

Future directions for implicit learning: Toward a clarification of issues associated with knowledge representation and consciousness

ANDREW NEAL

University of Queensland, Brisbane, Queensland, Australia

and

BERYL HESKETH

Macquarie University, Sydney, New South Wales, Australia

Areas of agreement and disagreement regarding knowledge representation and consciousness within implicit learning research are reviewed. It is argued that further progress in the field requires more precise definitions of abstract and episodic knowledge, and of conscious and unconscious forms of cognition. In particular, we argue that implicit learning research should be informed by more general theories of memory and performance, and that concepts such as consciousness are unlikely to act as explanatory constructs within this approach.

The target articles and commentaries reveal a wide diversity of opinion within the implicit learning field. Although a number of issues have been raised in this symposium, we believe that the two key issues dividing opinion in the field remain knowledge representation and consciousness. Although the symposium has revealed a number of areas of general agreement regarding knowledge representation, researchers still disagree over the flexibility of episodic models, the status of neuropsychological evidence, the computational efficiency of episodic models, and the scope of the episodic account. In relation to consciousness, researchers now appear to agree that task dissociations no longer represent convincing evidence for the existence of qualitative differences between implicit and explicit learning. However, the responses to this problem are varied. Dienes and Berry (1997), Stadler (1997), and Mathews (1997) continue to argue for the existence of such qualitative differences. Neal and Hesketh (1997) and Reber (1997) raise problems with the concept of qualitative differences but reach no firm conclusions regarding the status of such distinctions. Whittlesea and Dorken (1997) and Perruchet, Vinter, and Gallego (1997) specifically argue that the concept of qualitative differences is not necessary to explain the results from implicit learning studies. In order to resolve this issue, it is essential that the different metaphors used by researchers to characterize the relationship between implicit and explicit processes be clarified.

REPRESENTATION

From our perspective, one of the more interesting outcomes of the symposium is that it has clarified the precise nature of the commonalities and differences among the various approaches to studying implicit learning. In particular, it is worth noting that there appears to be agreement regarding several of the key assumptions of the episodic account. It now seems widely accepted that knowledge representation is dependent on the nature of the encoding and retrieval operations used by the individual, and that abstraction does not occur automatically whenever one attends to a structured stimulus display (Reber, 1997; Whittlesea & Dorken, 1997). Furthermore, critics of the episodic account (e.g., Mathews, 1997) have not disputed evidence suggesting that task performance is sensitive to the match between encoding and retrieval operations. However, critics of the episodic account have argued that pure episodic models are limited because they lack sufficient power to account for performance on complex tasks characteristic of learning in the "real world" (Mathews, 1997; Perruchet et al., 1997). These concerns are reviewed below, followed by an attempt to resolve the issue by developing a clearer specification of the distinction between abstract and episodic representations.

The major disagreement between supporters and critics of the episodic account concerns the flexibility of episodic models. Mathews (1997) argues that if individuals simply rely on prior instances to perform everyday tasks, generalization will be very limited. While Mathews agrees that learning in artificial grammar learning studies is often particularized and rigid, he argues that learning in the real world is far more flexible. There are at least two

Correspondence should be addressed to A. Neal, Department of Psychology, University of Queensland, Brisbane, QLD 4072, Australia (e-mail: andrew@psy.uq.edu.au).

senses in which knowledge could be "flexible." Knowledge could be flexible either because it allows the learner to respond to unanticipated task demands, or because it allows the learner to transfer his or her past experience to new situations. In contrast to Mathews, we believe that storing knowledge episodically may enhance both types of flexibility. Storing knowledge episodically may enhance an individual's ability to adapt to unanticipated changes in task demands, because it allows him or her to delay abstraction until it is actually needed and preserves the idiosyncratic details associated with individual learning episodes, which may subsequently prove useful for the task at hand. Performing abstraction at encoding or during storage (e.g., by modifying chunk strength) may limit the ability of the learner to respond to unanticipated task demands because it results in the loss of these seemingly irrelevant idiosyncratic details (see Whittlesea & Dorken, 1997).

