ELENI ZIORI

IMPLICIT KNOWLEDGE REPRESENTATION

ΠΕΡΙΛΗΨΗ

Παρά τον πλούτο ερευνών της διάκρισης άδηλου-έκδηλου, οι ερευνητές δεν έχουν συμφωνήσει μέχρι σήμερα στο είδος των πληροφοριών που κωδικοποιούνται στα έργα άδηλης μάθησης. Συγκεκριμένα, υπάρχει μια διαμάχη ανάμεσα σε θεωρίες που προτείνουν τη γνώση που προκύπτει από μια διαδικασία αφαίρεσης και σε θεωρίες που προτείνουν τη συγκεκριμένη γνώση (γνωστές ως θεωρίες των παραδειγμάτων). Το παρόν άρθρο παρουσιάζει τις διαφορετικές θεωρίες αναπαράστασης της άδηλης γνώσης που έχουν προταθεί μέχρι σήμερα και υποστηρίζει ότι δεν είναι απαραίτητο να ψάχνουμε για αποδείξεις υπέρ μιας συγκεκριμένης θεωρίας. Αντίθετα, είναι πιο λογικό να συμπεράνει κανείς ότι το ποιο είδος αναπαράστασης θα επικρατήσει καθορίζεται από τις απαιτήσεις του συγκεκριμένου έργου μάθησης.

Λέξεις κλειδιά: αναπαράσταση γνώσης, άδηλη γνώση, έκδηλη γνώση, αφηρημένη γνώση, συγκεκριμένη γνώση.

The representation of implicitly acquired knowledge.

Many theories of learning rely on the existence of two different modes of knowledge acquisition and representation. A dichotomy that has, recently, received an increasing interest cognitive psychologists is the one between explicit and implicit learning. According to a traditional definition of implicit learning offered by Reber (1993, p.5), "implicit learning is the acquisition of knowledge that takes place largely independently of conscious attempts to learn and largely in the absence of explicit knowledge about what was acquired". Implicit learning, by contrast to explicit learning, is generally thought to occur when: 1) the acquired knowledge cannot be reported verbally, and 2) participants learn without resource to conscious codebreaking strategies, such as hypotheses testing. Thus, implicit learning, by contrast to explicit learning, occurs when people acquire knowledge they are not fully conscious of.

A fundamental issue that has engaged researchers' attention in their attempt to throw light on implicit knowledge is how this knowledge is represented. Some have described implicit learning as 'abstract' in that they have equated this type of learning with a process that is associated with an abstraction of rules (e.g. Reber, 1967, 1989; Reber & Lewis, 1977). By contrast, others view implicit learning as the encoding and storage of whole exemplars (e.g. Brooks, 1978; Brooks & Vokey, 1991; Vokey & Brooks, 1992) or as the storage of fragments of exemplars (e.g. Dulany, Carlson, & Dewey, 1984; Perruchet & Pacteau, 1990; Servan-Schreiber & Anderson, 1990).

The implicit learning task that has been most extensively used in the debate about implicit knowledge representation is the *artificial grammar learning* (AGL) task introduced by Reber in the 1960s (see Reber, 1989; 1993 for an overview). In a typical AGL experiment, participants first study a list of letter strings generated by a finite state grammar and are asked to memorize them. After the learning phase, they are informed that the strings followed a complex set of rules, but no specific information is provided to them regarding the nature of the grammar. Then, they are presented with new strings of letters, some of which are grammatical and some are not, and are asked to classify them. In an early study, Reber (1967) found that patricipants could classify new strings 60-70% correctly on average, althoug, they were unable to describe the underlying rules. These findings were taken as evidence of implicitly acquired knowledge. Since Reber's early work, researchers have used a variety of experimental paradigms in order to investigate implicit learning, including artificial grammars (e.g. Dienes, Broadbent, & Berry, 1991; Mathews et al., 1989), concept learning tasks (e.g. Frick & Lee, 1995; Roberts & MacLeod, 1995), the control of complex systems (e.g. Berry & Broadbent, 1984, 1988; Stanley, Mathews, Buss, & Kotler-Cope, 1989), and sequence learning (e.g. Nissen & Bullemer, 1987; Lewicki, Czyzewska, & Hoffman, 1987). However, the presentation of the different views on the representation of implicit knowledge in the present article will focus mainly on the AGL task, since this is the task that has generated the most research on implicit learning.

