

Implicit learning shapes new conscious percepts and representations

PIERRE PERRUCHET and ANNIE VINTER
Université de Bourgogne, Dijon, France

and

JORGE GALLEGO
Université de Paris VII, Paris, France

We present here the lineaments of a new account of implicit learning, an account that does not rely on the notion of "implicit knowledge." In this account, improved performance depends on the action of unconscious mechanisms that structure the phenomenal, conscious experience of the world. This integrative view makes groundless the search for dissociations between conscious and unconscious influences that has been at the core of the research on implicit learning and memory. We contrast this view, on the one hand, to Dienes and Berry's (1997) proposal, which defines implicit learning by analogy with subliminal perception, and, on the other, to Neal and Hesketh's (1997) episodic account, in which subjective experience is a starting point for inquiry, rather than the phenomenon requiring explanation.

The basic postulate of implicit learning research is that performance improvement testifies to the acquisition of knowledge, and a large part of the literature in this area is devoted to the question of deciding whether or not this knowledge is implicit. Most authors argue that it is, thus admitting that implicit learning provides unconscious knowledge, at least in some sense. The papers in this issue by Dienes and Berry (1997; henceforth D&B) and Neal and Hesketh (1997; N&H) provide a new illustration of this dominant position. In contrast, some researchers (e.g., Perruchet & Pacteau, 1990) have been critical of this idea. They argue that the pieces of knowledge assumed to be implicit either are not necessary for performance or can be shown to be conscious when proper tests are administered (for a review, see Shanks & St. John, 1994). Our objective here is to continue beyond this critical view, by suggesting that the very question of the implicit or explicit nature of the knowledge that emerges from a learning episode may not be meaningful. We will present, first, the outlines of an alternative view of implicit learning (described in more detail in Perruchet & Gallego, in press, and Perruchet & Vinter, in press); then, we will show that this alternative view undermines the validity of the research aimed at providing positive or negative evidence for dissociations, whether between tasks or between processes.

Preparation of this article was supported by the University of Bourgogne and the Centre National de la Recherche Scientifique (ESA CNRS 5022). We would like to thank Henry L. Roediger III and three anonymous reviewers for their careful reviews and helpful suggestions on a previous draft of this article. Correspondence should be addressed to P. Perruchet, LEAD, Faculté des Sciences, Université de Bourgogne, 6 Boulevard Gabriel, 21000 Dijon, France (e-mail: perruche@satie.u-bourgogne.fr).

Implicit Knowledge as Optional Inference

As pointed out by N&H, "virtually all information processing models appear to assume that complex forms of information processing occur outside of awareness" (p. 34). However, this is a tacit assumption, and the possibility that this mode of processing differs from the conscious one is not even discussed. Information processing models construe (unconscious) mental activities as mimicking the organization of our conscious modes of thought. In the same way that we solve a problem by performing intentional operations on explicit representations, the adaptive behavior that occurs outside awareness is assumed to be due to unconscious operations acting on implicit representations.

Let us consider an example from auditory perception to illustrate that another view of mental life is possible. Humans are able to state the spatial orientation of a sound-emitting object, thanks to the processing of binaural cues. The binaural listener makes use of the physical differences in stimulation that arise between the two ears because of their separation in space. For instance, a sound source can be located, under some circumstances, on the basis of differences in the time of arrival of stimulation to the two ears of as little as .0001 sec (see, e.g., Handel, 1989). The mechanism responsible is the precedence effect, in which the neural activity produced by the first tone partially inhibits the response to the later one (Walach, Newman, & Rosenzweig, 1949). This mechanism is efficient because it exploits a specific property of sound waves—namely, their speed of propagation. However, everybody agrees, we presume, that it would make no sense to endow laypeople with some knowledge about the speed of propagation of sound. Insofar as the notion of knowledge makes sense here, it is only meaningful from the observer's (or scientist's) point of view. Now let

us assume that instead of being due to innate mechanisms the neural substrate of which is at least partially known, the ability to detect the source of a sound is acquired through practice. Our guess is that, in this case, improved sound localization, because it testifies to an improvement in subjects' abilities to exploit the speed of propagation of sound waves, might lead some developmental and cognitive psychologists to postulate the acquisition of some implicit knowledge about this property.

