Transfer of implicit knowledge across domains: How implicit and how abstract?

Authors: Zoltán Dienes and Gerry Altmann, University of Sussex

Editor: Dianne Berry

Book Title: How implicit is implicit learning?

Oxford University Press.

A Introduction

One of the putative advantages of conscious knowledge is its flexibility: It can be applied in novel ways to novel situations (e.g. Baars, 1988). Unconscious knowledge on the other hand is commonly supposed to be routinized, inflexible, and triggered by specific perceptual cues (e.g. Berry & Dienes, 1993; Schacter, 1987; Shiffrin & Schneider, 1977). For example, consider knowledge of directions for travelling a certain route. If the directions are held consciously, any route can be readily followed . On the other hand, once a route has been travelled many times, the directions do not have to be held consciously in mind; but in this case, the knowledge is inflexible in that only one particlar route will tend to be followed.

There are cases where conscious knowledge appears not to be flexibly deployed across domains, as illustrated by the literature on analogical transfer. People often do not notice that the known solution to a problem they have just come across could be used to solve a new problem in an analogous way. For example, Gick and Holyoak (1980, 1983) presented subjects with the classic radiation problem: Subjects imagined being a doctor whose task was to destroy a tumour with rays but without destroying the surounding healthy tissue. The solution was to present numerous low intensity rays at different angles such that they converged on the tumour with sufficient intensity to destroy tissue at just that point. Subjects were no more likely to find the solution to this problem if they previously had been exposed to an analogous military problem and its solution. The conscious knowledge was not flexibly deployed. The key difficulty for subjects in this experiment appeared to be noticing the relevance of the first problem for the second. When subjects were given hints about the relevance of the first problem for the second, their performance improved dramatically.

In some cases, people informed of the relevance of a prior problem may still fail to make an appropriate mapping between the domains. For example, novices often cannot see the common underlying structure of different maths problems, instead they see problems as similar if they have similar superficial detail (Chi, Feltovich, & Glaser, 1981). But at least one can say that when two problems have been appropriately parsed into corresponding elements with a transparent mapping, and subjects have been informed of the relevance of one problem for the other, the conscious knowledge can be applied in a completely flexible way. Davies (1989) suggested that in general, when people are asked to inferentially combine two pieces of conscious knowledge, they can in principle do so; on the other hand, they cannot combine belief with unconscious knowledge to yield further belief.

Another area where the flexibility of conscious and unconscious knowledge has been explored is implicit and explicit memory. Explicit memory refers to concsious recollection. In a typical implicit memory experiment, exposure to a written word leads to faster or more accurate processing of that word at a later time. Because this facilitation occurs regardless of whether the word is explicitly recognized as old or not, it may be regarded as unconscious memory. Explicit memory can be affected by peceptual cues, as in the case of context dependent memory. For example, divers who study a list of words in one of two contexts (either under water or on ground), will recall more words when tested in the same rather than the different context (Godden & Baddeley, 1975). However, explicit memory appears to be more affected by contexts that affect meaning rather than contexts that just affect superficial sensory details. When cued recall or recognition are used, contexts that affect only perceptual details and not meaning no longer influence retrieval (Baddeley, 1982). That is, direct cues as to the identity of the stimulus or its meaning override cues merely about perceptual detail. On the other hand, implicit memory remains sensitive to perceptual cues even when the identity of the stimulus is specified (Schacter, 1987). That is, unconcsious memory is more inflexible and tied to specific perceptual cues than conscious memory.

Expert knowledge is often thought to be relatively rigid and domain specific. Is this an example of rigid conscious knowledge? However, an expert's knowledge may not be all conscious (Berry & Dienes, 1993; Singley & Anderson, 1989). Singley and Anderson developed a model of how expert's acquire knowledge and how this knowledge could be transferred across different situations. They suggested that people start with an initial conscious declarative knowledge base that can in principle be used flexibly (albeit slowly) for different problems. With practice, specific productions (i.e. condition-action rules) relevant to the problems solved are built up and stored - these productions constitute the subjects unconscious knowledge. The unconscious knowledge transfers to a new problem only to the extent that the problem solution requires the same productions - the greater the number of productions in common to the solutions for the two problems, the greater the transfer. According to the theory, if the knowledge is to be put to a quite different use (e.g. contrast understanding a language to generating a language), there will be no transfer of unconscious knowledge, because productions are use specific. On the other hand, declarative knowledge can be put to many different uses. If the solution to an initial problem involves considering perceptual details, so that the conditions of the productions refer to perceptual detail, then there will be limited transfer of unconscious knowledge across different perceptual domains.