The Abstractive view.

According to this view, the knowledge that is acquired in implicit learning tasks is represented in an abstract way. The abstractive view is based on AGL findings that patricipants demonstrate better than chance performance when asked to make grammaticality judgements on novel letter strings (Reber, 1967, 1989, 1993). The term *abstract* is rather vague and might lead to confusion if it is not specifically defined. According to Reber (1993, p. 120-121), for example, "an abstract representation is assumed to be derived, yet separate from the original instantiation. Abstract codes contain little, if any, information pertaining to the specific stimulus features from which they were derived; the emphasis is on structural relationships among stimuli". Moreover, Reber claims that what participants encode when presented with stimuli is the relational features of these stimuli and not their superficial physical forms. In that sense, when people are asked to classify novel stimuli, they are assumed to compare the abstract codes of these stimuli with the previously acquired 'deep' knowledge about the rules. Another argument that Reber (1989, 1993) has proposed in order to show that implicit knowledge is represented in an abstract way is that the abstractive view can account for the transfer of knowledge across stimulus domains. This argument was further supported by findings that knowledge acquired in AGL tasks transfers to strings generated by the same grammar but instantiated with a different letter set.

An early study that investigated transfer to stimuli with different perceptual surface features was the one carried out by Reber (1969). In this study, participants had to memorize letter strings that were generated by a finite-state grammar. In a following session, one group of participants had to memorize novel strings that were constructed using the same letters but a different grammar, whereas another group of participants memorized novel strings that were constructed from a different letter set, but from the same grammar. Reber found that changing the grammar led to decrements in performance, whereas changing the perceptual characteristics (i.e., the vocabulary) had little or no effects at all.

These results were replicated by an experiment conducted by Mathews et al. (1989). In this experiment, which was run over a 4-week period, half of the patricipants were presented each week with test strings that were constructed from a different letter set, but from the same grammar, whereas the other half saw strings that consisted of the same letters as in the study phase. Subsequently, participants were given a forced-choice test, in which they had to discriminate between grammatical and non-grammatical strings. It was found that participants who had received a novel letter set performed as well as participants that were presented with the same letter set. This effect, however, was not found in a rule-instructed group. This suggests that participants learned the grammar in a way that allowed them to perform significantly above chance even when the surface features were changed.

Another study that provides compelling evidence in favour of the abstractive knowledge representation was the one carried out by Reber and Lewis (1977). They asked participants to initially memorize a set of grammatical strings and then perform an anagram task. Thus, Reber and Lewis assessed participants' knowledge of the grammar by their ability to solve anagram problems. They correlated the frequencies of different bigrams that appeared in participants' anagram solution with the frequencies of bigrams a) in the grammar as a whole and b) in the study strings. In the PTTTVV string, for example, the bigrams PT, TV and VV appear only once, whereas the bigram TT appears twice. They found that the frequency with which participants used the bigrams correlated more with how frequently they appeared in the whole set of strings that the grammar could generate rather than with how frequently they appeared in the learning strings. Thus, Reber and Lewis concluded that participants made their judgements relying on the deep structure of the grammar rather than on the specific study strings (but see Perruchet, Gallego, & Pacteau, 1992, for evidence that these results might have been attributed to several artefacts).

Another form of abstract representation is one that is instantiated and based on *prototypes*. According to this version of the abstractionist view, learning leads to a representation of the central tendencies of the features of previously encountered examples, and classification of novel stimuli is achieved by measuring the similarity between the novel stimuli and the prototypes (i.e., central tendencies). The two forms of abstract representation differ in the way they use the term abstract. In particular, the prototype representation involves encoding the central tendencies of specific stimuli properties and not feature covariations of stimulus types. However, a prototypes is considered abstract in that it involves the representation of only the central tendencies of features across a set of exemplars and not the representation of individual exemplars.

