

Time, Action, and Consciousness

Axel Cleeremans
Cognitive Science Research Unit
Université Libre de Bruxelles CP 191
Avenue F.-D. Roosevelt, 50
1050 Bruxelles — BELGIUM
Email: axcleer@ulb.ac.be

Jean-Christophe Sarrazin
Cognitive Science Research Unit
Université Libre de Bruxelles CP 191
Avenue F.-D. Roosevelt, 50
1050 Bruxelles — BELGIUM
Email: jcsarraz@ulb.ac.be

August, 2006

To appear in Human Movement Science (EWOMS special issue)

Abstract

Time plays a central role in consciousness, at different levels and in different aspects of information processing. Subliminal perception experiments demonstrate that stimuli presented too briefly to enter conscious awareness are nevertheless processed to some extent. Implicit learning, implicit memory, and conditioning studies suggest that the extent to which memory traces are available for verbal report and for cognitive control is likewise dependent on the time available for processing during acquisition. Differences in the time available for processing also determine not only the extent to which one becomes conscious of action, but also provides the basis for making attributions of authorship to experienced acts. In this paper, we offer a brief overview of these different findings and suggest that they can all be understood based on the fact that consciousness takes time. From this perspective, the availability of representations to conscious awareness depends on the quality of these representations — the extent to which they are strong, stable in time, and distinctive. High-quality representations occur when processes of global competition have had sufficient time to operate so as to make the system settle into the best possible interpretation of the input. Such processes implement global constraint satisfaction and critically depend on reentrant processing, through which representations can be further enriched by high-level constraints. We discuss these ideas in light of current theories of consciousness, emphasizing the fact that consciousness should be viewed as a process rather than as a static property associated with some states and not with others.

1. Introduction

When one thinks about consciousness, one often thinks about it as though it were a property associated either with entire organisms (e.g., humans are conscious, paramecia are not), with specific global states (e.g., I am conscious when awake, unconscious when asleep), or with particular regions of the brain (e.g., activity in primary visual cortex does not necessarily correlate with conscious awareness, but activity in some regions of frontal cortex do). The current focus on the “search for the neural correlates” of consciousness (Chalmers, 2000; Frith, Perry, & Lumer, 1999; Metzinger, 2000; Noë & Thompson, 2004), which relies heavily on imaging methods (e.g., fMRI) that have good spatial but rather poor temporal resolution, further emphasize localization over dynamics. But this eschews the fact that consciousness is foremost a process—not a static property associated with some neural states and not with others. Processes, are, by definition, dynamical, and hence take time to unfold. Appreciating the dynamics of consciousness requires sophisticated methods such as electroencephalography or magnetoencephalography (MEG) and conceptual tools such as dynamical systems theory (Freeman, 1999), the integration of which were pioneered by the likes of the late Francisco Varela in his neurophenomenological approach (Varela & Thompson, 2003) and by other authors such as Tononi and Edelman (1998).

We will not review these many exciting developments here, but instead simply focus on the notion that consciousness takes time and review available evidence in a

number of domains ranging from perception to action. As we will see in the rest of this article, numerous converging findings indeed suggest that time available for processing is a major determinant of whether one becomes conscious of some percept, memory trace, or even action. This important point also has significant implications for our very notion of how to best characterize conscious awareness. Indeed, as Dennett (1991) emphasized in his excellent “Consciousness Explained”, how the brain deals with the temporal relationships between different internal and external events offers one of the most compelling critiques against the notion that conscious awareness depends on activation in a single unified set of neuronal circuits—a global workspace. As Dennett puts it: “There is no single, definitive ‘stream of consciousness’, because there is no central Headquarters, no Cartesian Theater where ‘it all comes together’ for the perusal of a Central Meaner. Instead of such a single stream (however wide), there are multiple channels in which specialist circuits try, in parallel pandemoniums, to do their various things, creating Multiple Drafts as they go”. (Dennett, 1991, pp. 253-254).

Before we begin this overview, however, it is worth reflecting upon the question of why consciousness takes time. Answering this question involves further questioning about the functions that consciousness plays in information processing. While there is continuing puzzlement insofar as the putative functions of conscious experience—so-called phenomenal consciousness (Block, 1995)—are concerned, there is, in contrast, an emerging consensus about the function of so-called access consciousness (A-consciousness), which refers to our ability to report and act on our experiences. For a person to be in an A-conscious state entails that there is a representation in that person’s brain whose content is available for verbal report and for high-level

processes such as conscious judgment, reasoning, and the planning and guiding of action. There is wide agreement around the idea that conscious representations differ from unconscious ones in terms of such global accessibility (Baars, 1988): Conscious representations are informationally available to multiple systems in a manner that unconscious representations are not. Accessibility is in turn viewed as serving the function of making it possible for an agent to exert flexible, adaptive control over action. Briefly put, thus (we will return to these ideas in the discussion section), consciousness takes time because it takes time for any dynamical system to settle upon a state that best satisfies the many constraints imposed by diverse sources of such constraints, such as the environment with which the agent interacts, the goals and intentions of the agent, the properties of the action systems that it can control, and so on. Many computational mechanisms have been proposed to account for differences between conscious and unconscious information processing. It would too long to provide a principled overview here (but see Cleeremans, 2005, forthcoming; Maia & Cleeremans, 2005 for such overviews), but a salient point of agreement shared by several of the most popular current theories is that all such models, regardless of whether they assume specialized or non-specialized mechanisms, and regardless of whether they focus primarily on representational vehicles or on processes, converge towards assuming the following: Conscious representations differ from unconscious representations in that the former are endowed with certain properties such as their stability in time, their strength, or their distinctiveness. Cleeremans (forthcoming; Cleeremans & Jiménez, 2002) proposes the following definitions for these properties:

Stability in time refers to how long a representation can be maintained active during processing. O'Brien and Opie (1999), as well as Mathis and Mozer (1995), both suggested that stable representations are associated with conscious experience. There are many indications that different neural systems involve representations that differ along this dimension. For instance, the prefrontal cortex, which plays a central role in working memory (Baddeley, 1986), is widely assumed to involve circuits specialized in the formation of the enduring representations needed for the active maintenance of task-relevant information (Frank, Loughry, & O'Reilly, 2001; Norman & O'Reilly, 2001). Stability of representation is clearly related to availability to consciousness. For instance, the brief stimuli associated with subliminal presentation will result in weaker representations than supraliminal presentation does.

Strength of representation simply refers to how many processing units are involved in a given representation, and to how strongly activated these units are. Strength can also be used to characterize the efficiency of a an entire processing pathway, as in the Stroop model of Cohen, Dunbar & McClelland (1990). Strong activation patterns exert more influence on ongoing processing than weak patterns, and are most clearly associated with automaticity, to the extent that they dominate ongoing processing.

Finally, distinctiveness of representation refers to the extent of overlap that exists between representations of similar instances. Distinctiveness, or discreteness, has been hypothesized as the main dimension through which cortical and hippocampal representations differ (McClelland, McNaughton, & O'Reilly, 1995; O'Reilly & Munakata, 2000), with the latter becoming active only when the specific conjunctions of features that they code for are active themselves.

Strong, stable, and distinctive representations are thus explicit representations, at least in the sense put forward by Koch (2004): They indicate what they stand for in such a manner that their reference can be retrieved directly through processes involving low computational complexity (see also Kirsh, 1991, 2003). Conscious representations, in this sense, are explicit representations that have come to play, through processes of learning, adaptation, and evolution, the functional role of denoting a particular content for a cognitive system. Importantly, quality of representation should be viewed as a graded dimension.

Stability, strength, or distinctiveness can be achieved by different computational mechanisms. They can result, for instance, from the simultaneous top-down and bottom-up activation involved in so-called “reentrant processing” (Lamme, 2004), from processes of “adaptive resonance” (Grossberg, 1999), from processes of “integration and differentiation” (Edelman & Tononi, 2000), or from contact with the neural workspace, brought about by “dynamic mobilization” (Dehaene & Naccache, 2001). It is important to realize that the ultimate effect of any of these putative mechanisms is to make the target representations stable, strong, and distinctive. Hence an important computational principle through which to distinguish between conscious and unconscious representations is the following:

Availability to consciousness depends on quality of representation, where quality of representation is a graded dimension defined over stability in time, strength, and distinctiveness.

In the following, we overview empirical evidence for this principle in different domains. We will focus in particular on findings that suggest that the extent to which a representation is available to consciousness depends on time available for processing.