It is worth stressing that this inference ought to be paradoxical. Location detection is too fundamental to survival to be delayed in one's lifetime. However, the fact that the mechanisms of sound localization are subsequently adjusted to accommodate for the increasing distance between the two ears, presumably thanks to concurrent visual information, supports the view that this kind of ability is not beyond the scope of learning mechanisms. Another indication is provided by the remarkable use of auditory information in spatial localization by blind individuals who have had early visual orientation experience prior to losing their sight (Veraart & Wanet-Defalque, 1987). If the end product of the learned ability is the same as the hard-wired innate implementation, why should we draw different inferences about the likelihood of assuming the existence of an implicit knowledge base? We suggest that the common inference for implicit knowledge in cognitive science is no more justified than would be the claim that our ability to localize sounds in space testifies to the implicit knowledge of the properties of sound waves, or of any other pieces of implicit knowledge.

The Structure of Conscious Experience as the End Product of Implicit Learning

Our preceding argument does not mean, however, that implicit learning has no relation to any form of knowledge. Let us again consider the case of sound localization. The product of the mechanisms implemented in the auditory system is nothing less than our conscious experience of sound localization. We suggest that the neural modification following an implicit learning episode affects our conscious experience of the world in the same way as the hard-wired mechanisms located in the auditory system shape our conscious representation of the sound space.

In the framework that we propose, implicit learning forms the basis of conscious experience; that is to say, it shapes both subjects' perceptions and subjects' internal representations of the world. Because subjects have conscious perception and representations before starting any implicit learning tasks, *implicit learning may be thought of as allowing subjects to pass from conscious perceptions and representations to other, generally better structured, conscious perceptions and representations, through the action of intrinsically unconscious mechanisms.* For instance, after exposure to a corpus generated by a finite state grammar including VXT as a frequent recursion, subjects no longer perceive "V," "X," and "T" as three familiar but separate entities, but perceive "VXT" as an in-

creasingly familiar unit, with various consequences for behavior. The way in which the data are coded is changed, and this change directly affects our phenomenal experience. The phenomenon is obvious in any natural situation in which implicit learning is assumed to operate, whether this situation concerns the acquisition of a first or second language, natural categories, reading and writing abilities, or sensitivity to musical structure. It is hardly defensible that our subjective experience of that part of the environment with which we interact in each of these cases remains unchanged while training progresses. Our argument is that these changes in the way we consciously perceive and interact with the environment are at the core of implicit learning.

Note that, in our view (see Dulany, 1996, for a related perspective), neither the conscious nor the unconscious aspects of mind are negated. But they characterize different elements of mental life: processes and mechanisms responsible for the elaboration of knowledge are intrinsically unconscious, and the resulting mental representations and knowledge are intrinsically conscious.¹ No other components are needed, including the notion of implicit knowledge.

Dissociating Conscious From Unconscious Influences

A major consequence of the conception outlined above is that the search for dissociations between conscious and unconscious forms of thought, which is the primary objective of research in the implicit learning and memory fields, loses any meaning. This objective is equivalent to attempting to oppose the physiological excitatory and inhibitory processes operating in the auditory system to the conscious experience of sound space.

The concept of dissociation (and association as well) between conscious and unconscious influences presupposes that conscious and unconscious phenomena constitute two parallel domains, each of them having its own organization and its own influences on behavior. The "process dissociation procedure" proposed by Jacoby (e.g., 1991) provides an excellent illustration of this dualistic conception. In this method, performances in a situation where conscious and unconscious processes act together are compared with performances in a situation where the processes act in opposition. The resulting performance is expressed as a function of an additive combination of conscious (or intentional) and unconscious (or automatic) influences. Several researchers have questioned the postulate of independence between conscious and unconscious influences, which is a precondition for the validity of the procedure of estimation of the two kinds of influences (e.g., Curran & Hintzman, 1995). However, a prerequisite for assessing the level of dependence or independence is the idea that distinguishing between the two kinds of influences is meaningful. It is clear that, given our framework, the method itself is not valid, because the distinction between conscious and unconscious influences is groundless (hence we are far more critical than N&H, who propose abandoning the estima-

tion procedure, given the problems it raises, while keeping the fundamental insight of the process opposition procedure).