One of the major paradigms for investigating unconscious knowledge has been artificial grammar learning (e.g. Reber, 1967). In a typical study, subjects first memorize grammatical strings of letters generated by a finite-state grammar. Then, they are informed of the existence of the complex set of rules that constrains letter order (but not what they are), and are asked to classify grammatical and nongrammatical strings. Subjects can classify substantially above chance, indicating that they have learned some of the underlying grammatical structure. Reber (1989) argued that the knowledge was both largely unconscious and could be abstract in the sense of not being bound to any particular perceptual features. He argued that the degree of abstraction depended on the conditions of learning: if subjects are asked to keep distinct memories for separate items they acquire representations that are tied to the specifics of each training exemplar. If however, subjects are simply asked to observe a set of exemplars, then they acquire abstract knowledge. That is, according to Reber, unconscious knowledge, at least in some cases of artificial grammar learning, is not as inflexible as is commonly supposed. This chapter will explore the issue, first, by evaluating the criteria by which knowledge of an artificial grammar could be unconscious; second, by evaluating evidence for the flexibility of knowledge of artificial grammars; and third, by presenting new evidence directly testing the unconsciousness of knowledge not bound to specific perceptual features.

B Methodological criteria for testing whether knowledge is unconscious

The question of whether knowledge of artificial grammars can be unconscious has aroused considerable controversy (e.g.Dienes & Perner, 1995; Dulany, Carlson, & Dewey, 1984; Perruchet & Pacteau, 1990; Reber, 1989; Shanks & St John, 1994). The resolution of this controversy depends on which of several possible criteria is used for defining the divide between the conscious and the unconscious (Dienes & Perner, 1995). Cheesman and Merikle (1984) argued that the unconscious can usefully be defined in terms of a subjective threshold, and this is the criterion that we will be using in this chapter. According to a subjective threshold criterion, people have unconscious knowledge if they lack metaknowledge about their knowledge. That is, we could regard knowledge as unconscious if people do not know that they know it. A lack of metaknowledge can be revealed in two ways: A lack of correlation between confidence and accuracy (Chan, 1992; we will call this the zero correlation criterion); and accurate performance when subjects believe that they are guessing (Cheesman & Merikle, 1984; we will call this the guessing criterion). Both criteria are possible ways of demarcating a subjective threshold.

Both the guessing and zero correlation criteria tap only one of two possible forms of metaknowledge. Dienes and Perner (1996) argued that full metaknowledge requires both representing oneself as being in the possession of certain propositional content (they called this type of representation 'content explicit'); and also representing an appropriate propositional attitude towards that content - namely, representing the content as knowledge and not, for example, just as confabulation (they called the second type of representation 'attitude explicit'). Knowledge could be unconscious according to the metaknowledge criterion either because content explicitness was lacking or because attitude explicitness was lacking. Finding that confidence is unrelated to accuracy, or that subjects are accurate when they believe that they are guessing, shows that subjects' representations lack attitude explicitness. But they may have content explicitness, as in the possible case of accurate verbal reports offered as guesses. Reber (1993; page 136) suggested that subjects in artificial grammar learning experiments can show the converse; that is, they may know that they know something, even though they may not know what it was that they know. Similarly, in natural language we may be very confident in our judgements of grammaticality, and so the judgements per se would be attitude explicit, but be unable to describe the bases of those judgements, and so the bases would be content implicit. Both types of lack of metaknowledge - content implicitness and attitude implicitness - are possible, and even plausible in different circumstances. We will be exloring only whether subjects' representations can lack attitude explicitness, as assessed by the zero correlation and guessing criteria.

