Encyclopedia of Cognitive Science —IMPLICIT LEARNING (Article 97)

Implicit learning

Axel Cleeremans

Université Libre de Bruxelles, Brussels, Belgium

Headword: Implicit learning

Keywords: Implicit learning, consciousness, computational modeling, process dissociation procedure, representation

Contents list:

Article definition:

Implicit learning is the process through which we become sensitive to certain regularities in the environment (1) in the absence of intention to learn about those regularities (2) in the absence of awareness that one is learning, and (3) in such a way that the resulting knowledge is difficult to express.

Implicit learning — broadly construed, learning without awareness — is a complex, multifaceted phenomenon that defies easy definition. Frensch (1998) lists as many as eleven definitions in a recent overview — a diversity that is undoubtedly symptomatic of the conceptual and methodological challenges that continue to pervade the field 35 years after the term first appeared in the literature (Reber, 1967). According to Berry and

Dienes (1993), learning is implicit when we acquire new information without intending to do so, and in such a way that the resulting knowledge is difficult to express. In this, implicit learning thus contrasts strongly with explicit learning (e.g., as when learning how to solve a problem or learning a concept), which is typically hypothesis-driven and fully conscious.

Over the last decade or so, the field of implicit learning has come to embody ongoing questioning about three fundamental issues in the cognitive sciences, namely (1) consciousness (how we should conceptualize and measure the relationships between conscious and unconscious cognition); (2) mental representation (in particular the complex issue of abstraction); and (3) modularity and the architecture of the cognitive system (whether one should think of implicit and explicit learning as being subtended by separable systems of the brain or not). Computational modeling plays a central role in addressing these issues, most importantly by offering principled ways of deconstructing early characterizations of implicit learning as involving the unconscious acquisition of abstract knowledge.

1. Implicit cognition: The phenomena

Everyday experience suggests that implicit learning is a ubiquitous phenomenon. For instance, we often seem to know more than we can tell. Riding a bicycle, playing tennis or driving a car all involve mastering complex sets of motor skills that we find very difficult to describe verbally. These dissociations between our ability to report on cognitive processes and the behaviors that involve these processes are not limited to action but also extend to high-level cognition. Most native speakers of a language are unable to articulate the grammatical rules they nevertheless follow when uttering expressions of the language. Likewise, expertise in domains such as medical diagnosis or chess, as well as social or aesthetic judgments, all involve intuitive knowledge that one seems to have little introspective access to.

We also often seem to tell more than we can know. In a classic article, social psychologists Nisbett and Wilson (1977) reported on many experimental demonstrations that verbal reports on our own behavior often reflect reconstructive and interpretative processes rather than genuine introspection. While it is generally agreed that cognitive processes are not in and of themselves open to any sort of introspection, Nisbett and Wilson (1977) further claimed that we can sometimes be "(a) unaware of the existence of a stimulus that importantly influenced a response, (b) unaware of the existence of the response, and (c) unaware that the stimulus has affected the response". (p. 231).

Demonstrations of dissociations between subjective experience and various cognitive processes have now been reported in many domains of cognitive science. For instance, dissociations have been reported between conscious awareness and memory. Memory for previous events can be expressed explicitly, as a conscious recollection, or implicitly, as automatic, unconscious influences on behavior. Numerous studies have demonstrated dissociations between implicit and explicit memory, both in normal participants (see Schacter, 1987) as well in special populations. Amnesic patients, for instance, who exhibit severe or total loss in their ability to explicitly recall previous experiences (conscious recollection) nevertheless retain the ability to learn novel procedural skills or to exhibit sensitivity to past experiences of which they are not conscious.

Findings of “learning without awareness” have also been reported with normal subjects (Cleeremans, Destrebecqz, and Boyer, 1998). It is Arthur Reber, in a classic series of studies conducted in 1965 (Reber, 1967), who first suggested that learning might be "implicit" to the extent that people appear to be able to learn new information without intending to do so and in such a way that the resulting knowledge is difficult to express. Implicit learning contrasts with implicit memory in that it typically involves sensitivity to relationships between events rather than sensitivity to single events, and with subliminal perception in that it typically involves supra-liminal stimuli.

Implicit learning research has essentially been focused on three experimental paradigms: Artificial grammar learning, dynamic system control, and sequence learning. Additional paradigms that will not be discussed further include probability learning, hidden covariation detection, acquisition of invariant characteristics, or visual search in complex stimulus environments.

