First and third-person approaches in implicit learning research

Vinciane Gaillard, Muriel Vandenberghe, Arnaud Destrebecqz, & Axel Cleeremans

Address correspondence to:

Axel Cleeremans (*axcleer@ulb.ac.be*)
Cognitive Science Research Unit
50 av. FD Roosevelt (CP 191)
1050 Brussels
BELGIUM

To appear in *Consciousness and Cognition*

Abstract

How do we find out whether someone is conscious of some information or not? A simple answer is "We just ask them"! However, things are not so simple. Here, we review recent developments in the use of subjective and objective methods in implicit learning research and discuss the highly complex methodological problems that their use raises in the domain.

Introduction

Contemporary approaches to consciousness, best exemplified by what has been dubbed “the search for the neural correlates of consciousness” (Crick & Koch, 1990) emphasize localization over dynamics: We will know about the neural bases of conscious experience once we know which regions of the brain are specifically involved in conscious processing. While all agree that consciousness stems not from activity in a specific region but rather from the involvement of entire networks, the contrastive approach (Frith, Perry, & Lumer, 1999) required by most brain imaging studies nevertheless mandates that we set up experimental conditions in such a manner that we can compare what happens when processing takes place with or without consciousness. This strategy will often result in highlighting activity (or lack thereof) in a restricted set of regions. An unfortunate side effect of this otherwise commendable and necessary approach, however, is that our conception of the differences between conscious and unconscious processing has tended to become oversimplified. What it means to be conscious is reduced to, say, my ability to report on the absence or a presence of a particular visual experience, such as for instance a word presented under subliminal conditions (Dehaene et al., 2001), or a target in an attentional blink protocol (Sergent & Dehaene, 2004). In other words, consciousness is viewed as a binary dimension of information processing: I am either consciously aware of some state of affairs, or not, but there are no definite intermediate states of consciousness worth describing.

This dichotomous perspective also subtends contemporary theorizing about consciousness. Dehaene’s neural workspace model, for instance, itself rooted in Baars’ global workspace theory (Baars, 1988), predicts that all measures of conscious awareness one can come up with should be correlated (Dehaene, personal communication, June 11th, 2004). This makes sense based on the central idea that subtends the theory. Indeed, once a processor is “mobilized” into the global workspace, its contents become globally available through a process aptly dubbed as “ignition”. Ignition here means the highly non-linear, positive feedback loop through which the relevant representations are selected, emphasized, maintained, and made available, as a result, to a wide variety of other processors.

What is wrong with this picture? Beyond the obvious fact that consciousness is much, much more than mere perceptual awareness of some simple stimulus—an overly reductive characterization that one can hardly hold models such as the neural workspace accountable for, since one indeed has to start somewhere when thinking of something as complicated as consciousness[1]—, what is wrong with this picture is that it does not hold in the face of numerous demonstrations that there are in fact substantial dissociations between different measures of conscious awareness. When such dissociations occur as a result of differences in sensitivity, as we will discuss later, they are not particularly problematic for models that assume that once a representation has connected with the global workspace, it is conscious. Other dissociations, however, are more problematic, for they cannot be reconciled with a concept of consciousness that views it essentially as an all-or-none phenomenon.

In this paper, we first document these dissociations through several illustrative cases. Next, we reflect upon the often intricate methodological issues involved when attempting to establish differences between information processing with and without consciousness, focusing in particular on the domain of implicit learning. Finally, we close by suggesting that the time is now ripe for our concept of consciousness to be augmented—consciousness, far from being a "single thing", is instead a process—dynamic, extended through time, fragmented over different neural systems—a point emphasized by Dennett in his "Consciousness Explained" (Dennett, 1991).

What is knowledge?

At first sight, exploring consciousness empirically seems simple enough : Just ask people what the contents of their subjective experience is at some point in time, jot it down carefully, and correlate these reports with some objective measure of behaviour or of neural activity. This approach, which amounts to correlating first-person and third-person data, forms the basic gist of many contemporary studies, and it would

---

[1] The situation here is analogous to that of the proverbial drunkard looking for his keys not where he lost them, but where there is light : Since there is no hope, with our current methods, to find the neural correlates of a complex experience, then let us begin by trying to find the neural correlates of really, really simple experiences. Sound as this might be as a research strategy, let it not obscure the fact that it involves a severe reduction.

indeed be hard to imagine some other way of answering questions about consciousness. Many complications arise almost instantly, however (Cleeremans, 2001). These complications range from definitional (what does one mean by "consciousness"?) and methodological issues (how do I best assess conscious awareness?) to conceptual issues (how do I best interpret dissociation findings?).

A recent debate illustrates these points rather well. Bechara and colleagues (Bechara, Damasio, Tranel, & Damasio, 1997) explored performance in a simple gambling task (the "Iowa gambling task") while simultaneously measuring skin conductance. In this situation, participants, after being given $2000 in play money, were asked, on each of a series of trials, to choose a card from one of four decks. Each choice resulted in a win or in a loss. Subjects were told to play so as to maximize gains. Unknown to participants, decks differed in their overall ultimate yield, with some decks being disadvantageous and others being advantageous. Subjects were free to choose cards from any of the four decks and did not know how many trials had to be performed before the experiment stopped. Subjects were probed, through questions, about their knowledge of the situation at regular intervals during the game.

The results of these experiments indicated (1) that subjects started selecting cards from the advantageous decks before they were able to verbally motivate and explain their choices, and (2) that they exhibited a larger skin conductance response (SCR) just before taking a card from a disadvantageous deck. Thus, differential SCR responses to advantageous and disadvantageous decks appeared before subjects were able to motivate their decisions, as if their body knew which decks are risky before the relevant knowledge was available for verbal reports. This is what Bechara and colleagues dubbed the "somatic marker hypothesis". These results were taken to indicate implicit learning: Participants appear able to learn about which decks are "good" vs. "bad" and to use this knowledge to make the best decisions before being able to motivate their choices in answer to direct questions.