Storing knowledge episodically may also enhance transfer. Traditionally, it has been thought that transfer required the use of abstract knowledge. However, Brooks and Vokey (1991) found that individuals can discriminate grammatical from nongrammatical changed letter set strings by the use of analogy. For example, a subject could correctly classify the test item BDCCCB as grammatical because the relations between the letters resemble those in the training item MXVVVM. Whittlesea and Dorken (1993a) even found that subjects could perform different letter set transfer when there were no grammatical rules to abstract at all. It could, however, be argued that analogical transfer involves the use of abstract relationships between the items and therefore can be classified as a form of abstraction. In reply, we would argue that analogical transfer involves the use of retrieval-time abstraction, rather than the retrieval of precomputed abstractions from memory. The analogical transfer literature demonstrates that individuals rarely use precomputed abstractions, and instead, that transfer generally involves the retrieval of prior examples, the mapping of the correspondences between the items, and the abstraction of principles or procedures for solving the problem when needed (see Reeves & Weisberg, 1994, for a review). Even though individuals may ultimately use abstract knowledge to perform the task, these abstractions are often computed from episodic representations stored in memory. As described above, storing information episodically may facilitate these types of postcomputational strategies. This conclusion is consistent with findings demonstrating that transfer is facilitated by supplementing abstract principles with multiple examples (Gick & Holyoak, 1983).

In passing, it should also be noted that the use of postcomputational strategies in transfer tasks compounds the problems associated with the use of transfer as evidence for implicit abstraction. The reason for the exclusion of the transfer studies from our earlier paper (Neal & Hesketh, 1997) was that subjects might not use abstract knowledge to perform the transfer task (Brooks & Vokey, 1991;

Whittlesea & Dorken, 1993a). However, even if one demonstrates that subjects are using abstract knowledge to perform the transfer task, one still has to isolate whether this abstraction is occurring at retrieval or earlier, and whether this abstraction is implicit or explicit. The types of postcomputational abstraction studied in the analogical transfer literature appear to be explicit, rather than implicit.

A second disagreement between supporters and critics of the episodic account concerns the use of neuropsychological evidence. Reber (1997) argues that neuropsychological findings represent an important source of evidence regarding implicit learning. In particular, Reber argues that dissociations between performance on implicit and explicit tasks in amnesics questions the viability of episodic models. However, as argued in our previous paper, episodic models can account for these findings. Briefly, if it is assumed that amnesia makes the representations in memory less distinct by reducing the similarity gradient, this may impair recognition, while possibly even facilitating categorization. The precise nature of the relationship between recognition and categorization will depend on the similarity relationships between the stimuli.

A third disagreement concerns the computational efficiency of episodic models. Perruchet et al. (1997) argued that it is not realistic to assume that every episode encountered over the course of one's lifetime is stored as a separate trace (see also Mathews & Roussel, 1993). Unfortunately, these arguments appeal to intuitive theories of mind, rather than to known properties of the human information processing system. At present, it is simply not known whether concepts such as storage capacity should act as constraints on model development (see Whittlesea & Dorken, 1993b).

A final disagreement is related to the scope of the episodic account. Perruchet et al. (1997) argue that the episodic account is limited to microlevel phenomena, because it cannot explain how subjects actually process episodes in the first place. We agree with the argument that the account is limited, because it does not explain how episodes are processed in the first place. The problem is that virtually any set of results can be explained by an episodic account on a post hoc basis, by arguing that the task encouraged processing of one type or another. The theory needs a set of constraints to generate testable predictions regarding task processing. Perruchet et al.'s model provides one approach for developing such a set of constraints. However, we do not agree with the claim that the episodic account is limited to microlevel phenomena. The major advantage of the episodic account is that it provides an explanation of implicit learning phenomena that can accommodate the different mechanisms proposed by different research groups (similarity to whole strings, recognition of fragments, and abstraction of rules) within the context of a broader theory of human memory. As we demonstrated in our previous paper, this approach can

provide a principled account of the relationships between different tasks, such as categorization and recognition, that the other accounts cannot provide.