In Reber's seminal study of artificial grammar learning (Reber, 1967), subjects were asked to memorize meaningless letter strings generated by a simple set of rules embodied in a finite-state grammar (Figure 1). After this memorization phase, subjects were told that the strings followed the rules of a grammar, and were asked to classify novel strings as grammatical or not. In this experiment and in many subsequent replications, subjects were able to perform this classification task better than chance despite remaining unable to describe the rules of the grammar in verbal reports. This dissociation between classification performance and verbal report is the finding that prompted Reber to describe learning as implicit, for subjects appeared sensitive to and could apply knowledge that they remained unable to describe and had had no intention to learn.

INSERT FIGURE 1 HERE

In a series of studies that attracted renewed interest in implicit learning, Berry and Broadbent (Berry and Broadbent, 1984) showed that success in learning how to control a simulated system (e.g., a "sugar factory") so as to make it reach certain goal states was independent from ability to answer questions about the principles governing subject's inputs and the system's output: Practice selectively influenced ability to control the system, whereas verbal explanations about how the system works selectively influenced ability to answer questions.

Today, another paradigm — sequence learning — has become dominant in the study of implicit learning. In sequence learning situations (Clegg et al., 1988), participants are asked to react to each element of a sequentially structured visual sequence of events in the context of a serial reaction time task. On each trial, subjects see a stimulus that appears at one of several locations on a computer screen and are asked to press as fast and as accurately as possible on the key corresponding to its current location. Nissen and Bullemer (1987) first demonstrated that subjects progressively learned about the sequential structure of the stimulus sequence in spite of showing little evidence of being aware that the material contained structure. Numerous subsequent studies have indicated that subjects can learn about complex sequential relationships despite remaining unable to fully deploy this knowledge in corresponding direct tasks.

2. Demonstrating implicit cognition

These findings all suggest that unconscious influences on behavior are pervasive. This raises the question of how to best characterize the relationships between conscious and unconscious processes, and in particular whether one should consider that mental representations can be unconscious. Because there is no accepted operational definition of what it means for an agent to be conscious of something, three challenges need to be overcome when attempting to contrast conscious and unconscious cognition: A definitional challenge (what is it that we want to measure?), a methodological challenge (which measure offers an appropriate index of awareness?), and a conceptual challenge (what are the implications of dissociation findings?).

The definitional challenge. Addressing this first challenge involves delineating which aspects of consciousness "count" when assessing whether a subject is aware or not of a particular piece of information: Awareness of the presence or absence of a stimulus, conscious memory for a specific previous processing episode, awareness of one's intention to use some information, or awareness that one's behavior is influenced by some previous processing episode. Consciousness is not a single process or phenomenon, but rather encompasses many dimensions of experience. Different aspects of conscious processing are thus engaged by different paradigms. In subliminal perception studies, for instance, one is concerned with determining whether stimuli that have not been consciously encoded can influence subsequent responses. In contrast, implicit memory research has been more focused on retrieval processes, that is, on the unintentional, automatic effects that previously consciously perceived stimuli may exert on subsequent decisions. In studies of implicit learning, it is the relationships between ensembles of consciously processed stimuli that remain purportedly unconscious.

These subtle differences in which specific aspects of the situation are available to awareness emphasize the need to carefully distinguish between awareness during encoding and awareness during retrieval of information. Further, both encoding and retrieval can concern either individual stimuli or relationships between sets of stimuli, and both can either be intentional or not. Even so, there is continuing disagreement about exactly which criteria for conscious awareness should be used. In a recent review for instance, Butler and Berry (2001) found little evidence for implicit memory if one takes "implicit" to imply that performance (1) results from an unintentional retrieval strategy and (2) is not accompanied by conscious recollection. Likewise, in a landmark review article dedicated to implicit learning, Shanks and StJohn (1994) failed to identify convincing evidence for learning without awareness and therefore concluded that "Human learning is almost invariably accompanied by conscious awareness" (p. 394).

The methodological challenge. This second challenge consists of devising an appropriate measure of awareness. Most experimental paradigms dedicated to exploring the relationships between conscious and unconscious processing have relied on a simple quantitative dissociation logic aimed at comparing the sensitivity of two different measures to some relevant information: A measure C of subjects' awareness of the information, and a measure P of behavioral sensitivity to the same information in the context of some task. Unconscious processing, according to the quantitative dissociation logic, is demonstrated whenever P exhibits sensitivity to some information in the absence of correlated sensitivity in C. There are several important pitfalls with this reasoning, however.