These results were replicated by Maia and McClelland in a recent study (Maia & McClelland, 2004). However, and crucially, Maia and McClelland also found that the dissociation between performance and verbal reports observed by Bechara and colleagues vanishes when participants are probed using more detailed and elaborate questioning. For instance, whereas Bechara and colleagues simply asked participants "Tell me all you know about what is going on this game" and "Tell me how you feel about this game", Maia and McClelland used an elaborate structured questionnaire involving detailed probing about each deck of card, direct questions about expected gains and losses associated to each, as well as direct questions about participants own appraisal of their knowledge. Using these more elaborate questions, Maia and McClelland found that participants in fact possessed conscious knowledge of the good and bad decks as well as of the best strategies to deploy in the game. Congruently, the authors concluded that there was in fact little evidence to support the notion that learning was subtended by unconscious knowledge in this situation. The debate is now continuing (Bechara, Damasio, Tranel, & Damasio, 2005; Maia & McClelland, 2005). We will not discuss it further, but simply point out that it is highly illustrative of the methodological challenges that have characterized the field of implicit learning over the 30 years of its history. Indeed, the very same issues have come up repeatedly over the years in this domain. Reber's early studies (Reber, 1967, 1969) on artificial grammar learning, for instance, were later criticized by Dulany and colleagues (Dulany, Carlson, & Dewey, 1984), who, like Maia and McClelland, showed that using more elaborate probing of how people came to decide whether a string was grammatical or not revealed much more conscious knowledge than Reber thought was involved in performance. Similar reinterpretations of existing dissociations were subsequently proposed by authors such as Perruchet (Perruchet, Gallego, & Savy, 1990; Perruchet & Vinter, 2002) and Shanks (Shanks, 2005; Shanks & St John, 1994) and the debate continues unabated today.

It is interesting to reflect upon the nature of these debates. First, they highlight the fact that the genuine challenges come not from how to best apply or interpret third-person, objective data, but rather from how to best design first-person, subjective methods. Second, as one digs deeper into these debates, one realizes they often boil down to fundamental differences in our conception of what constitutes knowledge. Fu, Fu, and Dienes (Fu, Fu, & Dienes, submitted) have provided an insightful analysis of this state of affairs. They point out that objective measures (e.g., forced-choice recognition or discrimination) require subjects to indicate whether they are able to discriminate between features of the world ("worldy discrimination") whereas subjective measures (e.g., a confidence judgment) require subjects to discriminate between their own mental states ("mental state discrimination"). Interestingly, Fu, Fu and Dienes further suggest that researchers who advocate the use of objective measures to assess conscious mental states often consider that such measures measure not just knowledge, but conscious knowledge. Thus, if I show objective ability to discriminate, say, between the presence or absence of a stimulus, that should be taken as evidence not only that I have some representations that make it possible for me to perform the discrimination, but also as evidence that these representations are conscious. In other words, such authors tend to be sceptic about the possibility for unconscious knowledge. On the other hand, researchers who advocate the use of subjective measures to assess conscious mental states consider instead that if I claim to be guessing, for instance, yet nevertheless show evidence of ability to discriminate, then the relevant knowledge should be taken to be unconscious: People perform correctly, or at least better than chance would predict, yet also indicate that they think they have no conscious basis for their decisions. There are two ways to interpret such dissociation findings. Either one considers that participants were in fact conscious of the knowledge used to make the discrimination, but our measure of conscious awareness of such knowledge is flawed, insensitive, or otherwise biased, or one considers that our measure of conscious awareness was in fact appropriate, and thus that the dissociation should be interpreted as reflecting the influence of unconscious knowledge.

A recent study by Haynes and Rees (2005) is also relevant in this context, for it demonstrates that there can be neural representations that are (1) indicative of the presence or absence of a stimulus, yet (2) undetectable neither through objective nor through subjective reports. In other words, subjects are both unaware that they possess such representations at some point in time and unable to express them on an objective test! 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 in a first experiment, Haynes and Rees used visible gratings, in a second they used subliminal gratings rendered invisible by masking. In the latter case, subjects were shown unable to detect the orientation of the grating 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 then when they had not. V2 and V3, on the other hand, were only predictively active when subjects had consciously perceived the stimulus.

Such findings are important because they make us question what should be considered to constitute knowledge. Is a neural pattern that fails to influence behavior and that subjects do not know they possess knowledge? Is it even mental in any tractable sense? We will leave detailed discussion of this tricky issue to better thinkers than we are (Dennett, 1991; Kirsh, 1991; Mandler, 1983; Perner, 1991), but are convinced that it lies at the heart of many current debates concerning the extent to which information processing without consciousness is possible.

This being said, how do we proceed with an empirical approach to exploring dissociations and associations between objective and subjective data? In the next section, we overview historical approaches to this issue.

A brief historical sketch

Anyone interested in the debate concerning how to best collect data on mental processes might find it particularly relevant to go back in time, at the very beginnings of psychological science. According to the French philosopher Auguste Comte (1798-1857), psychology should be tied to introspection only. Despite his rather strong view, he had to abandon verbal reports, because the act of introspecting might distort the "object" observed, e.g. the conscious experience itself. As we will see in the remainder of this section, this argument still holds whenever authors want to invalidate the use of verbal reports as data (but see Ericsson & Simon, 1980; Nisbett & Wilson, 1977).

A few decades later, Titchener, one of Wilhelm Wund's former students, trained participants to become skilled introspectionists. At the time, introspection was taken as observation of mental phenomena, just as inspection is observation of physical phenomena. For the very same reason, i.e., the reports are just another type of objective behavior, but with an opposite result, Watson and the behaviorists sought to eliminate introspection from the psychological field. However, it is somewhat ironical, given that "even behaviorist experiments that limited subjects' responses to giving yes/no answers, marking number scales, or just pushing buttons are best viewed as involving verbal reports about the subjects' experiences" (Nahmias, 2002, p.6). In a third phase, Newell and Simon, in their quest for a computer program (GPS, or the "General Problem Solver") that would "parallel, hence explain, the human behavior" (Newell & Simon, 1972, p. 885), rehabilitated the use of concurrent verbal reports. Newell and Simon's work on GPS relied extensively on protocol analysis in problem solving, by means of 'talk aloud' and 'think aloud' methodologies.

What this brief historical sketch highlights is that the use of introspection was largely debated throughout the years—and still is today. Some researchers view the direct verbalization of the subjects' cognitive processes as the only way toward a better understanding of our mental states. Other authors criticize the method, whilst they agree on the fact that verbal reports can bring some insights on cognition. Yet another approach simply rejects recourse to any verbal data. In order to catch the essence of the debate, let us first examine what the term 'introspection' means.

As Nahmias suggests, we refer to Külpe's basic definition of "attentively experiencing a mental process" (1909, p. 9, cited by Nahmias, 2002) in our understanding of introspection. Hence, subjects do have to pay attention to their inner experience, sensations, thoughts, and to report them verbally. Verbal reports are often called subjective measures, as compared to objective measures "because they measure what states of knowledge the subject thinks he or she is in" (Ziori & Dienes, 2006, p.107). Overgaard (2001) also describes the first-person perspective as having a perspective on anything that someone is conscious of at a certain moment. He adds: "Third-person information is, according to this understanding, what we normally would understand by "objective science—anything that could be studied by an external observer" (Overgaard, 2001). The underlying assumption is that yes/no responses, keypresses, and response latencies are objective or third-person measures, while verbal reports are considered as first-person measures. However, the problem with this perspective is that introspectionists regard verbal reports as objective data as well!