It may be argued that an approach in terms of the dissociation of conscious and unconscious phenomena has largely demonstrated its heuristic value, in the sense that an overwhelming amount of experimental data has been generated and interpreted within this framework. This argument is not conclusive, however, because an alternative interpretation of the experimental data is possible. In line with the view outlined above, we suggest that all the methods devised to show dissociations between conscious and unconscious phenomena actually contrast various forms of conscious experience, all of which are due to unconscious mechanisms. Let us consider Jacoby's fame judgment experiments, a situation in which conscious and unconscious processes are assumed to act in opposition (e.g., Jacoby, Woloshyn, & Kelley, 1989). In these experiments, names that had previously been read within a list of nonfamous people were sometimes declared as famous in a subsequent test. This result has been interpreted as an automatic effect of the prior exposure to these names. Our account of this result is far simpler: subjects explicitly remembered that they had read these names in the past, but did not remember the relevant aspects of the spatiotemporal context of encoding—namely, that they read these names in a list of nonfamous people. The false fame effect in Jacoby's experiments reveals not the nonintentional effect of an earlier episode, but the intentional use of false or impoverished memory (we note that N&H do not reject this interpretation).²

Along the same lines, we suggest that all the measures taken from an implicit learning experimental setting reveal the unconscious formation of conscious percepts and representations, whether these measures are obtained in the test or transfer phase (the "performance") or in subsequent "explicit" tests. However, the content of the conscious knowledge base tapped by each of the two categories of measures is partly different from that tapped by the other. Let us suppose that a subject is shown material generated by a finite state grammar in which the trigram VXT occurs frequently. This trigram will become a subjective perceptual and representational unit. This (explicit) knowledge may be captured through subsequent performance in transfer tasks. For instance, the fact that VXT is a subjective unit is indicated by the result that new letter strings including VXT would tend to be judged as grammatical in the transfer test, or would tend to be perceived more fluently. Likewise, this knowledge may be captured by the result that VXT is generated in a subsequent recall test. However, the recall test also implies that the subject is able to recollect the spatiotemporal context of initial encoding. The difference between the two methods of measuring the effects of exposure to the material does not lie in the implicit versus explicit status of the underlying knowledge, but instead in the specific pieces of explicit knowledge that are required in order to perform the task. The "explicit" tests involve some memory about the spatiotemporal context

in which knowledge has been acquired, whereas this knowledge is not needed in the transfer task.

This framework makes it possible to account for the frequently observed parallel between knowledge revealed through transfer tasks and through explicit memory tasks in normal subjects. Indeed, both kinds of measures imply the formation of the same units of coding. But this framework also accounts for the result of improved performance in the transfer tasks unaccompanied by the memory of the spatiotemporal context in which the new ability was acquired, as has been observed, for instance, in amnesic patients. Simply, these dissociations are no longer accounted for in terms of the manifestation of implicit knowledge, which would be intact in amnesics. They are due to the fact that the transfer task does not require certain elements of explicit knowledge which are required in the explicit memory task. Above chance performance during the test phase in the absence of explicit remembering for having viewed the material during the study phase is not indicative of implicit knowledge, but of impoverished explicit knowledge. Each of us has explicit knowledge about the fact that London is the capital of England, and, presumably, some of us are able to remember the spatiotemporal contexts in which they read or heard this information, whereas others are not. This does not imply that the representation of London as England's capital is explicit for the former and implicit for the latter: it is fully conscious for everyone. Simply, the context in which this explicit knowledge was acquired either is still available in memory or has been forgotten. This is exactly the same trivial phenomenon that is captured by what is commonly presented as a dissociation between performance (or implicit knowledge) and explicit knowledge. We have no need to remember when and how we have learned that London is the capital of England to be successful in a test assessing our knowledge of the world's capitals. Likewise, we have no need to remember the spatiotemporal conditions in which we have been faced with VXT in order to perceive a letter string comprising VXT more fluently.

The Dienes and Berry Subjective Threshold Criterion

D&B's analysis starts from the analogy that they suggest between implicit learning and subliminal perception. This analogy is strikingly opposed to our framework. First, it leads one to confer on implicit learning a characteristic of subliminal perception—namely, that of being a marginal phenomenon, the very existence of which is sometimes questioned and the adaptive role of which remains unclear (whatever its theoretical interest). In our opinion, implicit learning does not possess such features: implicit learning is at the root of our conscious perception and representation of the world, and its importance for adaptive behavior is crucial. But the analogy with subliminal perception is flawed, in our opinion, for a more fundamental reason.