B.1 The zero correlation criterion

Chan (1992) investigated whether the incidentally acquired knowledge of an artificial grammar was implicit in the sense that subjects' confidence in their decisions predicts their accuracy. Chan initially asked one group of subjects (the incidentally trained subjects) to memorize a set of grammatical exemplars. Then in a subsequent test phase, subjects gave a confidence rating for their accuracy after each classification decision. Chan found that these subjects were just as confident in their incorrect decisions as they were in their correct decisions, and this was so despite the fact that in general subjects acquired almost completely veridical knowledge about the grammar (Reber, 1989). When this relation between confidence and accuracy was expressed as a correlation, the correlation was .20 (Chan's experiments six and seven) . He asked another group of subjects (the intentionally trained subjects) to search for rules in the training phase. For these subjects, confidence was strongly related to accuracy in the test phase - the correlation between confidence and accuracy was .54 (significantly greater than the value for incidental subjects) . That is, incidentally trained subjects lacked metaknowledge, intentionally trained subjects did not. Chan also found that the magnitude of the correlation depended on the type of stimuli used. For example, if incidentally trained subjects were asked to rate bigrams for grammaticality, then confidence was correlated with accuracy (r = .47), consistent with the idea that subjects acquired conscious knowledge of bigrams (Perruchet & Pacteau, 1990). Chan concluded that the correlation between confidence and accuracy could be a useful criterion of consciousness.

Dienes, Altmann, Kwan, & Goode (1995) replicated the lack of relationship between confidence and accuracy, but only under some conditions: Metaknowledge was low particularly when strings were longer than three letters and presented individually. It may be difficult to find conditions under which Chan's measure of metaknowledge is literally zero, but so long as it varies in meaningful ways (when classification performance is constant) it may provide a useful measure of degree of metaknowledge and hence degree of consciousness.

In summary, when subjects are trained in the normal incidental way, they have low, even if sometimes statistically significant, levels of metaknowledge about their knowledge of the artificial grammar.

B.2 The guessing criterion

Cheesman and Merikle (1984, 1986) introduced this criterion to determine when a stimulus was subliminally perceived. They initially gave subjects a sequence of trials in which one of four stimuli was presented. For each trial subjects said which of the stimuli was presented on that trial. Subjects also provided an estimate of their detection performance, varying from claimed random guessing to complete confidence. The subjective threshold was defined as occurring at the level of discriminative responding for which subjects claimed not to be able to detect perceptual information (that is, they claimed to be literally guessing). Cheesman and Merikle found that when subjects believed they were guessing, they were actually discriminating significantly above chance. Thus, subjects did not know that they knew what the stimulus was. Further, stimuli presented under these conditions also produced reliable semantic priming. That is, when subjects believed that they did not know what a stimulus was, the identity of the stimulus nonetheless affected their behaviour, providing an example of unconscious knowledge according to this criterion.

Dienes, Altmann, Kwan, & Goode (1995) also found subjects could lack metaknowledge about their knowledge of an artificial grammar according to the criterion of Cheesman & Merikle (1984): Even when subjects believed that they were literally guessing, they were still classifying test strings substantially above chance. Based on this criterion, about a third of subjects responses could be regarded as being based on purely unconscious knowledge.

B.3 conclusion

According to criteria of unconscious knowledge defined in terms of metaknowledge, subjects' knowledge of an artificial grammar can be unconscious, supporting Reber's (1989) claim. Now we consider the extent to which knowledge of artificial grammars is flexible and not bound to perceptual features. The question of flexibility has been investigated by exploring the extent to which subjects' knowledge can transfer between perceptually different domains.

C. Transfer across domains

Reber (1969) asked subjects to memorize strings of letters generated by a finite state grammar. The more strings subjects had previously studied, the easier they found it to memorize new strings generated by the grammar. There was also a benefit when the new strings were constructed from a different letter set, but the same grammar. That is, subjects could apply their knowledge of the grammar to a completely different letter set, leading Reber to claim that the knowledge was abstract and not perceptually bound. Subsequent work by Mathews, Buss, Stanley, Blanchard-Fields, Cho, & Druhan (1989), Brooks and Vokey (1991), Whittlesea and Dorken (1993), Gomez and Schvaneveldt (1994), and Manza and Reber (this volume) exposed subjects to strings constructed from one letter set, and later tested subjects' ability to classify as grammatical or not strings constructed from another letter set. As well as showing transfer performance significantly above chance, these studies have also all demonstrated a transfer decrement, that is, an advantage of same domain performance over transfer performance. Altmann, Dienes, and Goode (1995) extended these findings to transfer between different modalities. In each of four experiments, involving transfer between letters and music, and between graphic symbols and nonsense syllables, despite significant transfer, subjects performed better when tested in the same domain as they were trained in rather than the different domain.

