons become spontaneously activated first, and selectively amplify a possibly very small sensory signal; and a bottom-up mode (**F**) in which initially unattended sensory signals carry sufficient strength as to cause a reorienting of top-down amplification towards them.

**Figure 3.** Main paradigms used in neuroimaging studies of subliminal priming. **A**, word repetition paradigm, in which a word prime is flashed for 29 ms, hidden by forward and backward masking shapes (Dehaene et al., in press; Dehaene et al., 2001). Repetition of the same physical stimulus as prime and target leads to feature-based priming in occipital

cortex, while cross-case word repetition leads to case-independent priming in the left fusiform “visual word form area”. **B**, number comparison paradigm, in which a numerical prime is flashed for 43 ms, hidden by forward and backward letter strings (Dehaene, Naccache et al., 1998; Naccache et al., 2002; Naccache & Dehaene, 2001a, 2001b). Repeating the same quantity as prime and target leads to quantity-based priming in the left and right intraparietal sulci, while repeating the same motor response leads to response priming in the left and right motor cortex. Altogether, those results indicate that a subliminal prime can proceed through an entire series of visual, semantic and motor stages without becoming conscious.

**Figure 4.** Anatomical connectivity (left) and functional states of activity (right) in a recent simulation of a subset of the proposed workspace system (for details, see Dehaene et al., 2003). Two hierarchies of pyramidal neurons code for two input stimuli T1 and T2 at successively higher levels. Higher-level workspace neurons send long-distance axons in a top-down manner all the way down to the earliest cortical area, as well as to many other workspace areas (reportability). At the higher-level, cortical representations of T1 and T2 also inhibit each other via inhibitory interneurons. Simulated peri-stimulus time histograms (right) illustrate that a very brief presentation (40 ms) of either T1 or T2 leads to an initial bottom-up activation through the hierarchy, then a top-down amplification in the reverse direction, finally resulting in a long-lasting self-sustained state of coherent activity. When T2 is presented within a short period after T1, however, such “ignition” fails to occur and only the first stages of bottom-up propagation are seen. This is proposed to constitute the neural basis of the attentional blink.

## References


- Allman, J. M., Hakeem, A., Erwin, J. M., Nimchinsky, E., & Hof, P. (2001). The anterior cingulate cortex. The evolution of an interface between emotion and cognition. *Ann N Y Acad Sci*, 935, 107-117.
- Baars, B. J. (1989). *A cognitive theory of consciousness*. Cambridge, Mass.: Cambridge University Press.
- Beck, D. M., Rees, G., Frith, C. D., & Lavie, N. (2001). Neural correlates of change detection and change blindness. *Nature Neuroscience*, 4, 645-650.
- Bianchi, L. (1922). *The mechanism of the brain and the functions of the frontal lobes*. New York: W. Wood.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18(2), 227-287.
- Block, N. (2001). Paradox and cross purposes in recent work on consciousness. *Cognition*, 79(1-2), 197-219.
- Changeux, J. P. (1983). *L'homme neuronal*. Paris: Fayard.
- Changeux, J. P., & Dehaene, S. (1989). Neuronal models of cognitive functions. *Cognition*, 33, 63-109.
- Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *J Exp Psychol Hum Percept Perform*, 21(1), 109-127.
- Cohen, L., Dehaene, S., Naccache, L., Lehéricy, S., Dehaene-Lambertz, G., Hénaff, M. A., et al. (2000). The visual word form area: Spatial and temporal characterization of an initial stage of reading in normal subjects and posterior split-brain patients. *Brain*, 123, 291-307.
- Crick, F. (1994). *The Astonishing Hypothesis: The Scientific Search for the Soul*. New York: Charles Scribner's Sons.
- Crick, F., & Koch, C. (1995). Are we aware of neural activity in primary visual cortex? *Nature*, 375, 121-123.
- Crick, F., & Koch, C. (1998). Constraints on cortical and thalamic projections: the no-strong-loops hypothesis. *Nature*, 391(6664), 245-250.
- Crick, F., & Koch, C. (2003). A framework for consciousness. *Nat Neurosci*, 6(2), 119-126.
- Debner, J. A., & Jacoby, L. L. (1994). Unconscious perception: attention, awareness, and control. *J Exp Psychol Learn Mem Cogn*, 20(2), 304-317.
- Dehaene, S., Artiges, E., Naccache, L., Viard, A., & Martinot, J. L. (submitted). Conscious and subliminal conflicts in normal and schizophrenic subjects: The role of the anterior cingulate.
- Dehaene, S., & Changeux, J. P. (1989). A simple model of prefrontal cortex function in delayed-response tasks. *Journal of Cognitive Neuroscience*, 1, 244-261.
- Dehaene, S., & Changeux, J. P. (1991). The Wisconsin Card Sorting Test: Theoretical analysis and modelling in a neuronal network. *Cerebral cortex*, 1, 62-79.