Verbal reports as data?

A famous theoretical controversy on the status of verbal reports as data illustrates this issue. On the one hand, Nisbett and Wilson (1977) argued against the possibility of accurate introspection: "there is by now enough evidence discrediting introspective reports to allow us to ignore any argument based on introspection" (p. 255). On the other hand, Ericsson and Simon (1984) point out that verbal reports constitute a rich source of data—one that can be combined with other data and that can be of greatest value in providing an integrated and full account of cognitive processes and structures. According to them, the anti-introspectionists have mistakenly assumed that yes/no responses, eye fixations or response latencies, which have been acceptable to psychology, are not themselves types of verbal responses. "Once we recognize that such responses are in fact verbal responses, then the question becomes not whether verbal responses should be accepted but rather under what circumstances they should be accepted." (Ericsson & Simon, 1980, p. 216).

There are different types of verbalization procedures as a function of time of verbalization. Concurrent verbal reports are collected as the subject is engaged in the task. Just as Auguste Comte himself had been thinking, these types of reports have been criticized for interfering with normal cognitive processes (Nisbett & Wilson, 1977). Ericsson and Simon argued on their part that "Thinking-aloud, as distinguished from explanation, will not change the structure and course of the task process, although it may slightly decrease the speed of the task performance" (Ericsson & Simon, 1980, p. 221).

In contrast to concurrent verbalization, retrospective verbal reports are collected after the subject has completed the task. As will be discussed in the next section, these types of postexperimental reports have been widely used in implicit learning experiments to assess whether subjects had been aware of the relevant features of the task (Cleeremans, Destrebecqz, & Boyer, 1998). This procedure has been criticized because subjects may forget or inaccurately recall the relevant features of the experimental situation, which of course cannot be interpreted as implying that they were unaware when engaged in the task.

The same problem is evidenced by Hannula, Simons, and Cohen (2005), who recently commented on the weakness of neuroimaging studies on implicit perception (e.g., Dehaene et al., 2001), pointing out that the conclusions of such studies often rely exclusively on subjective measures. Because response biases and interindividual differences influence confidence thresholds, Hannula et al. (2005) concluded that consciousness cannot be assessed based on verbal reports only: “they measure what is reported, not what is reportable” (p. 249).

As another illustration, let us cite Nisbett and Wilson’s (1977) article showing that subjects verbalizing retrospectively in a variety of settings about the motives for their behavior were no more accurate than observers were in identifying the important situational factors that actually determined their behavior. The authors reviewed a large sample of social psychology studies, where subjects were found unable to accurately describe their own inner processes. Nisbett and Wilson concluded that subjects apparently changed their attitudes in the absence of any subjective experience of change and that this observation was a significant argument against the reliability of introspection.

Nevertheless, according to Ericsson and Simon (1980), general retrospective questions such as “How did you do this task?” or “What did you think the experiment was about?” implicitly or explicitly request from the subject an interpretation of how she performed the task at hand. In other words, subjects have to speculate and theorize about their own cognitive processes, with the consequence that verbal reports are extremely difficult to interpret or to use as behavioural data: “Such probing for hypothetical states can never tap subjects' memories for their cognitive processes, since the information was never in memory" (Ericsson & Simon, 1980, p. 246).

To summarize, Ericsson and Simon (1980) delineated three central questions that must be addressed if verbal reports are to be used as data in psychological research. Firstly, do the instructions to verbalize and the probes to respond in either a general or a specific manner have any effect on the cognitive processes? The use of a structured questionnaire or some other means of limiting the subject's answers to a few possible categories such as "yes" or "no" might indeed not be the best way to proceed, according to Overgaard (2001). Secondly, are the verbal reports complete, i.e., do they consist of all the information the subjects have about their own cognitive processes? Finally, are the verbal reports consistent with other empirical data on behavior? We will see in the next sections that these questions are indeed of primary importance in the study of consciousness in general, and in the study of implicit learning in particular.

First- and third-person approaches are complementary

Looking closer at the relationships between first- and third-person approaches, we can classify them along three dimensions. Firstly, the procedure can be categorized as involving the first- or the third-person perspective depending on who is considered to be the <u>observer</u> of the relevant mental processes. Are subjects taken to be authoritative about what their self-observations mean, or is it that the experimenter is taken to be the only 'valid' observer? The first-person perspective is assumed in the former case, whereas the third-person perspective is assumed in the latter.

The second dimension concerns the <u>mode</u> of the response (e.g., verbal reports, keypresses, mouse clicks…). As we mentioned earlier, verbal reports are considered to be subjective measures, whereas response latencies, eye fixations, reaction times and so on are taken to constitute objective measures.

Finally, the third dimension is about the <u>object</u> of the response. As Fu, Fu, and Dienes suggest, the object of the response may either be objective or 'outer' (i.e., concerning features of the world, external to the subject's perspective) or it may be subjective, meaning ‘inner’ (i.e., concerning subject’s internal states). In a similar vein, Dienes and Perner (2004) distinguish between worldly subjectivity and experiential subjectivity. “Objective discrimination of stimuli entails that you are aware of the stimuli, but not that you are consciously aware of them. Subjective measures, which directly test for the existence of second order states, thus (according to the theory) directly test for the existence of conscious awareness.” (Dienes & Perner, 2004, p. 174). Thus, ‘outer’ responses concern the state of the world at a certain time, while ‘inner’ responses concern <u>the mental attitude</u> that the subject entertains *vis-à-vis* his or her ‘outer’ response.

The main implication of this perspective is that low confidence is no longer a means by which relevant conscious knowledge is excluded from measurement; rather confidence itself becomes an object of study and can be directly assessed on every trial (Ziori & Dienes, 2006).

To sum up, authors widely agree on the validity and on the reliability of verbal data as a source of information about cognitive processes as long as they are elicited with care and interpreted with full understanding of the circumstances under which they were obtained. To cite Ericsson and Simon (1980, p. 243): “When clear probes are used for specific retrospective memory and when reports are requested immediately after the last trial(s), informative verbal reports can usually be obtained […]. The failure of subjects to report some information does not demonstrate the uselessness of verbal protocols. Incompleteness of reports may make some information unavailable, but is does not invalidate the information that is present”. Nevertheless, consciousness is not equivalent to reportability, hence, “free reports can be a useful tool of detecting unconscious knowledge only to the degree that they are used as a complementary measure of conscious knowledge and as a means of testing the validity of other subjective measures” (Ziori & Dienes, 2006, p.119). The association of objective and subjective measures thus constitutes a promising avenue of research toward a better understanding of consciousness, as will be demonstrated in the next section.