We believe that the major limitation of the episodic account, at present, is that it draws only on very general principles of human memory, such as the encoding specificity principle. As a result, it has proved very difficult to distinguish between episodic and abstractive accounts empirically (see, e.g., Stadler, 1997). Further progress will depend on more clearly specifying the nature of abstract and episodic representations. In our previous paper, we briefly raised the possibility of distinguishing between abstract and episodic representations in terms of the binding between the item, retrieval cue, and context. This proposal was based on a theory of data structures and memory access processes developed by Humphreys, Wiles, and Dennis (1994). The theory provides an abstract computational-level description of memory processes in terms of the goal of computation and the inputs and outputs required for the task (Marr, 1982). In the theory, inputs to memory tasks consist of sets of items and include words, relations, and contexts. In artificial grammar learning studies, "words" can correspond to individual letters, fragments (bigrams or trigrams), or whole letter strings, depending on the manner in which the individual encodes the information. Outputs of memory tasks are either items or decisions. For example, in a recall task, the output is a word, whereas in recognition, the output is a decision (yes/no). Information is stored in data structures in the form of bindings. Common types of bindings include bindings between items, bindings between items and context, and bindings between items and relations. A simplified description of two common memory tasks (lexical decision and list-specific recognition) is provided below, followed by an application of the technique to the artificial grammar learning task.

Lexical decision tasks require the subject to identify whether a given target item is a word or not. Humphreys et al. (1994) identify the inputs into the task as the target item and the data structure (i.e., all items and associated bindings in memory). The functional specification of the task is as follows. First, the set of all known words is retrieved from the data structure. The intersection between the set of all known words and the target item is then computed. Finally, if the target item matches one or more of the items in the set of all known words, the subject identifies the item as a word.

List-specific recognition tasks require the subject to indicate whether a given target item has occurred on a previously studied list. The inputs into the task are the target item, the context identified by the experimenter's instructions (the study list), and the data structure. There are two possible specifications for list-specific recognition. First, the subject could retrieve the set of all items that are bound to the study context from memory (using context as the retrieval cue), and compute the intersection between this set and the target item. If there is a match between the target item and the items retrieved from memory, the subject identifies the item as old. Al-

ternatively, the subject could retrieve the set of all contexts that are bound to the target item (using the target item as the retrieval cue) and compute the intersection between this set and the study context. If there is a match between the study context and the contexts retrieved from memory, the subject identifies the item as old.

The functional specifications provided by Humphreys et al. (1994) may provide some insights into the nature of abstract and episodic approaches to artificial grammar learning. A functional specification for abstractive approaches to artificial grammar learning may resemble the specification for lexical decision tasks. The inputs into the string discrimination task may simply be the target item and the data structure. The subject could retrieve the set of nonword letter strings from memory and compute the intersection between this set and the target item. If there is a match between the target item and the letter strings retrieved from memory, the subject would identify the item as grammatical. This specification represents an abstractive strategy, because the subject does not make use of bindings between the letter strings and the study context. Indeed, abstractive approaches to artificial grammar learning may predict that contextual bindings are not stored in memory in the first place, or are lost as the individual encodes additional items.

An episodic specification for the artificial grammar learning task may resemble that for list-specific recognition. The inputs into the task may be the test item, the study context, and the data structure. Either the subject could retrieve the set of all items that are bound to the study context and compute the intersection with the target item, or the subject could retrieve the set of all contexts that are bound to the target item and compute the intersection with the "study context." If there is a match between the context(s) and item(s), the subject would identify the letter string as grammatical. This specification represents an episodic strategy, because it involves the use of contextual bindings.