- Dehaene, S., & Changeux, J. P. (1997). A hierarchical neuronal network for planning behavior. *Proc Natl Acad Sci U S A*, *94*, 13293-13298.
- Dehaene, S., & Changeux, J. P. (2000). Reward-dependent learning in neuronal networks for planning and decision making. *Prog Brain Res*, *126*, 217-229.
- Dehaene, S., Jobert, A., Naccache, L., Ciuciu, P., Poline, J. B., Le Bihan, D., et al. (in press). Letter binding and invariant recognition of masked words: Behavioral and neuroimaging evidence. *Psychological Science*.
- Dehaene, S., Kerszberg, M., & Changeux, J. P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proc Natl Acad Sci U S A*, *95*(24), 14529-14534.
- Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, *79*, 1-37.
- Dehaene, S., Naccache, L., Cohen, L., Le Bihan, D., Mangin, J. F., Poline, J. B., et al. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience*, *4*, 752-758.
- Dehaene, S., Naccache, L., Le Clec'H, G., Koechlin, E., Mueller, M., Dehaene-Lambertz, G., et al. (1998). Imaging unconscious semantic priming. *Nature*, *395*, 597-600.
- Dehaene, S., Sergent, C., & Changeux, J. P. (2003). A neuronal model linking subjective reports and objective neurophysiological data during conscious perception. *Proc Natl Acad Sci U S A*, in press.
- Di Lollo, V., Enns, J. T., & Rensink, R. A. (2000). Competition for consciousness among visual events: the psychophysics of reentrant visual processes. *J Exp Psychol Gen*, *129*(4), 481-507.
- Di Virgilio, G., & Clarke, S. (1997). Direct interhemispheric visual input to human speech areas. *Hum Brain Mapp*, *5*, 347-354.
- Edelman, G. M. (1989). *The remembered present*. Basic Books: New York.
- Edelman, G. M. (1993). Neural Darwinism: selection and reentrant signaling in higher brain function. *Neuron*, *10*, 115-125.
- Forster, K. I., & Davis, C. (1984). Repetition priming and frequency attenuation in lexical access. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 680-698.
- Goldman-Rakic, P. S. (1988). Topography of cognition: Parallel distributed networks in primate association cortex. *Annual Review of Neuroscience*, *11*, 137-156.
- Granon, S., Faure, P., & Changeux, J. P. (2003). Executive and social behaviors under nicotinic receptor regulation. *Proc Natl Acad Sci U S A*, submitted.
- Greenwald, A. G., Draine, S. C., & Abrams, R. L. (1996). Three cognitive markers of unconscious semantic activation. *Science*, *273*(5282), 1699-1702.
- Jasper, H. H. (1998). Sensory information and conscious experience. *Adv Neurol*, *77*, 33-48.
- 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.
- Kovacs, G., Vogels, R., & Orban, G. A. (1995). Cortical correlate of pattern backward masking. *Proc Natl Acad Sci U S A*, *92*, 5587-5591.