First-and third-person methods in implicit learning research

While the overview we provided above suggests that first-person methodologies constitute an essential way of understanding and measuring consciousness

(Overgaard, 2001), it is important to keep their limitations in mind. Such methods do capture gross and simple features of conscious experience (Chalmers, 1999), but they might not actually be sensitive enough to provide an accurate measurement, neither for dissociating between conscious and unconscious knowledge nor for differentiating between features of conscious experience. For instance, while there is no doubt that I should conclude that a subject is conscious of some information if she can describe it to me, the reverse does not necessarily hold. Thus, poor performance on a questionnaire or verbal report task does not necessarily imply that I am unaware of some information.

These issues come out most pointedly in the domain of implicit learning. Implicit learning—broadly construed, learning without awareness—is a complex, multifaceted phenomenon that defies easy definition (Cleeremans et al., 1998). Learning is implicit when we acquire new information without intending to do so, and in such a way that the resulting knowledge is difficult to express. In this, implicit learning thus contrasts strongly with explicit learning (e.g., as when learning how to solve a problem or learning a concept), which is typically hypothesis-driven and fully conscious. Implicit learning contrasts with implicit memory in that it typically involves sensitivity to relationships between events rather than sensitivity to single events, and with subliminal priming in that it typically involves supra-liminal stimuli. Implicit learning research has essentially been focused on two experimental paradigms: Artificial grammar learning and sequence learning. Additional paradigms include dynamic system control, probability learning, hidden covariation detection, acquisition of invariant characteristics, or visual search.

In Reber's seminal study of artificial grammar learning, subjects were asked to memorize meaningless letter strings generated by a finite-state grammar. After memorization, subjects are asked to classify novel strings as grammatical or not. In this and many subsequent experiments, subjects can classify strings better than chance despite remaining unable to describe the rules of the grammar. This dissociation between classification and verbal report is the finding that prompted Reber to describe learning as implicit. Subjects indeed appeared sensitive to and could apply knowledge that they remained unable to describe and had had no intention to learn.

Today, another paradigm—sequence learning—has become dominant. Subjects 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, a stimulus appears at one of several locations on a computer screen and subjects are asked to press as fast as possible on a corresponding key. Nissen and Bullemer (1987) first demonstrated that subjects 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 have indicated that subjects can learn about complex sequential relationships despite remaining unable to deploy this knowledge in direct tasks.

In light of these dissociation findings, Shanks and St John (1994) pointed out that verbal reports do not always satisfy two criteria that they consider to be critical; the *information* and *sensitivity* criteria. According to the information criterion, the task used to measure conscious knowledge must tap into the same knowledge base upon which learning is based. Otherwise, learning could be described as unconscious not because participants are unable to access their knowledge consciously but simply because they are probed about irrelevant features of the training material that they did not need to process in order to perform the task. For instance, this might be the case in sequence learning studies in which participants are asked to report first- or second-order sequential regularities between successive elements while zero-order information, such as variations in the frequencies of the sequence elements, are sufficient to account for RT performance. In artificial grammar learning tasks, participants may have consciously learned some fragments of the training strings but not report this information when asked about the generating rules.

According to Shanks and St John's (1994) sensitivity criterion, the test used to measure conscious knowledge must be sensitive to all of the relevant information. If this criterion is not met, unconscious influences on performance might be overestimated because some conscious knowledge remains undetected by the awareness test. This issue is at the heart of the debate between Bechara et al. and Maia and McClelland that we discussed in the introduction. There are several reasons to argue that verbal reports fail the sensitivity criterion and therefore do not constitute an adequate measure of awareness. Firstly, subjects might fail to report fragmentary knowledge held with low confidence. Secondly, in the case of verbal reports, performance and awareness tests may implement very different retrieval contexts (Shanks & St John, 1994). With respect to the sequence learning paradigm, for instance, none of the contextual cues that are available to support performance in the RT task (such as the visual presentation of the stimuli, the requirement of motor responses, and a sustained pace of responding) are available to support verbal reports.

It is possible to improve the sensitivity of the awareness test by using questionnaires that involve more specific queries about the relevant knowledge—this was the strategy adopted by Maia and McClelland in their attempts to replicate Bechara et al. findings. However, the use of questionnaires imposes conditions on the test phase that are rather different from those in the training phase. Moreover, as Ericsson and Simon (1980) had already pointed out, participants may have forgotten some of the knowledge their learning performance was based on at the time of the verbal report or questionnaire.

Many authors have therefore suggested that valid tests of awareness should involve forced-choice tasks such as recognition tasks. It has been argued that these are able to detect conscious knowledge left undetected by verbal reports or questionnaires (Frensch, Buchner, & Lin, 1994; Perruchet & Amorim, 1992; Shanks & Johnstone, 1998, 1999; Shanks & St. John, 1994; Willingham, Nissen, & Bullemer, 1989).

Under the assumption that consciousness allows recollection of the acquired knowledge, forced-choice tests have generally taken the form of generation or recognition tasks at the end of sequence learning. In a typical generation task, participants are requested to reproduce the training sequence themselves by pressing the key corresponding to the location of the next stimulus instead of reacting to the current target. In the recognition task, they are presented with a sequence fragment, and after reacting to each element as they did in the RT task (i.e. by pressing as fast and as accurately as possible on the corresponding key), they are asked to identify whether or not the fragment was part of the training sequence. In sequence learning studies, forced-choice tests of conscious knowledge that prompt reproduction of the training sequence or differentiation of old and new sequence fragments have quite systematically indicated that participants were able to express a great deal of the knowledge they acquired in the SRT task (e.g. Perruchet & Amorim, 1992).

In artificial grammar learning tasks, forced-choice measures of learning have also taken the form of strings continuation tasks. In such completion tasks, participants have to generate the end of stems consisting of a few letters so as to produce a grammatical string (Dienes, Broadbent, & Berry, 1991). Other authors have used tasks in which participants had to underline fragments that made a sequence grammatical or ungrammatical (Dulany et al., 1984). All such studies have indicated that participants actually know more about the training material that was initially thought.

These results have often been interpreted as an indication that implicit learning is only implicit because the methods used to probe awareness are insensitive, and have thus led many to question the very existence of an implicit form of learning. Others, however, have challenged the assumption that generation or recognition performance depends solely on conscious knowledge (e.g., Jiménez, Méndez, & Cleeremans, 1996). It is indeed true that these tasks involve the same type of retrieval conditions as the SRT task: participants have to react to visual stimuli by giving motor responses and, at least in the case of recognition tasks, at the same pace of responding. There is therefore little reason to believe that SRT tasks and awareness tests tap into different knowledge bases, and thus that implicit knowledge does not contribute to generation or recognition performance.