We referred above to the traditional information processing perspective, in which presumably unconscious

mental activities are conceived as mimicking the organization of conscious modes of thought. The D&B analogy provides a specific instantiation of this general framework. Let us consider the distinction that the authors draw at the very beginning of their paper, between solving an algebraic problem and speaking one's native language. In the former case, behavior is governed by explicit rules and strategy, whereas in the latter, explicit rule knowledge may be lacking. D&B suggest that behavior is grounded on a knowledge base in both cases, the difference being that this knowledge base is situated either above or below a threshold of consciousness. The implicit or explicit nature of knowledge is conceived as a continuum, using the model of the energy or the signal-to-noise ratio of an external stimulus in subliminal perception. In our conception, the representation of the rules underlying the task structure is not below a certain threshold; insofar as implicit learning is concerned, such a representation does not exist at all. Improved performance after an implicit learning episode is consecutive to a change in the conscious perception and representation of the environment, owing to the involvement of unconscious mechanisms. Obviously, the rules structuring the environment can be discovered, through scientific investigations based on hypothetico-deductive reasoning and logical inference. In the latter case, they become explicitly represented in subjects' minds. However, as we illustrated in the analogy of sound localization above, our point is that there is no continuity between the (intrinsically unconscious) mechanisms that exploit some structural properties of the environment to ensure immediate behavioral adaptations and the (intrinsically conscious) knowledge of these properties. To resume, the former is not a "subliminal" version of the latter.

D&B propose two criteria to assess whether or not knowledge is below the subjective threshold. The guessing criterion posits that subjects believe they are guessing when they are above chance, and the zero correlation criterion asserts that confidence is unrelated to the accuracy of responding. The justification of the authors is that "these characteristics might correspond to an important aspect of what the layperson means by unconscious learning" (p. 5). We have no privileged insight into the layman's view of unconscious learning, but we believe that the D&B criteria are not compatible with the role conferred on implicit learning by most contributors in the field. For instance, we fail to see how these criteria can be reconciled with the fact of conceiving the acquisition of one's mother tongue as a prototypical instance of implicit learning, as D&B do themselves at the very beginning of their paper. It seems obvious that a sentence that we judge as grammatical with a high degree of confidence is more likely to actually be grammatical than a sentence for which a judgment of well-formedness is associated with low confidence. The very same line of reasoning can be pursued with other real-life manifestations of implicit learning, whether they are instantiated in the physical, so-

cial, musical, or any other domain. D&B are seemingly aware of this contradiction, and they attempt to account for the case of natural language by introducing a distinction between "content explicit" representation (in which subjects represent themselves as being in the possession of a certain propositional content) and "attitude explicit" representation (in which subjects represent this content as knowledge, and not just as guessing). Language, they argue, is a case of attitude explicitness without content explicitness. Although content explicitness without attitude explicitness may adequately qualify some representations, we fail to grasp the meaning of the reverse combination, according to which we represent as knowledge, and not as guessing, a certain propositional content that we do not represent as being in our possession.

We acknowledge that, in some sense, the explicit knowledge gained through implicit learning processes is a "guess," because no rational, analytical justifications are normally available. The well-formedness of a linguistic statement or of a musical piece, or the natural processes of categorization, cannot be justified, except, in some cases, by a posteriori considerations. At the end of an experimental training session, participants may perceive a letter string as familiar, or expect that the next occurrence of a target will be in a given location, without being able to justify their feelings. In the same way, we perceive a sound as coming from a given location, without being able in any way to specify the cues that we are using. In this sense, all these experiences are "guesses." We have no introspective access to the mechanisms that shape our phenomenal experience. However, this lack of rational justifications is not related to the confidence that we may have in the correctness of our experiences, except, of course, when the task implies that such justifications should be provided. Just as we may be quite confident (or not) about the location of a sound-emitting object, we may also be quite confident (or not) about the adequation of our subjective experience as shaped by implicit learning processes, whether this experience consists in perceiving a sentence as grammatical, a musical piece as well-formed, or a given animal as a dog. A low level of confidence is in no way a characteristic of the end product of implicit learning.

The Future of Implicit Learning Research

We suggest that the widespread notion of implicit knowledge, and hence today's dominant trend toward establishing dissociations between implicit and explicit forms of thought, should be rejected. This critical recommendation must not be misunderstood. We are *not* suggesting that one should abandon research into the kind of phenomena that are captured in the implicit learning experiments. On the contrary, we believe that our perspective opens up a far more fruitful and exciting field of research than that offered by the current approach. Rather than conceive unconscious and conscious phenomena as two parallel components of mind, we propose a fully in-

tegrated perspective in which each act of thought involves conscious representations, shaped by unconscious processes during prior experiences.