First, the measures C and P cannot typically be obtained concurrently. This "retrospective assessment" problem (Shanks and St. John, 1994) entails that finding that C fails to be sensitive to the relevant information need not necessarily imply that information was processed unconsciously during encoding, but that, for instance, it might have been forgotten or otherwise distorted before retrieval. It is therefore important, but often impossible — short of resorting to online measures of awareness such as made possible through brain imaging techniques — to obtain concurrent measures of learning and awareness, partly because measuring awareness concurrently to processing entails a form of “observer paradox” in which what is measured comes to be influenced by the measurement itself.

A second issue is to ensure that the information revealed through C is in fact relevant to perform the task. As Shanks and St. John (1994) suggested, many studies of implicit learning have failed to respect this "information" criterion. For instance, successful classification in an artificial grammar learning task need not necessarily be based on knowledge of the rules of the grammar, but can instead involve knowledge of the similarity relationships between training and test items. Subjects asked about the rules of the grammar would then understandably fail to offer relevant explicit knowledge, not because they have no awareness of the relevant rules, but simply because these rules are not necessary to perform the classification task successfully.

A third issue is to ensure that C and P are both equally sensitive to the relevant information (the “sensitivity” criterion; Shanks & StJohn, 1994). At first sight, verbal reports and other subjective measures such as confidence ratings would appear to offer the most direct way through which to assess the contents of subjective experience. The use of subjective measures to assess awareness was first advocated by Cheesman and Merikle (1984), who also introduced the notions of subjective and objective thresholds. Performance on a given task (i.e., identification) is said to be below the subjective threshold if one can show that performance is better than chance while subjects indicate they are guessing through subjective measures such as confidence judgments. Performance is said to be below the objective threshold if it fails to differ from chance on objective measures of awareness such as forced-choice tests (e.g., recognition, presence-absence decisions, or identification). Unconscious perception, for instance, would thus be demonstrated whenever performance is below the subjective threshold and above the objective threshold. This logic can also be applied to the domain of implicit learning, and several studies have now applied these ideas in the domains of artificial grammar learning and sequence learning. Overall, these studies indicate that the knowledge acquired by participants in these empirical situations can indeed be implicit to the extent that it is "below the subjective threshold".

Even if the different criteria briefly overview above are fulfilled, however, it might be elusive to hope to be able to obtain measures of awareness that are simultaneously <u>exclusive</u> and <u>exhaustive</u> with respect to knowledge held consciously. In other words, finding null sensitivity in <u>C</u>, as required by the dissociation paradigms for unconscious processing to be demonstrated, might simply be impossible because no such absolute measure exists. A significant implication of this conclusion is that, at least with normal participants, it makes little sense to assume that conditions exist where awareness can simply be "turned off".

It might therefore instead be more plausible to assume that any task is always sensitive to both conscious and unconscious influences. In other words, no task is <u>process-pure</u>. Hence, according to this logic, even performance on forced-choice, objective tests such as recognition might be influenced by unconscious influences. These tests could therefore <u>overestimate</u> explicit knowledge. This is the <u>contamination</u> problem. Two methodological approaches that specifically attempt to overcome the conceptual limitations of the dissociation logic have been developed.

The first approach was introduced by Reingold and Merikle (1988), who suggested that the search for absolute measures of awareness should simply be abandoned in favor of approaches that seek to compare the sensitivity of <u>direct</u> measures and <u>indirect</u>

measures of some discrimination. Direct measures involve tasks in which the instructions make explicit reference to the relevant discrimination, and include objective measures such a recognition or recall. In contrast, indirect measures, such as stem completion in memory tasks, make no reference to the relevant discrimination. By assumption, direct measures should exhibit greater or equal sensitivity than indirect measures to consciously held task-relevant information, for subjects should be expected to be more successful in using conscious information when instructed to do so than when not. Hence, demonstrating that an indirect task is more sensitive to some information than a comparable direct task can only be interpreted as indicating unconscious influences on performance. This “relative sensitivity approach” has been successfully applied in the study of subliminal perception and implicit memory. Jiménez et al. (1996) first applied this logic to the study of implicit learning by comparing direct and indirect measures of the knowledge acquired by participants trained on an SL task. Using detailed correlational analyses, Jiménez et al. were able to show that some sequence knowledge tended to be expressed exclusively through choice reaction decisions.