Indeed, it has been shown that participants are able to reproduce the training sequence in a generation task even when they claim to guess the location of the next sequence element (Shanks & Johnstone, 1998). Furthermore, in a recognition task, subjects may tend to respond faster to old sequence fragments than to novel ones. Recognition ratings may therefore reflect this improved feeling of perceptual and motor fluency rather than explicit recollection of the training material (see Cohen & Curran, 1993; Perruchet & Amorim, 1992; Perruchet & Gallego, 1993; Willingham, Greeley, &

Bardone, 1993 for relevant discussion). Performance in both recognition and generation tasks, rather than depending exclusively on conscious knowledge, is thus likely to depend on both implicit and explicit influences. By the same token, it must also be emphasized that sequence acquisition during the SRT task is itself likely to involve both implicit and explicit components.

To summarize, forced-choice tasks are more likely to meet information and sensitivity criteria than free reports or questionnaires. This improvement in sensitivity, however, comes at a cost, namely that one is forced to endorse the so-called *exclusiveness assumption* (Reingold & Merikle, 1988), according to which the test of awareness must be sensitive *only* to the relevant conscious knowledge. Unfortunately, the most sensitive tests of awareness are also the most likely to be contaminated by implicit knowledge (Neal & Hesketh, 1997). The logic of quantitative dissociation has therefore been questioned through the argument that no task can be used as an absolute test of awareness that would simultaneously be sensitive to <u>all</u> of a subject's conscious knowledge and <u>only</u> to the relevant conscious knowledge. In other words, it is highly implausible that any task can be considered as "process-pure". To further improve awareness tests, different solutions have been proposed to overcome this so-called "contamination" problem. These procedures, which we discuss in the next sections, were initially proposed in the fields of subliminal perception and implicit memory, and later applied to the domain of implicit learning.

Subjective measures of awareness

Cheesman & Merikle (1984) have introduced the notion of subjective and objective thresholds in subliminal perception. In a typical experiment, the task simply consists in identifying a series of visual targets that are briefly flashed on a computer screen. Perception is said to be under the *subjective* threshold when participants are able to identify the target at above chance performance whilst stating that they did not perceive it consciously. Perception is said to be under the *objective* threshold when identification is at chance. Hence, perception is under the subjective threshold when a subject does not know that he knows the identity of the target, or in other words, when he has no *metaknowledge*. Perception is under the objective threshold when the target has simply not been perceived. According to Cheesman & Merikle, perception is unconscious when it is under the subjective threshold.

Dienes & Berry (1997) since suggested the application of the same threshold criteria to the study of implicit learning. In this framework, the acquired knowledge would be over the objective threshold when performance in a forced-choice task is above baseline. Learning could be described as unconscious if knowledge remains under the subjective threshold at the same time, i.e. if participants claim to respond at chance in the forced-choice task used to measure conscious knowledge. This procedure can indeed be extremely fruitful when attempting to disentangle conscious and unconscious knowledge given that, as discussed above, both types of knowledge can subtend performance in a forced-choice task.

Dienes, Altmann, Kwan, & Goode (Dienes, Altmann, Kwan, & Goode, 1995; see also Dienes & Perner, 1999) have described two criteria that make it possible to demonstrate unconscious knowledge acquisition. The first —the guessing criterion — corresponds to the criterion used by Cheesman & Merikle: knowledge is unconscious when performance is above chance while participants claim to perform at chance. The second criterion — the zero-correlation criterion — is met when confidence levels and performance rates are uncorrelated. Several studies have now applied these ideas in the domains of artificial grammar learning (Dienes & Altmann, 1997) and sequence learning (Shanks & Johnstone, 1998). Overall, these studies indicate that the knowledge acquired by participants in these empirical situations can indeed be implicit to the extent that it is “below the subjective threshold”. In a sequence learning experiment, using the guessing criterion after the generation task, Shanks & Johnstone (1998) asked participants to indicate whether they felt that they were reproducing the training sequence or that they responded randomly. Only 3 out of 15 participants felt that some fragments of the sequence were familiar and part of the training sequence. When these 3 subjects were excluded from the analysis, generation performance was still above chance level. In another experiment from the same study, these authors applied the zero-correlation criterion by asking participants to rate how confident they were in their generation performance on a scale ranging from 0 to 100. They observed that experimental subjects reproduced more of the training sequence than a control group presented with a random sequence during training. However, some of the experimental subjects failed to show a higher level of confidence than the control participants.

These results seem to suggest that learning was, at least partly, unconscious. However, Reingold & Merikle (1990) have insisted that subjective measurement of unconscious knowledge must be interpreted with caution given that subjective measures of awareness depend on participants' interpretation of the task instructions. Indeed, participants might essentially give a different interpretation of the term "guess" than the experimenter. Accordingly, Shanks & Johnstone have argued that, in their first study, participants' responding in the generation task was based on fragmentary (but, crucially, nevertheless conscious) knowledge held with low confidence. A similar interpretative problem may arise in Shanks and Johnstone's second study which used the zero-correlation criterion; in this study participants were able to evaluate their confidence level based not only on the conscious accessibility of the acquired knowledge but also on how much they believed was expected from them. For instance, a given subject might have underestimated his level of confidence because he assigned a high level of expectancy to the experimenter. This problem highlights why the zero-correlation criterion should not involve comparing confidence levels across different subjects, but rather within subjects.

To summarize, the criticisms previously laid against verbal reports may also hold for subjective measures because, in both cases, subjects have to decide themselves whether they have access to some knowledge or whether they were able to perform a task efficiently.

Is there some way in which our current first-person methods might be improved upon? In the next section, we examine a recent proposal to use Signal Detection Theory on subjective data, and briefly report on an experiment that makes use of this method. The method the basic premise of subjective approaches literally: What people have to do when reporting on subjective states is to perform signal detection on their own mental representations—signal detection on the mind, so to speak.

Signal detection on the mind

Tunney and Shanks (2003) have recently proposed to implement Signal Detection Theory within an artificial grammar learning task. As an alternative to the guessing criterion discussed above, Tunney and Shanks suggested exploring whether participants know when they are correct and when they are incorrect. To achieve this, they asked participants to report on their phenomenal states by means of confidence ratings. Thus, participants first have to discriminate and are then asked whether they are confident that their judgment was correct or not. Using the artificial grammar learning paradigm, the authors asked their participants to classify new strings of letters according to whether they follow the grammatical rules or not. After each classification, participants also had to express a confidence judgment regarding their performance. The idea is the following: if participants are aware of the knowledge they use to classify sequences, they should be more confident when they make correct decisions than when they make incorrect ones. Conversely, if participants have no awareness of the information they are using, high and low confidence responses would be randomly assigned to correct and incorrect discriminations. This method is thus built upon—and indeed identical with—the *zero-correlation criterion* methodology proposed by Dienes et al. (1995), according to which confidence should predict accuracy when knowledge can be accessed consciously.