2. Temporal effects in subliminal perception and priming

The abundant literature dedicated to subliminal priming (Marcel, 1983) makes it clear that stimuli that are presented in such a way that they fail to be consciously perceived — for instance because they are presented for a very brief duration and subsequently masked — can nevertheless undergo substantial processing. The exact extent to which such stimuli can be processed semantically remains a significant area of debate (Hannula, Simons, & Cohen, 2005; Holender, 1986), but there is no question that such stimuli are indeed represented at some level. As a case in point, a recent study by Haynes and Rees (2005) suggests that sometimes, subjects' brains seem to know more than they can tell. Haynes and Rees used sophisticated signal processing techniques to demonstrate that it is possible to predict (albeit not perfectly) what stimulus (gratings oriented to the left to the right) a subject has been exposed to based on a single fMRI image of activity in their visual cortex—a method they dubbed “mindreading”. Strikingly, this was the case regardless of whether or not subjects had actually consciously perceived the stimulus. Indeed, whereas a first experiment, Haynes and Rees used visible gratings, in a second they used subliminal gratings rendered invisible by masking. Subjects were shown unable to detect the orientation of the latter despite prolonged exposure lasting up to 15 seconds. Interestingly, Haynes and Rees found graded differences in V1, V2 and V3 activity between the two conditions.

V1's activity, for instance, was predictive of the stimulus regardless of whether or not subjects had perceived it, but more so when they had than when they had not. V2 and V3, on the other hand, were only active when subjects had consciously perceived the stimulus. Haynes and Rees concluded by pointing out that "*Whether to be represented in conscious experience information has to cross a threshold level of activity, or perhaps needs to be relayed to another region of the brain, is an interesting question for further research*". (p. 689).

Earlier work by Naccache and collaborators (Dehaene et al., 2001) also indicated that while invisible words elicited activity in the left fusiform gyrus, visible words, in contrast, elicited activity in a much wider network of cortical areas that includes prefrontal and parietal regions. Hence in both cases, the differences between visible and invisible stimuli brought about by varying the duration of presentation (combined with masking) are (1) that the areas specialized for processing the stimulus tend to be more active when the stimulus has been consciously perceived, and (2) that consciously perceived stimuli result in widespread global processing that engage numerous areas distributed throughout the cortex. A further important point is that stimuli that completely fail to be perceived nevertheless undergo some processing. To summarize, the findings briefly reviewed above suggest the more time is available for processing, the deeper processing can be, and the more likely that the resulting representations are conscious.

3. Temporal effects in implicit learning

Similar temporal effects are also found in the very different domain of implicit learning (Cleeremans, Destrebecqz, & Boyer, 1998). In implicit learning situations, participants are shown to learn about the regularities contained in a stimulus environment in a manner that is independent of intentions and that results in knowledge that is difficult or impossible to verbalize. Even though the domain is also contentious (Shanks & St. John, 1994), the debate concerns not so much the existence of implicit learning per se but rather how one should interpret existing dissociation findings between ability to improve one's performance through learning and ability to verbalize or otherwise express the relevant knowledge. This is a debate that we will not address in this article, for it is rather intricate (Cleeremans, 1997). Instead, we will rather focus on the differences that result, in typical implicit learning paradigms, from varying the duration between events during training.

One of the most widely used paradigms through which to study implicit learning is sequence learning. In a typical sequence learning situation (see Clegg, DiGirolamo, & Keele, 1998), participants are asked to react to each element of a sequentially structured visual sequence of events in the context of a serial reaction time task. On each trial, subjects see a stimulus that appears at one of several locations on a computer screen and are asked to press as fast and as accurately as possible on the key corresponding to its current location. Nissen and Bullemer (1987) first demonstrated that subjects asked to respond as fast and as accurately as possible to a series of visual events progressively learned about the sequential structure of the stimulus sequence in spite of showing little evidence of being aware that the material contained structure.

Numerous subsequent studies of this effect have indicated that subjects can learn about complex sequential relationships despite remaining unable to fully elicit this knowledge in corresponding direct, explicit tasks (e.g. Cleeremans & McClelland, 1991; Cohen, Ivry, & Keele, 1990; Curran & Keele, 1993; Reed & Johnson, 1994; Willingham, Nissen, & Bullemer, 1989), thus suggesting that learning can occur without corresponding changes in our ability to elicit the acquired knowledge explicitly.

In a recent sequence learning study, Destrebecqz & Cleeremans (2001) aimed at manipulating the extent to which learning was explicit or implicit. To do so, they explored how changes in the response-to-stimulus interval (RSI; the amount of time that elapses between the response and the onset of the next stimulus) in the SRT task influenced the extent to which participants learned about the sequence. Their main hypothesis was that reducing the value of the RSI to 0 ms might selectively impair the development of conscious expectations about the location of the next stimulus.

To find out about how much participants consciously know about the sequence, one uses so-called direct tests such as free generation, which simply consists of asking participants to generate a sequence of stimuli after training is complete. To assess the relative contribution of implicit and explicit learning, Destrebecqz and Cleeremans adapted the process dissociation procedure (Debnar & Jacoby, 1994; Jacoby, 1991) to such a generation task. The process dissociation procedure is based on comparing performance in two generation tasks—inclusion and exclusion—that differ only by the instructions given to participants. In the inclusion task, participants are told to try to reproduce the training sequence. Performance in this task can depend either on

explicit recollection or on guessing based on intuition or familiarity. Hence, both implicit and explicit knowledge can contribute to inclusion performance. By contrast, in the exclusion task, participants are told to produce a sequence of stimuli that differs as much as possible from the training sequence. Implicit and explicit influences are thus now set in opposition, for to successfully avoid reproducing the sequence, one has to consciously retrieve its regularities so as to be able to produce some other sequential transitions. Of course, one can also exclude based on a mere feeling of familiarity, which may not be associated with explicit recollection of the sequence. In other words, one may estimate that a given response would reproduce the training sequence without remembering that response explicitly, and, based on this (explicit) feeling of familiarity, conservatively decide to generate a different transition. This is precisely why applying the process dissociation procedure to a free generation task makes it possible to dissociate implicit and explicit knowledge acquisition: If participants keep generating the training sequence against exclusion instructions, one can safely assume that such performance only reflects the automatic influence of unconscious knowledge.

Destrebecqz and Cleeremans's hypothesis that reducing the RSI to zero during training would result in qualitative differences in the knowledge acquired was confirmed. Indeed, while all participants were able to learn about the sequential regularities contained in the material and to project this knowledge in the generation task performed under inclusion instructions, only participants trained with an RSI of 250 ms were able to successfully exclude their knowledge when the generation task was performed under exclusion instructions. Participants trained with an RSI of 0ms instead tended to continue to preferentially reproduce the regularities of the training

sequence under such exclusion instructions. Destrebecqz and Cleeremans suggested that these participants lacked conscious control over their sequential knowledge. This impression was further strengthened by the results of a recognition task showing that these participants were unable to correctly discriminate between old and new sequence fragments. In contrast, participants trained with an RSI of 250 ms were perfectly capable of discriminating between novel and familiar sequence fragments.

Based on these results, they therefore concluded (1) that direct tasks such as generation or recognition may indeed be influenced by implicit knowledge, (2) that their data offered clear evidence that learning can be unconscious to the extent that the relevant knowledge may influence performance yet remain outside conscious control and recollection, and (3) that time available for processing each stimulus during the SRT task is critical in determining the extent to which sequence knowledge is available to conscious awareness. This third conclusion, — that explicit sequence learning depends on the time available to process each event over the course of the SRT task — was somewhat speculative since existing studies do not offer a coherent picture of the influence of temporal factors on sequence learning.

Indeed, the results of the Destrebecqz & Cleeremans (2001) study are consistent with some, but not with all of the previous studies that explored the importance of temporal factors on sequence learning (e.g., Hsiao & Reber, 2001). Such temporal factors were initially proposed as a way of understanding the effects of a secondary task on sequence learning performance. In the “dual-task” version of the SRT experiment, (Nissen & Bullemer, 1987) either a low- or a high-pitched tone is produced during the RSI. Instructions require participants not only to respond to each

stimulus location (the primary task) but also to keep a running count of how many low-pitched tones have occurred during each block (the secondary task). In their original study, Nissen & Bullemer (1987) argued that a secondary tone-counting task impairs sequence learning because it exhausts participants' attentional resources.

Other authors have instead suggested that the detrimental effect of the tone-counting task is due to scheduling conflicts between performing the main and secondary tasks, rather than to attentional load. Stadler (1995), for instance, argued that a secondary tone-counting task impairs sequence learning, not because it divides attention, but because it introduces variability in the RSI. Stadler pointed out that the secondary task lengthens the RSI only for the target trials in which the tone count must be updated. This incidental lengthening of the RSI would then have effects similar to those resulting from actually inserting a pause between those trials and the next. The pauses would disrupt participants' ability to parse the sequence into consistent chunks — a process which, according to Stadler, is essential to sequence learning.

Frensch & Miner (1994), in contrast, attribute the detrimental effects of the secondary task to short-term memory limitations: Secondary tasks impair sequence learning not so much because they make it hard for participants to chunk the sequence consistently, but simply because they lengthen the response-to-stimulus interval (RSI), and that this lengthening makes it more difficult for participants to link together the memory traces corresponding to successive elements of the sequence in short-term memory. Consistently, Frensch and Miner (1994) reported that sequence learning is impaired when the RSI is increased to the unusually high value of 1500

ms.