The second approach — Larry Jacoby's "Process Dissociation Procedure" (Jacoby, 1991) is based on the argument that, just as direct measures can be contaminated by unconscious influences, indirect measures can likewise be contaminated by conscious influences: Particular tasks can simply not be identified with particular underlying processes. The process dissociation procedure thus aims to tease apart the relative contributions of conscious and unconscious influences on performance. To do so, two conditions are compared in which conscious and unconscious influences either both contribute to performance improvement, or act against each other. For instance, subjects might be asked to memorize a list of words and then, after some delay, to perform a stem completion task in which word stems are to be completed either so as to form one of the words memorized earlier (the inclusion condition) or so as to form a different word (the exclusion condition). If the stems nevertheless tend to be completed by memorized words under exclusion instructions, then one can only conclude that memory for these words was implicit, since if subjects had been able to consciously recollect them, they would have avoided using them to complete the stems. Numerous experiments have now been designed using the process dissociation procedure. They collectively offer convincing evidence that performance can be influenced by unconscious information in the absence of conscious awareness. In the context of implicit learning research, Destrebecqz and Cleeremans (2001) adapted the process dissociation procedure to sequence learning, asking trained participants to either generate a sequence that resembled the training sequence (inclusion) or a sequence that was as different as possible from the training sequence (exclusion). Results indicated that under certain conditions, participants were unable to exclude familiar sequence fragments, thus suggesting that they had no control over the knowledge acquired during training. Destrebecqz and Cleeremans (2001) concluded that this knowledge was best described as implicit, for its expression is not under conscious control.

Despite the considerable methodological advances achieved over the past decade or so, assessing awareness in implicit learning and related fields remains particularly challenging. The central issue of the extent to which information processing can occur in the absence of conscious awareness remains as controversial today as it was 35 years ago.

The conceptual challenge. The third, conceptual challenge is to determine how to best interpret existing dissociation results. Dunn and Kirsner (1988) pointed out that even crossed double dissociations between two tasks do not necessarily indicate the involvement of separable, independent processes. Other authors have appealed to theoretical and simulation work to call the dissociation logic into question. Many authors have described non-modular architectures that can nevertheless produce double dissociations. Plaut (1995) explored these issues in the context of cognitive neuropsychology. In a compelling series of simulation studies, Plaut not only showed that lesioning a single connectionist network in various ways could account for the double dissociations between concrete and abstract word reading exhibited by deep dyslexic patients, but also that lesions in a single site produced both patterns of dissociations observed with patients. In other words, the observed dissociations can clearly not be attributed to architectural specialization, but can instead be a consequence of functional specialization (functional modularity) in the representational system of the network. These issues are also debated in the context of implicit learning research.

3. The role of computational modeling

Computational modeling has played a central role in deconstructing early verbal theories of implicit learning (1) by offering "proof of existence" demonstrations that elementary, associative learning processes (as opposed to rule-based learning) are in fact often sufficient to account for the data, (2) by making it possible to cast specific predictions that can then be contrasted with those of competing models, and (3) by making it possible to explore how specific computational principles can offer novel, unitary accounts of the data.

Detailed computational models have now been proposed for all three main paradigms of implicit learning. Two families of models are currently most influential: Neural network models, and fragment-based models. Neural network models typically consist of simple auto-associator models (Dienes, 1992) or of networks capable of processing sequences of events, such as the Simple Recurrent Network (SRN, see Figure 2) introduced by Elman (1990) and applied to sequence learning by Cleeremans and McClelland (1991). Such models assume that over the course of training, information about the statistical structure of the stimulus material is stored in the connection weights between the model's processing units. This information is thus subsymbolic to the extent that it is embedded in the same structures that are used to support processing itself. Fragment-based models (e.g., Perruchet and Vinter, 1998), in contrast, are variants of exemplar-based models which assume that learning results in the acquisition of memory structures such as whole exemplars or fragments thereof.

INSERT FIGURE 2 HERE

While no type of model can currently claim generality, both approaches share a number of central assumptions: (1) learning involves elementary association or recoding processes that are highly sensitive to the statistical features of the training set, (2) learning is viewed essentially as a mandatory by-product of ongoing processing, (3) learning is based on the processing of exemplars and produces distributed knowledge, and (4) learning is unsupervised and self-organizing.

More recently, hybrid models that specifically attempt to capture the relationships between symbolic and subsymbolic processes in learning have also been proposed. Sun (2001), for instance, has introduced models that specifically attempt to link the subsymbolic, associative, statistics-based processes characteristic of implicit learning with the symbolic, declarative, rule-based processes characteristic of explicit learning.