Interestingly, Tunney and Shanks (2003) suggested to categorize confidence responses in terms of signal detection theory (SDT). According to them, correct discriminations made with high confidence may be categorized as *Hits*, while incorrect discriminations made with high confidence may be categorized as *False Alarms*. Indeed, when participants are aware of the knowledge used to perform a discrimination, they presumably believe themselves to be correct and should respond with high confidence. On the other hand, when participants believe themselves to be correct and respond with high confidence when they are, in fact, incorrect, they presumably do not possess knowledge that can be accessed consciously, so that these responses can be considered as false alarms. The authors then proposed to compute a *d'*, which represent participants' awareness of their own performance. The main advantage of the SDT procedure in this context is that it ensures that the sensitivity of the measure is unbiased, that is, unaffected by the participants' own report criteria (i.e., where participants place the criterion for making high and low confidence judgments).

In addition, Tunney and Shanks (2003, Exp. 4) showed that a binary method can detect lower levels of awareness for a given sample size than a rating technique can. Within the same group of participants, they observed that the acquired knowledge was categorized as “explicit” when the participants had to provide binary confidence judgments (by choosing one of the following answers: *yes conforms to rules – more confident*, *yes conforms to rules – less confident, no does not conform to rules – less confident,* and *no does not conform to rules – more confident*), while it was categorized as “implicit” when participants had to provide confidence judgment on a continuous scale (from 50% - complete guess - to 100% - absolutely certain). The authors concluded that measures based on a binary scale provide a more sensitive measure of awareness.

We thus applied the methodology suggested by Tunney and Shanks (2003) to the sequence learning paradigm. Participants performed a Serial Reaction Time task including 18 training blocks, each containing eight repetitions of a 12-elements sequence. They were trained with either a 0 or 1000 ms RSI (Response-Stimulus Interval, i.e., the time interval between the response of the participant and the occurrence of the following stimulus in the sequence). Both groups exhibited learning effects during the learning phase, as evidenced by the RTs decrease with practice and the RTs increase when the training sequence was modified. Subsequently, participants were confronted to a recognition task, in which they were presented with fragments consisting of three trials. Half of the items were part of the training sequence; the other half was not. Participants were asked to react to the stimuli as fast and as accurately as possible, just like in the learning phase, and then to answer the two questions “*Have you seen this short sequence before?*” and “*Are you confident in your response?*” by pressing one of two buttons marked *Yes* or *No*.

For each participant, two d' values were calculated and are reported in Table 1. The first, d'_d, indexed participants’ ability to discriminate between triplets from the training sequence (“old triplets”) and triplets from the transfer sequence (“new triplets”). Hits were calculated from *yes* responses to training triplets, and false alarms were calculated from *yes* responses to transfer triplets. Table 1 shows that participants reliably discriminated between old and new triplets in both the 0ms RSI group [mean d'_d = 0.54, t (18) = 2.58, p < .05] and in the 1000ms RSI group [mean d'_d = 0.83, t (19) = 2.77, p < .05]. Moreover, there was no significant difference between the groups insofar as their discrimination ability is concerned [t (37) = -.814, p > .1]. This suggests that RSI duration did not influence ability to discriminate between old and new triplets.

Insert Table 1 about here

Importantly, the second value, d'_c, measured participants’ ability to discriminate between correct and incorrect responses. Hits and false alarms can be computed as follows in this situation: Considering participants’ responses to the question “*Are you confident in your response?*”, hits are defined as *yes* responses to correct discrimination decisions (including both endorsements and rejections), and false alarms are defined as *yes* responses to incorrect discrimination decisions (again including both endorsements and rejections). As proposed by Tunney and Shanks (2003), if participants are aware of the information used to discriminate between old and new triplets, they should be more confident in correct than in incorrect discrimination responses. On the other hand, if participants are unaware of the information used to discriminate between old and new triplets, then *yes* and *no* confidence responses should be distributed equally between correct and incorrect discriminations.

Thus, explicit knowledge would result in d'_c values greater than zero while implicit knowledge would result in d'_c values close to zero. Table 1 presents the obtained d'_c, values and shows that participants trained with a 1000ms RSI had more explicit knowledge than would be expected by chance [mean d'_c, = 0.471, t (19) = 2.1, p < .05]. Importantly, participants trained with a 0ms RSI were no more confident in their correct discriminations than in their errors [mean d'_c, = 0.266, t (18) = 1.59, p > .1]. According to Tunney and Shanks’ method, we can thus conclude that the knowledge acquired by participants in the 0ms RSI group can be defined as implicit. In other words, these participants did not know that they were able to discriminate between old and new triplets, showing that they had no relevant *metaknowledge*.

This recognition task, and the manner in which we analyzed it, fits well with the perspective developed by Fu, Fu and Dienes (submitted). Indeed, we obtained both *objective* measures (the yes/no recognition responses), which require subjects to indicate whether they are able to discriminate between features of the world, and *subjective* measures (the yes/no confidence judgments), which require subjects to discriminate between their own mental states. Our results indicate a dissociation between these two measures of awareness, as the objective measure is the same for both groups of participants while the subjective measure differs between groups. While both groups were shown able to discriminate between old and new triplets. suppressing the RSI during training subsequently led participants to be less confident in their responses when compared with participants who had been trained with a 1000 ms RSI. This suggests implicit knowledge in the 0ms RSI group. The demonstration is particularly convincing since we used a binary rather than continuous method for the recognition task, —a method that Tunney and Shanks (2003) suggest is more sensitive to explicit knowledge.

Dissociating objective and subjective measures of consciousness therefore appears to be a fruitful way to obtain more precise information regarding the nature of the knowledge acquired during the learning phase of implicit learning experiments. This new procedure illustrates how systematically combining first- and third-person data offers us interesting data for the study of consciousness.

Conclusions

In some cases, we know more than we can tell. In others, we actually tell more than we can know. Finding out what is the case in any specific situation clearly requires a combination of first- and third-person approaches—indeed, this appears to be the only way to go about these issues. Implicit learning offers a rich domain through which to pursue these explorations for it doesn’t coerce us into reducing consciousness to a static, dichotomous property associated with some mental states and not with others, but instead makes it possible to approach it as what it is: A complex, multifaceted phenomenon.