Implicit learning has been characterized above as the process that allows subjects to shape new conscious percepts and representations. In human adults, these percepts and representations are highly structured. One perceives and represents unitary and meaningful objects, and not a set of juxtaposed features. One codes the discourse of a speaker as a succession of words and structured sentences, not as a stream of individual auditory signals. We consider that the primary aim of a theory of implicit learning is to account for these sophisticated abilities. Because many of these abilities are acquired during childhood, a theory of implicit learning must be relevant to the traditional laboratory situations in adults as well as to developmental phenomena. Researchers have, among other tasks, to explain how children become able to parse sensory input into physically and linguistically relevant units, and how they become increasingly sensitive to the structural constraints embedded in their environment. Future research should focus on understanding how unconscious learning may proceed in a way that endows human subjects with such a powerful perceptual and representational system.

Is N&H's episodic account a candidate theory for fulfilling this objective? The interest of the N&H account resides in its highlighting the fact that behavior may be affected by individual episodes rather than simply by large amounts of training. The scope of this account seems tightly limited, however. Indeed, this account posits as a premise the very fact that we are attempting to explain—namely, that subjects are able to process meaningful episodes. In consequence, its scope of relevance is limited to certain microlevel phenomena. An episodic account appears to be a quite improbable explanation as soon as we consider learning on the scale of the life span. Extending an episodic account of learning as a developmental theory, for instance, leads to a quite unrealistic view of development, in which people store a linearly increasing number of separate episodes from early infancy to old age. Below we present the outlines of an account whose objective is to explain the formation of the most fundamental abilities, in keeping with the general framework described above.

Our account, which has been applied to conventional implicit learning phenomena (Perruchet & Gallego, in press) and to certain developmental data (Perruchet & Vinter, in press), is based on the idea that the formation of subjective, conscious units of perception and representation is due to the action of unconscious and mandatory associative mechanisms. This account has evolved from the view that N&H designate as the “fragmentary knowledge account” of artificial grammar learning (see, e.g., Perruchet & Pacteau, 1990). We acknowledge, however, that the notion of “fragments” implies several questionable assumptions. It suggests that, owing to certain

constraints inherent to the processing system, subjects encode only a part of a primitive entity. Moreover, it suggests that this part is a randomly selected sample of this primitive entity. In our view, the primitive entity (for instance the set of three to eight letters forming an item generated by a finite-state grammar) exists only for the experimenter. From the subject's point of view, the primitives are in fact the letters, or perhaps some familiar combinations of a few letters. Our proposal is that the “fragments” are in fact subjective units that are built by linking together the primitives of the system. Let us posit now that these units are the end products of associative processes, such as those which are described in the traditional literature on associative learning. These processes endow the system with considerable power. Indeed, it can be shown in the context of artificial grammar that the resulting units are not random samples of the material. Instead, because they are formed from the most frequently associated letters in the displayed strings, they tend to match the units that are meaningful from an analytical, structural standpoint. For instance, if the grammar includes a recursive loop, the set of letters forming this loop will tend to form a cognitive unit, as a simple application of ubiquitous associative laws (for a demonstration, see Perruchet & Gallego, in press; and for a different account of the same phenomenon, see Servan-Schreiber & Anderson, 1990).

This account provides a powerful explanation of the main result of artificial grammar studies—namely, that after being exposed to a set of grammatical items, participants become able to discriminate between grammatical and nongrammatical *new* items. The new letter strings displayed during the test tend to be a recombination of small sets of letters displayed during the familiarization phase. Because associative processes entail a match between subject's cognitive units and the structurally relevant units, the small sets of letters which are recombined tend to be those that subjects have previously encoded as a single unit. Moreover, a conception framing implicit learning as the formation of cognitive units through the action of associative processes can be easily expanded to encompass phenomena that go beyond the complexity of laboratory situations. This is mainly due to the possibility of hierarchical processing which it offers. Hierarchical processing occurs whenever the units built from a given set of primitives themselves become the primitives for higher level units. Hierarchical processing is useful, not only as a tool for apprehending the complexity of the world with limited capacities, but because many domains that we have to penetrate, such as natural language, are themselves hierarchically organized.

Of course, this account is still at an early stage. Many aspects need to be made more precise, and many important questions have not even been raised. We hope that we have achieved our aim of convincing the reader that the future of research should not lie in endless discussions about whether knowledge should be qualified as explicit

or implicit, but instead will consist in accounting for the formation of conscious perceptions and representations through implicit learning mechanisms.