All these models have been essentially directed at addressing two questions: (1) What knowledge do people acquire in implicit learning situations? and (2) To what extent should demonstrated dissociations be interpreted as reflecting the involvement of separable learning systems?

Rules vs. Statistics. Early characterizations of implicit knowledge have tended to describe it as abstract, based essentially on findings that subjects exhibit better-than-chance transfer performance, as when asked to make grammaticality judgments on novel strings in the context of artificial grammar learning situations (Reber, 1989). Likewise, it has often been assumed that the reaction time savings observed in sequence learning tasks reflect the acquisition of "deep" knowledge about the rules used to generate the stimulus material (Lewicki et al., 1987). These abstractionist accounts have generally left it unspecified what the exact form of the acquired knowledge may be, short of noting that it must somehow represent the structure of the stimuli and their relationships, and be independent of the surface features of the material. The latter claim was further substantiated by findings that artificial grammar learning knowledge transfers to strings based on the same grammar but instantiated with a different letter set, or even across modalities, as when training involves letter strings but transfer involves tone sequences.

However, there now is considerable evidence that non-abstractionist mechanisms are largely sufficient to account for the data. Brooks (1978) first suggested that subjects in artificial grammar learning experiments were classifying novel strings based not on abstract knowledge of the rules, but simply based on the extent to which novel grammatical or ungrammatical strings are similar to whole exemplars memorized during training. Perruchet and colleagues (Perruchet and Pacteau, 1990) showed that the knowledge acquired in both artificial grammar learning and sequence learning tasks might consist of little more than explicitly memorized short fragments or chunks of the training material such as bigrams or trigrams, or simple frequency counts. Both learning and transfer performance can then be accounted for by the extent to which novel material contains memorized chunks. Figure 3 illustrates some of the computational possibilities that have been suggested in the context of artificial grammar learning tasks, ranging from purely exemplar-based approaches to neural network models.

INSERT FIGURE 3 HERE

More recently, accounts that assume separate memory systems for representing general or specific knowledge in artificial grammar learning tasks have been proposed based on evidence that significant sensitivity to grammaticality remains even when similarity and fragment overlap is carefully controlled for.

Overall, while it is clear that the knowledge acquired in typical implicit learning situations need not be based on the unconscious acquisition of symbolic rules, significant areas of debate remain about the extent to which unitary, fragment-based mechanisms are sufficient to account for sensitivity to both the general and specific features of the training material. Simulation models have generally been suggestive that such mechanisms are in fact sufficient to account simultaneously for both grammaticality and similarity effects.

Based on these properties of successful models of implicit learning, it is appealing to consider it as a complex form of priming whereby experience continuously shapes memory, and through which stored traces in turn continuously influence further processing. Such priming, far from involving the sort of passive and automatic acquisition of abstract structure that were previously assumed to lie at the heart of implicit learning, is in fact highly dependent on task demands and attentional processing during acquisition, as well as on the congruence between learning and transfer conditions (Whittlesea and Dorken, 1993).

Finally, while both fragment-based and neural network models make it clear how sensitivity to the distributional properties of an ensemble of stimuli can emerge out of the processing of exemplars, they differ in whether they assume that the shared features of the training materials are represented as such or merely computed when needed. This locus of abstraction issue is a difficult one that is unlikely to be resolved by modeling alone. Overall thus, it appears that the knowledge acquired through implicit learning is best described as lying somewhere on a continuum between purely exemplar-based representations and more general, abstract representations —a characteristic that neural network models are particularly apt at capturing.

Separable systems? Dissociations between implicit and explicit learning or processing have often been interpreted as suggesting the existence of separable memory systems. For instance, Squire and collaborators have shown that artificial grammar learning is largely preserved in amnesia (e.g., Knowlton et al., 1992), to the extent that amnesic patients perform at the same level as normal controls when asked to classify strings as grammatical or not, but are severely impaired when asked to discriminate between familiar and novel instances (or fragments) of the strings. These results suggest that the processes that subtend declarative and non-declarative memory depend on separable brain systems respectively dedicated to representing either information about the specific features of each encountered exemplar on the one hand (the hippocampus and related structures), and information about the features shared by many exemplars on the other hand (the neocortex).

In this case also however, computational modeling often casts the empirical findings in a different light. For instance, Kinder and Shanks (2001) were able to simulate the observed dissociations by tuning a single parameter (the learning rate) in a Simple Recurrent Network trained on the same material as used in the behavioral studies, and therefore concluded that a single-system account is in fact sufficient to account for the data.

