From the Strüngmann Forum Report: Better Than Conscious?
Decision Making, the Human Mind, and Implications For Institutions
Edited by Christoph Engel and Wolf Singer
© MIT Press. ISBN-10: 0-262-19580-1 ISBN-13: 978-0-262-19580-5
http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=11498

2

Conscious and Nonconscious Processes

Distinct Forms of Evidence Accumulation?

Stanislas Dehaene[1,2,3,4]

[1] INSERM, U562, Cognitive Neuroimaging Unit, 91191 Gif/Yvette, France
[2] CEA, DSV/I2BM, NeuroSpin Center, 91191 Gif/Yvette, France
[3] Univ Paris-Sud, IFR49, 91191 Gif/Yvette, France
[4] Collège de France, 75005 Paris, France

Abstract

Among the many brain events evoked by a visual stimulus, which ones are associated specifically with conscious perception, and which merely reflect nonconscious processing? Understanding the neuronal mechanisms of consciousness is a major challenge for cognitive neuroscience. Recently, progress has been achieved by contrasting behavior and brain activation in minimally different experimental conditions, one of which leads to conscious perception whereas the other does not. This chapter reviews briefly this line of research and speculates on its theoretical interpretation. I propose to draw links between evidence accumulation models, which are highly successful in capturing elementary psychophysical decisions, and the conscious/nonconscious dichotomy. In this framework, conscious access would correspond to the crossing of a threshold in evidence accumulation within a distributed "global workspace," a set of recurrently connected neurons with long axons that is able to integrate and broadcast back evidence from multiple brain processors. During nonconscious processing, evidence would be accumulated locally within specialized subcircuits, but would fail to reach the threshold needed for global ignition and, therefore, conscious reportability.

An Experimental Strategy for Exploring Consciousness

Although the nature of consciousness remains a formidable problem, Lionel Naccache and I argue that it can be approached through behavioral and brain-imaging methods:

> 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 (Dehaene and Naccache 2001).

In that respect, identifying the neural bases of consciousness need not be any more difficult than, say, identifying that of other states of mind (e.g., face perception or anger). Bernard Baars (1989) outlined a simple *contrastive method* which, in his own terms, consists simply in contrasting pairs of similar events, where one is conscious but the other is not. Baars noted that in the last forty years, experimental psychology and neuropsychology have identified dozens of contrasts relevant to consciousness. Examples include normal vision versus blindsight; extinguished versus seen stimuli in patients with hemineglect; masked versus nonmasked visual stimuli; habituated versus novel stimuli; accessed versus nonaccessed meanings of ambiguous stimuli; distinctions within states of consciousness (sleep, coma, wakefulness, arousal); voluntary versus involuntary actions; or even explicit problem solving versus implicit "incubation."

In this chapter, I focus on the masking paradigm, perhaps the simplest and most productive situation in which to study conscious access in normal subjects. During masking, a target visual stimulus is flashed briefly on a computer screen. It can be followed or preceded by a "mask": another visual stimulus presented at the same screen location or just nearby. Under the right conditions, presentation of the mask erases the perception of the target stimulus, and subjects report that they are no longer able to see it. Yet the target stimulus still induces behavioral priming effects and brain activation patterns which correspond to nonconscious or "subliminal" (below threshold) processing. Focusing on what types of processing can occur under subliminal masking conditions, and what additional processes unfold once the stimulus is unmasked, can thus shed considerable light on the nature of conscious access.

How Do We Measure Whether Conscious Access Occurred?

As mentioned above, once an appropriate paradigm such as masking is available, studying the cerebral correlates of conscious access need not be more difficult than, say, studying face perception. In both cases, one correlates brain activity with the presence or absence of the relevant aspect of the stimulus (face vs. nonface stimulus, or conscious vs. nonconscious perception). What is special about conscious access, however, is that it is defined solely in subjective terms. Thus, Lionel Naccache and I have argued:

> The first crucial step is *to take seriously introspective phenomenological reports*. Subjective reports are the key phenomena that a cognitive neuroscience of consciousness purports to study. As such, they constitute primary data that need to

> be measured and recorded along with other psychophysiological observations (Dehaene and Naccache 2001).

Increasingly, therefore, consciousness research relies on subjective reports as a defining criterion. Ideally, one should measure the extent of conscious perception on every single trial, possibly using a graded scale to capture even fine nuances of the percept (Del Cul et al. 2007; Sergent et al. 2005; Sergent and Dehaene 2004). For an identical objective stimulus, one may then contrast the brain activation observed when it is or is not subjectively seen.

The emphasis on subjective reporting goes against a long tradition in psychophysics and experimental psychology, which has emphasized the need for objective criteria based on signal-detection theory. According to this tradition, a masked stimulus is accepted as being subliminal only if performance on some direct task of stimulus perception falls to chance level (zero d-prime). There are several difficulties associated with this objective definition, however. First, it tends to overestimate conscious perception, as there are many conditions in which subjects perform better than chance, and yet deny perceiving any stimulus. Second, it requires accepting the null hypothesis of chance-level performance; usually d-prime never quite drops to zero, and whether it is significant or not depends merely on the number of trials dedicated to its measurement. Finally, performance can be at chance level for some tasks, but not others. Does above-chance performance on the former tasks count as evidence of conscious perception, or merely of subliminal processing? The issue seems unsolvable unless we have a good theory of which tasks can only be performed at a conscious level, and thus constitute appropriate objective measures of conscious access, and which tasks can operate under subliminal conditions.

By focusing first and foremost on subjective reports, we can avoid this somewhat Byzantine discussion of what constitutes a good subliminal stimulus. It is an empirical fact that, when subjects rate a stimulus subjectively as having been seen consciously, a major transition occurs such that the stimulus also becomes available for a variety of objective tasks. For instance, Figure 2.1 shows data from a masking paradigm (Del Cul et al. 2007) where subjects were asked, on every trial, to perform two tasks on a masked digit: (a) a subjective task of rating the stimulus visibility; (b) an objective, forced-choice task of deciding whether the stimulus was larger or smaller than five. As the interval between the target and mask increased, both subjective and objective performance increased in a nonlinear sigmoidal manner. Both sigmoids allowed for the definition of a threshold (placed at the inflection point). We found that these subjective and objective definitions of the consciousness threshold were virtually identical and highly correlated between subjects. Furthermore, both were degraded jointly in patients with schizophrenia or multiple sclerosis (Del Cul et al. 2006; Reuter et al. 2007). Interestingly, below this threshold, the objective and subjective tasks could be dissociated, as there was a proportion of

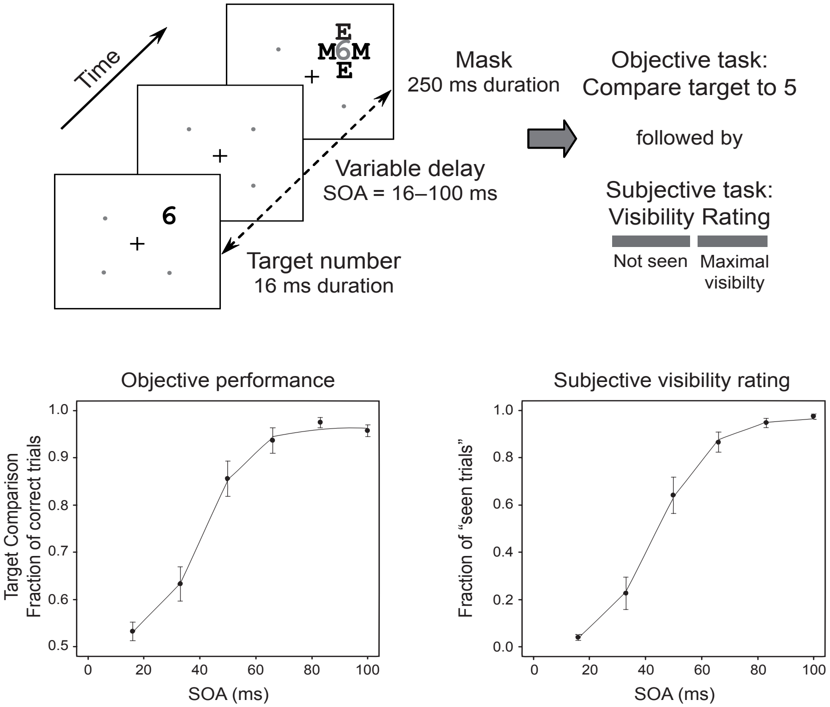

Figure 2.1 Example of a masking paradigm where objective and subjective measures concur to define a threshold for perceptual consciousness (after Del Cul et al. 2007). A digit is flashed at one of four parafoveal locations and is followed after a variable delay by a surrounding letter mask (top left panel). On each trial, participants are asked to perform an objective task (decide if the digit is larger or smaller than 5) and a subjective task (rate the stimulus visibility). Both measures concur: performance is low at short delays, but suddenly jumps to a high value above a threshold delay (around 50 ms). This method thus defines a range of "subliminal" (below-threshold) stimuli. SOA = stimulus onset asynchrony.

trials in which objective performance remained higher than chance, although subjects denied subjective perception.