Another hypothesis about the effects of the RSI on sequence learning was put forward by Willingham et al. (1997), who argued that lengthening the RSI does not impair sequence learning *per se*, but only the expression of knowledge about the sequence. Willingham et al. reported that participants trained with a 1500ms RSI showed impaired sequence learning as compared to participants trained with a 500 ms RSI. However, when transferred to a shorter RSI, the former group showed the same level of sequence learning as the latter.

At first sight, the study of Destrebecqz and Cleeremans (2001) appears to contradict these results. Indeed, they found that a higher value of the RSI tends to improve explicit sequence learning. They argued that people trained with an RSI of 250 ms are given more opportunities to link together high-quality memory traces and to develop stronger representations of the sequential constraints of the training material — an account that is totally inconsistent with the findings of Frensch and Miner (1994), for instance.

In this controversial context, to further explore and clarify the role of temporal factors on sequence learning, they therefore conducted a new experiment in which they manipulated the RSI over three different values (RSI = 0, 250, or 1500 ms), and in which performance was assessed through a wider array of objective and subjective measures (Destrebecqz & Cleeremans, 2003). In this experiment, the authors replicated the results of their previous study (Destrebecqz & Cleeremans, 2001): in the RSI 0 ms condition, participants who were denied preparation to the next stimulus

in the SRT task learned the sequence but were unable (1) to refrain from expressing their knowledge under exclusion instructions and (2) to project this knowledge in a recognition task. These findings confirm that sequence learning can occur unconsciously (but see Wilkinson & Shanks, 2004 for a failure to replicate and for a different interpretation). They also found new evidence suggestive of increasing degrees of explicitness as the RSI is lengthened. When the RSI was set to a standard 250 ms, participants acquired explicit sequence knowledge that they could both control in the exclusion task as well as recollect in the recognition task. Importantly, further increasing the RSI to 1500 ms further allowed participants to acquire meta-knowledge about their knowledge—participants's confidence ratings in their generation performance was now correlated with their actual performance.

To summarize, the results of the Destrebecqz & Cleeremans (2001, 2003) experiments show that manipulating the RSI influences the extent to which sequence learning is explicit. Beyond the striking fact that such a small change in the pacing of the stimulus material can result in large qualitative changes in the extent to which the acquired knowledge is available to conscious awareness, these results also suggest that the differences between implicit and explicit learning might in fact be best viewed as resulting from continuous, gradual changes along a single dimension—“quality of representation”. According to this perspective, representations that are more stable, stronger, and more distinctive are more available to conscious control.

These ideas were confirmed by a recent imaging study (Peigneux et al., 2005) using essentially the same design as Destrebecqz and Cleeremans (2001). Participants were scanned using PET during exclusion after training on an SRT task using either

and RSI of 0 or 250 msec. Results showed that activity in the striatum subtends the implicit component of performance during recollection of the sequence, whereas the anterior cingulate/mesial prefrontal cortex (ACC/MPFC) supports the explicit component. Importantly, it was found that the ACC/MPFC exerts control on the activity of the striatum during retrieval of the sequence after explicit learning, that is, when the RSI had been set to 250 msec, whereas the activity of these regions was uncoupled when learning had been essentially implicit (RSI = 0 msec). These data thus suggest that implicit learning processes can be successfully controlled by conscious knowledge when learning is essentially explicit. They also supply further evidence for a partial dissociation between the neural substrates supporting conscious and nonconscious components of performance during recollection of a learned sequence, and suggest that increasing the RSI makes it for possible for richer representations of the sequential material to be developed. Developing such high-quality representations, crucially, takes time. Interestingly, similar effects of temporal delays are found in the conditioning literature, which we review in the next section.

4. Temporal effects in conditioning

The SRT findings described above find an interesting echo in conditioning research. Classical conditioning is a simple form of associative learning that has been studied extensively in vertebrate and invertebrate animals (Carew & Sahley, 1986; Rescorla & Wagner, 1972; Squire & Kandel, 1999). The best understood example of classical conditioning in vertebrates is conditioning of the eyeblink response (Gormezano, Schneiderman, & Fuentes, 1962; J. Kim & Thompson, 1997). In eyeblink classical conditioning a neutral conditioned stimulus (CS), such as a tone, is presented just

before an unconditioned stimulus (US), such as a mild puff of air to the eye. After repeated pairings of the CS and the US, the CS elicits a learned or conditioned eyeblink response (CR) in advance of the US. In this view, the two most commonly studied forms of eyeblink conditioning are delay and trace conditioning. In delay conditioning, the CS is presented and remains on until the US is presented. The two stimuli then overlap and co-terminate. In trace conditioning, an empty interval separates the CS and the US.

This review of classical conditioning and awareness focuses on delay and trace eyeblink classical conditioning essentially because the neural substrates sustaining eyeblink conditioning are better understood than any other form of conditioning in the vertebrate. Indeed, work with rabbits first demonstrated a clear distinction between delay and trace eyeblink conditioning (Thompson & Krupa, 1994). In this study, authors found that the acquisition and retention of delay eyeblink conditioning require the cerebellum and associated brainstem structures; no forebrain structures, as the hippocampus, are required. Further, rabbits made decerebrate after removal of the forebrain tissue (including cerebral cortex, basal ganglia, limbic system, thalamus, and hypothalamus) exhibited normal retention of delay eyeblink conditioning (Mauk & Thompson, 1987). Findings in humans are consistent with the animal work. Thus delay eyeblink conditioning was impaired in patients with cerebellar (Daum et al., 1993; Lye et al., 1988; Topka et al., 1993) or brainstem lesions (Solomon, Stowe, & Pendlebury, 1989), but was intact in amnesic patients with damage that includes the hippocampus (Daum, Channon, & Canavan, 1989; Gabrieli et al., 1995). In this context, because delay conditioning is independent of the forebrain, and intact in amnesia, it appears to be a typical example of nondeclarative

memory (Squire, 1992). As delay conditioning, trace eyeblink conditioning requires the cerebellum (Woodruff-Pak, Lavond, & Thompson, 1985). However, since acquisition and retention of trace conditioning were severely disrupted in mammals when the hippocampus and neocortex were damaged (J. J. Kim, Clark, & Thompson, 1995; Powell et al., 2001; Weible, McEchron, & Disterhoft, 2000), they differ in that trace conditioning also requires these structures. Here again, findings in human are consistent. In amnesic patients with damage that include the hippocampus, trace eyeblink conditioning was severely impaired (McGlinchey-Berroth et al., 1997).

Some studies that have explored the importance of awareness in eyeblink classical conditioning have used both delay and trace conditioning in differential conditioning or single-cue conditioning. More precisely, one CS (designated CS+) is paired with the US and a second CS (designated CS-) is paired with the absence of the US. In a large series of studies Clark, Squire and collaborators (Clark, Manns, & Squire, 2001; Clark, Manns, & Squire, 2002; Clark & Squire, 1998, 1999, 2000; Manns, Clark, & Squire, 2000a, 2000b, 2001, 2002) have argued that delay and trace eyeblink conditioning involve fundamentally different kinds of learning. In these findings, trace conditioning is related to awareness (or declarative knowledge) about the stimulus contingencies, while delay conditioning is not. Participants who develop awareness of the stimulus contingencies (i.e., the awareness of the relationship between the CS and the US) are successful at trace conditioning because they have acquired conscious representations that allow them to expect the US when the CS is presented.

One of the most convincing demonstrations of dissociations between conscious knowledge and behaviour was obtained in a eye-blink delay conditioning situation (Perruchet, 1985). In this experiment, people were exposed to a series of identical tones (the CS), 50% of which could be followed after a short interval by an air puff directed to the left cornea (the US). Immediately after each tone was presented (and before the puff occurred in reinforced trials), people were asked to indicate the extent to which they expected the tone to be followed by an air puff on a 0-7 points scale. A trial-by-trial analysis of the results indicated that eye blink responses were increasingly more likely to occur after presentation of a tone if the corresponding trial had been preceded by a series of reinforced trials (i.e., trials during which the tone had indeed been followed by an air puff). In stark contrast, however, people's subjective expectancy of the occurrence of an air puff tended to decrease with the number of reinforced trials that preceded the trial under consideration. In other words, people's eye blink responses were completely dissociated from their conscious expectations about when each tone would be followed by an air puff! People's expectations followed the gambler's fallacy (Anderson, 1960; Burns & Corpus, 2004; Jarvik, 1951; Keren & Lewis, 1994), whereas their blink responses simply reflected the strength of the CS-US association.

Using the same method, Clark et al. (2001) observed that the dissociation between subjective expectancy and eyeblink responses disappeared and even turned into an association when the performance of the participants was no more measured in a delay but in a trace paradigm. That is, CR probability was high when expectancy of the US was high whereas the CR probability was low when expectancy of the US was low. These findings were interpreted as supporting the idea that delay conditioning

and trace conditioning are fundamentally different in their dependence on awareness. Further, the authors suggested, “in the case of trace conditioning, the cerebellum is responsible for the acquisition, storage, and generation of the CR, and the hippocampal formation provides input to the cerebellum that is essential for the acquisition process. One possibility is that the hippocampus provides temporally shifted information to the cerebellum so that the CS and US arrive at the cerebellum in the temporally overlapping fashion that the cerebellum can use.” (p. 530). This pattern of data was taken as indicative that consciousness is required for conditioning to occur under trace conditions but not under delays conditions.