It should also be noted that the abstractive and episodic specifications described above can accommodate the fragmentary knowledge account of artificial grammar learning (e.g., Perruchet & Pacteau, 1990). The abstractive and episodic specifications remain valid, regardless of whether the subject encodes the items as individual letters, fragments, groups of fragments, or whole items. Fragmentary knowledge, therefore, may be used either in an abstractive sense or in an episodic sense, depending on whether the individual makes use of contextual bindings.

Relational bindings may provide another basis for distinguishing between abstractive and episodic strategies. The abstractive and episodic strategies described above predominantly involve the use of associative information. However, subjects also appear to use higher order processes, such as analogical reasoning and hypothesis testing during artificial grammar learning studies. These types of higher order processes require the use of relational information (Halford et al., 1994). Halford et al. have demonstrated that the relational information needed

to perform complex tasks, such as analogical reasoning, can be represented in the Humphreys et al. (1994) model, using bindings of increasing complexity. The simplest level of relational information used by humans is the unary relation, which involves a binding between a single item and an attribute. The most complex level of relational information generally used by humans involves a binding between four items and a relation. Abstractive strategies involving the use of relational information probably represent a form of abstraction that is qualitatively different from abstractive strategies involving the use of associative information.

CONSCIOUSNESS

There appears to be agreement among participants in the symposium that dissociations between performance and self-reported declarative knowledge can no longer be taken as convincing evidence for the existence of implicit learning. The only substantive defense of the use of free report was provided by Reber (1997), who argued that free report should not be totally abandoned, since it captures much of the phenomenological significance of implicit learning. However, Reber did accept that other criteria, such as Dienes and Berry's (1997) subjective threshold criteria, are likely to be more empirically defensible. Although we agree that the dissociation between performance and self-reported declarative knowledge does capture the phenomenological significance of implicit learning, we believe that the variety of alternative explanations for dissociations limits their usefulness as an experimental tool. In particular, as long as measures of performance and declarative knowledge fail the exhaustiveness and exclusiveness criteria (Reingold & Merikle, 1988), task dissociations will never provide a resolution to the question of whether learning can occur outside of awareness.

One response to problems with the task dissociation paradigm has been to question the usefulness of attempting to find "pure" forms of implicit or explicit learning (Mathews, 1997; Reber, 1997). According to this argument, research is producing a biased picture of implicit learning, because the experimental procedures used to isolate the contribution of implicit and explicit processes change the nature of the phenomenon under study. Instead, it is argued that it makes more sense to focus on the interaction between implicit and explicit processes, than on isolating the contribution of both. However, the validity of this argument critically depends on what one actually means by "implicit" and "explicit" processes. If one believes that implicit and explicit processes reflect the operation of separate "systems," one needs to specify the independent properties of the systems, in order to understand the interaction between the two.

A similar problem emerges if one believes that implicit and explicit processes are separated by a threshold, and that processes operating above the threshold are

qualitatively different from processes that operate below the threshold (e.g., Dienes & Berry, 1997). If one adopts a threshold model, it is difficult to avoid the need to specify the nature of the qualitative distinction between performance above and below the cut-point. Inevitably, we tend to be drawn toward the concept of a cut-point that separates the two. Simply arguing that the two processes interact and that it is difficult to isolate the contributions of both does not avoid the problem.

As Reber (1997) acknowledges, the concept of a cut-point that dissociates conscious from unconscious processes is problematic. Perhaps the best illustration of the problem with the concept of a cut-point, and of threshold theories in general, has been provided by Dennett (1993). Dennett points out that if there is no single structure in the brain where consciousness occurs, it follows that there can be no single point in time when consciousness occurs, either. If consciousness occurs at a multitude of different sites, any stimulus must reach these sites at different times, and there can be no single point in time during an ongoing sequence of processing when the individual can be said to have become conscious of that stimulus. Conscious experience, therefore, has to be both spatially and temporally distributed across the brain. If so, there can be no single "dividing line" that separates conscious from unconscious processing.