In these studies, subjects were not told of the mapping between the domains, and they were not given accuracy feedback. So it is not clear whether the transfer decrement (i.e. the difference between transfer and same domain performance) reflects the problem of inducing a mapping across domains, or reflects the possibility that at least some of the knowledge is bound to particular perceptual features. There are three further studies that bear on this question. Whittlesea and Dorken (1993) trained subjects on exemplars from two grammars. The exemplars were distinguished by the task subjects performed on them: For one grammar, subjects pronounced the exemplars; for the other grammar, subjects spelled them. Subjects were later asked to pronounce or spell test items, and then classify these strings as either belonging to the 'pronouncing' grammar or to the 'spelling' grammar. Whittlesea and Dorken argued that discrimination between the grammars was achieved simply by feelings of familiarity induced by test experiences perceptually matching or mismatching representations of prior training experiences: When test strings were common to both grammars, subjects tended to classify them as belonging to the spelling grammar if they spelt them and as belonging to the pronouncing grammar if they pronounced them. Manza and Reber (this volume) found no significant transfer decrement between sequences of tones and lights when subjects were informed of the mapping. However, the same domain performance was 58% and transfer performance was only 54%, consistent with a substantial decrement (the upper limit of the 95% confidence interval on the decrement must be at least 2 *(58%-54%) = 8%). Finally, Wright (1993), using number stimuli, found significant decrements when format of the stimuli was changed from digits (e.g. 1836) during training to words (e.g. one-eight-three-six) rather than digits at test: Transfer performance was 57% and same domain performance 66%. However, in this case, the mapping between the domains may not have been so straight forward: 1836 could be encoded in a number of different ways (e.g. one-thousand-and-thirty-six, eighteen-thirty-six, etc).

In summary, despite the ability of subjects to transfer across domains, the transfer decrement leaves open the question of whether the knowledge is partly perceptualy bound, or if it is independent of specific perceptual features, as Reber (1989) claimed.

D. Can transfer be implicit?

The evidence reviewed above suggests that:

(1) Some knowledge of an artificial grammar may be perceptually bound, but this has not yet been directly tested. Experiment one will test this claim by looking at transfer between domains with known clear mappings.

Also:

(2a) Knowledge of an artificial grammar that applies to the same domain can be unconscious.

(2b) Some knowledge can transfer between domains. However, the question of whether unconscious knowledge can be flexibly deployed across domains remains unanswered. It may be that the unconscious knowledge - point (2a) - is not the same knowledge that allows transfer - point (2b). That is, we do not know whether transfer is based wholly on conscious knowledge, wholly on unconscious knowledge, or partly on both. Experiment two will address this question by measuring metaknowledge in a transfer paradigm.

D.1 Experiment one: Is knowledge of an artificial grammar perceptually-bound?

D.1.1 Introduction

Experiment one investigated transfer in a case where the mapping between the domains was transparent so any transfer decrement must arise because the knowledge was perceptually-bound. The domains were colours and names of colours.

In the classification phase, a counterbalancing procedure was used in which the grammatical items for half the subjects were the nongrammatical items for the other half of the subjects, and vice versa. This was achieved by training half the subjects on one grammar (grammar A), and the other half of the subjects on another grammar (grammar B). Half the test items were garmmatical according to grammar A but not B, and the other half were grammatical according to grammar B but not A. Thus, classification performance could only be above 50% on average if subjects had learned to apply some part of the grammar from the training phase to the test phase. This procedure obviates the need for a control group and thus makes more efficient use of subjects. Even if subjects could detect the two grammar types by simply examining test items (Redington & Chater, in press), population performance would still only be above 50% if subjects had learned something from the training phase.

D.1.2 Method

D.1.2.1 Design

The factors of theoretical interest were training domain (words vs colours) and testing domain (same vs transfer). Training domain was a between-subjects factor, and testing domain was a within-subjects factor. In addition, there were the counterbalancing between-subjects factors of order (same domain tested first vs transfer domain tested first) and grammar trained on (grammar A vs grammar B). A fully factorial design was used.

Forty-eight volunteers from the University of Sussex were used as subjects. Equal numbers of subjects were randomly assigned to each of the cells of the design. No subject had participated in any other artificial grammar learning experiment.