It is important to note that these findings—in particular the notion that associative conditioning can occur without consciousness in humans — are contentious. In fact, evidence for a separate conditioning mechanism that is independent of higher cognitive processes has been remarkably difficult to obtain (Brewer, 1974; Dawson & Schell, 1985; Lovibond & Shanks, 2002). Conditioning in humans does appear to be closely tied to attention, consciousness, and language. Nevertheless, the findings of Perruchet (1985) provide strong evidence for the possibility that conditioned responses can dissociate from subjective expectancy, and the findings of Squire and colleagues are at least suggestive that the pattern of relationships between subjective expectancy and behavior may be dependent on the presence of a delay between the CS and the US. An intriguing question remains, however, for the data do not indicate whether it is the case that conscious awareness is required so as to make it possible for subjects to bridge the temporal gap that separates the CS from the US, or whether it is the case that the presence of the delay is responsible for the onset of conscious awareness of the relationship under trace conditions. Congruently with the SRT

findings reviewed above, we would suggest that the latter holds, that is, that the presence of delays, during which processing of the stimuli can be enriched, result in higher-quality representations of the association between the CS and the US.

In the preceding sections, we reviewed several studies emphasizing the role that time plays in making the results of information processing available to consciousness in different domains—subliminal perception, implicit learning, and conditioning. Considered together, these studies all suggest that time available for processing is a central determinant of the extent to which the relevant representations are conscious. We suggest that this is so because (1) the degree to which a given representation is conscious depends on its quality (stability, strength, distinctiveness), and (2) that increasing time available for processing increases a representation's quality. In the remainder of this review, we ask whether such temporal effects are also at play in the rather different domain of action. Under which conditions are we aware of our own actions? Is our sense of self and agency—central aspects of conscious experience—dependent on temporal correlations between our intentions, our actions, and their consequences? We examine these issues in the following sections.

5. Temporal effects and awareness of action

Many studies have now demonstrated that motor control can be achieved independently of conscious awareness, even for intentional actions. For example, Goodale et al. (1986) report on a pointing experiment in which the target occasionally jumped several degrees, unnoticed by the subjects. Nevertheless, subjects were able to adjust the trajectory of their moving hand to the target position. In this case subjects

were aware neither of the sensory information that elicited the movement correction nor of the change of the motor program that resulted.

In another experiment, Castiello et al. (1991) found that awareness of an unexpected target jump occurred more than 200 ms after the motor system had initiated an appropriate movement correction. A third example is provided by a study by Fournieret and Jeannerod (1998). By giving false visual feedback about the trajectory of hand movements, Fournieret and Jeannerod demonstrated that subjects (who could not see their hand) were nevertheless able to voluntarily achieve the desired result of drawing a straight line on a computer screen. The necessary corrective movements were produced in spite of the fact that subjects remained unaware of having produced them, at least so long as the corrections remained under a certain (but surprisingly large) threshold.

Intending to provide further insight in these processes, Varraine et al. (2002) studied the on-line dialogue between high level controlled processing and low level automatic processing in motor control. This was studied in locomotion because it is an automatic behaviour that can be intentionally modulated. Subjects walked continuously on a treadmill with the instruction to maintain either a constant walking speed (compensation condition) or constant propulsive forces (no-intervention condition); they were sometimes faced with slow variations in resistance that they had to detect. This experiment produced two main results. First, subjects remained unaware of their force increase (in compensation) or of their walking velocity decreases (in no-intervention) for a long time, although these modifications went largely beyond the variability range in which they were able to intentionally control

their force (in no-intervention) or their velocity (in compensation). Second, detection of the increase in resistance occurred at the same time in both conditions. The authors concluded that the sudden awareness of a movement pattern produced at a low level was found to emerge from the interaction between a top down mechanism where intentional control of goal feedback delays the conscious perception of the other sensory sources, and a bottom up mechanism where high-level mechanisms of sensorimotor integration come into play beyond a discrepancy threshold between different sources of sensory information.

These and other studies suggest that many aspects of motor control take place unconsciously: We are not aware of the precise details of the motor commands that generate our actions, nor of the way in which immediate sensory information is used to fine-tune these commands. As Haggard (2005) puts it: “Much of the information processing underlying our conscious actions is ‘automatic’. We are aware only of the tip of the action iceberg.” (p. 290). What we know about our own movements may thus be a filtered, delayed version of the rapid operation of the action circuits embodied by the dorsal stream (Milner & Goodale, 1998), and may partly be a representation of what we thought we would do (rather than what we actually did). In other words, awareness of action might be influenced both by reconstructive and by preconstructive processes rather than simply consisting in a faithful, transparent appraisal of our actions.

These ideas about the relationships between motor control and awareness of action have recently been embodied in computational models called “forward models”. (Wolpert, 1997; Wolpert, Ghahramani, & Jordan, 1995). Such models (see also

Grush, 2004; Jordan & Rumelhart, 1992; Theofilou, Destrebecqz, & Cleeremans, 2004) aim at addressing the fact that many control problems (and acting adaptively is the best example of such problems) are difficult because they require solving two separate sub-problems: (1) learning about the effects of particular actions on the environment, which involves developing a model of the system one is attempting to control (the so-called “forward” model), and (2) learning which particular actions are most appropriate to achieve a desired goal, that is, learning how to control the system (the “inverse” problem). Forward models make it possible to solve both problems simultaneously. To do so, they generally consist of two interconnected networks. The first (“inverse” model) takes as input a goal and a description of the current state, and produces actions. The second, “forward”, model takes the response of the first network (an action) and a description of the current state as input, and produces a prediction of how the to-be-controlled system (the “plant”, in control theory parlance) would change if the produced action were carried out.

Crucially, such models provide a mechanism through which fast, accurate control of action can be achieved. Indeed, such control would in fact be all but impossible if one were to wait for the results of sensory feedback about the consequences of our actions to be computed whenever a motor command is issued. The forward, predictive model bypasses these delays by anticipating the sensory consequences of our actions even before the latter are carried out, based on a copy of the efference signal sent by the motor system to the muscles, so making it possible to achieve the open-loop control characteristic of skilled action.

As Haggard (2005) points out, while such models were not initially developed to explain conscious experience, they provide a useful framework to think about which aspects of action control are associated with conscious experience and which are not. Based on our hypothesis that consciousness takes time, it makes sense that motor control, as mediated by the interaction between the inverse and forward components of the model, is not associated with conscious experience. This forms the basis for the dissociations between action and perception described above and makes sense from an evolutionary point of view. As Frith et al. (2000) note “*The changes in representation that result from our own movements are entirely predictable on the basis of those movements and do not require any attention. It seems plausible that to be aware of representations which change every time we moved our body, or even our eyes, would be a positive disadvantage. (p. 1775).* Unlike the motor system, which has to provide fast and efficient adjustments to our intended movements, the mechanisms that underlie conscious perception appear to have evolved so as to maintain stability and to bring focus on unexpected events.

We can now ask what happens when more time lapses between the time a stimulus is shown and the time an action is directed towards it? A recent study by Hu & Goodale (2000) is particularly illustrative in this context.

Hu and Goodale compared, in a series of three experiments, the effect of relative and absolute size on manual prehension and manual conscious estimates of perceived size. In each experiment, right-handed subjects were presented with two different sized 3-D objects in a virtual display, and were instructed to pick-up or estimate the size of one of them.

In the first experiment, subjects were requested to pick up the smaller of two virtual objects in one condition, and the larger one in the other condition. In reality, the target objects were identical on all trials; they were simply paired with a smaller object on some trials and with a larger one on other trials, so producing a size-contrast illusion.

In the second experiment, one of the virtual objects was marked with a red dot on its top surface. From trial to trial, the marked object was paired with a larger, smaller, or same-sized object. Subjects were instructed to always pick up the marked object on each trial. In both the first and the second experiments, half the subjects were tested in delayed grasping, which involved a 5 seconds delay between viewing the objects and initiating the grasp, and half were tested in real-time grasping, that is, without a delay.

In the third experiment, using the same display of virtual objects as in the second experiment, subjects were requested to estimate the size of the marked object using their index finger and thumb (i.e., they showed how big the object looked to them). After estimating the target object's size, they picked it up. All subjects gave their estimates either immediately or after a delay.

Recording of hand movements revealed that when subjects in the first and the second experiments picked up the target object in real time, their grip aperture in flight was not significantly affected, regardless of whether the object was accompanied by a larger object or a smaller one. However, when subjects picked up the target object after a delay, their grip aperture in flight was larger when the target

object was accompanied by a smaller object than when it was accompanied by a larger object.

A similar size-contrast effect was also observed in the third experiment in which subjects gave manual estimates of the perceived size of the target object. The effect of conscious perception was observed both when the estimates were given immediately or after a 5 seconds delay. The authors suggested that real-time visuomotor control relies on absolute metrics, whereas delayed grasping operates the relative metrics produced by conscious perception. Thus, with respect to the metrics and frame of reference that each system uses and with respect to the time-scale over which each system operate, there is strong evidence that control of action and conscious perception can be dissociated.