An alternative response to the problem of distinguishing implicit and explicit processes is to redefine the nature of the distinction between conscious and unconscious cognition. Several commentators have adopted this position (Perruchet et al., 1997; Reber, 1997; Whittlesea & Dorken, 1997). Although the arguments put forth by these commentators all differ in important respects, they all share the belief that searching for dissociations between implicit and explicit processes is meaningless. Reber suggests that implicit and explicit processes should be thought of as a continuum, rather than a dichotomy. The function of consciousness is thought to be control; hence processes that operate toward the implicit end of the continuum, such as conditioned taste aversion, operate automatically, whereas processes that operate at the explicit end of the continuum are open to self-reflection. Concepts such as cut-points and systems are redundant within this view.

Whittlesea and Dorken (1997) adopt a slightly different approach. They argue that implicit knowledge is an unnecessary concept because experimental findings can be explained by reference to the processing operations and retrieval cues used in the tasks, rather than the operation of conscious and unconscious forms of information processing. Standard implicit learning instructions encourage subjects to encode specific learning experiences, which can then be used in a variety of ways, often unanticipated at encoding. For this reason, there is no need to store implicit knowledge about the task. The dissociation between performance and declarative knowledge can be explained simply on the basis that individu-

als make false attributions regarding the basis for their behavior. Even in supposedly "explicit" tasks, individuals do not have direct introspective access to the cognitive processes involved in task performance. If self-report is an attribution in both implicit and explicit learning tasks, there is no need to assume that the conscious status of task performance is qualitatively different in the two situations, and the concept of a threshold distinguishing the two is redundant. Within this approach, interesting research questions can still be framed regarding the dissociation between performance and self-report. However, the focus of this research should be on the types of conditions that facilitate or impair correct attributions regarding task performance, rather than the types of conditions that influence the operation of conscious and unconscious learning systems.

In many respects, Perruchet et al. (1997) make a very similar argument to that of Whittlesea and Dorken (1997). Both argue that implicit knowledge is an unnecessary concept and that individuals do not have introspective access to the basic processes underlying learning. Specifically, Perruchet et al. argue that learning is an associative process, reflecting changes that occur at the neural level. These changes, which can be thought of as analogous to the changes in the connection weights within a connectionist network, are argued to be intrinsically unconscious. However, these associative changes are argued to shape our perceptual awareness and conscious internal representation of the world. Continuing the analogy with connectionist systems, our perceptual awareness and internal representation of the world is thought to reflect the current pattern of activation within the network. While individuals are conscious of these representations, they are assumed to have no introspective access to the underlying processes that generate them. Perruchet et al. argue that the concept of implicit knowledge is meaningless, because the "knowledge" consists of a pattern of neural interconnections that can never be made conscious. Concepts such as a cut-point that dissociates conscious and unconscious processing, or separate implicit and explicit systems, are meaningless within this approach, because individuals have access only to the products of these processes, and not to the processes themselves. For example, there is no threshold which these associative processes can exceed in order to become conscious.

In many respects, the arguments developed by Whittlesea and Dorken (1997) and Perruchet et al. (1997) are consistent with that of Dennett (1993). Like Perruchet et al., Dennett argues that patterns of neural activation within the brain are intrinsically conscious and contribute to our perceptual awareness and internal representation of the world. However, whereas Perruchet et al. claim that there can be no threshold which associative changes can exceed to become conscious, Dennett claims that there can be no threshold which patterns of activation need to exceed in order to become conscious. Whether a certain pattern of activation appears conscious is simply a func-

tion of how long that pattern remains active in the brain. This approach can explain why a subject cannot report the identity of a backwardly masked stimulus, even though the stimulus can influence performance on subsequent forced choice tasks. If a stimulus is displayed for a brief period of time, and backwardly masked, the mask will disrupt the representation of the stimulus. If the representation is disrupted, it will exert only a transient influence on the system, there will be no record of its existence left, and the stimulus will be unavailable for verbal report. The concept of "perception without awareness" is not needed in order to explain the findings.