D.1.2.2 Materials

The two grammars used were taken from Reber (1969) and Dienes et al (1985). Eighteen strings between length 5 and 9 items (words or colours) were randomly sampled from each grammar to form the two sets of training strings, a set for grammar A and a set for grammar B. A different set of 29 strings (with the same length constraints as the training strings) were sampled from each grammar, and randomly combined to make the test set (that is, the test set was comprised of 58 strings). The colour stimuli used the colours brown, green, blue, red, and mauve, and the word stimuli used the same colour names. Colours were presented as adjacent solid-filled squares of side 1cm. Colour names were displayed with letters occupying a square of side 0.3 cm, and there was 0.8 cm separation between each word in the same sequence. The training and test items are shown in the Appendix (in their letters embodiment as used in Experiment two).

D.1.2.3 Procedure

In the training phase, subjects were given instructions to study and remember as much as possible about the stimuli, but subjects were not informed that the stimuli were rule-governed. The 18 training strings were presented individually with a 5 sec exposure for each presentation. The set of 18 was cycled through 6 times, with a different random order each time. Subjects made no overt response during the training phase.

In the test phase, subjects were informed that there was a complex set of rules that governed the order of items in each string. They were told that half the strings they were about to see would obey the rules and half would disobey the rules, and their task was to classify the strings. Each string remained on the screen until the subject made a decision.

D.1.3 Results

Note that the counterbalancing factors (grammar, order) produced no significant effects, and so analyses are reported without these factors.

Table 1 displays the mean classification performance for the different conditions.
A 2 X 2 (training group (colours vs words) by testing domain (same vs transfer)) mixed model analysis of variance indicated only a significant effect of testing domain, $F(1,46) = 11.53$, $p = .0014$. That is, subjects performed better in the same domain rather than the transfer domain. Despite this transfer decrement, subjects still performed above chance in the transfer conditions, both ps < .001, by t-test.

D.1.4 Discussion

Experiment one found a significant transfer decrement even though the mapping between the domains was a previously well-learned one. That is, the decrement could not be due to the problems of determining the mapping between the domains, but it must be due to the fact that the knowledge was perceptually bound. Performance in the transfer domain was

76% of performance in the same domain (i.e. (63-50)/(67-50)), representing a drop of 24% due to the change in perceptual embodiment. This is the type of drop found in the implicit memory literature after a change in perceptual conditions: For example, Bassili, Smith, and McLeod found an average drop of 38% when modality was changed; Blaxton (1989) found an average drop of 22% when visual format was changed; and Roediger and Blaxton (1987) found an average drop of 33% when modality was changed. That is, in contrast to the claim of Reber (1989), incidentally acquired knowledge of artificial grammars is not completely abstract, but is perceptually bound to about the same extent as implicit memory.

Even though implicit learning and implicit memory are both called 'implicit', this does not mean that they rely on the same mechanisms or have the same properties. In particular, implicit memory reveals itself on tests of specific previously exposed items; implicit learning, on the other hand, shows subjects ability to generalize to new items. Perhaps a mechanism that allows generalization might have been expected to be less context bound than one that produces sensitivity to exact matches. However, generalization and sensitivity to exact matches can be just different aspects of the same type of mechanism (e.g. McClelland & Rumelhart, 1985), and the current results show that both generalization and sensitivity to exact matches can be context bound to the same extent.

D.2 Experiment two: Metaknowledge and transfer

D.2.1 Introduction

Experiment two tested the relationship between consciousness and transfer. On the one hand, Reber (1989) claimed that transfer was based on unconscious knowledge. On the other hand, Shanks and St John (1994) argued that transfer could be based on conscious knowledge. Experiment two tested whether the knowledge underlying transfer was unconscious according to the criteria of metaknowledge: the zero correlation criterion and the guessing criterion. Transfer was tested between the domains of letters and colours. That is, in this experiment there was no a priori relation between the domains. This provides a stronger test of whether the type of transfer typically seen in the literature can really be unconscious.

For example, transfer may occur unconsciously between domains with well-learned mappings but not between arbitrary domains (which present a more complex problem).

D.2.2 Method

Exactly the same procedure and materials were used as for experiment one, with two exceptions First, instead of using words as one of the domains, experiment two used letters. The letter stimuli used the letters M, T, V, R, and X. Each letter occupied a square of side 1 cm. Second, in the testing phase, after making a decision, the subject was asked for a confidence judgement on a scale from 50 to 100, where 50 indicated a complete guess and 100 indicated complete certainty.