In view of such results, the following operational definitions of conscious and nonconscious processing may be proposed. First, on a single-trial basis, priority should be given to subjective reports in defining what constitutes a "conscious" trial. Second, when averaging across trials, the threshold for conscious access may be identified with the major nonlinearity that occurs in both subjective and objective performance as the stimulus is progressively unmasked. Third, the presence of nonconscious processing can be inferred whenever objective performance departs from subjective reports; for instance, by remaining above-chance in a region of stimulus space where subjective reports fall to zero.

The latter hypothesis lies at the heart of the *dissociation method*, which has been used by many others to separate conscious and nonconscious processing. For instance, masking conditions can be found that create a U-shaped curve for subjective perception as a function of target-mask interval. Other aspects of performance, such as response time and brain activity patterns, vary monotonically with the same stimulus parameter, thus clearly reflecting nonconscious stimulus processing (Haynes et al. 2005; Vorberg et al. 2003).

Subliminal Processing and Evidence Accumulation Models

A broad array of research has focused on the issue of the depth of subliminal processing of masked visual stimuli: to what extent is a masked stimulus that is reported subjectively as "not seen" processed in the brain? Here I present only a brief overview of this line of research (for a broader review, see Kouider and Dehaene 2007). The main goal is to examine these data in relation to models of decision making by evidence accumulation, which have proven highly successful in mathematical modeling of chronometric and neurophysiological data from simple psychophysical decisions (e.g., Laming 1968; Link 1992; Smith and Ratcliff 2004; Usher and McClelland 2001; Shadlen, this volume).

For the sake of concreteness, one such accumulation model is presented in Figure 2.2. This particular model was shown to capture much of what is known about simple numerical decisions and their neural bases (Dehaene 2007). While many variants can be proposed (see Smith and Ratcliff 2004), this model incorporates mechanisms that are generic to a variety of psychophysical tasks. To illustrate this, consider the task of deciding if a number, presented either as a set of dots or as an Arabic numeral, is smaller or larger than 10. The model assumes the following steps:

1. Visual perception of the stimulus.
2. Semantic coding along the appropriate dimension (here, numerosity).
3. Categorization of the incoming evidence in relation to the instructions. This is achieved by separating this continuum into pools of units, each favoring a distinct response (here, units preferring numbers larger than 10 and units preferring numbers smaller than 10).
4. Computation of a log likelihood ratio (logLR), a quantity which estimates the likelihood that response R1 or R2 is correct, given the sensory evidence.
5. Stochastic accumulation of the logLR over a period of time, until a threshold amount is obtained in one direction or the other.
6. Emission of a motor response when the threshold is exceeded.

Models of this form have been shown to capture the details of chronometric data, including the shape of response time (RT) distributions and speed-accuracy trade-offs. Within the context of conscious versus nonconscious

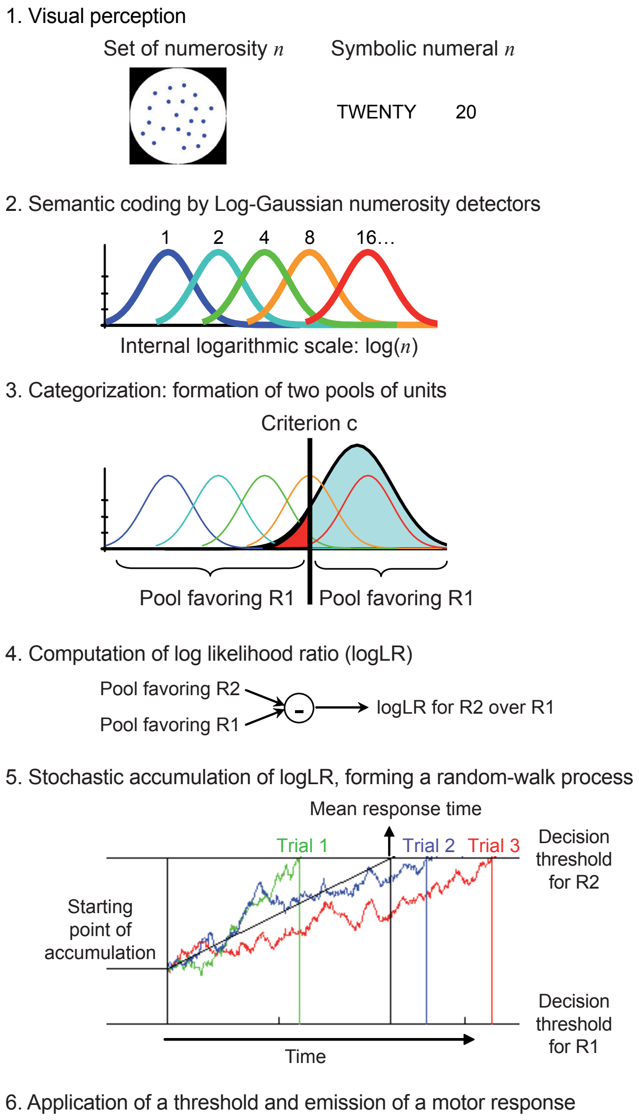

Figure 2.2 Proposed theoretical model of decision making in an objective numerical comparison task (for a full mathematical exposition, see Dehaene 2007). Subjects first encode each input number as a random variable on an internal continuum (top). The decision mechanism consists in accumulating evidence by adding up the log likelihood ratios (logLRs) for or against each of the two possible responses provided by successive samples of the random variable (middle). As a result, each trial consists in an internal random walk of the accumulated logLR (bottom). A response is emitted whenever the random walk reaches one of two response thresholds. Evidence reviewed in the present chapter suggests that all stages of the model can begin to operate in the absence of consciousness.

computation, a key question is: Which of the model's mechanisms can operate under subliminal conditions, and which cannot?

Subliminal Perception

Extensive research has demonstrated that a subliminal masked stimulus can be processed at the perceptual level. The main support comes from the *repetition priming* experiment, in which a subliminal prime is shown to facilitate the subsequent processing of an identical stimulus presented as a target. Priming is evidenced behaviorally as a reduction of response time on repeated trials compared to nonrepeated trials and neurally as a reduction in the amount of evoked brain activity (repetition suppression).

Repetition priming indicates that a subliminal stimulus can be registered perceptually. As illustrated in Figure 2.3, however, priming can occur at multiple levels. In extrastriatal cortex, priming is sensitive to the repetition of the exact same stimulus. In more anterior sectors of fusiform cortex, priming is more abstract and can resist a change in surface format, for example, when the same word is presented in upper case or lower case (Dehaene et al. 2001).

Subliminal Semantic Processing

At an even more abstract level, semantic subliminal priming has been observed, for example, in the left lateral temporal cortex for synonym words such as sofa/couch (Devlin et al. 2004) or for Japanese words presented in Kanji and Kana notations (Nakamura et al. 2005). Likewise, numerical repetition priming has been observed in bilateral intraparietal cortex when number words are presented in Arabic or word notations (Naccache and Dehaene 2001a). These observations have been confirmed by detailed behavioral studies (Naccache and Dehaene 2001b; Reynvoet et al. 2002). In terms of the model presented in Figure 2.2, they suggest that subliminal primes can partially bias the level of semantic coding.

The reality of subliminal semantic processing is confirmed by several empirical findings. Subliminal words can evoke an N400 component of the event-related potential, which depends on their semantic relation to a previously presented word (Kiefer 2002; Kiefer and Brendel 2006). Subliminal words that convey an emotion (e.g., rape, shark) can cause an activation of the amygdala (Naccache et al. 2005), and the threshold for their conscious perception is lowered, indicating that they receive distinct processing even prior to conscious access (Gaillard, Del Cul et al. 2006).

In many of these cases, brain activation evoked by a subliminal stimulus is much reduced compared to the activation evoked by the same stimulus under conscious perception conditions (Dehaene et al. 2001). However, there are

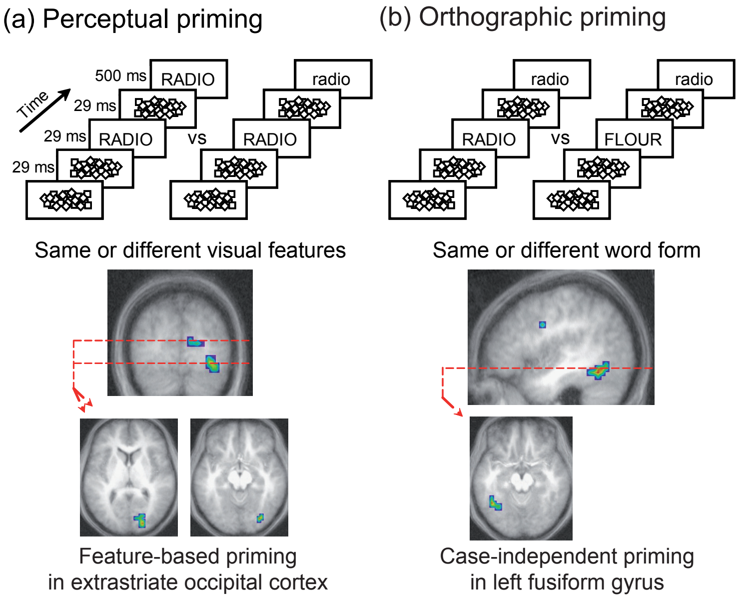

Figure 2.3 Brain imaging evidence for nonconscious processing at multiple levels of word and digit processing. All experiments rely on the priming method (Naccache and Dehaene 2001a), which consists in examining whether a subliminal prime can modulate the processing of a subsequent conscious target. The nature of the prime-target relation changes the site of modulation of brain activation: (a) shared low-level visual features cause perceptual priming in extrastriate occipital cortex; (b) case-independent orthographic priming of words occurs in the left occipito-temporal "visual word form area" (Dehaene et al. 2001).

some cases in which a full-blown activation can be observed in the absence of conscious perception. In early visual areas, even heavily masked stimuli can produce essentially unchanged event-related responses in both fMRI (Haynes et al. 2005) and ERPs (Del Cul et al. 2007). In higher visual areas, large nonconscious responses have been observed under conditions of light masking, where invisibility is due to distraction by a secondary task (e.g., the attentional blink paradigm). Even a late (~400 ms) and abstract semantic event such as the N400 can be largely (Sergent et al. 2005) or even fully (Luck et al. 1996) preserved during the attentional blink.