4. Conclusions

The study of differences between implicit and explicit processing is a major endeavor for the cognitive neurosciences. Indeed, as our knowledge of how the brain works accumulates, our knowledge about how the mind works is changing rapidly. In particular, many existing distinctions previously described in purely functional, binary terms, such as the implicit-explicit distinction, the controlled-automatic distinction, or the declarative-procedural distinction, are now being characterized anew in terms of graded characterizations of the computational problems that corresponding brain systems have been evolved to solve. In this respect, the study of implicit learning, from the perspective of computational cognitive neuroscience, has a bright future, for it is through the development of sensitive paradigms through which to explore the differences between conscious and unconscious cognition that one can best contribute to the search for the neural, behavioral, and computational correlates of consciousness.

Glossary

Auto-associator network. A neural network the task of which consists of producing its inputs as outputs.

Direct vs. indirect measures. Direct measures of some discrimination involve tasks in which the instructions make explicit reference to the relevant discrimination, and include objective measures such a recognition or recall. In contrast, indirect measures, such as stem completion in memory tasks, make no reference to the relevant discrimination.

Double dissociation. A double dissociation between two measures A and B obtains whenever there exists two independent variables, each of which exclusively influences one of measures A and B. The double dissociation is said to be "crossed" when the two variables have opposite effects on measures A and B.

Finite-state grammar. A finite-state grammar is a simple directed graph consisting of nodes connected by labeled arcs. Sequences of symbols can be generated by entering the grammar through an "in" node, and by moving from node to node until an "out" node is reached. Each transition between a node and the next produces the label associated with the arc linking the two nodes.

References

Berry DC and Broadbent DE (1984) On the relationship between task performance and associated verbalizable knowledge. *Quarterly Journal of Experimental Psychology* **36**: 209-231.

Berry DC and Dienes Z (1993) *Implicit learning: Theoretical and empirical issues*. Hove, UK: Lawrence Erlbaum Associates.

Butler LT and Berry DC (2001) Implicit memory: Intention and awareness revisited. *Trends in Cognitive Sciences* **5**:192-197.

Brooks LR (1978) Non-analytic concept formation and memory for instances. In: Rosch E and Lloyd B (eds) *Cognition and Concepts*, pp. 16-211. Mahwah, N.J.: Lawrence Erlbaum Associates.

Cheesman J and Merikle PM (1984). Priming with and without awareness. *Perception and Psychophysics* **36**: 87-395.

Cleeremans A and McClelland JL (1991) Learning the structure of event sequences. *Journal of Experimental Psychology : General* **120**: 235-253.

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

Clegg BA, DiGirolamo GJ and Keele SW (1998) Sequence learning. *Trends in Cognitive* Sciences **2**:75-281.

Destrebecqz A and Cleeremans, A (2001) Can sequence learning be implicit? New evidence with the Process Dissociation Procedure. *Psychonomic Bulletin & Review* **8**: 343-350.

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

Dunn JC and Kirsner K (1988) Discovering functionally independent mental process: The principle of reversed association. *Psychological Review* **95**: 91-101.

Elman JL (1990) Finding structure in time. *Cognitive Science* **14**: 179-211.

Frensch PA (1998) One concept, multiple meanings: On how to define the concept of implicit learning. In: Stadler MA and Frensch PA (eds) *Handbook of implicit learning*, pp. 47-104. Thousand Oaks, CA: Sage Publications.

Jacoby LL (1991) A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language* **30**: 513-541.

Jiménez L, Mendéz C and Cleeremans A (1996) Comparing direct and indirect measures of sequence learning. *Journal of Experimental Psychology□Learning, Memory and Cognition* **22**: 948-969.

Kinder A and Shanks DR (2001) Amnesia and the Declarative/Nondeclarative distinction: A recurrent network model of classification, recognition, and repetition priming. *Journal of Cognitive Neuroscience* **13**: 648-669.

Knowlton BJ, Ramus SJ and Squire LR (1992) Intact artificial grammar learning in amnesia: Dissociation of classification learning and explicit memory for specific instances. *Psychological Science* **3**:172-179.

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

Nisbett RE and Wilson TD (1977) Telling more than we can know: Verbal reports on mental processes. *Psychological Review* **84**: 231-259.

Nissen MJ and Bullemer P (1987) Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology* **19**:1-32.