Crucially, whether action and conscious perception are congruent with each other again appears to depend on the presence of a delay. This is interestingly confirmed by the observation that the presence of a temporal delay between perception and action completely reverses the dissociation patterns exhibited by patients suffering from visual agnosia. Thus, patient DF, who has been extensively studied by Milner and Goodale (Milner & Goodale, 1998), typically exhibits a striking dissociation between her (preserved) ability to accurately perform grasping movements and her (impaired) ability to report on the size and orientation of objects. This and other findings formed the basis for the suggestion that the visual system is organized in functionally and anatomically separate ventral and dorsal pathways corresponding roughly to “vision for perception” and to “vision for action”. Now, Goodale et al.

(1994) found that inserting a 2 seconds delay between stimulus presentation and initiation of action resulted in a complete breakdown of grasping performance by DF.

This suggests that accurate delayed action requires active maintenance of conscious perceptual representations, which in normal participants results in increasing congruence between action and conscious perception, and is simply impossible in the case of DF for she has no conscious perception of visual objects. There is an interesting parallel between these findings and the conditioning findings we reviewed above. In both cases indeed, the presence of a delay results in increased congruence between responses and reportable subjective experience, so that responses are influenced by conscious representations as long as sufficient time is available for the latter to emerge, but not otherwise.

6. Temporality of conscious experience and the sense of agency

Strikingly, temporal factors also appear to play a crucial role in our sense of agency, but in a different manner than the studies we have reviewed so far suggest. In a nutshell, the central notion is that we derive a sense of authorship for our own actions by (unconsciously) comparing the predicted consequences of our own actions and the actual consequences of these actions, as forward models suggest. In most situations, especially those that are routine, the actual state of the motor system will closely correspond to the state predicted before the action was completed. In such highly automatized cases, one would thus expect *attenuated* conscious awareness of action based on the fact that the underlying representations have become, through

prior training, very strong, and can thus occur without necessarily resulting in strong associated phenomenal experience.

Relevant evidence comes from studies on tickling. It is well known that the intensity of the tactile conscious experience is greatly reduced in comparison with the sensation when someone else tickles us (Weiskrantz, Elliott, & Darlington, 1971). Interestingly, a recent study by Blakemore et al. (1999) suggests that temporal factors play a central role in explaining why it is impossible to tickle oneself. In their experiment, a tactile stimulus on the palm of the right hand was either externally produced by a robot, or self-produced by the subject. In the conditions in which the tactile stimulus was self-produced, subjects moved the arm of a robot with their left hand to produce the tactile stimulus on their right hand via a second robot. Subjects were asked to rate the intensity of the tactile sensation, and they consistently rated self-produced tactile stimuli as less tickly, less intense, and less pleasant than externally produced tactile stimuli.

Crucially, using their robotic setup, Blakemore et al. were able to manipulate the correspondence between the action of the subjects' left hand and the tactile stimulus on the right hand, by varying both the temporal and the spatial relationships between the movements of the left hand and the resulting effects on the right hand. Results showed that perceived ticklishness increased as a function of the extent of temporal and spatial decorrelation between action and sensation. The authors suggested that in the case of self-produced movements, the resultant tactile sensation is attenuated because of the correspondence in space and time between the tactile stimulus and its causal motor command — just as forward models would predict. Conscious

experience of being tickled is tightly dependent on the error between the predicted sensory feedback of the motor system and the actual sensory feedback produced by the movement. It is tempting to suggest that ongoing assessment of such (time-dependent) discrepancies between actions and their effects forms the basis for representations of the acting self.

Thus, (1) it takes some time to become aware of an action that we have just carried out, as the phenomenology of reflex action illustrates particularly well (“I pulled my hand out of the fire before realizing what was happening”), and (2) our brain uses the (lack of) time that lapses between the occurrence of specific neural states (e.g. SMA activation) and subsequent action to attribute the action to itself or not.

If this is the case, one would thus expect to be able to influence the perceived temporal relationships between events by manipulating the extent to which people think they caused the event or not. It is well known that if an event B systematically follows an event A in close temporal proximity, one will conclude, given other appropriate conditions, that A caused B (Michotte, 1954). But is it the case that manipulating our sense of authorship (i.e., did my actions cause this effect?) also changes our perception of the temporal relationships between our actions and their consequences? A particularly illustrative study of this relationship has been carried out by Haggard, Clark and Kalogeras (2002) on the judgment of the perceived time of either a voluntary action or an involuntary movement.

Haggard et al.’s study indicates that when we believe our actions to have caused an event, the event is perceived to occur earlier than if we think we did not cause it. In

one condition, subjects judged the timing of an auditory tone by reporting on the corresponding position of a rapidly moving clock hand. In the second, crucial condition, subjects pressed a key that caused a tone to follow 250 ms later. Again, subjects judged the time of the tone. Comparing the data across conditions, the perceived times of the tone when keypress and tone were causally linked were compared with the conditions in which the tone occurred by itself. Remarkably, when the tone was causally linked to the subjects' keypress, subjects judged the tone to have occurred 46 ms earlier than if these events had occurred alone. In a second experiment with different subjects, the delay between the keypress and subsequent tone was varied (it was set to be 250, 450, or 650 ms), and subjects judged the time of the tone. Haggard et al. found that the further apart the keypress and tone, the less the "temporal attraction" of the tone to the keypress. Interestingly, comparable involuntary movements caused by TMS reversed this "temporal attraction" effect.

Haggard et al. concluded that intentional actions and their effects are bound together through a process they dubbed "intentional binding", and whose function would be to make it possible for agents to attribute to themselves the authorship of events their own actions caused.

Thus, it seems that temporality of the conscious experience of an action and our sense of agency are reciprocally linked. On the one hand, temporal factors impinge on our sense of agency: Whether we attribute an action to ourselves or not depends in part on whether the actual consequences of a movement can be matched against its internally predicted effects, and whether such consequences occur in close temporal proximity to relevant neural states. On the other hand, "intentional binding" illustrates

the effect that our sense of agency has on conscious experience of temporal relationships, and again that *the phenomenology of intentional action requires an appropriate predictive link between intentions and effects* (Haggard & Clark, 2003; p. 695): Occurrence of intentions causes binding when the intention produces a voluntary action, but not when the intention is interrupted by TMS (Haggard & Clark, 2003).

Finally, the work of Libet is of course also highly relevant in this context. In a well-known and highly puzzling study, Libet et al. (1983) explored the relationships between conscious intention and its neural correlates by asking subjects to perform a voluntary movement while simultaneously recording EMG activity as well as the readiness potential (RP) — the EEG signal that indicates preparation of a movement. Subjects also had to indicate the moment at which they first experienced the urge to move their finger by noting the position of a rapidly rotating hand on the face of a clock. This protocol thus yields three types of data that Libet was able to align with each other in time: (1) the moment subjects had actually lifted their finger, (2) the moment they had reported first experiencing the urge to move (Libet's "W" judgment) and (3) the moment that the readiness potential began to rise. Libet found that while subject's self-reported moment of first experiencing the urge to move preceded actual movement by about 200 ms, RP onset itself preceded conscious experience of the urge to move by about 350 ms. This strikingly reverses our folk-psychological notion of what it means to perform a voluntary action. Libet's data suggest that intention to act does not precede the brain activity that causes movement, but rather that unconscious brain activity produces both conscious experience of the intention to act as well as the action itself. Wegner (2002) interestingly elaborated on

these ideas to frame his “theory of apparent mental causation”, according to which both our actions and our experience of intention to act are caused by the same unconscious neural processes, so that conscious will turns out to be nothing more than an entirely illusory attribution of causality through which conscious experience of intention to act is linked to action.

Beyond the challenging philosophical issues involved (which we will not address in this article), Libet’s work is particularly relevant in suggesting that conscious experience of willing an action is delayed with respect to the onset of the neural activity that caused it. This is again congruent with our hypothesis that consciousness takes time, and is also consistent with our experience (or rather, lack thereof) of automaticity, as we briefly discuss in the next section.

7. Consciousness as global constraint satisfaction

The findings we have reviewed so far all converge to suggesting that time available for processing is a central determinant of the extent to which a representation is conscious. While this does not explain conscious experience—it could be the case that consciousness is just “fame in the brain”, as Dennett (2001) would have it, or it could be that further explanation is needed (see Block, 2005, for further discussion)—it helps identify putative neural mechanisms and computational principles associated with its onset. Here, our suggestion, in line with other proposals, is conscious states are associated with the existence of what we have called high-quality representations, and that such high-quality representations take some time to emerge out of the recurrent processing so characteristic of cortical computation. In this context, many

authors have recently emphasized that prefrontal cortex (PFC) play a specific role not only in integrating information from different regions of the brain, but also in being specifically involved in tasks that require active maintenance of information over temporal delays (Funahashi, Chafee, & Goldman-Rakic, 1993; Rao, Rainer, & Miller, 1998). As Koch (2004) puts it, “Prefrontal regions are (...) in an eminent position to integrate information from all sensory and motor modalities.” (p. 130). According to Hu and Goodale (2000), “the ventrolateral and dorsolateral prefrontal regions not only have rich connections with the ventral and dorsal visual pathways, respectively (Baizer, Ungerleider, & Desimone, 1991; Cavada & Goldman-Rakic, 1998), but they are also reciprocally connected with each other and with other cortical areas that receive input from both regions (Barbas & Pandya, 1989; Watanabe-Sawagushi, Kubota, & Arikuni, 1991).” (p. 865).