Subliminal Accumulation of Evidence towards a Decision

Dehaene et al. (1998) and Leuthold (Leuthold and Kopp 1998) first showed that a subliminal stimulus can bias a decision all the way down to the response programming level. The paradigm used by Dehaene et al. (1998) is illustrated in Figure 2.3c, d. Subjects had to categorize numbers as being larger or smaller

(c) Semantic priming

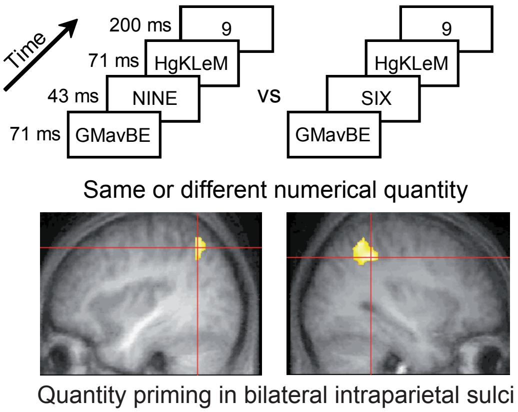

(d) Motor congruity priming

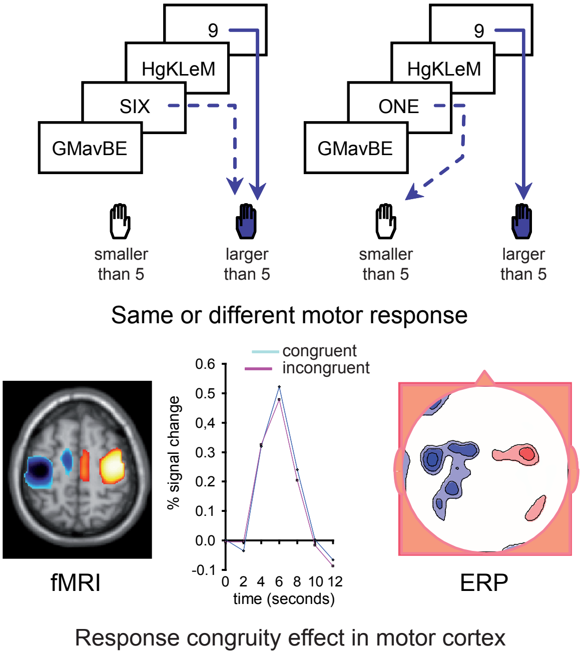

Figure 2.3 (continued) (c) Repetition of a number, in Arabic or verbal notation, causes semantic priming in bilateral intraparietal sulci (Naccache and Dehaene 2001a); (d) congruence of the motor responses associated with the prime and target modulates motor cortex activity, as if motor representations accumulate partial evidence from the prime before accumulating the main evidence arising from the target (Dehaene et al. 1998).

than five by pressing a right- or left-hand button (the response mappings were assigned randomly and switched in the middle of the experiment). Unknown to the participants, a subliminal number was presented prior to each target. A congruity effect was observed: on congruent trials, where the prime fell on the same side as the target (e.g., 9 followed by 6, both being larger than 5), responses were faster than on incongruent trials where they fell on different sides of 5 (e.g., 1 followed by 6). This effect could be measured by fMRI and ERP recordings of the motor cortex as a partial accumulation of motor bias towards the response side elicited by the prime.

Thus, activation evoked by an unseen prime can propagate all the way down to the motor level. Within the context of the model presented in Figure 2.2, this implies that semantic coding of the stimulus, categorization by application of arbitrary instructions, and response selection by evidence accumulation can all proceed, at least in part, without conscious perception. Research by Vorberg et al. (2003) supports this conclusion well. Using primes shaped as arrows pointing left or right, Vorberg et al. showed that the behavioral priming effect increased monotonically with the time interval separating the prime from the mask (while conscious prime perception was either absent or followed a non-monotonic, U-shaped curve). Those results, presented in Figure 2.4, can be captured mathematically using an evidence accumulation model similar to the one presented in Figure 2.2. Vorberg et al's model supposes that the various response alternatives are coded by leaky accumulators which receive stochastic input: first from the prime, then from the target. The accumulators add up sensory evidence until a predefined threshold is reached, after which a response is emitted. Mathematical analysis and simulations show that this model can reproduce the empirical observation of a bias in response time. At long SOAs, the model predicts that primes can also induce a high error rate, especially if the response threshold is relatively low—a prediction which is empirically supported by the data.

Role of Instruction and Attention in Subliminal Processing

Subliminal processing was previously thought to be automatic and independent of attention. In recent years, however, several effects from top-down modulation on subliminal processing have been identified.

Modulation by Instructions

Task instructions readily alter the fate of subliminal stimuli. As just described, masked primes can elicit instruction-dependent activation in the motor cortex (Dehaene et al. 1998; Eimer and Schlaghecken 1998; Leuthold and Kopp 1998; Vorberg et al. 2003). Even details of the instructions provided to subjects, such as whether they are told that the targets consist of all numbers 1

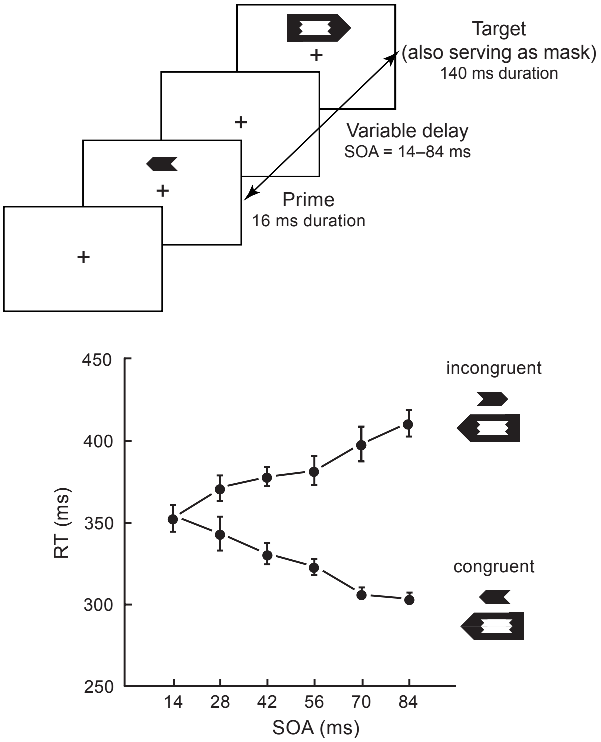

Figure 2.4 Evidence suggesting a partial accumulation of evidence from a nonconscious prime during a simple sensorimotor task (after Vorberg et al. 2003). Subjects classify target arrows as pointing right or left, while a masked prime also points left or right. A linear priming effect is seen: as the prime-target delay increases, congruent primes induce a monotonic speed-up of response times, while incongruent primes cause a monotonic slowing down. The slope of the effect is such that the difference in response time (RT) is essentially equal to the prime-target delay (SOA = stimulus onset asynchrony), suggesting that evidence is being continuously accumulated, first from the prime, then from the target.

through 9 or just the numbers 1, 4, 6 and 9, can affect subliminal priming (Kunde et al. 2003). Though still debated, those results suggest that the arbitrary stimulus-response mappings conveyed by conscious instructions can also apply to nonconscious stimuli. As noted above, within the framework of evidence accumulation models, this implies that an entire instruction set, reflected in how the stimulus is categorized and mapped onto responses, can be partially applied to a nonconscious stimulus.

Modulation by Executive Attention

Within-task changes in executive attention also seem to impact on subliminal processing. For instance, Kunde et al. (2003) studied the “Gratton effect,” a strategic increase in executive control which follows Stroop interference trials. The effect is such that, if on trial *n–1* subjects experience a cognitive conflict due to a Stroop-incongruent trial, then on trial *n* the Stroop effect is reduced, as if subjects somehow regain stronger control over the task (perhaps by focusing attention more tightly around the time of the target). Kunde et al. manipulated the consciousness of the conflict by presenting, on each trial, a subliminal or supraliminal prime followed by a conscious target. They observed that the Gratton effect could only be induced by a conscious trial (i.e., the conflict at trial *n–1* had to be a conscious conflict). Once established, however, the increase in control applied to both subliminal and supraliminal trials: the effect of conflict at trial *n* was diminished, whether or not this conflict was consciously perceived. This suggests that executive attention, once modified by a conscious stimulus, can have an impact on subsequent subliminal processing.