In summary, there is now a substantial body of opinion arguing that the metaphors traditionally used to characterize conscious and unconscious processes should be discarded. Taken to its logical extreme, the current argument suggests that since (1) individuals never have introspective access to the basic processes involved in learning, and (2) self-report consists of attributions based on our (conscious) internal representations of the world, all learning has to be implicit. According to this logic, there can be no explicit learning. This position is interesting, because it represents the reverse of the standard criticism of implicit learning research—namely, that so-called implicit learning may actually be explicit. While Reber (1997) may be right that much of the controversy over implicit learning in the past has been generated by the claim that complex cognitive operations may occur outside of awareness, it appears that much of the current reappraisal of the implicit learning literature is motivated by difficulties associated with claims that any form of cognitive operations could occur in consciousness at all.

An alternative response to the problems with the task dissociation paradigm is to operationalize the distinction between conscious and unconscious processes in terms of intention (Neal & Hesketh, 1997; Stadler, 1997). As Stadler points out, refocusing attention onto the intention/nonintention distinction does not simply represent a linguistic sleight of hand (Reber, 1997), because it involves shifting from an essentialist definition of consciousness to an operational definition. If the intention/nonintention distinction is to prove useful within an operational approach, research must demonstrate the existence of qualitative differences between intentional and nonintentional learning, or between intentional and nonintentional performance. At present, it is unclear whether we have research paradigms sufficiently powerful to demonstrate these differences, or even whether the concept of such qualitative differences makes sense in the first place. There are currently two approaches for investigating the existence of qualitative differences between intentional and nonintentional processes—namely, the process dissociation procedure (Neal, Hesketh, & Andrews, 1995), and the retrieval intentionality criterion (Stadler, 1997). As pointed out in our previous paper (Neal & Hesketh, 1997), the process dissociation procedure cannot provide strong evidence for the existence of qualitative distinctions be-

tween intentional and nonintentional processes, because the effects of experimental manipulations can be explained with reference to encoding and retrieval operations, rather than concepts such as intention. For example, the existence of false alarms to exclusion items may reflect difficulties in list discrimination (see also Perruchet et al., 1997). The retrieval intentionality criterion may suffer from the same problems. Stadler reports findings demonstrating that an indirect task (sequence learning) was more sensitive than a direct task (recognition). The two tasks were identical, except for the instructions to predict the location of the targets (indirect task) or to identify whether the sequence had been previously seen (direct task). Using Reingold and Merikle's (1988) logic, Stadler argued that these findings qualified as a demonstration of implicit learning. However, these findings may also be explained with reference to encoding and retrieval operations. Using the Humphreys et al. (1994) model, the direct and indirect tasks can be seen to differ in the use of contextual bindings and the response required. The recognition task requires contextual bindings and a yes/no decision. The prediction task does not require contextual bindings, and it requires an item for an output. It is likely that recognition performance will be particularly poor in the sequence learning paradigm, because there is so little contextual information supplied during practice. The greater sensitivity of the prediction task, therefore, may not be attributable to differences in intentional control.

CONCLUSION

In summary, there appears to be substantial agreement that dissociations between conscious and unconscious forms of learning or knowledge should no longer form the central focus of implicit learning research, and that the very concept of such a dissociation may lack meaning. Instead, most commentators appear to agree that we need to focus on the underlying processes involved in implicit learning studies. The approach that we favor is to embed this research within the context of more general theories of memory and reasoning. Regardless of the specific theoretical model adopted, we believe that a major advantage of focusing on the underlying processes is that consciousness will no longer need to be used as an explanatory construct. Although this position reflects a form of "hardware independent functionalism," it should not be misinterpreted. Like Perruchet et al. (1997), we are not suggesting that consciousness should be ignored. We believe that consciousness is a phenomenon to be explained by cognitive theory, rather than an explanatory construct in its own right, and we strongly support Reber's (1997) claim that consciousness still represents the constitutive problem in psychology.