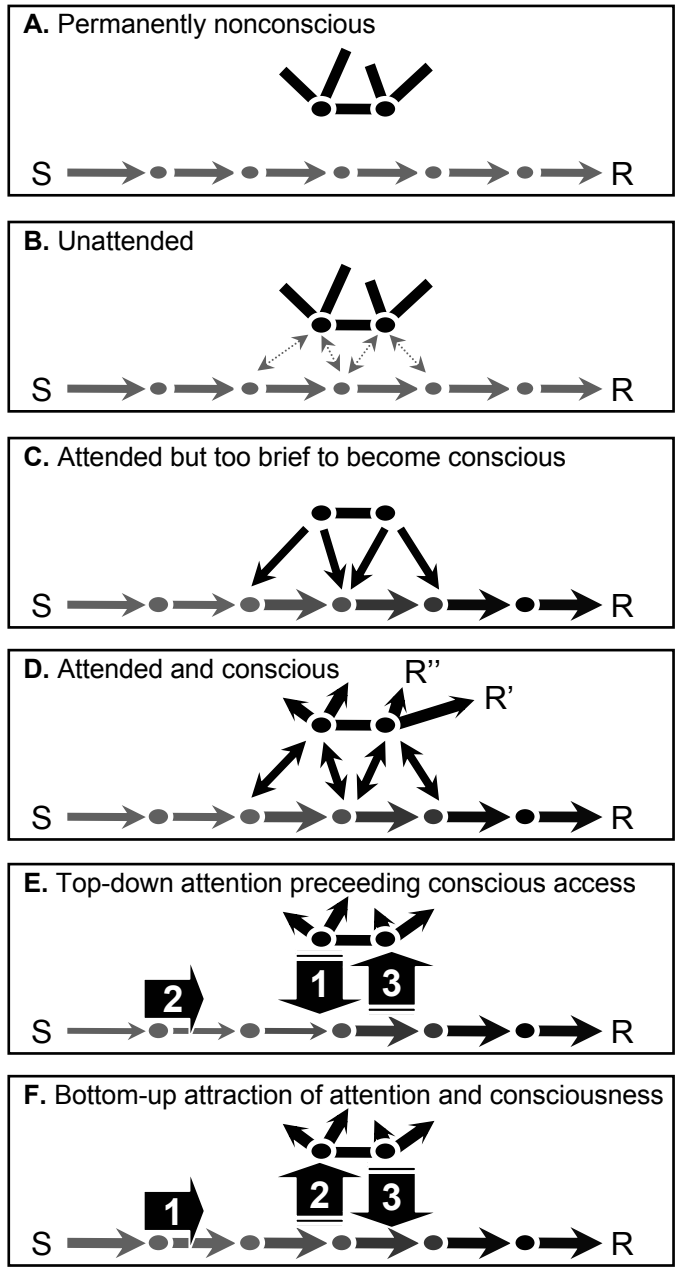
- Lamme, V. A. (2003). Why visual attention and awareness are different. *Trends Cogn Sci*, 7(1), 12-18.
- Lamme, V. A., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci*, 23(11), 571-579.
- Lamme, V. A., Zipser, K., & Spekreijse, H. (2002). Masking interrupts figure-ground signals in V1. *J Cogn Neurosci*, 14(7), 1044-1053.
- Llinas, R., Ribary, U., Contreras, D., & Pedroarena, C. (1998). The neuronal basis for consciousness. *Philos Trans R Soc Lond B Biol Sci*, 353(1377), 1841-1849.
- Lumer, E. D., & Rees, G. (1999). Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proc Natl Acad Sci U S A*, 96(4), 1669-1673.
- Naccache, L., Blandin, E., & Dehaene, S. (2002). Unconscious masked priming depends on temporal attention. *Psychological Science*, 13, 416-424.
- Naccache, L., & Dehaene, S. (2001a). The Priming Method: Imaging Unconscious Repetition Priming Reveals an Abstract Representation of Number in the Parietal Lobes. *Cereb Cortex*, 11(10), 966-974.
- Naccache, L., & Dehaene, S. (2001b). Unconscious semantic priming extends to novel unseen stimuli. *Cognition*, 80, 215-229.
- Neely, J. H., & Kahan, T. A. (2001). Is semantic activation automatic? A critical re-evaluation. In H. L. Roediger, J. S. Nairne, I. Neath & A. M. Surprenant (Eds.), *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 69-93). Washington D.C.: American Psychological Association.
- Norman, D. A., & Shallice, T. (1980). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz & D. Shapiro (Eds.), *Consciousness and self-regulation* (Vol. 4). New York: Plenum Press.
- Paus, T. (2000). Functional anatomy of arousal and attention systems in the human brain. *Prog Brain Res*, 126, 65-77.
- Pins, D., & ffytche, D. (2003). The neural correlates of conscious vision. *Cereb Cortex*, 13, 461-474.
- Posner, M. I., & Rothbart, M. K. (1998). Attention, self-regulation and consciousness. *Philos Trans R Soc Lond B Biol Sci*, 353(1377), 1915-1927.
- Revonsuo, A. (2001). Can Functional Brain Imaging Discover Consciousness In the Brain? *J Consciousness Studies*, 8, 3-50.
- Rolls, E. T., Tovee, M. J., & Panzeri, S. (1999). The neurophysiology of backward visual masking: information analysis. *J. Cogn. Neurosci.*, 11(3), 300-311.
- Super, H., van der Togt, C., Spekreijse, H., & Lamme, V. A. (2003). Internal state of monkey primary visual cortex (V1) predicts figure-ground perception. *J Neurosci*, 23(8), 3407-3414.
- Thompson, K. G., & Schall, J. D. (1999). The detection of visual signals by macaque frontal eye field during masking. *Nature Neuroscience*, 2, 283-288.
- Tononi, G., & Edelman, G. M. (1998). Consciousness and complexity. *Science*, 282(5395), 1846-1851.
- Vogel, E. K., Luck, S. J., & Shapiro, K. L. (1998). Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *J Exp Psychol Hum Percept Perform*, 24(6), 1656-1674.