Modulation of Subliminal Priming by Temporal Attention

An impact of temporal attention on subliminal processing was demonstrated by Naccache et al. (2002) in a numerical masked priming paradigm. They showed that subliminal priming was present when subjects could attend to the time of presentation of the prime-target pair, but vanished when stimuli could not be temporally attended. Kiefer and Brendel (2006) observed a similar effect in an experiment investigating the N400 potential elicited by masked words. Unseen masked words elicited a much larger N400 when they were temporally attended than when they were not.

In terms of evidence accumulation models, temporal attention effects may relate to the little-studied issue of how the accumulators are reset and opened. To operate optimally, the accumulators must be emptied before each trial, and evidence must only be accumulated once the stimulus is actually present. The above effects can be interpreted as showing that semantic and decision-related evidence arising from subliminal primes fails to be accumulated whenever it is presented outside of the temporal window when the target is expected. Alternatively, if accumulated, it is reset to zero and therefore cannot bias target processing.

Modulation by Spatial Attention

Kentridge et al. (1999, 2004) first reported that blindsight patient GY could use consciously perceived cues to enhance unconscious processing of visual targets. When a target was presented in his scotoma region, patient GY responded more quickly and accurately when it was validly cued by a consciously perceptible arrow pointing to it, than when he was invalidly cued. In both cases, he still claimed that he could not see the target. Modulation of subliminal priming by spatial attention was also observed in normal subjects (Lachter et al. 2004; Marzouki et al. 2007).

In summary, task preparation includes many different components, including attention to the relevant stimulus parameter (e.g., number) and to the likely location and presentation time of the stimulus, as well as preparation of a stimulus-response mapping and setting of executive-level parameters (e.g., response threshold). Evidence suggests that essentially all of these task-preparation components, once prepared for a conscious target, apply as well to a nonconscious target.

Recent Evidence for Extended Subliminal Processing

Recently, subliminal research has gone one step further and asked whether task-preparation processes themselves can be primed subliminally. The central issue is whether processes traditionally associated with a "central executive" system can also unfold in the absence of consciouness.

Pessiglione et al. (2007) demonstrated that one aspect of task setting—motivation—could be cued subliminally. Prior to each trial of a force-generation task, subjects were presented with conscious information about the amount of money they could earn on the subsequent trial: one penny or one pound. Unknown to them, each conscious monetary cue was preceded by a subliminal image which could be congruent or incongruent with the conscious image. This subliminal information modulated the subject's motivation, as evidenced by a modulation of both the applied force and the amount of activation of a bilateral ventral pallidal region known to convey reward anticipation information.

In a similar line of research, Mattler (2003) presented a series of experiments in which a square or diamond shape successively cued increasingly abstract aspects of the task: response finger, response hand, stimulus modality (auditory or visual), or the requested task (pitch or timbre judgment). For instance, in one experiment, subjects heard a variable sound which, if preceded by a square, had to be judged for its timbre and, if preceded by a diamond, had to be judged for its pitch. Unknown to the subject, each instruction cue was preceded by a masked prime which could be congruent or incongruent with the cue. Response times were systematically shorter on congruent trials and this effect increased with the prime-mask interval in a manner which was dissociated from the U-shaped curve for conscious perception. Thus, even task selection seemed to be biased by a subliminal cue.

Unfortunately, Mattler's (2003) results could also be interpreted as a conflict at a purely visual level of cue identification; that is, the measure response time included components of cue identification, task selection, and task execution, and the observed priming might have arisen from the perceptual component alone. To demonstrate firmly that a subliminal prime could affect task selection, Lau and Passingham (2007) resorted to functional imaging. They selected tasks of phonological versus semantic judgment on visual words that are associated with broadly different cortical networks. Using a design similar to Mattler's, they then showed that not only the response time but the entire task-related network was modulated up or down as a function of whether the subliminal prime was congruent or incongruent with the task information provided by the visible cue. This subliminal task-cueing effect was not sufficient to reverse the conscious task cue, but it did yield an increase in subjects' error rate.

One last paradigm of relevance to the present discussion was developed by van Gaal et al. (2007). They showed that a subliminal cue could fulfill the role of a "stop signal" requiring subjects to interrupt their ongoing response to a main task. Unconscious stop signals yielded a minuscule but still significant slowing down of response time and increase in errors. Thus, subliminal stimuli can trigger the first hints of a task interruption.

How can one interpret such high-level priming effects? One possibility is that, even at the "central executive" level, task selection and task control processes continue to operate according to rules of evidence accumulation, which can be biased by subliminal priming. According to the model illustrated in Figure 2.2, subjects select a motor response by forming two pools of units: those accumulating evidence for response R1 and those accumulating evidence for response R2. Perhaps the "central executive" consists of nothing but similar decision mechanisms organized in a control hierarchy (Koechlin et al. 2003). At a higher level, similar evidence accumulation processes would be involved in the selection of one of two tasks, T1 and T2. Those accumulators would accrue evidence provided by conscious cues, but also by subliminal cues. Sigman and Dehaene (2006) presented precisely such a model of task selection in a dual-task context. They showed how the time to select which task to perform added a variable duration to the overall response time which could be captured well by an accumulator model. It remains to be seen whether these ideas can be extended to an entire hierarchy of interacting decision systems, as proposed, for example, by Koechlin et al (2003).

Limits to Subliminal Processing

Given this wealth of evidence which indicates that subliminal processing can extend to a high cognitive level, one may reasonably ask if there are any limits to subliminal processing. Are there mental processes that can be executed only once conscious perception has occurred? This question naturally arises in relation to the evolutionary role of consciousness. Although the evidence remains fragmentary, several mental operations can be associated speculatively with conscious-level processing.

Durable and Explicit Information Maintenance

Priming experiments show that subliminal information tends to be short-lived: after about 500 ms, priming effects typically cease to be detectable (Greenwald et al. 1996; Mattler 2005). To bridge delays of a few seconds, information is thought to be stored in working memory by active populations of neurons, particularly in prefrontal cortex. When information reaches this working memory stage, Dehaene and Naccache (2001) have suggested that it is always consciously accessible. Kunde et al.'s (2003) work, reviewed above, fits nicely with this conclusion, since it shows that only the conscious variables of trial *n*–1 can be carried out to trial *n*. Similar evidence is provided by the trace-conditioning paradigm, in which conditioning across a temporal gap only occurs if subjects report being aware of the relations among the stimuli (Clark et al. 2002). Additional supporting data has been reviewed by Dehaene and Naccache (2001). Altogether, the evidence points to a crucial role of consciousness in bridging information across a delay.

Global Access and Novel Combinations of Operations

Consciousness has been suggested to play an essential role in the expression of novel behaviors that require putting together evidence from multiple sources (e.g., by confronting evidence spread across several trials). For instance, Merikle et al. (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 occurred: 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 take advantage of the predictability of the target from the prime to become faster on incongruent trials than on congruent trials, thus inverting the Stroop effect. Crucially, this strategic inversion occurred only 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 and Joordens 1997). Here, 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.

This conclusion may need to be qualified in the light of recent evidence, reviewed above, that is, task switching or task stopping can be modulated partially by subliminal cues (Lau and Passingham 2007; Mattler 2003; van Gaal et al. 2007). Note, however, that this evidence was always obtained under conditions of highly routinized performance. Subjects performed hundreds of trials with consciously perceived task cues before the same cues, presented subliminally, began to affect task choice. This is very different from the rapid deployment of novel strategies that, presumably, can only be deployed under conscious conditions.

Intentional Action

As noted by Dehaene and Naccache (2001), the spontaneous generation of intentional behavior may constitute a third property specifically associated with conscious perception. It is noteworthy that, in all of the above priming tasks, although subliminal primes modulate the response time to another conscious stimulus, they almost never induce a full-blown behavior in and of themselves. Only on a very small proportion of trials do subliminal primes actually cause overt responses. When they do, such trials are typically labeled as unintended errors by the subject (and by the experimenter).

As a related example, consider the case of blindsight patients (Weiskrantz 1997). Some of these patients, even though they claim to be blind, show an excellent performance in pointing to objects. As noted by Dennett (1992) and Weiskrantz (1997), a fundamental difference with normal subjects, however, is that 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.

In summary, nonconscious stimuli do not seem to reach a stage of processing at which information representation enters into a deliberation process that supports voluntary action with a sense of ownership. If they do reach this stage, it is only with a trickle of activation that modulates decision time but does not determine the decision outcome.