Overall, the abstractive view has not been very conclusive. Even Reber (1993) himself argues that it is not very clear how people encode the various components of the novel stimuli and compare them with their deep representations. So, do the above mentioned results unequivocally support the argument that implicit learning is equated with the abstraqction of unconscious rules? According to the exemplar view, the answer to this question is negative.

The exemplar view.

The basis of the exemplar view is that stimuli are encoded and stored in memory as separate instances (e.g. Brooks, 1978; Medin & Schaffer, 1978; Neal & Hesketh, 1997; Vokey & Brooks, 1992). In this approach, implicit learning involves the storage of whole exemplars and not patterns of covariation among features. Consequently, participapants can use analogy to stored exemplars to classify new exemplars at above chance levels without making any inductions or changing the instatiated memories in any way. This is a major difference between the exemplar models and the abstractive view in which recoding the stimuli and revising the existing representations are considered essential.

Thus, in the above-mentioned studies on AGL, transfer performance could be based on the similarity between the transfer string and a specific training string. In other words, the exemplar view argues that patricipants in AGL tasks are classifying novel strings relying not on abstract knowledge of rules, but on the similarity of grammatical or ungrammatical strings to 'whole exemplars' memorized during training. For example, the strings TXXVPT and BLLMKB can be seen as similar because both strings begin and end with the same letter and have identical letters at the same positions. Brooks and Vokey (1991; see also Gomez, Gerken, & Schvaneveldt, 2000; Tunney & Altmann, 1999) found that such abstract similarities (or *abstract analogies* as they named them) had a significant effect on transfer performance even when different letter sets were used. This finding casts serious doubt on Reber's (1989) argument that knowledge is abstract in that it can show positive transfer between different stimulus domains.

In order to show that transfer is based on the similarity between a test instance and a study instance, Vokey and Brooks (1992) manipulated the similarity between transfer strings and study strings by constructing grammatical or nongrammatical test strings that were either similar or dissimilar to whole individual study strings. They found that both grammaticality and similarity had a large effect on participants' transfer performance. Thus, grammatical strings were classified better than nongrammatical strings, and similar strings were judged as grammatical more often than dissimilar strings. However, these findings may not be interpreted solely as evidence that people memorized whole letter strings. Alternatively, it has been argued that the similarity between test strings and study strings may be confounded with the frequency with which partial strings (e.g. bigrams) in the test items occur in the study items. Thus, when grammatical and ungrammatical test strings were equated for fragment knowledge, no exemplar effects were found (Knowlton & Squire, 1994; Shanks, Johnstone, & Staggs, 1997).

The fragmentary view.

This view is similar to the exemplar view in that they both rely on instantiated representations of stimuli. The difference between the two positions is that, according to the fragmentary view, only fragments of stimuli are stored and not whole stimuli. Thus, when participants are called to classify novel stimuli, they are assumed to compare the fragments that are already represented in their memory with the chunks or fragments that the test stimuli contain and then base their judgements on the similarity or match of these codings. Several versions of this view have been proposed so far. It has been argued, for instance, that in a typical AGL task, patricipants store units of information consisting of 'chunks' of the training stimuli, such as bigrams and trigrams or simple frequency counts (Dulany et al., 1984; Perruchet & Pacteau, 1990). Similarly, Servan-Schreiber and Anderson (1990) proposed the so-called competitive chunking model[1], according to which

---

1. Apart from the competitive chunking model, several other compuratational models of implicit learning have been proposed so far, including exemplar models (Dienes, 1992), classifier systems (Druhan & Mathews, 1989), and connectionist models (Cleeremans & McClelland, 1991; Dienes, 1992). However, the present article will only briefly refer to some connectionist models (see following section), which have been characterized as the most successful candidate models of implicit learning (for an extensive description of the various models, see Cleeremans, 1993; Dienes, 1992, 1993).

learning consists of forming and applying increasingly higher-order chunks. Each time an item is encoded as a single chunk, its memory representation is strengthened, and is perceived as maximally familiar. For instance, the string PTVPXVPS may first be chunked as (PT), (VPX), and (VPS). After repeated exposure, the more complex string ([PT] [VPX]) may be created, and finally the whole store may be formed, wihch will seem more familiar when later encountered. The chunks that Servan-Schreiber and Anderson (1990) proposed involve higher-order representations than those that both Dulany et al. (1984) and Perruchet and Pacteau (1990) suggested.