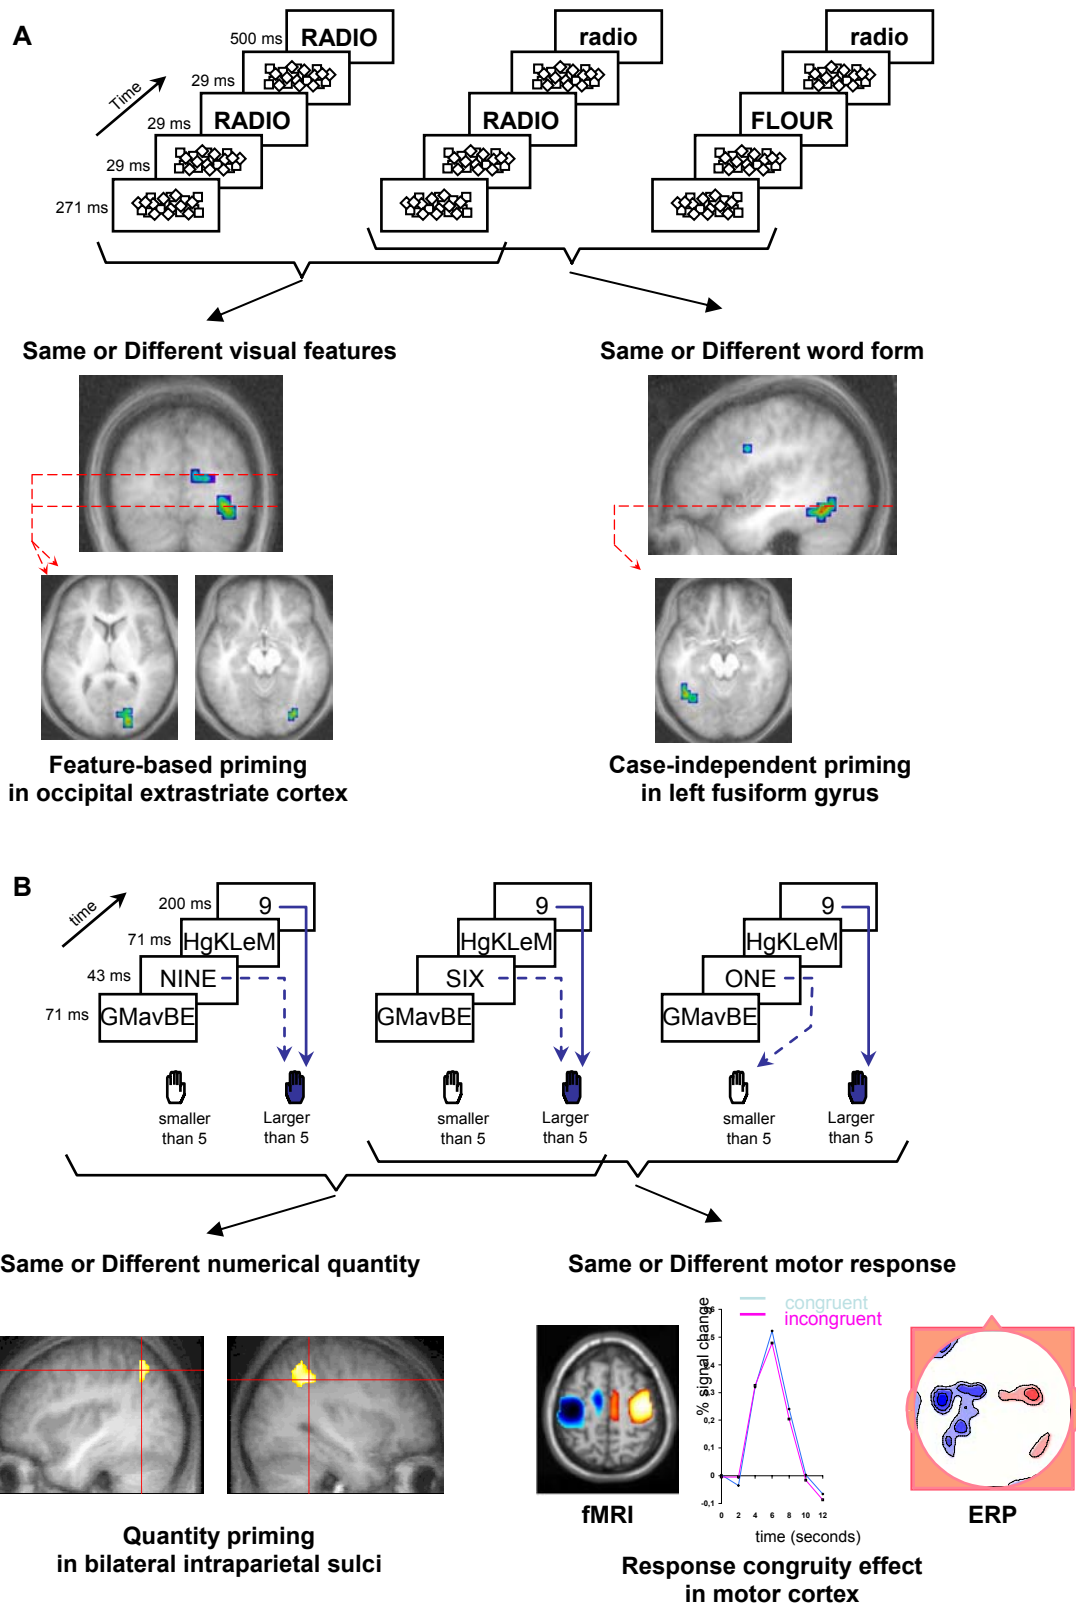
- Vuilleumier, P., Armony, J., Clarke, K., Husain, M., Driver, J., & Dolan, R. (2002).  
Neural response to emotional faces with and without awareness: event-related  
fMRI in a parietal patient with visual extinction and spatial neglect.  
*Neuropsychologia*, 40(12), 2156.
- Weiskrantz, L. (1997). *Consciousness lost and found: A neuropsychological exploration*.  
New York: Oxford University Press.

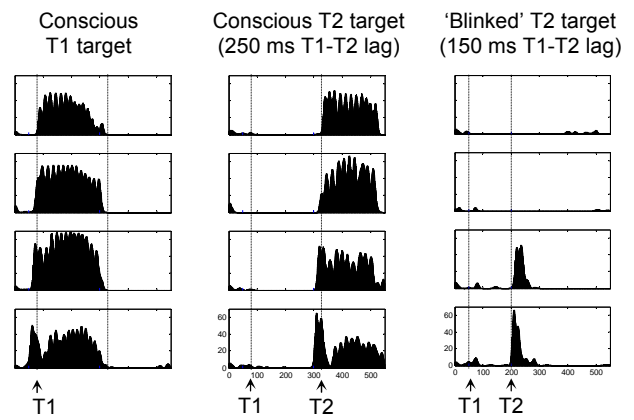



Dehaene & Changeux – Figure 1



Dehaene & Changeux – Figure 2





Dehaene & Changeux – Figure 4