O’Reilly and collaborators (Frank, Loughry, & O’Reilly, 2001; Norman & O’Reilly, 2001) have further suggested, in their “tripartite” theory, that prefrontal cortex has evolved specialized recurrent circuits dedicated to the active maintenance of the enduring representations characteristic of working memory tasks—a computational objective that is incompatible with those that posterior cortex and hippocampus address.

It therefore seems reasonable to speculate that PFC neurons could be part of a recurrent cortical network that biases organized perceptual information for the production of actions, and that operates at a larger time-scale compared to the motor control feedforward pathways. In fact, as soon as feedforward activation has reached a particular area, horizontal connections start to connect distant cells within that area,

and feedback connections start sending information from higher level areas back to lower levels. Together, these connections provide what is called recurrent processing.

The importance of recurrent connections for conscious experience has recently been emphasized by several authors (Dehaene, Sergent, & Changeux, 2003; Edelman, 2003; Lamme, 2004). From this perspective, recurrent processing is necessary for conscious experience (Lamme & Roelfsema, 2000), whereas purely feedforward processing would not be associated with conscious awareness. The notion that recurrent interactions at a nearly global scale sustain the emergence of conscious experience makes sense in light of the idea that dynamical neural networks implement global constraint satisfaction. The idea is that a network with recurrent connexions arrives at an interpretation of a given input by settling into a stable state (Mathis & Mozer, 1996; Rumelhart et al., 1986). In a recent review, Maia and Cleeremans (2005) have provided an integrated view of attention, working memory, cognitive control, and consciousness based on a single mechanism, namely global competition between representations, biased by prefrontal cortex. They suggested that these functions are not distinct functions implemented by separate brain systems, but that they should instead be understood emerging out of the dynamics of global competition. Thus, instead of arguing whether a particular brain area or group of neurons contributes to consciousness or not, this approach aims at characterizing the kinds of computational principles that might account for key properties of conscious experience, and suggests that one focuses on the mechanisms that drive change, either during processing of individual stimuli or over the longer periods that characterize skill acquisition.

Congruently, changes in specific aspects of conscious experience correlate with changes in activity in specific brain areas widely distributed over both posterior (occipital and temporal) and anterior or frontal regions. Changes in the degree to which neural activity is distributed within the brain also accompany the transition between unconscious, conscious and controlled performance, and automatic performance (Poldrack et al., 2005). When tasks are novel, brain activation related to the task is widely distributed; when the task has become automatic, activation is more localized and may shift to a different set of areas (Petersen et al., 1998). In animal studies, neural activity related to sensory stimuli can be recorded in many brain regions before habituation. After habituation sets in (a time when humans report that stimuli tend to disappear gradually from consciousness), the same stimuli evoke neural activity exclusively along their specific sensory pathways (Horel et al., 1967). These observations suggest that when tasks are automatic and require little or no conscious control, the spread of signals that influence performance involves a more restricted and dedicated set of circuits that become “*functionally insulated*” (Tononi & Edelman, 1998, p. 1847). This, in turn produces a gain in speed and precision, but a loss in context-sensitivity, accessibility, and flexibility (Baars, 1988).

7. Conclusions

In this paper we have explored a simple idea—the idea that consciousness takes time. We have suggested that the reason this is so is that the extent to which representations become available to conscious awareness depends on their quality, by which we mean their stability, their strength, and their distinctiveness. High-quality representations emerge as the result of global competition biased by top-down

modulation, which implements global constraint satisfaction. The contents of conscious experience at some point in time thus reflect the application of the brain's knowledge on the current situation so as to yield the most adapted representations in the service of action. Such processes in turn critically depend on recurrent, or reentrant processing, by which high-level, global knowledge can come to influence and modulate the processing conducted in the lower areas of a hierarchy of modular processors. Such modulatory influences, however, cannot occur if there isn't sufficient time to fully process the input, nor if action is quick enough.

To support these ideas, we have reviewed converging evidence that even relatively small temporal delays between (1) the onset of two stimuli, (2) between a stimulus and a response or (3) between an action and its effects, all profoundly modify subjective experience, reportability, and awareness of action. The findings we have reviewed suggest (1) that aspects of information processing—from perception to action—can occur without conscious awareness, thus resulting in dissociations between subjective reports and actual behavior, (2) that allowing more time for processing to take place tends to reduce or eliminate such dissociations, so that subjective reports and behavior become congruent with each other. Strikingly, this applies to very different domains, from subliminal perception and subliminal priming to implicit learning and conditioning, from the influence of perception on action to our sense of agency.

This suggests that common factors are at play in all these cases. Our suggestion is that all such findings can be understood based on (1) the notion that consciousness depends on high-quality representations, and on (2) the notion that developing such

high-quality representations takes time—time for mechanisms of global constraint satisfaction to operate through recurrent processing so as to make it possible for the brain to integrate current input, prior knowledge and expectations about future states into unified, rich representations.

One might of course quarrel with the conclusion that the dissociations we have reviewed reflect a genuine distinction between conscious and unconscious representations. Perruchet and Vinter (2002), for instance, propose instead that the very notion of “unconscious representation” should simply be abandoned, and that information processing would be best characterized as involving unconscious associative processes acting on (necessarily) conscious representations—however fleeting and weak the latter might be. We note that this characterization of conscious information processing is in fact not inconsistent with our main claim, namely that time available for processing determines the extent (the degree) to which the relevant representations are conscious, and that the differences are brought about through processes of global constraint satisfaction. Whether one can actually distinguish between very weak conscious representations and genuinely unconscious representations remains a very, very tough empirical issue. Since the two might be functionally equivalent, perhaps the dispute is simply terminological...

Acknowledgements

Axel Cleeremans is a Research Director with the National Fund for Scientific Research (Belgium). Jean-Christophe Sarrazin is a Marie Curie Post-doctoral fellow (contract MEIF-CT-2005-515499) at the Université Libre de Bruxelles. We thank Pierre Perruchet and an anonymous referee for their valuable comments on a previous draft of this article. This work was supported by an institutional grant from the Université Libre de Bruxelles and by the European Commission.

References

- Anderson, N. H. (1960). Effect of first-order conditional probability in a two-choice learning situation. *Journal of Experimental Psychology*, *56*, 73-93.
- Baars, B. J. (1988). *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.
- Baddeley, A. D. (1986). *Working Memory*. New York, NY: Oxford University Press.
- Baizer, J. S., Ungerleider, L. G., & Desimone, R. (1991). Organization of visual inputs to the inferior temporal and posterior parietal cortex in macaques. *Journal of Neuroscience*, *11*, 168-190.
- Barbas, H., & Pandya, D. N. (1989). Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. *Journal of Comparative Neurology*, *256*, 353-375.
- Blakemore, S. J., Frith, C. D., & Wolpert, D. M. (1999). Spatiotemporal prediction modulates the perception of self-produced stimuli. *Journal of Cognitive Neuroscience*, *11*(5), 551-559.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, *18*, 227-287.
- Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences*, *9*(2), 46-52.
- Brewer, W. F. (1974). There is no convincing evidence for operant or classical conditioning in adult humans. In W. B. Weimer & D. S. Palermo (Eds.), *Cognition and the symbolic processes*. Mahwah, N.J.: Lawrence Erlbaum Associates.
- Burns, B. D., & Corpus, B. (2004). Randomness and inductions from streaks: "Gambler's fallacy" versus "Hot hand". *Psychonomic Bulletin and Review*, *11*, 179-184.
- Carew, T. J., & Sahley, C. L. (1986). Invertebrate learning and memory: from behavior to molecules. *Annual Review of Neuroscience*, *9*, 435-487.
- Castiello, U., Paulignan, Y., & Jeannerod, M. (1991). Temporal dissociation of motor responses and subjective awareness. *Brain*, *114*, 2639-2655.
- Cavada, C., & Goldman-Rakic, P. S. (1998). Posterior parietal cortex in rhesus monkey: II. Evidence for segregated corticocortical networks linking sensory and limbic areas with the frontal lobe. *Journal of Comparative Neurology*, *287*, 422-445.
- Chalmers, D. J. (2000). What is a neural correlate of consciousness? In T. Metzinger (Ed.), *Neural Correlates of Consciousness. Empirical and Conceptual Questions*. (pp. 17-39). Cambridge, MA: MIT Press.
- Clark, R. E., Manns, J. R., & Squire, L. R. (2001). Trace and delay eyeblink conditioning: Contrasting phenomena of declarative and nondeclarative memory. *Psychological Science*, *12*, 304-308.
- Clark, R. E., Manns, J. R., & Squire, L. R. (2002). Classical conditioning, awareness, and brain systems. *Trends in Cognitive Sciences*, *6*(12), 524-531.
- Clark, R. E., & Squire, L. R. (1998). Classical conditioning and brain systems: The role of awareness. *Science*, *280*, 77-81.