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

Perruchet P and Vinter A (1998) PARSER: A model for word segmentation. *Journal of Memory and Language* **39**: 246-263.

Plaut DC (1995) Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropschology* **17**: 291-326.

Reber AS (1967) Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior* **6**: 855-863.

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

Reingold EM and Merikle PM (1988) Using direct and indirect measures to study perception without awareness, *Perception and Psychophysics* **44**: 563-575.

Shanks DR and St. John MF (1994) Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences* **17**: 367-447.

Schacter DL (1987) Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **13**: 501-518.

Sun R (2001) *Duality of the Mind*. Mahwah, NJ: Lawrence Erlbaum Associates.

Whittlesea BWA and Dorken MD (1993) Incidentally, things in general are particulary determined: An episodic-processing account of implicit learning. *Journal of Experimental Psychology: General* **122**: 227-248.

Bibliography

Berry DC (1997) *How implicit is implicit learning?* Oxford: Oxford University Press.

Berry DC and Dienes, Z (1993) *Implicit learning: Theoretical and empirical issues*. Hove, UK: Lawrence Erlbaum Associates.

Cleeremans A (1993) *Mechanisms of implicit learning: Connectionist models of sequence processing*. Cambridge: MIT Press.

French RM and Cleeremans A (2002). *Implicit learning and consciousness: An empirical, computational and philosophical consensus in the making*. Hove, UK: Psychology Press.

Reber AS (1993) *Implicit learning and tacit knowledge: An essay on the cognitive unconscious*. Oxford, UK: Oxford University Press.

Stadler MA and French PA (1998) *Handbook of implicit learning*. Thousand Oaks, CA: Sage Publications.

Figure captions

Figure 1: A finite-state grammar is a simple directed graph consisting of nodes connected by labeled arcs. Sequences of symbols can be generated by entering the grammar through an “in” node, and by moving from node to node until an “out” node is reached. Each transition between a node and the next produces the label associated with the arc linking the two nodes. Concatenating the symbols together produces strings of symbols, in this case, letters of the alphabet.

Figure 2: The Simple Recurrent Network (SRN) introduced by Elman (1990). The network takes the current element of a sequence as input, and is trained to predict the next element using back-propagation. Context units, which on time step contain a copy of the activation pattern that existed over the network’s hidden units on the previous time step, enable previous information to influence current predictions.

Figure 3: A representation of different computational approaches to artificial grammar learning.