REFERENCES

- BROOKS, L. R., & VOKEY, J. R. (1991). Abstract analogies and abstracted grammars: Comments on Reber (1989) and Mathews et al. (1989). *Journal of Experimental Psychology: General*, **120**, 316-323.
- DENNETT, D. C. (1993). *Consciousness explained*. Harmondsworth, U.K.: Penguin.
- DIENES, Z., & BERRY, D. C. (1997). Implicit learning: Below the subjective threshold. *Psychonomic Bulletin & Review*, **4**, 3-23.
- GICK, M. L., & HOLYOAK, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, **15**, 1-38.
- HALFORD, G. S., WILSON, W. H., GUO, J., GAYLER, R. W., WILES, J., & STEWART, J. E. M. (1994). Connectionist implications for processing capacity limitations in analogies. In K. J. Holyoak & J. Barnden (Eds.), *Advances in connectionist and neural computational theory: Vol. 2. Analogical connections* (pp. 363-415). Norwood, NJ: Ablex.
- HUMPHREYS, M. S., WILES, J., & DENNIS, S. (1994). Toward a theory of human memory: Data structures and access processes. *Behavioral & Brain Sciences*, **17**, 655-692.
- MARR, D. (1982). *Vision: A computational investigation into human representation and processing of visual information*. San Francisco: W. H. Freeman.
- MATHEWS, R. C. (1997). Is research painting a biased picture of implicit learning? The dangers of methodological purity in scientific debate. *Psychonomic Bulletin & Review*, **4**, 38-42.
- MATHEWS, R. C., & ROUSSEL, L. G. (1993). Automatic abstraction of stimulus structure from episodes: Comment on Whittlesea and Dorken (1993). *Journal of Experimental Psychology: General*, **122**, 397-400.
- NEAL, A., & HESKETH, B. (1997). Episodic knowledge and implicit learning. *Psychonomic Bulletin & Review*, **4**, 24-37.
- NEAL, A., HESKETH, B., & ANDREWS, S. (1995). Instance-based categorization: Automatic versus intentional forms of retrieval. *Memory & Cognition*, **23**, 227-242.
- 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., & GALLEGO, J. (1997). Implicit learning shapes new conscious percepts and representations. *Psychonomic Bulletin & Review*, **4**, 43-48.
- REBER, A. S. (1997). Implicit ruminations. *Psychonomic Bulletin & Review*, **4**, 49-55.
- REEVES, L. M., & WEISBERG, R. W. (1994). The role of content and abstract information in analogical transfer. *Psychological Bulletin*, **115**, 381-400.
- REINGOLD, E. V., & MERIKLE, P. M. (1988). Using direct and indirect measures to study perception without awareness. *Perception & Psychophysics*, **44**, 563-575.
- STADLER, M. A. (1997). Distinguishing implicit and explicit learning. *Psychonomic Bulletin & Review*, **4**, 56-62.
- WHITTLESEA, B. W. A., & DORKEN, M. D. (1993a). Incidentally, things in general are particularly determined: An episodic-processing account of implicit learning. *Journal of Experimental Psychology: General*, **122**, 227-248.
- WHITTLESEA, B. W. A., & DORKEN, M. D. (1993b). Potential power of coding particular experiences: Reply to Mathews and Roussel (1993). *Journal of Experimental Psychology: General*, **122**, 401-404.
- WHITTLESEA, B. W. A., & DORKEN, M. D. (1997). Implicit learning: Indirect, not unconscious. *Psychonomic Bulletin & Review*, **4**, 63-67.