- Clark, R. E., & Squire, L. R. (1999). Human eyeblink conditioning: Effects of manipulating awareness of the stimulus contingencies. *Psychological Science*, *10*, 14-18.
- Clark, R. E., & Squire, L. R. (2000). Awareness and the conditioned eyeblink response. In D. Woodruff-Pak & J. E. Steinmetz (Eds.), *Eyeblink classical conditioning. Vol. 1: Applications in humans* (pp. 229-251). Boston, MA: Kluwer.
- Cleeremans, A. (1997). Principles for implicit learning. In D. C. Berry (Ed.), *How implicit is implicit learning?* (pp. 195-234). Oxford: Oxford University Press.
- Cleeremans, A. (2005). Computational correlates of consciousness. In S. Laureys (Ed.), *Progress in Brain Research* (Vol. 150, pp. 81-98). Amsterdam: Elsevier.
- Cleeremans, A. (forthcoming). *Being Virtual*. Oxford, UK: Oxford University Press.
- Cleeremans, A., Destrebecqz, A., & Boyer, M. (1998). Implicit learning: News from the front. *Trends in Cognitive Sciences*, *2*, 406-416.
- Cleeremans, A., & Jiménez, L. (2002). Implicit learning and consciousness: A graded, dynamic perspective. In R. M. French & A. Cleeremans (Eds.), *Implicit Learning and Consciousness: An Empirical, Computational and Philosophical Consensus in the Making?* (pp. 1-40). Hove, UK: Psychology Press.
- Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology : General*, *120*, 235-253.
- Clegg, B. A., DiGirolamo, G. J., & Keele, S. W. (1998). Sequence learning. *Trends in Cognitive Sciences*, *2*, 275-281.
- Cohen, A., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*, *97*, 332-361.
- Cohen, A., Ivry, R. I., & Keele, S. W. (1990). Attention and structure in sequence learning. *Journal of Experimental Psychology : Learning, Memory and Cognition*, *16*, 17-30.
- Curran, T., & Keele, S. W. (1993). Attentional and nonattentional forms of sequence learning. *Journal of Experimental Psychology : Learning, Memory and Cognition*, *19*, 189-202.
- Daum, I., Channon, S., & Canavan, A. G. M. (1989). Classical conditioning in patients with severe memory problems. *J. Neurol. Neurosurg. Psychiatry*, *59*, 47-51.
- Daum, I., Schugens, M. M., Ackermann, H., Lutzenberger, W., Dichgans, J., & Birbaumer, N. (1993). Classical conditioning after cerebellar lesions in humans. *Behavioral Neuroscience*, *107*, 748-756.
- Dawson, M. E., & Schell, A. M. (1985). Information processing and human autonomic classical conditioning. In P. K. Ackles (Ed.), *Advances in psychophysiology, Vol. 1* (pp. 89-165). Greenwich, CT: JAI Press.
- Debner, J. A., & Jacoby, L. L. (1994). Unconscious perception : Attention, awareness and control. *Journal of Experimental Psychology : Learning, Memory and Cognition*, *20*, 304-317.
- 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*(7), 752-758.
- Dehaene, S., Sergent, C., & Changeux, J.-P. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences of the U.S.A.*, *100*(14), 8520-8525.
- Dennett, D. C. (1991). *Consciousness Explained*. Boston, MA.: Little, Brown & Co.
- Dennett, D. C. (2001). Are we explaining consciousness yet? *Cognition*, *79*, 221-237.
- Destrebecqz, A., & Cleeremans, A. (2001). Can sequence learning be implicit? New evidence with the Process Dissociation Procedure. *Psychonomic Bulletin & Review*, *8*(2), 343-350.
- Destrebecqz, A., & Cleeremans, A. (2003). Temporal effects in sequence learning. In L. Jiménez (Ed.), *Attention and Implicit Learning* (pp. 181-213). Amsterdam: John Benjamins.
- Edelman, G. M. (2003). Naturalizing consciousness: A theoretical framework. *Proceedings of the National Academy of Sciences of the U.S.A.*, *100*, 5520-5524.
- Edelman, G. M., & Tononi, G. (2000). *Consciousness. How matter becomes imagination*. London: Penguin Books.
- Fourneret, P., & Jeannerod, M. (1998). Limited conscious monitoring of motor performance in normal subjects. *Neuropsychologia*, *36*, 1133-1140.
- Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective & Behavioral Neuroscience*, *1*(2), 137-160.
- Freeman, W. (1999). *How brains make up their minds*. London: Wiedenfeld & Nicholson.
- Frensch, P. A., & Miner, C. S. (1994). Effects of presentation rate and individual difference in short-term memory. *Memory and Cognition*, *22*, 95-110.
- Frith, C. D., Perry, R., & Lumer, E. (1999). The neural correlates of conscious experience : An experimental framework. *Trends in Cognitive Sciences*, *3*, 105-114.
- Funahashi, S., Chafee, M. V., & Goldman-Rakic, P. S. (1993). Prefrontal neuronal activity in rhesus monkeys performing a delayed anti-saccade task. *Nature*, *365*, 753-756.
- Gabrieli, J. D., McGlinchey-Berroth, R., Gluck, M. A., Cermak, L. S., & Disterhoft, J. F. (1995). Intact delay-eyeblick classical conditioning in amnesia. *Behavioral Neuroscience*, *109*, 819-827.
- Goodale, M. A., Jakobson, L. S., & Keillor, J. M. (1994). Differences in the visual control of pantomimed and natural grasping movements. *Neuropsychologia*, *32*, 1159-1178.
- Goodale, M. A., Pelisson, D., & Prablanc, C. (1986). Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature*, *320*, 748-750.
- Gormezano, I., Schneiderman, D. E., & Fuentes, I. (1962). Nictitating membrane: Classical conditioning and extinction in the albino rabbit. *Science*, *138*, 33-34.
- Grossberg, S. (1999). The link between brain learning, attention, and consciousness. *Consciousness and Cognition*, *8*, 1-44.

- Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27(3), 377-396.
- Haggard, P. (2005). Conscious intention and motor cognition. *Trends in Cognitive Sciences*, 9(6), 290-295.
- Hannula, D. E., Simons, D. J., & Cohen, N. J. (2005). Imaging implicit perception: Promise and Pitfalls. *Nature Reviews Neuroscience*, 6(247-255).
- Haynes, J.-D., & Rees, G. (2005). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nature Neuroscience*, 8(5), 686-691.
- Holender, D. (1986). Semantic activation without conscious activation in dichotic listening, parafoveal vision, and visual masking : A survey and appraisal. *Behavioral and Brain Sciences*, 9, 1-23.
- Horel, J. A., Vierck, C. J. J., Pribram, K. H., Spinelli, D. N., John, E. R., & Ruchkin, D. S. (1967). Average evoked responses and learning. *Science*, 158, 394-395.
- Hsiao, A. T., & Reber, A. (2001). The dual-task SRT procedure: Fine-tuning the timing. *Psychonomic Bulletin & Review*, 8, 336-342.
- Hu, Y., & Goodale, M. A. (2000). Grasping after a delay shifts size-scaling from absolute to relative metrics. *Journal of Cognitive Neuroscience*, 12(5), 856-868.
- Jacoby, L. L. (1991). A process dissociation framework : Separating automatic from intentional uses of memory. *Journal of Memory and Language*, 30, 513-541.
- Jarvik, M. E. (1951). Probability learning and a negative recency effect in the serial anticipation of alternative symbols. *Journal of Experimental Psychology*, 41, 291-297.
- Jordan, M. I., & Rumelhart, D. E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive Science*, 16, 307-354.
- Keren, G., & Lewis, C. (1994). The two fallacies of gamblers: Type I and Type II. *Organizational Behavior and Human Decision Processes*, 60, 75-89.
- Kim, J., & Thompson, R. F. (1997). Cerebellar circuits and synaptic mechanisms involved in classical eyeblink conditioning. *Trends in Neurosciences*, 20, 177-181.
- Kim, J. J., Clark, R. E., & Thompson, R. F. (1995). Hippocampectomy impairs the memory of recently, but not remotely, acquired trace eyeblink conditioned responses. *Behavioral Neuroscience*, 109, 195-203.
- Kirsh, D. (1991). When is information explicitly represented? In P. P. Hanson (Ed.), *Information, Language, and Cognition*. New York, NY: Oxford University Press.
- Kirsh, D. (2003). Implicit and explicit representation. In L. Nadel (Ed.), *Encyclopedia of Cognitive Science* (Vol. 2, pp. 478-481). London, UK: Macmillan.
- Koch, C. (2004). *The quest for consciousness. A neurobiological approach*. Englewood, CO: Roberts & Company Publishers.
- Lamme, V. A. F. (2004). Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness. *Neural Networks*, 17(5-6), 861-872.
- Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, 23(11), 571-579.
- Lovibond, P. F., & Shanks, D. R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology-Animal Behavior Processes*, 28(1), 3-26.