Forty volunteers from the University of Sussex were used as subjects. Equal numbers were randomly assigned to each of the cells of the design.

D.2.3 Results

Note that the counterbalancing factors (grammar, order) produced no significant effects, and so analyses are reported without these factors.

D.2.3.1 Classification performance

Table 2 displays the mean classification performance for the different conditions.
A 2 X 2 (training group (colours vs letters) by testing domain (same vs transfer)) mixed model analysis of variance indicated only a significant effect of testing domain, $F(1,38) = 64.86$, $p < .0001$. That is, subjects performed better in the same domain rather than the transfer domain. Despite this transfer decrement, subjects still performed above chance in the transfer conditions: For colours to letters, $t(19) = 3.43$, $p < .01$; and for letters to colours, $t(19) = 2.62$, $p < .02$.

D.2.3.2 Zero correlation criterion

Table 3 displays the mean difference in confidence between when subjects made a correct decision and when subjects made an incorrect decision. Note that Table 3 does not display correlations, because, consistent with our previous application of the criterion, the relationship between confidence and accuracy can be most simply expressed as a difference rather than a correlation ; see Dienes, Altmann, Kwan, & Goode, 1995. The difference score allows the degree of metaknowledge to be assessed on a meaningful scale (i.e. the confidence scale, which is the subjects estimate of their actual classification performance). A 2 X 2 (training group (colours vs letters) by testing domain (same vs transfer)) mixed model analysis of variance indicated only a significant effect of testing domain, F(1,38) = 36.75, p < .0001. That is, subjects had more metaknowledge in the same domain rather than the transfer domain. In fact, there was no evidence that subjects possessed any metaknowledge in the transfer domain: the mean score for the transfer conditions did not differ significantly from zero, t(38) = 1.06, p > .30. The upper limit of the 95% Confidence Interval of metaknowledge in the transfer conditions was 2.2, indicating that whatever metaknowledge subjects may have had, it was very small. Even in the same domain the amount of metaknowledge was small - the difference score was 5.5 on a scale that goes up to 50.

D.2.3.3 Guessing criterion

Table 4 displays the mean percentage correct classification when subjects believed that they were literally guessing (i.e. gave a confidence rating of 50). A 2 X 2 (training group (colours vs letters) by testing domain (same vs transfer)) mixed model analysis of variance indicated no significant effects. Subjects classified significantly above chance in all conditions - all ps < .02 by t-test.

D.2.4 Discussion

Experiment two replicated the finding that subjects can transfer their knowledge across arbitrarily different domains without prior knowledge of the mapping. The aim of experiment two was to determine the extent to which subjects possessed metaknowledge about this knowledge; i.e. the extent to which the knowledge was conscious. According to Reber (1989), the transferable knowledge is unconscious. According to Shanks & St John

(1994) it should be conscious. Also, because transfer presumably requires more sophisticated types of processes than same domain application of knowledge, determining how conscious transfer knowledge is has implications for how sophisticated the unconscious is. For example, Redington & Chater (in press) showed how transfer could be produced by subjects engaging in an elaborate series of code breaking strategies, matching up memorized fragments from the training phase with the test stimuli, and pooling information across test items and memorized fragments. Dienes, Altmann, & Gao (1995) presented a connectionist model of transfer which settles into the correct mapping between the domains, but also uses the results from successive test stimuli to hone in on a correct mapping.

Experiment two found the normal mixture of implicit and explicit knowledge in the same domain. Subjects were more confident when they made a correct rather than an incorrect decision, indicating that they had some awareness of, say, when they were applying knowledge and when they were just guessing (see Dienes & Perner, 1995, for further discussion of the meaning of this criterion). However, there was also some purely implicit knowledge: When subjects believed that they were literally guessing, they were classifying at a rate substantially above chance.

Crucially, in the transfer conditions, there was no evidence of explicit knowledge to within the limits of the sensitivity of the experiment - subjects were no more confident in correct than in incorrect decisions. On the other hand, there was clear evidence of implicit knowledge - when subjects believed they were literally guessing, they were classifying at a rate substantially above chance. That is, as far as we can tell, transfer is based largely on implicit knowledge.[1]

E General Discussion

The aim of this chapter was to explore the extent to which unconscious knowledge could be applied flexibly. Reber (1989) claimed that when subjects incidentally learn artificial grammars, they acquire knowledge that is both abstract and unconscious; that is, contrary to the claims made in the implicit memory (Schacter, 1987) and automaticity

(Shiffrin & Schneider, 1977) literatures, Reber claimed that unconscious knowledge could be flexibly deployed.