Cerebral Bases of Conscious and Nonconscious Computations

The hypothesis that conscious information is associated with a second stage of processing that cannot be deployed fully for subliminal stimuli meshes well with recent experiments that have directly compared the brain activation evoked by conscious versus nonconscious stimuli. Many such experiments have been performed with fMRI, and they converge to suggest that, relative to a masked stimulus, an unmasked stimulus is amplified and gains access to high levels of activation in prefrontal and parietal areas (Dehaene et al. 2006; Dehaene et al. 2001; Haynes et al. 2005; for review and discussion, see Kouider et al. 2007). Most relevant to the present discussion are time-resolved experiments using ERPs or MEG that have followed the processing of a stimulus in time as it crosses or does not cross the threshold for conscious perception. My colleagues and I have performed such experiments under conditions in which invisibility was created either by masking (Del Cul et al. 2007; see also Koivisto et al. 2006; Melloni et al. 2007; van Aalderen-Smeets et al. 2006) or by inattention during the attentional blink (Gross et al. 2004; Kranczioch et al. 2003; Sergent et al. 2005). In both cases, we were able to analyze a subset of trials in which the very same stimulus was presented, but was or was not consciously perceived according to subjective reports.

The results were highly convergent in coarsely separating two periods of stimulus processing. During the first 270 ms, brain activation unfolded in an essentially unchanged manner whether or not the stimuli were consciously perceived. Strong visual activation was seen, quickly extending to the ventral temporal visual pathway. In the case of the attentional blink, the nonconscious activation extended even further in time, with very strong left lateral temporal activity around 400 ms plausibly associated with semantic-level processing (see also Luck et al. 1996). However, around 270 ms, an important divergence occurred, with a sudden surge of additional activation being observed on conscious trials only. Over a few tens of milliseconds, activation expanded into bilateral inferior and dorsolateral frontal regions, anterior cingulate cortex, and posterior parietal cortex. As shown in Figure 2.5, this activity was reduced drastically on nonconscious trials: only short-lived activation was seen, quickly decaying towards zero about 500 ms after stimulus presentation. By contrast, activation seemed to be amplified actively on conscious trials.

The parsing of brain activation into two stages— early activation by subliminal stimuli, followed by late global amplification and reverberation—seems to be a generic phenomenon that can be observed in various stimulus modalities, by a variety of methods, and in multiple species. Thus Victor Lamme and collaborators (2002), using electrophysiological recordings in macaque area V1, have distinguished early feed-forward versus late feedback responses. They found that only the latter were sensitive to attention and reportability. Using intracranial electrodes in human epileptic patients, my team has obtained evidence for a similar division in human subjects during subliminal versus conscious word reading (Gaillard, Naccache et al. 2006; Naccache et al. 2005). In many electrodes, subliminal words evoked only a first peak of activation whereas conscious words evoked a similar but magnified peak followed by a sustained period of activation.

To give yet a third example, Nieuwenhuis et al. (2001) used ERPs in humans to track error detection and compensation processes. When subjects made an undetected erroneous saccade, an early error-related negativity was observed over mesial frontal electrodes, presumably reflecting a nonconscious triggering of an anterior cingulate system for error detection. However, only when the error was detected consciously was this early waveform amplified and followed by a massive P3-like waveform associated presumably with the expansion of activation into a broader cortical and subcortical network.

A Global Workspace Model of Conscious Access

Jean-Pierre Changeux and I have suggested that these global self-amplifying properties of brain activation during conscious access can be accounted for by the concept of a "global workspace" (Dehaene and Changeux 2005; Dehaene and Naccache 2001; Dehaene et al. 2003). This model, which has been backed

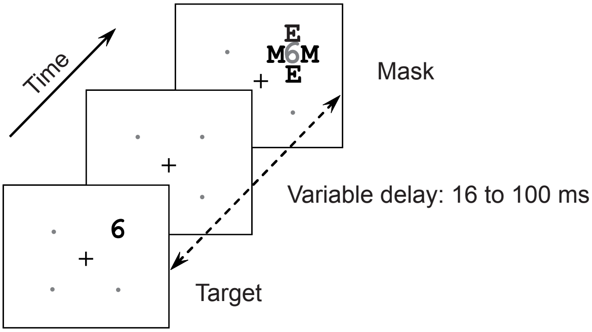

1. Effects of target-mask delay

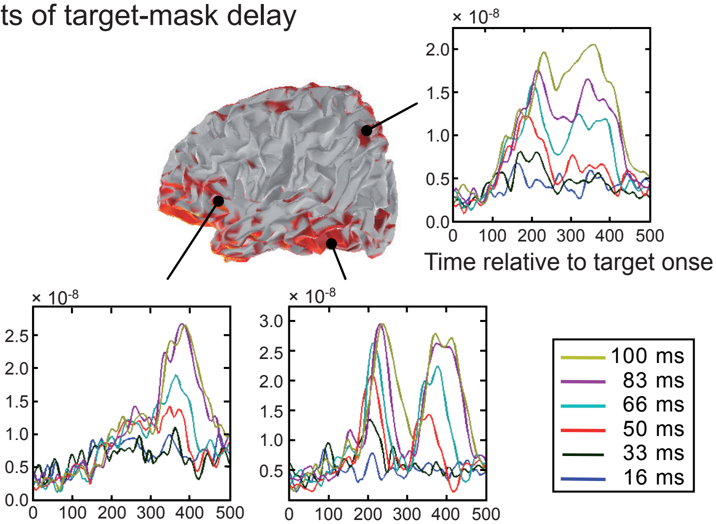

2. Effects of target visibility at a fixed delay (50 ms)

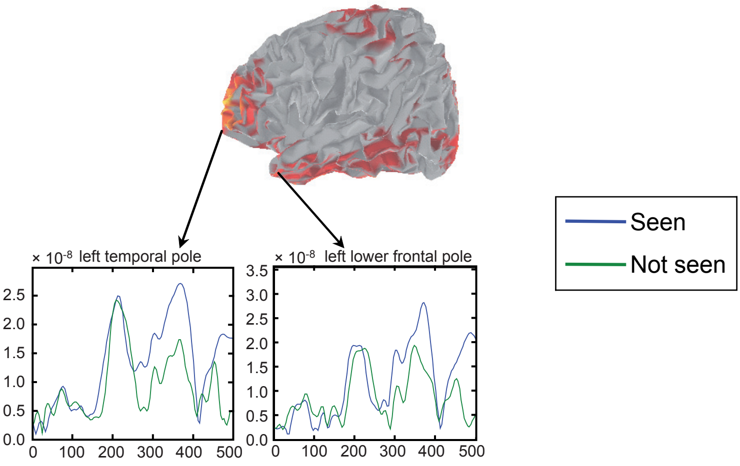

up with explicit computer simulations of realistic thalamo-cortical networks, supposes that access to consciousness again corresponds to a form of accumulation of activation within a recurrently connected network. However, this accumulation is postulated to occur, not just locally, but within a highly distributed set of columns coding for the same object within distinct brain areas. These columns are interconnected in a reciprocal manner by distinct cortical "workspace neurons" with long-distance axons. As a result, an entire set of distributed brain areas can function temporarily as a single integrator, with a strong top-down component such that higher association areas send supportive signals to the sensory areas that first excited them.

Computer simulations show that such a network, when stimulated by a brief pulse of activation, presents complex dynamics with at least two distinct stages. In the first stage, activation climbs up the thalamo-cortical hierarchy in a feed-forward manner. As it does, the higher levels send increasingly stronger top-down amplification signals. If the incoming signal is strong enough, then at a certain point a dynamic threshold is crossed and activation becomes self-amplifying and increases in a nonlinear manner. During this second stage, the whole distributed assembly coding for the stimulus at multiple hierarchical levels then "ignites" into a single synchronously activated state. In peripheral neurons, this creates a late second peak of sustained firing. The corresponding brain state is illustrated schematically in Figure 2.6.

Why would this global brain state correspond to conscious access? Computer simulations show that once stimulus-evoked activation has reached highly interconnected associative areas, two important changes occur:

1. The activation can now reverberate for a long time period, thus holding information on-line for a duration essentially unrelated to the initial stimulus duration.
2. Stimulus information represented within the global workspace can be propagated rapidly to many brain systems.

Figure 2.5 Changes in brain activity associated with crossing the threshold for conscious perception during masking (after De Cul et al. 2007). The paradigm is described in Figure 2.1, and involves varying the delay between a digit and the subsequent mask. Event-related potentials are recorded with a 128-channel electrode net and reconstructed on the cortical surface with BrainStorm software. As the delay increases, thus rendering the stimulus increasingly visible, activation increases monotonically in posterior areas, then a threshold effect is seen. The late part of the activation (beyond 270 ms) suddenly increases nonlinearly in a sigmoidal manner once the delay exceeds a critical value which coincides with the threshold value for conscious perception. This nonlinear activation is highly global and occurs simultaneously in inferior and anterior prefrontal cortex as well as in posterior parietal and ventral occipito-temporal cortices. Even when the delay is fixed, the same results are seen when sorting the individual trials into seen versus not-seen (bottom panel): there is a clear separation between an initial period where activation is identical for seen and not-seen trials, and a later period (>270 ms) where activation suddenly re-increases globally on seen trials.