- Lye, R. H., O'Boyle, D. J., Ramsden, R. T., & Schady, W. (1988). Effects of a unilateral cerebellar lesion on the acquisition of eye-blink conditioning in man. *Journal of Physiology (London)*, *403*, 58.
- Maia, T. V., & Cleeremans, A. (2005). Consciousness: Converging insights from connectionist modeling and neuroscience. *Trends in Cognitive Sciences*, *9*(8), 397-404.
- Manns, J. R., Clark, R. E., & Squire, L. R. (2000a). Awareness predicts the magnitude of single-cue trace eyeblink conditioning. *Hippocampus*, *19*, 181-186.
- Manns, J. R., Clark, R. E., & Squire, L. R. (2000b). Parallel acquisition of awareness and trace eyeblink conditioning. *Learning and Memory*, *7*, 267-272.
- Manns, J. R., Clark, R. E., & Squire, L. R. (2001). Single-cue delay eyeblink conditioning is unrelated to awareness. *Cognitive, Affective & Behavioral Neuroscience*, *2*, 192-198.
- Manns, J. R., Clark, R. E., & Squire, L. R. (2002). Standard delay eyeblink conditioning classical conditioning is independent of awareness. *Journal of Experimental Psychology: Animal Behavior Processes*, *28*, 32-37.
- Marcel, A. J. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology*, *15*, 197-237.
- Mathis, W. D., & Mozer, M. C. (1995). On the computational utility of consciousness. In G. Tesauro & D. S. Touretzky (Eds.), *Advances in neural information processing systems* (Vol. 7, pp. 10-18). Cambridge: MIT Press.
- Mathis, W. D., & Mozer, M. C. (1996). Conscious and unconscious perception: A computational theory. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 324-328). Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Mauk, M. D., & Thompson, R. F. (1987). Retention of classically conditioned eyelid responses following acute decerebration. *Brain Research*, *403*, 89-95.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419-457.
- McGlinchey-Berroth, R., Carrillo, M. C., Gabrieli, J. D. E., Brawn, D. M., & Disterhoft, J. F. (1997). Impaired trace eyeblink conditioning in bilateral, medial-temporal lobe amnesia. *Behavioral Neuroscience*, *100*, 243-252.
- Metzinger, T. (2000). *Neural Correlates of Consciousness. Empirical and Conceptual Questions*. Cambridge, MA.: MIT Press.
- Michotte, A. (1954). *The perception of causality*. New York: Basic Books.
- Milner, A. D., & Goodale, M. A. (1998). *The visual brain in action*. Oxford, England: Oxford University Press.
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirement of learning: Evidence from performance measures. *Cognitive Psychology*, *19*, 1-32.
- Noë, A., & Thompson, E. (2004). Are there neural correlates of consciousness? *Journal of Consciousness studies*, *11*(1), 3-28.
- Norman, K., & O'Reilly, R. C. (2001). *Modeling hippocampal and neocortical contributions to recognition memory: A complementary learning systems approach* (No. Technical Report 01-02): Institute of Cognitive Science — University of Colorado, Boulder.
- O'Brien, G., & Opie, J. (1999). A connectionist theory of phenomenal experience. *Behavioral and Brain Sciences*, *22*, 175-196.

- O'Reilly, R. C., & Munakata, Y. (2000). *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*. Cambridge, MA: MIT Press.
- Peigneux, P., Destrebecqz, A., Laureys, S., Degueldre, C., Delfiore, G., Aerts, J., et al. (2005). The neural correlates of implicit and explicit sequence learning: Interacting networks revealed by the process dissociation procedure. *Learning and Memory*, 12(5), 480-490.
- Perruchet, P. (1985). A pitfall for the expectancy theory of human eyelid conditioning. *Pavlovian Journal of Biological Science*, 20, 163-170.
- Perruchet, P., & Vinter, A. (2002). The self-organizing consciousness. *Behavioral and Brain Sciences*, 25(3), 297-330.
- Petersen, S. E., van Mier, H., Fiez, J. A., & Raichle, M. E. (1998). The effects of practice on the functional anatomy of task performance. *Proceedings of the National Academy of Sciences USA*, 95, 853-860.
- Poldrack, R. A., Sabb, F. W., Foerde, K., Tom, S. M., Asarnow, R. F., Bookheimer, S. Y., et al. (2005). The neural correlates of motor skill automaticity. *Journal of Neuroscience*, 25(22), 5356-5364.
- Powell, D. A., Skaggs, H., Churchwell, J., & McLaughlin, J. (2001). Posttraining lesions of the medial prefrontal cortex impair performance of Pavlovian eyeblink conditioning but have no effect on concomitant heart rate changes in rabbits (*Oryctolagus cuniculus*). *Behavioral Neuroscience*, 115, 1029-1038.
- Rao, S. C., Rainer, G., & Miller, E. K. (1998). Integration of what and where in the primate prefrontal cortex. *Science*, 276, 821-824.
- Reed, J., & Johnson, P. (1994). Assessing implicit learning with indirect tests: Determining what is learned about sequence structure. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 20, 585-594.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical Conditioning II: Current Theory and Research* (pp. 64-99): Appleton-Century-Crofts.
- Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. (1986). Schemata and sequential thought processes in PDP models. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel Distributed Processing: Explorations in the microstructure of cognition. Volume 2: Psychological and Biological Models* (pp. 7-57). Cambridge, MA: MIT Press.
- Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, 17, 367-447.
- Solomon, P. R., Stowe, G. T., & Pendlebury, W. P. (1989). Disrupted eyelid conditioning in a patient with damage to cerebellar afferents. *Behavioral Neuroscience*, 103, 989-902.
- Squire, L. R. (1992). Declarative and nondeclarative memory: multiple brain systems supporting learning and memory. *Journal of Cognitive Neuroscience*, 4, 232-243.
- Squire, L. R., & Kandel, E. R. (1999). *Memory: From mind to molecules*. New York: W.H. Freeman.
- Stadler, M. A. (1995). Role of attention in implicit learning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 21, 674-685.
- Theofilou, D., Destrebecqz, A., & Cleeremans, A. (2004). Applying forward models to sequence learning: A connectionist implementation. In H. Bowman & C.

- Labouise (Eds.), *Proceedings of the Eight Neural Computation and Psychology Workshop: Connectionist Models of Cognition and Perception II* (pp. 51-61). London: World Scientific.
- Thompson, R. F., & Krupa, D. J. (1994). Organization of memory traces in the mammalian brain. *Annual Review of Neuroscience*, 20, 519-550.
- Tononi, G., & Edelman, G. M. (1998). Consciousness and complexity. *Science*, 282(5395), 1846-1851.
- Topka, H., Valls-Sole, J., Massaquoi, S. G., & Hallett, M. (1993). Deficit in classical conditioning in patients with cerebellar degeneration. *Brain*, 116, 961-969.
- Varela, F. J., & Thompson, E. (2003). Neural synchrony and the unity of mind: A neurophenomenological perspective. In A. Cleeremans (Ed.), *The Unity of Consciousness* (pp. 266-287). London: Oxford University Press.
- Varraine, E., Bonnard, M., & Pailhous, J. (2002). The top down and bottom up mechanisms involved in the sudden awareness of low level sensorimotor behavior. *Cognitive Brain Research*, 13(3), 357-361.
- Watanabe-Sawagushi, K., Kubota, K., & Arikuni, T. (1991). Cytoarchitecture and intrafrontal connections of the frontal cortex of the brain of the hamadryas baboon (*Papio Hamadryas*). *Journal of Comparative Neurology*, 311, 108-133.
- Wegner, D. M. (2002). *The Illusion of Conscious Will*. Cambridge, MA: Bradford Books, MIT Press.
- Weible, A. P., McEchron, M. D., & Disterhoft, J. F. (2000). Cortical involvement in acquisition and extinction of trace eyeblink conditioning. *Behavioral Neuroscience*, 114, 1058-1067.
- Weiskrantz, L., Elliott, J., & Darlington, C. (1971). Preliminary observations of tickling oneself. *Nature*, 230(598-599).
- Wilkinson, L., & Shanks, D. R. (2004). Intentional control and implicit sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 354-359.
- Willingham, D. B., Greenberg, A. R., & Cannon Thomas, R. (1997). Response-to-stimulus interval does not affect implicit motor sequence learning, but does affect performance. *Memory & Cognition*, 25, 534-542.
- Willingham, D. B., Nissen, M. J., & Bullemer, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology : Learning, Memory and Cognition*, 15, 1047-1060.
- Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in Cognitive Sciences*, 1(6), 209-216.
- Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, 269, 1180-1182.
- Woodruff-Pak, D. S., Lavond, D. G., & Thompson, R. F. (1985). Trace conditioning: abolished by cerebellar nuclear lesions but not lateral cerebellar cortex aspirations. *Brain Research*, 348, 249-260.