Perruchet and Pacteau (1990), for instance, presented a group of participants with a set of letter strings and another group only with bigrams that were contained in those strings. They found that both groups could distinguish between grammatical and nongrammatical strings significantly above chance, even though the group that saw the whole strings performed better than the group that saw only the bigrams. However, when Perruchet and Pacteau eliminated all strings that were nongrammatical simply because they contained some non-permissible bigrams, which were easier for the group that saw the whole strings to classify, performance of the two groups was identical. In another experiment, participants were asked to remember isolated bigrams. It was found that participants' memory of isolated bigrams could account for their accuracy in classifying grammatical and nongrammatical items (although not completely — not for permissible bigrams in non — permissible locations). Perruchet and Pacteau concluded that people acquire knowledge of permissible bigrams independent of knowledge of whole exemplars or knowledge of positional dependencies of bigrams, and that it is the former knowledge that determines participants' classification performance.

The argument, however, that people do not acquire any knowledge about the positional dependencies of bigrams has been challenged by a number of researchers (e.g. Dienes, Broadbent, & Berry, 1991; Dulany et al., 1984; Johnstone & Shanks, 1999). Dienes et al. (1991), for example, found that when participants were asked to classify permissible bigrams as grammatical or not, they performed significantly above chance only if the bigrams appeared at a grammatically legal position. Thus, participants could acquire some knowledge about the positions of bigrams within whole strings.

Moreover, there have been studies that provided evidence that participants use both rule and fragment knowledge when they classify test stimuli (e.g. Knowlton & Squire, 1994, Experiment 2b; Meulemans & Van der Linden, 1997). For instance, Meulemans & Van der Linden (1997) showed that after training on few letter strings, participants classified the test strings relying on fragment knowledge (Experiment 2a), whereas after prolonged training, they used knowledge about the rules of the grammar (Experiment 2b). However, Johnstone and Shanks (1999) showed that in Experipment 2b information about rules and fragment locations was confounded. Thus, using a biconditional grammar, which allowed them to unconfound rule and fragment knowledge, they provided evidence that classification judgements can be explained in terms of fragment knowledge (that includes positional information about fragments).

To summarize, as it has become evident from the discussion so far, whether stored exemplars or chunks rather than a more abstract knowledge base best account for knowledge representation in implicit learning tasks still remains an open question. It has been suggested, however, that we should view implicit learning as lying "somewhere on a continuum of exemplar-based to abstract" (Dienes, 1993, p. 167). There may be cases, in which both abstract and specific types of knowledge may be acquired as well as cases, where it is the requirements of the task that determines which type dominates. It has been argued (Cleeremans, 1994) that such a representational continuum from the specific to the abstract is best captured by the way connectionist models represent knowledge.

The connectionist view.

Connectionist systems represent inputs in a parallel-distributed fashion by a large number of elementary computational units or nodes and patterns of interconnectedness between nodes. A number of such connectionist models have, recently, been applied in implicit learning tasks with considerable success (e.g. Dienes, 1992; Cleeremans 1994; Cleeremans & McClelland, 1991). For instance, Dienes (1992) used the AGL paradigm to test 9 versions of two exemplar models that had previously been used in concept learning tasks, namely the model of Medin and Schaffer (1978) and the model of Hintzman (1986), as well as 16 different types of connectionist models. He found that only a number of the connectionist models were adequate models of AGL (see Dienes, 1992, 1993 for a detailed description of all the models).

Connectionist networks are especially suitable to model implicit learning, which is a process that relies on similarities and patterns in the input. More specifically, what makes connectionist models rather promising models of implicit learning is the fact that they propose certain mechanisms that are compatible with some generally accepted characteristics of implicit learning. In particular, implicit learning is generally thought to proceed through the detection of complex covariations in the environment in an unselective way. Moreover, implicit learning is independent of participants' intentions or conscious control, and results in knowledge that is difficult to verbalize. Similarly, connectionist models learn in a self-organizing way determined by changes in connection weights, and are sensitive to statistical covariations of stimuli. Most importantly, connectionist networks represent knowledge without resorting to specific rules, but in a distributed fashion that results from patterns of interconnectedness among nodes, and explains their rule-governed behaviour. Finally, these distributed representations allow successful performance, but are also, as Cleeremens (1993) argues, "holistic" in that they cannot be readily analyzed. This characteristic of connectionist representations, by analogy, fits with the finding that in many implicit learning tasks, participants acquire knowledge that is not available to consciousness.