Contrary to Reber's claims, Experiment one demonstrated that knowledge of artificial grammars was actually relatively inflexible in its application, in that there was limited transfer between two domains with obvious mappings (i.e. colours and the names of those colours). In fact, the degree of inflexibility was about the same as that found in the implicit memory literature. Experiment two investigated the extent to which transferable or flexible knowledge was unconscious, defining consciousness in terms of metaknowledge. The results indicated that the knowledge had an unconscious component in both the same and different domains. In fact, consistent with Reber's claim, but contrary to the more traditional view, there was no evidence that the knowledge that could transfer had an explicit component. That is, the <u>flexible</u> knowledge appeared to be quite <u>unconscious</u>.

In summary, the knowledge that is incidentally acquired about artificial grammars is partly perceptually bound, but the remaining knowledge that can apply across domains is at least partly unconscious. Although unconscious knowledge has some flexibility, it remains possible, in keeping with the traditional view, that entirely conscious knowledge would be <u>more</u> flexible under some conditions - in that it might not show a decrement in transfer across domains where there is a transparent mapping (as in Experiment one). In contrast, it might show more of a decrement across domains in which there is a completely arbitrary mapping (as in Experiment two). In general, the current studies show that artificial grammar learning is a useful paradigm for exploring the links between consciousness and flexibility.

References

Altmann, G. T. M., Dienes, Z., & Goode, A. (1995). On the modality independence of implicitly learned grammatical knowledge. <u>Journal of Experimental Psychology: Learning, Memory, and Cognition</u>.

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

Baddeley, A. D. (1982). Domains of recollection. Psychological Review, 89, 708-729.

Bassili, J. N., Smith, M. C., & MacLeod, C. M. (1989). Auditory and visual word stem completion: Separating data-driven and conceptually driven processes. Quarterly Journal of Experimental Psychology, 41, 439-453.

Berry, D. C. and Dienes, Z. (1993). Implicit learning: Theoretical and empirical issues. Hove: Lawrence Erlbaum.

Blaxton, T. A. (1989). Investigating dissociations among memory measures: Support for a transfer-appropriate processing framework. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 657-668.

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.

Chan, C. (1992). Implicit cognitive processes: theoretical issues and applications in computer systems design. Unpublished D.Phil thesis, University of Oxford.

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

Cheesman, J., & Merikle, P. M. (1986). Distinguishing conscious from unconscious perceptual processes. Canadian Journal of Psychology, 40, 343-367.

Chi, M. T. H., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, 5, 121-152.

Davies, M. (1989). Tacit knowledge and subdoxastic states. In A. George (Ed.), Reflections on Chomsky (pp. 131-152). Oxford: Basil Blackwell.

Dienes, Z., Altmann, G., & Gao, S-J (1995). Mapping across domains without feedback: A neural network model of transfer of implicit knowledge. In L.S. Smith and P.J.B. Hancock (Eds), Neural Computation and Psychology: Proceedings of the 3rd Neural Computation and Psychology Workshop (NCPW3), Stirling, Scotland, 31 August - 2 September 1994, pp 19-33. Springer-Verlag: London.

Dienes, Z., Altmann, G., Kwan, L, Goode, A. (in press) Unconscious knowledge of artificial grammars is applied strategically. Journal of Experimental Psychology: Learning, Memory, & Cognition

Dienes, Z., & Perner, J. (1995) Implicit knowledge in people and connectionist networks. In G. Underwood (Ed), Implicit cognition, Oxford University Press.

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

Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. Cognitive Psychology, 12, 306-355.

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. Cognitive Psychology, 15, 1-38.

Godden, D., & Baddeley, A. D. (1975). Context dependent memory in two natural environments: On land and under water. British Journal of Psychology, 66, 325-331.

Gomez, R. L., & Schvaneveldt, R. W. (1994). What is learned from artificial grammars? Transfer tests of simple association. Journal of Experimental Psychology: Learning, Memory, and Cognition, 20, 396-410.

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 in 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.

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 Behaviour, 6, 855-863.