REFERENCES

- CLEEREMANS, A. (1993). *Mechanisms of implicit learning: A connectionist model of sequence processing*. Cambridge, MA: MIT Press, Bradford Books.
- CURRAN, T., & HINTZMAN, D. L. (1995). Violations of the independence assumption in process dissociation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 531-547.
- DIENES, Z., & BERRY, D. (1997). Implicit learning: Below the subjective threshold. *Psychonomic Bulletin & Review*, **4**, 3-23.
- DULANY, D. E. (1996). Consciousness in the explicit (deliberative) and implicit (evocative). In J. D. Cohen & J. W. Schooler (Eds.), *Scientific approaches to the study of consciousness* (pp. 179-212). Mahwah, NJ: Erlbaum.
- ELMAN, J. L. (1990). Representation and structure in connectionist models. In G. Altman (Ed.), *Cognitive models of speech processing* (pp. 345-382). Cambridge, MA: MIT Press.
- HANDEL, S. (1989). *Listening: An introduction to the perception of auditory events*. Cambridge, MA: MIT Press.
- JACOBY, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory & Language*, **30**, 513-541.
- JACOBY, L. L., WOLOSHYN, V., & KELLEY, C. M. (1989). Becoming famous without being recognized: Unconscious influences of memory produced by dividing attention. *Journal of Experimental Psychology: General*, **118**, 115-125.
- MANDLER, G. (1991). Your face looks familiar but I can't remember your name: A review of dual process theory. In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 207-225). Hillsdale, NJ: Erlbaum.
- NEAL, A., & HESKETH, B. (1997). Episodic knowledge and implicit learning. *Psychonomic Bulletin & Review*, **4**, 24-37.
- PERRUCHET, P., & GALLEG0, J. (in press). A subjective unit formation account of implicit learning. In D. Berry (Ed.), *How implicit is implicit learning*. Oxford: Oxford University Press.
- PERRUCHET, P., & PACTEAU, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General*, **119**, 264-275.
- PERRUCHET, P., & VINTER, A. (in press). Implicit learning and development: The implicit knowledge assumption reconsidered. In M. Stadler & P. Frensch (Eds.), *Handbook of implicit learning*. Thousand Oaks, CA: Sage Publications.
- SERVAN-SCHREIBER, D., & ANDERSON, J. R. (1990). Learning artificial grammars with competitive chunking. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 592-608.
- SHANKS, D. R., & ST. JOHN, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral & Brain Sciences*, **17**, 367-447.
- VERAART, C., & WANET-DEFALQUE, M.-C. (1987). Representation of locomotor space by the blind. *Perception & Psychophysics*, **42**, 132-139.
- WALLACH, H., NEWMAN, E. B., & ROSENZWEIG, M. R. (1969). The precedence effect in sound localization. *American Journal of Psychology*, **62**, 315-336.

NOTES

1. In this conception, in contrast with the information processing tradition, processes *generate* representations, instead of manipulating (storing, transforming, comparing, retrieving, etc.) built-in representations. A similar distinction emerges in connectionist networks. The "representations" provided through cluster or factor analysis of the pattern of activation in the studies by Cleeremans (e.g., 1993) or Elman (e.g., 1990) are the by-product of algebraic operations. However, there is no direct manipulation of these representations as such. Indeed, algebraic operations only concern individual connections.

2. This general reasoning is also relevant for the distinction between familiarity and recollection. In the conventional view, the feeling of familiarity and the judgment of recollection tap, respectively, unconscious and conscious phenomena (e.g., Mandler, 1991). In our view, both familiarity and recollection are conscious experiences. In the first case, the phenomenal experience may be "I have seen this face somewhere in the past"; in the second case, the phenomenal experience may be "I saw this face in the corridor yesterday." The difference is not that the first experience is unconscious and the second one is conscious: what differs is the content, and notably the accuracy, of the phenomenal experience. Certain spatiotemporal contextual elements are lacking in the feeling of familiarity. Also, in our view, both familiarity and recollection are due to unconscious processes. There is a consensus about the fact that we have no access to the processes that generate the feeling of familiarity. However, we fail to see the sense in which we have a better conscious access to the mechanisms leading to explicit recollection judgments.

(Manuscript received September 10, 1996;
revision accepted for publication November 6, 1996.)