As McClelland and Rumelhart (1985) argue, connectionist networks may represent exemplar-based information or they can develop representations based on the shared properties of instances depending on the number of hidden units or on the structure of the training items. Similarly, Cleeremens (1994) argued that the sequence recursive network (SRN) of Elman (1990), used to model AGL and sequence learning, may develop internal representations that are like the abstract representation of the grammar. Thus, each state of a finite-state grammar is encoded by a pattern of activation over the hidden units of the network, resulting in an abstract representation that has captured relevand dimensions of the grammar. However, this representation is not abstract in that it cannot be transferred to a shifted letter set (but see Dienes, Altman, & Gao, 1999, for a description of a SRN that can transfer implicit knowledge across domains).

Moreover, Cleeremans (1994) used a simple feed-forward network to model human performance on a sequence prediction task. In brief, participants in this task have to predict the location of a target event, which is determined by a complex, biconditional rule (i.e., a rule that determines the mapping between corresponding events in the first and the second half of a sequence). Cleeremans found that his network, which simulated data quite closely, developed representations that were intermediate between instantiated and abstract. In particular, the network behaved as if its predictions were determined by similarity to stored exemplars rather than by the rules specified by the experimenters. However, the network also behaved in a way that yielded abstract representations in that it extracted the relevant properties of the training exemplars and relied only on them to make similarity judgements.

Thus, it seems that connectionist networks operate rather on a continuum that extends from general or abstract knowledge representations to specific or exemplar based ones. A more radical approach, proposed by Whittleasea (1997; Whittleases & Dorken, 1993), suggests that implicit learning cannot be adequately explained through the acquisition of a preferred type of knowledge. As will be shown below, this approach focuses on processing experience rather than on stimulus structures.

The episodic-processing account.

According to the episodic-processing account, the type of structural information that will be encoded (i.e., rules, exemplars or chunks) depends on the proscessing demands of the task (Whittleasea, 1997; Whittleasea & Dorken, 1993; see also Johnstone & Shanks, 2001). As Whittleasea and Dorken (1993, p. 229) argue, this approach "emphasizes the processing conducted within particular experiences as the primary explanatory mechanism for memory, and that the processing conducted depends on the particular demands and affordances of the encoding episode". More specifically, the episodic processing account assumes that a) participants encode processing information in addition to structural information, b) which structural aspects of test stimuli will be encoded as well as subsequent test performance is determined by the kind of processing that took part during training (transfer-appropriate processing; see Morris, Bransford, & Franks, 1977), and c) different tasks carrierd out on the same stimuli will result in different kinds of processing (encoding variability).

In a series of experiments, Whittleasea and Dorken (1993) presented participants with strings, such as ENRIGOB, that were generated from a finite-state grammar, and asked them to memorize these strings by either pronouncing or spelling them. In the test phase, participants were asked to classify the items about to be displayed by pronouncing half of them and spelling the remaining half. It was found that for items that were spelled during training and pronounced at test, or for items that were pronounced during training and spelled at test, classification performance was at chance levels. In contrast, when there was an overlap of study and test processing, classification was reliably above chance. Thus, it was shown that, during training patricipants encoded processing information in addition to the structural aspects of stimuli, and that performance in the test phase was successful only when the task demands of the prior processing experience were reinstated.