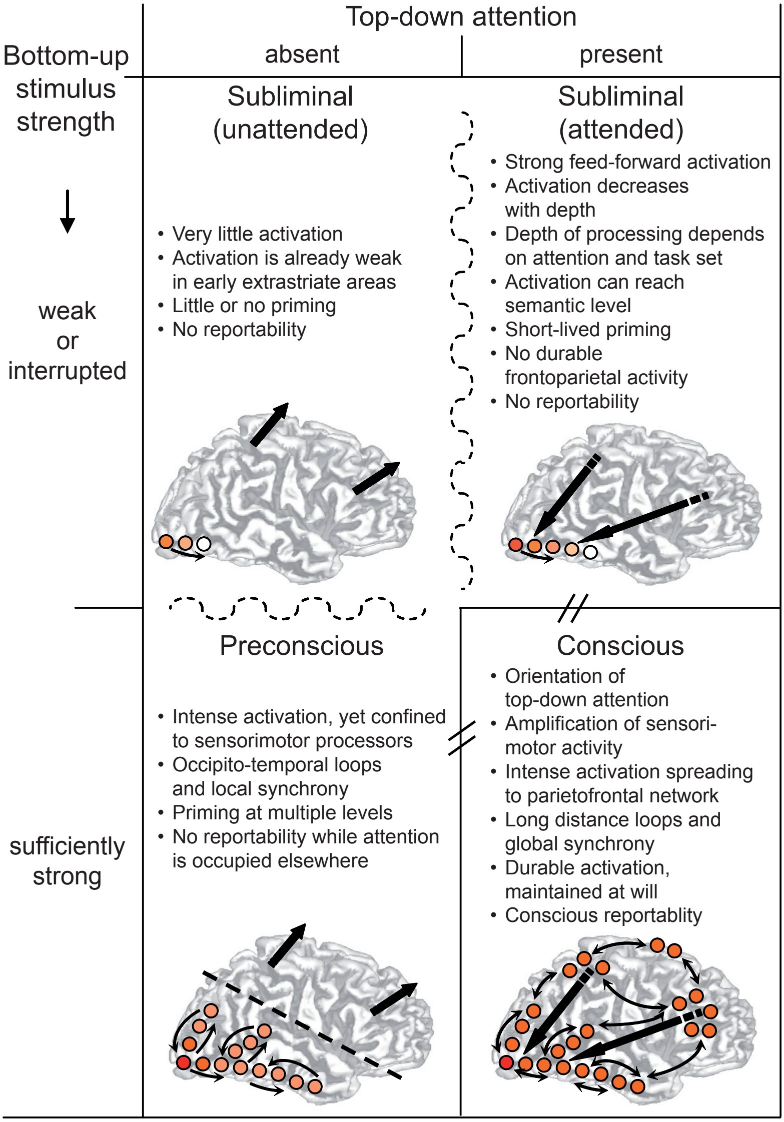

Figure 2.6 Theoretical proposal of a distinction between brain states of subliminal, preconscious, and conscious processing (after Dehaene et al. 2006). Conscious processing occurs when the accumulated stimulus-evoked activation exceeds a threshold and evokes a dynamic state of global reverberation ("ignition") across multiple high-level cortical areas forming a "global neuronal workspace," particularly involving prefrontal, cingulate and parietal cortices (bottom right). These areas can maintain the information on-line and broadcast it to a variety of other processors, thus serving as a central hub for global access to information—a key property of conscious states. Subliminal processing corresponds to a data-limited situation where only a trickle of

We argue that both properties are characteristic of conscious information processing. As noted above, the information can be maintained in time, buffered from fast fluctuations in sensory inputs, and can be shared across a broad variety of processes including evaluation, verbal report, planning, and long-term memory (Baars 1989).

Anatomically, the model postulates that workspace neurons are particularly dense in prefrontal, parietal, and anterior cingulate cortices, thus explaining why these regions are recurrently found to be associated with conscious access across various paradigms and modalities (Dehaene et al. 2006). However, according to the model, workspace neurons are also present to variable degrees in essentially all of the cortex, thus permitting essentially any active cortical contents to be brought together into a single brain-scale assembly. Indeed, it would seem likely that this long-distance network has been subject to a particular selective pressure in humans. A number of recent observations support this possibility, including (a) the disproportionate increase of prefrontal white matter volume in our species (Schoenemann et al. 2005), (b) the massive increase in dendritic branching and spine density in prefrontal cortex across the primate lineage (Elston 2003); and (c) the presence in anterior cingulate cortex of large projection neurons ("spindle cells") seemingly unique to humans and great apes (Nimchinsky et al. 1999).

Accounting for Subliminal Processing

The proposed workspace architecture separates, in a first minimal description, two computational spaces, each characterized by a distinct pattern of connectivity. Subcortical networks and most of the cortex can be viewed as a collection of specialized and automatized processors, each attuned to the processing of a particular type of information via a limited number of local or medium-range connections that bring to each processor the "encapsulated" inputs necessary to its function. On top of this automatic level, we postulate 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, thus allowing many different processors to exchange information.

---

partial evidence is able to propagate through specialized cerebral networks, yet without reaching a threshold for global ignition and thus without global reportability (top line). The orientation and depth of subliminal processing may nevertheless depend on the top-down state of attention (top right). A distinct nonconscious state, preconscious processing, corresponds to a resource-limited situation where stimulus processing is blocked at the level of the global neuronal workspace while it is temporarily occupied by another task. A preconscious stimulus may be temporarily buffered within peripheral sensory areas and later accessed by the fronto-parietal system once it is released by its distracting task. In this case, information switches from nonconscious to conscious.

Can this model explain observations on subliminal processing? According to the proposed model, subliminal processing corresponds to a condition of specialized processing without global information accessibility (see Figure 2.6). A subliminal stimulus is a stimulus that possesses sufficient energy to evoke a feed-forward wave of activation in specialized processors, but it has insufficient energy or duration to trigger a large-scale reverberating state in a global network of neurons with long-range axons. As explained above, simulations of a minimal thalamo-cortical network (Dehaene and Changeux 2005) indicate that such a nonlinear self-amplifying system possesses a well-defined dynamic threshold. While it has been observed that activation exceeding a threshold level grows quickly into a full-scale ignition, a slightly weaker activation propagates forward, sometimes all the way into higher areas. It, however, loses its self-supporting activation and dies out quickly. Subliminal processing would correspond to the latter type of network state.

Let us examine briefly how this schematic model may account for the data reviewed in the preceding sections. We have seen that a masked visual stimulus that is not consciously reportable is nevertheless processed at multiple levels, including visual but also semantic, executive, and motor levels. These observations mesh well with the notion of an ascending wave of feed-forward activation that begins to accumulate within decision systems, but does not lead to a full-blown activation crossing the response threshold. Recent theorizing suggests that local neural assemblies recurrently interconnected by glutamatergic synapses with a mixture of AMPA and NMDA receptors can operate as accumulators of evidence (Wong and Wang 2006). The global workspace model suggests that such multiple integrators can operate in parallel during subliminal processing, each integrating evidence for or against their preferred stimulus. In priming experiments, where a subliminal stimulus is followed by a supraliminal target, this partial accumulation of evidence evoked by the prime would shift the baseline starting level of these accumulators, thus creating priming effects in response time, determined primarily by the congruity of the prime and target.

As long as the prime-based accumulation remains subthreshold, and therefore fails to trigger a global recurrent assembly, there is nothing in the global workspace model that prevents subliminal processing from occurring at any cognitive level, including higher-level control processes. However, the model predicts that only the most specialized processors, tightly attuned to the stimulus, should be capable of activating strongly to a subliminal stimulus. This prediction meshes well with the narrow localized activation measured by fMRI and intracranial recordings in response to subliminal words and digits (Dehaene et al. 2001; Naccache and Dehaene 2001a; Naccache et al. 2005). Note that, under the model's hypotheses, subliminal processing is not confined to a passive spreading of activation, independent of the subject's attention and strategies, as previously envisaged. On the contrary, whichever task and attentional set are prepared consciously, it can serve to orient and amplify the processing of a subliminal stimulus, even if its bottom-up strength remains insufficient for global ignition. This aspect of our model agrees with the many top-down influences on subliminal processing that have been observed experimentally.

Finally, the model predicts correctly that subliminal activation may be very strong within the first 100–300 ms after stimulus presentation, but progressively dies out in the next few hundreds of milliseconds as time elapses and as the stimulus reaches higher levels of representation. Such a decay of subliminal activation, both in time and in cortical space, has indeed been observed experimentally with high-density recordings of event-related potentials (Del Cul et al. 2007; see Figure 2.5). It can explain why only small behavioral influences of subliminal stimuli are measurable at higher cognitive levels (van Gaal et al. 2007; Mattler 2003), and why most if not all subliminal priming effects decay to a nonmeasurable level once the prime-target interval exceeds 500 ms (Mattler 2005). Only very rarely are subliminal effects seen beyond the range of a few seconds. My colleagues and I have suggested that when they do (Gaillard et al. 2007), it may be because the subliminal stimulus has caused structural changes (e.g., changes in synaptic efficacy) rather than it being due to lingering brain activity.