In this paper, we have attempted to overview the challenges that combining first- and third-person methods entail in this domain. What comes out is a mixed bag. On the one hand, considerable, largely methodological, progress has been achieved over the years. The methods that are now typically used in implicit learning research, as we have illustrated, are far more sophisticated then they were 30 years ago. On the other hand, serious issues remain in our appreciation of phenomenology. In an excellent recent paper dedicated to introspection, Marcel (2003) enumerates some of these problems. Most are related to the “observer’s paradox”, namely the fact that the very act of asking a subject to introspect on his mental states changes the nature of these mental states. Hence, attention, according to Marcel, cannot only influence its object, but also distort it and even create it! Marcel amusingly continues by listing several “coping strategies” researchers can deploy to deal with such issues. For instance, one can attempt to minimize or otherwise manipulate the manner in which reports are produced; one can combine several different objective methods; one can explore what happens when reports are to be produced unexpectedly, and so on. The basic problem with all approaches to phenomenological content, however, remains that there is no guarantee that things would have been the same had people not been asked to produce a report. The possibility that introspective awareness is a distinct state of consciousness, however, is itself a question that can be approached empirically, as Overgaard and Soerensen (2004) suggest.

To close, let us also reflect upon the theoretical implications of these discussions. The use of subjective data to explore consciousness, which we fully endorse when used in combination with objective data, is not fully neutral theoretically. In fact, it rests on the assumption that conscious awareness depends on the existence of systems of meta-representations targeted to first-order representations and through which the latter are made conscious. This assumption, however, very much remains an assumption. For instance, one can assume that first-order representations only become conscious when targeted by a relevant meta-representation—a non-inferential “higher-order thought” (Dienes & Perner, 1999; Rosenthal, 1997). One can also think that representations are associated with phenomenal experience because of their properties, for instance, in virtue of their stability (O'Brien & Opie, 1999), their quality (Cleeremans, 2005) or in virtue of the fact that they have “won the competition” (Dennett, 2001). These are extremely difficult, fundamental questions that are perhaps best approached by combining the sophisticated behavioural methods we have overviewed with computational methods that force us to formalize our underlying theories (Atkinson, Thomas, & Cleeremans, 2000; Maia & Cleeremans, 2005).

Authors note

AC is a Research Director of the National Fund for Scientific Research (F.N.R.S., Belgium). VG and MV are Research Assistants with the same institution. This article is based in part on Destrebecqz and Peigneux (2005). This work was supported by an institutional grant to Axel Cleeremans. We would like to thank Zoltan Dienes and a anonymous referee for useful comments on previous versions of this manuscript.

References

Atkinson, A. P., Thomas, M. S. C., & Cleeremans, A. (2000). Consciousness: Mapping the theoretical landscape. Trends in Cognitive Sciences, 4(10), 372-382.

Baars, B. J. (1988). A cognitive theory of consciousness. Cambridge: Cambridge University Press.

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. (2005). The iowa gambling task and the somatic marker hypothesis: Some questions and answers. Trends in Cognitive Science, 9(4), 159-162.

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. Science, 275, 1293-1295.

Chalmers, D. J. (1999). First-person methods in the science of consciousness. In Consciousness bulletin: Center for Consciousness Studies, University of Arizona.

Cheesman, J., & Merikle, P. M. (1984). Priming with and without awareness. Perception and Psychophysics, 36, 387-395.

Cleeremans, A. (2001). Conscious and unconscious processes in cognition. In N. J. Smelser & P. B. Baltes (Eds.), International encyclopedia of the social and behavioral sciences (Vol. 4, pp. 2584-2589). London: Elsevier.

Cleeremans, A. (2005). Computational correlates of consciousness. In S. Laureys (Ed.), Progress in brain research (Vol. 150, pp. 81-98). Amsterdam: Elsevier.

Cleeremans, A., Destrebecqz, A., & Boyer, M. (1998). Implicit learning: News from the front. Trends in Cognitive Sciences, 2, 406-416.

Cohen, A., & Curran, T. (1993). On tasks, knowledge, correlations, and dissociations : Comment on perruchet and amorim (1992). Journal of Experimental Psychology : Learning, Memory and Cognition, 19, 1431-1437.

Crick, F. H. C., & Koch, C. (1990). Towards a neurobiological theory of consciousness. Seminars in the Neurosciences, 2, 263-275.

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.

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., & Peigneux, P. (2005). Methods for studying unconscious learning. In S. Laureys (Ed.), Progress in brain research (Vol. 150). Amsterdam: Elsevier.

Dienes, Z., & Altmann, G. T. M. (1997). Transfer of implicit knowledge across domains: How implicit and how abstract? In D. C. Berry (Ed.), How implicit is implicit learning? (Vol. 5, pp. 107-123): Oxford University Press.

Dienes, Z., Altmann, G. T. M., Kwan, L., & Goode, A. (1995). Unconscious knowledge of artificial grammars is applied strategically. Journal of Experimental Psychology: Learning, Memory and Cognition, 21, 1322-1338.

Dienes, Z., & Berry, D. C. (1997). Implicit learning : Below the subjective threshold. Psychonomic Bulletin & Review, 4, 3-23.

Dienes, Z., Broadbent, D. E., & Berry, D. C. (1991). Implicit and explicit knowledge bases in artificial grammar learning. Journal of Experimental Psychology : Learning, Memory and Cognition, 17, 875-887.

Dienes, Z., & Perner, J. (1999). A theory of implicit and explicit knowledge. Behavioral and Brain Sciences, 22, 735-808.

Dienes, Z., & Perner, J. (2004). Assumptions of a subjective measure of consciousness: Three mappings. In R. Gennaro (Ed.), Higher order theories of consciousness (pp. 173-199). Amsterdam: John Benjamins Publishers.

Dulany, D. E., Carlson, R. A., & Dewey, G. I. (1984). A case of syntactical learning and judgement : How conscious and how abstract? Journal of Experimental Psychology : General, 113, 541-555.

Ericsson, K. A., & Simon, H. A. (1980). Verbal reports as data. Psychological Review, 87, 215-251.

Ericsson, K. A., & Simon, H. A. (1984). Protocol analysis: Verbal reports as data. Cambridge, MA: MIT Press.

Frensch, P. A., Buchner, A., & Lin, J. (1994). Implicit learning of unique and ambiguous serial transitions in the presence and absence of a distractor task. Journal of Experimental Psychology : Learning, Memory and Cognition, 20, 567-584.

Frith, C. D., Perry, R., & Lumer, E. (1999). The neural correlates of conscious experience : An experimental framework. Trends in Cognitive Sciences, 3, 105-114.

Fu, Q., Fu, X., & Dienes, Z. (submitted). Implicit sequence learning and conscious awareness.