Moreover, Wright and Whittleasea (1998) presented participants with four digit-strings that were constructed so that they followed an odd-even-odd-even rule. One group of participants were asked to say each digit of each string aloud and judge whether it was a high number (i.e., greater than four) or a low number (i.e., lower than five), whereas another group had to pronounce the two digit pairs of each string. All test strings were novel: half of the test strings were constructed by reversing the order of the two digit pairs of strings that were presented in the training phase (and thus were more familiar), whereas the remaining half were completely novel strings. Next, participants were asked to discriminate between *old* items (previously seen) and *new* ones. It was found that patricipants who processed the training items as digit pairs were more likely to judge test strings consisting of familiar digit pairs as old and strings consisting of unfamiliar digit pairs as new than were participants who had processed the strings on a digit-by-digit basis. Thus, it was shown that the original processing experience that depended on the demands of the different tasks determined what level of stimulus structure was encoded.

Similar results were obtained by Johnstone & Shanks (2001), who used a biconditional grammar that allowed rule and fragment knowledge to be unconfounded. In the 'match' condition, participants were simply asked to memorize letter strings, whereas in the 'edit' condition participants were encouraged to test huypotheses in order to discover the rules of the grammar. The results showed that the processing demands of the two different tasks (edit and match) determined the structural form of the acquired knowledge, with the match group classifying on the basis of fragment knowledge and the edit group using only rule knowledge.

CONCLUSION

Researchers have suggested various types of information (i.e., rules, exemplars, or chunks) encoded in implicit learning tasks. Moreover, Whittleasea's (e.g. 1997) approach suggests that the types of structural information acquired depend on the processing demands of the tasks. Although the issue of implicit knowledge representation is far from resolved, it seems reasonable to conclude that both instantiated and abstract representation are possible candidate representation of implicit knowledge, and that, in any cases, either of them may dominate depending on the requirements of the learning task. Thus, what seems more important than deciding which type of knowledge is encoded in implicit learning tasks is to look for any specific exprerimental conditions or individual differences that provide evidence that implicit and explicit knowledge are qualitative different.

REFERENCES

Berry, D.C., & Broadbent, D.E. (1984). On the relatilonship between task performance and associated verbalisable knowledge. *Quarterly Journal of Experimental Psychology*, *36*, 209-231.

Berry, D.C., & Broadbent, D.E. (1988). Interactive tasks and the implicit-explicit distinction. *British Journal of Psychology*, *79*, 251-272.

Brooks, L. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization*, 169-211. Hillsdale, NJ: Erlbaum.

Brooks, L., & 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.

Cleeremans, A. (1993). *Mechanisms of implicit learning.* Cambridge, MA: MIT Press.

Cleeremans, A. (1994). The representation of structure in explicit prediction tasks. In C. Umiltà & M. Moscovitch (Eds.), *Attention and Performance XV: Conscious and non-conscious information processing*, Cambridge, MA: MIT Press, 783-809.

Cleeremans, A., & McClelland, J.L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General, 120*, 235-253.

Dienes, Z. (1992). Connectionist and memory-array models of artificial grammar learning. *Cognitive Science, 16*, 41-79.

Dienes, Z. (1993). Implicit concept formation. In Berry, D.C. and Dienes, Z. (Eds.), *Implicit learning: Theoretical and empirical issues*, 81-112. London: Lawrence Erlbaum Associates.

Dienes, Z., Altman, G.T.M., & Gao, S.J. (1999). Mapping across domains without feedback: A neural network model of transfer of implicit knowledge. *Cognitive Science, 23 (1)* 53-82.

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

Elman, J.L. (1990). Finding structure in time. *Cognitive Science, 14*, 179-211.

Druhan, B. B., & Mathews, R.C. (1989). THYIOS: A classifier system model of implicit knowledge of artificial grammars. *In Proceedings of the Eleventh Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Laurence Erlbaum Associates.

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

Frick, R.W., & Lee, Y.S, (1995). Implicit learning and concept learning. *The Experimental Psychology Society*, 762-782.

Gomez, R.L., Gerken, L., & Schvaneveldt, R.W. (2000). The basis of transfer in artificial grammar learning. *Memory & Cognition, 28*, 253-263.

Hintzman, D.L. (1986). Schema abstraction in a multiple-trace memory model. *Psychological Review, 93*, 411-428.

Johnstone, T., & Shanks, D.R. (1999). Two mechanisms in implicit artificial grammar learning? Comment on Meulemans and Van der Linden (1997). *Journal of Experimental Psychology: Learning, Memory and Cognition, 25*, 524-531.