A Distinct State of Preconscious Processing

Simulations of the global workspace have revealed that global workspace ignition can also be prevented in a different manner, suggesting a distinct state of nonconscious processing that we have proposed to call *preconscious* (or potentially conscious, or P-conscious). Contrary to subliminal processing, where the incoming stimulus itself does not have enough energy or duration to trigger a supra-threshold reverberation of activation, preconscious processing corresponds to a neural process that potentially carries enough activation for conscious access, but is temporarily blocked from activating the global workspace due to its transient occupancy by another stimulus. Simulations have shown that such a competitive interaction for global access can occur when two stimuli are presented in short succession, in a paradigm akin to the "attentional blink." The first target (T1) creates a global workspace ignition, but while this global state is occurring, lateral inhibition prevents a second target (T2) from entering the workspace. Essentially, the global workspace acts as a central bottleneck (Chun and Potter 1995; Pashler 1994) whose occupancy by T1 deprives the T2-evoked neural assembly from its top-down support. The corresponding postulated brain state is illustrated schematically in Figure 2.6.

Computer simulations (Dehaene and Changeux 2005) suggest that during preconscious processing, T2 activation is blocked sharply at the central level; it can, however, be quite strong at peripheral levels of processing. It may excite resonant loops within medium-range connections that may maintain the representation of the stimulus temporarily active in a sensory buffer for a few hundreds of milliseconds. As a result, a preconscious stimulus is literally on the brink of consciousness and can compete actively for conscious access with other stimuli, including the currently conscious one. Furthermore, although temporarily blocked, a preconscious stimulus may later achieve conscious access once the central workspace is freed. This aspect of the model may correspond to the empirical observation of a "psychological refractory period" in behavioral dual-task performance (Pashler 1984; Sigman and Dehaene 2005), in which one task is put on hold while another task is being processed. The model assumes that the key difference between the psychological refractory period and attentional blink phenomena is the possibility of a lingering of T2-induced activation in peripheral circuits. T2 may never gain access to conscious processing if its preconscious representation is erased prior to the orienting of top-down attention (as achieved by masking in the attentional blink paradigm).

At present, only a few studies have examined brain activity during states where conscious access is prevented by top-down attentional competition, such as the attentional blink (for review, see Marois and Ivanoff 2005). Time-resolved experiments suggest that the initial activation by an unseen T2 can be very strong and essentially indistinguishable from that evoked by a conscious stimulus during a time window of about 270 ms (Sergent et al. 2005). The attentional blink then creates a sudden blocking of part of the activation starting around 270 ms, particularly in inferior prefrontal cortex (Sergent et al. 2005), and a global state of fronto-parietal synchrony indexed by the scalp P3 and by evoked oscillations in the beta range is prevented from occurring (Gross et al. 2004; Kranczioch et al. 2003). Other fMRI experiments also point to a distributed prefronto-parietal network as the main locus of the bottleneck effect in competition paradigms, consistent with the global workspace model (Dux et al. 2006; Kouider et al. 2007).

Conclusion: Conscious Access as a Solution to von Neumann's Problem?

The purpose of this chapter was to survey the rich cognitive neuroscience literature on nonconscious processing and to establish links with evidence accumulation models. The main generalizations that I have proposed to draw from these observations are the following:

1. *Subliminal processing* corresponds to a state of partial accumulation of evidence within multiple sensory, semantic, executive, and motor networks, yet without reaching a full-blown decision threshold.
2. Nonconscious processing can also occur in a distinct state of *preconscious processing,* where evidence accumulation can proceed normally within posterior sensory and semantic networks while being blocked from

accessing anterior networks due to competition with another attended mental representation.

3. *Conscious access* is associated with the crossing of a dynamic threshold beyond which the stimulus activation reverberates within a global fronto-parietal network. The sensory representation of the stimulus can thus be maintained online and be used for higher-level executive processes, such as reasoning and decision making.

I end with a final speculative note on one of the possible functions of consciousness in evolution. In his 1958 book, *The Computer and the Brain*, von Neumann asked how a biological organ such as the brain, where individual neurons are prone to errors, could perform multistep calculations. He pointed out that in any analogical machine, errors accumulate at each step so that the end result quickly becomes imprecise or even useless. He therefore suggested that the brain must have mechanisms that discretize the incoming analogical information, much like the TTL or CMOS code of current digital chips is based on a distinction of voltages into high (between 4.95 and 5 volts) versus low (between 0 and 0.05 volts).

Tentatively, I surmise that the architecture of the "conscious workspace" may have evolved to address von Neumann's problem. In the human brain, one function of conscious access would be to control the accumulation of information in such a way that information is pooled in a coherent manner across the multiple processors operating preconsciously and in parallel, and a discrete categorical decision is reached before being dispatched to yet other processors. By pooling information over time, this global accumulation of evidence would allow the inevitable errors that creep up during analog processing to be corrected or at least to be kept below a predefined probability level. Many decision models already postulate such an accumulation of evidence within local brain systems such as the oculo-motor system (see, e.g., Shadlen, this volume). The role of the conscious global workspace would be to achieve such accumulation of evidence in a unified manner across multiple distributed brain systems and, once a single coherent result has been obtained, to dispatch it back to essentially any brain processor as needed by the current task. This architecture would permit the execution of a multistep mental algorithm through successive, consciously controlled steps of evidence accumulation followed by result dispatching. The latter proposal is consistent with recent findings from the "psychological refractory period" paradigm, where response time in a dual-task situation was shown to result from a temporal succession of multiple non-overlapping stochastic accumulation periods (Sigman and Dehaene 2005, 2006).

While clearly speculative and in need of further specification, the proposed architecture seems to combine the benefits of two distinct computational principles: massive parallel accumulation of evidence at a nonconscious level, followed by conscious broadcasting of the outcome permitting the operation of the human brain as a slow serial "Turing machine."

References

Baars, B. J. 1989. A Cognitive Theory of Consciousness. Cambridge, MA: Cambridge Univ. Press.

Chun, M. M., and M. C. Potter. 1995. A two-stage model for multiple target detection in rapid serial visual presentation. *J. Exp. Psychol.: Hum. Perc. Perf.* **21(1)**:109–127.

Clark, R. E., J. R. Manns, and L. R. Squire. 2002. Classical conditioning, awareness, and brain systems. *Trends Cogn. Sci.* **6(12)**:524–531.

Dehaene, S. 2007. Symbols and quantities in parietal cortex: Elements of a mathematical theory of number representation and manipulation. In: Sensorimotor Foundations of Higher Cognition, ed. P. Haggard and Y. Rossetti. Oxford: Oxford Univ. Press, in press.

Dehaene, S., and J. P. Changeux. 2005. Ongoing spontaneous activity controls access to consciousness: A neuronal model for inattentional blindness. *PLoS Biol.* **3(5)**:e141.

Dehaene, S., J. P. Changeux, L. Naccache, J. Sackur, and C. Sergent. 2006. Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends Cogn. Sci.* **10(5)**:204–211.

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

Dehaene, S., L. Naccache, L. Cohen, D. L. Bihan, J. F. Mangin et al. 2001. Cerebral mechanisms of word masking and unconscious repetition priming. *Nat. Neurosci.* **4(7)**:752–758.

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

Dehaene, S., C. Sergent, and J. P. Changeux. 2003. A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proc. Natl. Acad. Sci.* **100**:8520–8525.

Del Cul, A., S. Baillet, and S. Dehaene. 2007. Brain dynamics underlying the non-linear threshold for access to consciousness. *PLoS Biol.*, **5(10)**:e260.

Del Cul, A., S. Dehaene, and M. Leboyer. 2006. Preserved subliminal processing and impaired conscious access in schizophrenia. *Arch. Gen. Psychiatry* **63(12)**:1313–1323.

Dennett, D. C. 1992. Consciousness Explained. London: Penguin.

Devlin, J. T., H. L. Jamison, P. M. Matthews, and L. M. Gonnerman. 2004. Morphology and the internal structure of words. *Proc. Natl. Acad. Sci.* **101(41):**14,984–14,988.

Dux, P. E., J. Ivanoff, C. L. Asplund, and R. Marois. 2006. Isolation of a central bottleneck of information processing with time-resolved fMRI. *Neuron* **52(6)**:1109–1120.

Eimer, M., and F. Schlaghecken. 1998. Effects of masked stimuli on motor activation: Behavioral and electrophysiological evidence. *J. Exp. Psychol.: Hum. Perc. Perf.* **24(6)**:1737–1747.

Elston, G. N. 2003. Cortex, cognition and the cell: New insights into the pyramidal neuron and prefrontal function. *Cerebral Cortex* **13(11)**:1124–1138.

Gaillard, R., L. Cohen, C. Adam, S. Clémenceau, M. Baulac et al. 2007. Subliminal words durably affect neuronal activity. *Neuroreport* **18(15)**:1527–1531.

Gaillard, R., A. Del Cul, L. Naccache, F. Vinckier, L. Cohen et al. 2006. Nonconscious semantic processing of emotional words modulates conscious access. *Proc. Natl. Acad. Sci.* **103(19)**:7524–7529.

Gaillard, R., L. Naccache, P. Pinel, S. Clemenceau, E. Volle et al. 2006. Direct intracranial, fmri, and lesion evidence for the causal role of left inferotemporal cortex in reading. *Neuron* **50(2)**:191–204.

Greenwald, A. G., S. C. Draine, and R. L. Abrams. 1996. Three cognitive markers of unconscious semantic activation. *Science* **273(5282)**:1699–1702.