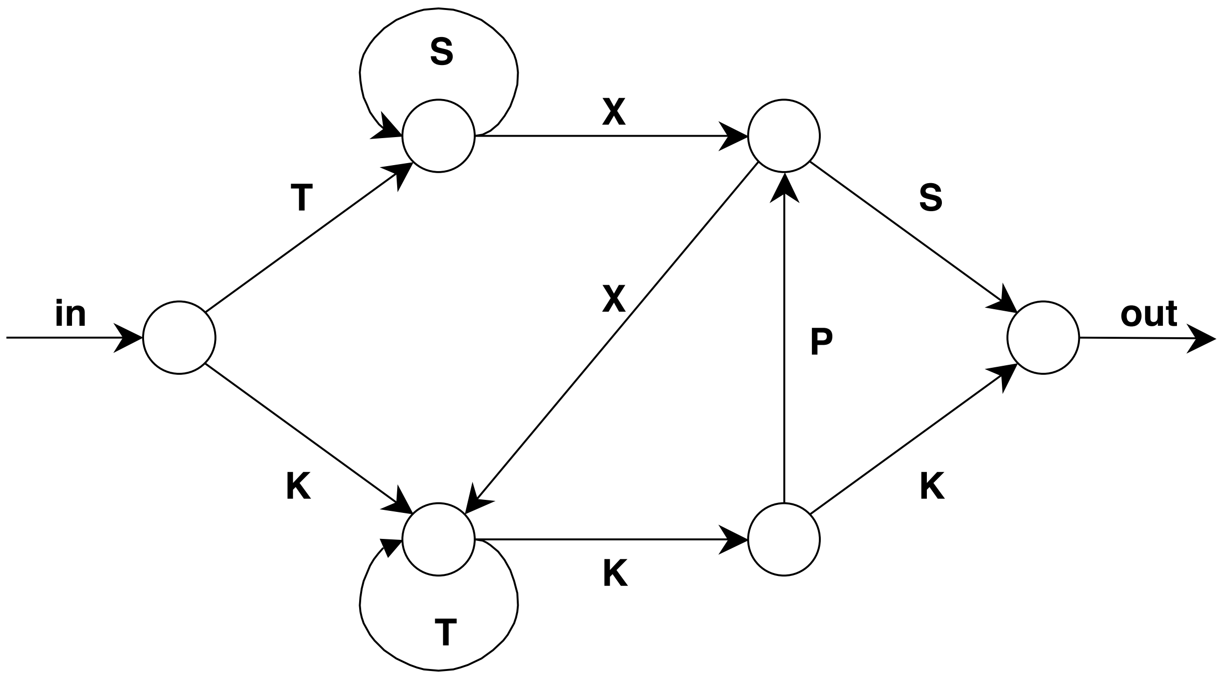
S
X
T
S
in
X
P
out
K
K
K
T

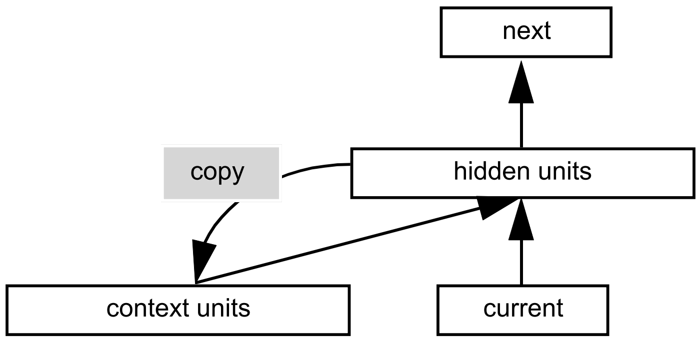
next
copy
hidden units
context units
current

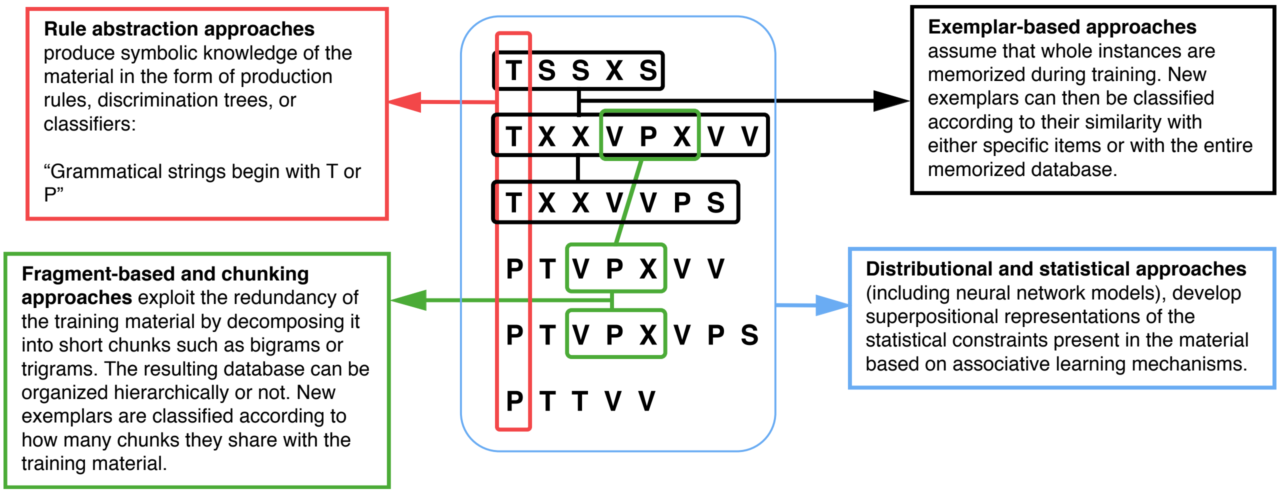
Rule abstraction approaches produce symbolic knowledge of the material in the form of production rules, discrimination trees, or classifiers:
"Grammatical strings begin with T or P"
T S S X S
T X X V P X V V
T X X V V P S
P T V P X V V
P T V P X V P S
P T T V V
Exemplar-based approaches assume that whole instances are memorized during training. New exemplars can then be classified according to their similarity with either specific items or with the entire memorized database.
Fragment-based and chunking approaches exploit the redundancy of the training material by decomposing it into short chunks such as bigrams or trigrams. The resulting database can be organized hierarchically or not. New exemplars are classified according to how many chunks they share with the training material.
Distributional and statistical approaches (including neural network models), develop superpositional representations of the statistical constraints present in the material based on associative learning mechanisms.