Johnstone, T., & Shanks, D.R. (2001). Abstractionist and processing accounts of implicit learning. *Cognitive Psychology, 42*, 61-112.

Kemler Nelson, D.G. (1984). The effect of intention on what concepts are acquired. *Journal of Verbal Learning and Verbal Behavior, 23*, 734-759.

Knowlton, B.J., & Squire, L.R. (1994). The information acquired during artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory and Cognition, 20*, 79-91.

Lewicki, P., Czyzewska, M., & Hoffman, H. (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition, 13*, 523-530.

Mathews, R.C., Buss, R.R., Stanley, W.B., Blanchard-Fields, F., Cho, J.R., & Druhan, B. (1989). The role of implicit and explicit processes on learning from examples: A synergistic effect. *Journal of Experimental Psychology: Learning, Memory and Cognition, 15*, 1083-1100.

McClelland, J.L., & Rumelhart, D.E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General, 114*, 159-188.

Medin, D.L., & Schaffer, M.M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207-238.

Meulemans, T., & Van der Linden, M. (1997). Associative chunk strength in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory and Cognition, 23*, 1007-1028.

Morris, C.C., Bransford, J.D., & Franks, J.J. (1977). Levels of processing versus transfer appropriate processing. *Journal of Verbal Learning and Verbal Behaviour, 16*, 519-533.

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

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

Perruchet, P., Gallego, J., & Pacteau, C. (1992). A reinterpretation of some earlier evidence for abstractiveness of implicitly acquired knowledge. *Quarterly Journal of Experimental Psychology, 44A*, 193-210.

Perruchet, P., & Pacteau, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General, 119*, 264-275.

Reber, A.S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, *6*, 317-327.

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

Reber, A.S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General, 118*, 219-235.

Reber, A.S. (1993). *Implicit learing and tacit knowledge: An essay on the cognitive unconscious.* New York: Oxford University Press.

Reber, A.S., & Lewis, S. (1977). Toward a theory of implicit learning: The analysis of the form and structure of a body of tacit knowledge. *Cognition, 5*, 333-361.

Roberts, P.L., & MacLeod, C. (1995). Representational consequences of two modes of learning. *Quarterly Journal of Experimental Psychology, 48A*, 296-319.

Servan-Schreiber, E., & Anderson, J.R. (1990). Learning artificial grammars with competitive chunking. *Journal of Experimental Psychology: Learning, Memory and Cognition, 16*, 592-608.

Shanks, D.R., Johnstone, T., & Staggs, L. (1997). Abstraction processes in artificial grammar learning. *Quarterly Journal of Experimental Psychology, 50A*, 216-252.

Smith J.D., & Shapiro J.H. (1989). The occurrence of holistic categorization. *Journal of Memory and Language, 28*, 386-399.

Stanley, W.B., Mathews, R.C., Buss, R.R., & Kotler-Cope, S. (1989). Insight without awareness: On the interaction of verbalisation, instruction, and practice in a simulated process control task. *Quarterly Journal of Experimental Psychology, 41*, 553-577.

Tunney, R.J., & Altman, G.T.M. (1999). The transfer effect in artificial grammar learning: Reappraising the evidence on the transfer of sequential dependencies. *Journal of Experimental Psychology: Learning, Memory and Cognition, 25*, 1322-1333.

Vokey, J.R., & Brooks, L.R. (1992). Salience of item knowledge in learning artificial grammars. *Journal of Experimental Psychology: Learning, Memory and Cognition, 18*, 328-344.

Whittlesea, B.W.A. (1994). The representation of general and particular knowledge. In K. Lamberts, & D. Shanks (Eds.), *Knowledge, Concepts, and Categories*, (pp. 335-370). Cambridge, MA: MIT Press.

Whittlesea, B.W.A., & Dorken, M.D. (1993). Incidentally, things in general are particularly determined: An episodic-processing account of implicit learning. *Journal of Experimental Psychology: General, 122*, 227-248.

Wright, R., & Whittlesea, B.W.A. (1998). Implicit learning of complex structures: Active adaptation and selective processing in acquisition and application. *Memory & Cognition, 26*, 402-420.