Gross, J., F. Schmitz, I. Schnitzler, K. Kessler, K. Shapiro et al. 2004. Modulation of long-range neural synchrony reflects temporal limitations of visual attention in humans. *Proc. Natl. Acad. Sci.* **101(35)**:13,050–13,055.

Haynes, J. D., J. Driver, and G. Rees. 2005. Visibility reflects dynamic changes of effective connectivity between V1 and fusiform cortex. *Neuron* **46(5)**:811–821.

Kentridge, R. W., C. A. Heywood, and L. Weiskrantz. 1999. Attention without awareness in blindsight. *Proc. Roy. Soc. Lond. B* **266(1430)**:1805–1811.

Kentridge, R. W., C. A. Heywood, and L. Weiskrantz. 2004. Spatial attention speeds discrimination without awareness in blindsight. *Neuropsychologia* **42(6)**:831–835.

Kiefer, M. 2002. The N400 is modulated by unconsciously perceived masked words: Further evidence for an automatic spreading activation account of N400 priming effects. *Brain Res. Cogn. Brain Res.* **13(1)**:27–39.

Kiefer, M., and D. Brendel. 2006. Attentional modulation of unconscious "automatic" processes: Evidence from event-related potentials in a masked priming paradigm. *J. Cogn. Neurosci.* **18(2)**:184–198.

Koechlin, E., C. Ody, and F. Kouneiher. 2003. The architecture of cognitive control in the human prefrontal cortex. *Science* **302(5648)**:1181–1185.

Koivisto, M., A. Revonsuo, and M. Lehtonen. 2006. Independence of visual awareness from the scope of attention: An electrophysiological study. *Cereb. Cortex* **16(3)**:415–424.

Kouider, S., and S. Dehaene. 2007. Levels of processing during non-conscious perception: A critical review of visual masking. *Phil. Trans. Roy. Soc. Lond. B,* **362(1481)**:857–875.

Kouider, S., S. Dehaene, A. Jobert, and D. Le Bihan. 2007. Cerebral bases of subliminal and supraliminal priming during reading. *Cereb. Cortex* 17(9):2019–2029.

Kranczioch, C., S. Debener, and A. K. Engel. 2003. Event-related potential correlates of the attentional blink phenomenon. *Brain Res. Cogn. Brain Res.* **17(1)**:177–187.

Kunde, W., A. Kiesel, and J. Hoffmann. 2003. Conscious control over the content of unconscious cognition. *Cognition* **88(2)**:223–242.

Lachter, J., K. I. Forster, and E. Ruthruff. 2004. Forty-five years after Broadbent (1958): Still no identification without attention. *Psychol. Rev.* **111(4)**:880–913.

Laming, D. R. J. 1968. Information Theory of Choice-Reaction Times. London: Academic Press.

Lamme, V. A., K. Zipser, H. and Spekreijse. 2002. Masking interrupts figure-ground signals in V1. *J. Cogn. Neurosci.* **14(7)**:1044–1053.

Lau, H. C., and R. E. Passingham. 2007. Unconscious activation of the cognitive control system in the human prefrontal cortex. *J. Neurosci*. **27(21)**:5805–5811.

Leuthold, H., and B. Kopp. 1998. Mechanisms of priming by masked stimuli: Inferences from event-related potentials. *Psychol. Sci.* **9**:263–269.

Link, S. W. 1992. The Wave Theory of Difference and Similarity. Hillsdale, NJ: Lawrence Erlbaum.

Luck, S. J., E. K. Vogel, and K. L. Shapiro. 1996. Word meanings can be accessed but not reported during the attentional blink. *Nature* **383(6601)**:616–618.

Marois, R., and J. Ivanoff. 2005. Capacity limits of information processing in the brain. *Trends Cogn. Sci.* **9(6)**:296–305.

Marzouki, Y., J. Grainger, and J. Theeuwes. 2007. Exogenous spatial cueing modulates subliminal masked priming. *Acta Psychol.* **126(1)**:34–45.

Mattler, U. 2003. Priming of mental operations by masked stimuli. *Perc. Psychophys.* **65(2)**:167–187.

Mattler, U. 2005. Inhibition and decay of motor and nonmotor priming. *Perc. Psychophys.* **67(2)**:285–300.

Melloni, L., C. Molina, M. Pena, D. Torres, W. Singer et al. 2007. Synchronization of neural activity across cortical areas correlates with conscious perception. *J. Neurosci.* **27(11)**:2858–2865.

Merikle, P. M., and S. Joordens. 1997. Parallels between perception without attention and perception without awareness. *Conscious Cogn.* **6(2–3)**:219–236.

Merikle, P. M., S. Joordens, and J. A. Stolz. 1995. Measuring the relative magnitude of unconscious influences. *Conscious. Cogn.* **4**:422–439.

Naccache, L., E. Blandin, and S. Dehaene. 2002. Unconscious masked priming depends on temporal attention. *Psychol. Sci.* **13**:416–424.

Naccache, L., and S. Dehaene. 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., and S. Dehaene. 2001b. Unconscious semantic priming extends to novel unseen stimuli. *Cognition* **80(3)**:215–229.

Naccache, L., R. Gaillard, C. Adam, D. Hasboun, S. Clémenceau et al. 2005. A direct intracranial record of emotions evoked by subliminal words. *Proc. Natl. Acad. Sci.* **102**:7713–7717.

Nakamura, K., S. Dehaene, A. Jobert, D. Le Bihan, and S. Kouider. 2005. Subliminal convergence of Kanji and Kana words: Further evidence for functional parcellation of the posterior temporal cortex in visual word perception. *J. Cogn. Neurosci.* **17(6)**:954–968.

Nieuwenhuis, S., K. R. Ridderinkhof, J. Blom, G. P. Band, and A. Kok. 2001. Error-related brain potentials are differentially related to awareness of response errors: Evidence from an antisaccade task. *Psychophys.* **38(5)**:752–760.

Nimchinsky, E. A., E. Gilissen, J. M. Allman, D. P. Perl, J. M. Erwin et al. 1999. A neuronal morphologic type unique to humans and great apes. *Proc. Natl. Acad. Sci.* **96(9)**:5268–5273.

Pashler, H. 1984. Processing stages in overlapping tasks: Evidence for a central bottleneck. *J. Exp. Psychol.: Hum. Perc. Perf.* **10(3)**:358–377.

Pashler, H. 1994. Dual-task interference in simple tasks: Data and theory. *Psychol. Bull.* **116(2)**:220–244.

Pessiglione, M., L. Schmidt, B. Draganski, R. Kalisch, H. Lau et al. 2007. How the brain translates money into force: A neuroimaging study of subliminal motivation. *Science* **316**:904–906.

Reuter, F., A. Del Cul, B. Audoin, I. Malikova, L. Naccache et al. 2007. Intact subliminal processing and delayed conscious access in multiple sclerosis. *Neuropsychologia* **45(12)**:2683–2691.

Reynvoet, B., M. Brysbaert, and W. Fias. 2002. Semantic priming in number naming. *Qtly. J. Exp. Psychol.* **55(4)**:1127–1139.

Schoenemann, P. T., M. J. Sheehan, and L. D. Glotzer. 2005. Prefrontal white matter volume is disproportionately larger in humans than in other primates. *Nat. Neurosci.* **8(2)**:242–252.

Sergent, C., S. Baillet, and S. Dehaene. 2005. Timing of the brain events underlying access to consciousness during the attentional blink. *Nat. Neurosci.* **8(10)**:1391–1400.

Sergent, C., and S. Dehaene. 2004. Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychol. Sci.* **15(11)**:720–728.

Sigman, M., and S. Dehaene. 2005. Parsing a cognitive task: A characterization of the mind's bottleneck. *PLoS Biol.* **3(2)**:e37.

Sigman, M., and S. Dehaene. 2006. Dynamics of the Central Bottleneck: Dual-task and task uncertainty. *PLoS Biol.* **4(7)**:e220.

Smith, P. L., and R. Ratcliff. 2004. Psychology and neurobiology of simple decisions. *Trends Neurosci.* **27(3)**:161–168.

Usher, M., and J. L. McClelland. 2001. The time course of perceptual choice: The leaky, competing accumulator model. *Psychol. Rev.* **108(3)**:550–592.

van Aalderen-Smeets, S. I., R. Oostenveld, and J. Schwarzbach. 2006. Investigating neurophysiological correlates of metacontrast masking with magnetoencephalography. *Adv. Cogn. Psychol,* **2(1)**:21–35.

van Gaal, S., K. R. Ridderinkhof, W. P. M. van den Wildenberg, and V. A. Lamme. 2007. Exploring the boundaries of unconscious processing: Response inhibition can be triggered by masked stop-signals. *J. Vision* **7(9)**:425.

von Neumann, J. 1958. The Computer and the Brain. London: Yale Univ. Press.

Vorberg, D., U. Mattler, A. Heinecke, T. Schmidt, and J. Schwarzbach. 2003. Different time courses for visual perception and action priming. *Proc. Natl. Acad. Sci.* **100(10)**:6275–6280.

Weiskrantz, L. 1997. Consciousness Lost and Found: A Neuropsychological Exploration. New York: Oxford Univ. Press.

Wong, K. F., and X. J. Wang. 2006. A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* **26(4)**:1314–1328.