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.

Jiménez, L., Méndez, C., & Cleeremans, A. (1996). Comparing direct and indirect measure of sequence learning. Journal of Experimental Psychology-Learning Memory and Cognition, 22, 948-969.

Kirsh, D. (1991). When is information explicitly represented? In P. P. Hanson (Ed.), Information, language, and cognition. New York, NY: Oxford University Press.

Macmillan, N. A., & Creelman, C. D. (1991). Detection theory: A user's guide. New York: Cambridge University Press.

Maia, T. V., & Cleeremans, A. (2005). Consciousness: Converging insights from connectionist modeling and neuroscience. Trends in Cognitive Science, 9(8), 397-404.

Maia, T. V., & McClelland, J. L. (2004). A reexamination of the evidence for the somatic marker hypothesis: What participants really know in the iowa gambling task. Proc Natl Acad Sci U S A, 101(45), 16075-16080.

Maia, T. V., & McClelland, J. L. (2005). The somatic marker hypothesis: Still many questions but no answers. Trends in Cognitive Science, 9(4), 162-164.

Mandler, J. M. (1983). Representation. In J. H. Flavell & E. M. Markman (Eds.), Cognitive development (Vol. 3, pp. 255-308). New York, NY: Wiley.

Marcel, A. J. (2003). Introspective report: Trust, self-knowledge and science. Journal of Consciousness studies, 10(9), 167-186.

Nahmias, A. (2002). Verbal reports on the contents of consciousness - reconsidering introspectionist methodology. Psyche, 8(21).

Neal, A., & Hesketh, B. (1997). Episodic knowledge and implicit learning. Psychonomic Bulletin & Review, 4, 24-37.

Newell, A., & Simon, H. A. (1972). Human problem solving. Englewood Cliffs, NJ: Prentice-Hall.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can do: Verbal reports on mental processes. Psychological Review, 84, 231-259.

Nissen, M. J., & Bullemer, P. (1987). Attentional requirement of learning: Evidence from performance measures. Cognitive Psychology, 19, 1-32.

O'Brien, G., & Opie, J. (1999). A connectionist theory of phenomenal experience. Behavioral and Brain Sciences, 22, 175-196.

Overgaard, M. (2001). The place for phenomenology in experiments on consciousness. Psycoloquy, 12(9), 1-13.

Overgaard, M., & Sœresen, T. A. (2004). Introspection distinct from first-order experiences. Journal of Consciousness studies, 11(7-8), 77-95.

Perner, J. (1991). Understanding the representational mind. Cambridge, MA: MIT Press.

Perruchet, P., & Amorim, M. A. (1992). Conscious knowledge and changes in performance in sequence learning : Evidence against dissociation. Journal of Experimental Psychology-Learning Memory and Cognition, 18, 785-800.

Perruchet, P., & Gallego, J. (1993). Associations between conscious knowledge and performance in normal subjects : Reply to cohen and curran (1993) and willingham, greeley, and bardone (1993). Journal of Experimental Psychology : Learning, Memory, & Cognition, 19, 1438-1444.

Perruchet, P., Gallego, J., & Savy, I. (1990). A critical reappraisal of the evidence for unconscious abstraction of deterministic rules in complex experimental situations. Cognitive Psychology, 22, 493-516.

Perruchet, P., & Vinter, A. (2002). The self-organizing consciousness. Behavioral and Brain Sciences, 25(3), 297-330.

Reber, A. S. (1967). Implicit learning of artificial grammars. Journal of Verbal Learning and Verbal Behavior, 5, 855-863.

Reber, A. S. (1969). Transfer of syntactic structure in synthetic languages. Journal of Experimental Psychology, 81, 115-119.

Reingold, E. M., & Merikle, P. M. (1988). Using direct and indirect measures to study perception without awareness. Perception & Psychophysics, 44, 563-575.

Reingold, E. M., & Merikle, P. M. (1990). On the inter-relatedness of theory and measurement in the study of unconscious processes. Mind and Language, 5, 9-28.

Rosenthal, D. (1997). A theory of consciousness. In N. Block, O. Flanagan & G. Güzeldere (Eds.), *The nature of consciousness: Philosophical debates*. Cambridge, MA: MIT Press.

Sergent, C., & Dehaene, S. (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. Psychological Science, 15(11), 720-728.

Shanks, D. (2005). Implicit learning. In K. Lamberts & R. Goldstone (Eds.), Handbook of cognition (pp. 202-220). London: Sage.

Shanks, D. R., & Johnstone, T. (1998). Implicit knowledge in sequential learning tasks. In M. A. Stadler & P. A. Frensch (Eds.), Handbook of implicit learning (pp. 533-572). Thousand Oaks, CA: Sage Publications.

Shanks, D. R., & Johnstone, T. (1999). Evaluating the relationship between explicit and implicit knowledge in a serial reaction time task. Journal of Experimental Psychology : Learning, Memory, & Cognition.

Shanks, D. R., & St John, M. F. (1994). Characteristics of dissociable human learning systems. The Behavioral and Brain Sciences, 17, 367-447.

Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. Behavioral and Brain Sciences, 17, 367-447.

Tunney, R. J., & Shanks, D. (2003). Subjective measures of awareness and implicit cognition. Memory and Cognition, 31(7), 1060-1071.

Willingham, D. B., Greeley, T., & Bardone, A. M. (1993). Dissociation in a serial response time task using a recognition measure : Comment on perruchet and amorim (1992). Journal of Experimental Psychology-Learning Memory and Cognition, 19, 1424-1430.

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.

Ziori, E., & Dienes, Z. (2006). Subjective measures of unconscious knowledge of concepts. Mind and Society, 5(1), 105-122.

Table 1
Measures of Discrimination Accuracy (d'd) and Binary Awareness (d'c) for both groups of participants

	RSI 0 ms				RSI 1000 ms			
	Discrimination (d'd)		Binary Awareness (d'c)		Discrimination (d'd)		Binary Awareness (d'c)	
Measure	*Mean*	*SE*	*Mean*	*SE*	*Mean*	*SE*	*Mean*	*SE*
p(H)	0.76	0.03	0.57	0.05	0.72	0.03	0.63	0.04
p(FA)	0.60	0.05	0.49	0.04	0.47	0.06	0.50	0.06
C	-0.66	0.17	-0.11	0.09	-0.23	0.17	-0.13	0.13
d'	0.53	0.21	0.27	0.17	0.83	0.30	0.47	0.22

Note—The C statistic is a measure of response bias. Values greater than zero indicate liberal bias, and those less than zero indicate conservatism; see Macmillan and Creelman (1991) and Tunney and Shanks (2003).