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

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

Redington, M., & Chater, N. (in press). Transfer in artificial grammar learning: A re-evaluation. Journal of Experimental Psychology: General

Roediger, H. L., & Blaxton, T. A. (1987). Effects of varying modality, surface features, and retention interval on priming in word-fragment completion. Memory & Cognition, 15, 379-388.

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

Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. Behavioural and Brain Sciences, 17, 367-448.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. Psychological Review, 84, 127-190.

Singley, M. K., & Anderson, J. R. (1989). The transfer of cognitive skill. Cambridge, MA: Harvard University 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. (1993). The role of 'explicit' variables on artificial grammar learning. Paper presented to the Third Annual Meeting of the Canadian Society for Brain, Behaviour, snd Cognitive Science and the Experimental Psychology Society, Toronto, July, 1993.

Author note

Preparation of this chapter was partly supported by a grant from the MRC. We would like to thank Alastair Goode for preparing stimuli and the writing the Turbopascal programs for experiments one and two, and also Elizabeth Ely, David Hoffbrand, Lee Hogarth, and Sairah Salim for running the subjects in experiments one and two.

Footnotes

1The total number of guess responses was no lower in the transfer rather than same domain conditions (in fact, 23 compared to 16, F(1,38) = 9.28, p = .004). Thus, it might be thought that the equivalent percentage correct performance in the transfer and same domain conditions indicated a complete transfer of implicit knowledge across domains. However, it may be that there was more implicit knowledge in the same domain condition than we measured but its expression was masked by the presence of explicit knowledge applying to an overlapping set of test strings.

Table 1 Experiment one: Classification performance

	Training group	
	Colours	Words
Testing:		
Same domain	67 (11)	67 (11)
Transfer domain	62 (11)	64 (14)

Note: Dependent variable is percentage correct classification. Standard deviations appear in parentheses.

Table 2 Experiment two: Classification performance

	Training group	
	Colours	Letters
Testing:		
Same domain	73 (11)	75 (11)
Transfer domain	58 (11)	55 (9)

Note: Dependent variable is percentage correct classification. Standard deviations appear in parentheses.

Table 3 Experiment two: Zero correlation criterion

	Training group	
	Colours	Letters
Testing:		
Same domain	5.5 (4.5)	5.5 (3.6)
Transfer domain	1.4 (3.1)	0.3 (4.0)

Note: Dependent variable is difference in confidence ratings between when subjects made a correct decision and when they made an incorrect decision. Standard deviations appear in parentheses.

Table 4 Experiment two: Guessing criterion

	Training group	
	Colours	Letters
Testing:		
Same domain	63 (19)	63 (24)
Transfer domain	63 (19)	60 (14)

Note: Dependent variable is percentage correct classification when subjects believed that they were literally guessing. Standard deviations appear in parentheses.

Appendix

Training items grammar A

XXRVM

VTTVM

VTTTVM

VTTTVTM

XXRTTVM

XXRTVTM

VTVTRVM

VTTVTRVM

XXRTTTVM

VTVTRVTM

VVTRTVTM

XMMXRTTVM

VTVTRTTVM

VTTVTRTVM

VVTRTTVTM

XMXRTTTVM

XXRTTTVTM

VTVTRTVTM

Grammar B

XXRRM

XMTRRM

XXRRRM

VVRMTM

XMVTRXM

VVRXRRM
XMVRXRM
XMVTRXRM
XMVTRMTM
VVRMTRRM
VVTTRMTM
VVTRXRRRM
XMVTTTRXM
XMVRXRRRM
XMVTTRXRM
VVTTTRMTM
VVTTRXRRM
VVRMVRXRM

Test items

Grammar A

VTTTVTM
VTTVTM
VTTTVTRVM
XXRVM
VTVTRVM
VVTRTTVTM
XMMMXRVTM
XXRTTVM
XMMXRVM
VTTVTRVM
XXRTTVTM
XXRTVTRVM

VVTRTTVTM

VTVTRTVM

XMXRTTTVM

VVTRTM

XMMXRTVM

XMXRTVM

XMXRVTRVM

VTVTM

XMMMXRTVM

XMMXRTVTM

XMXRTVM

XXRTTTVTM

VVTRVTRVM

XXRVTM

XXRVTRTVM

XXRVTRVM

VTTVTRTVM

Grammar B

VVTRMVRXM

VVTTTRXRM

VVTTRXRM

VVTRMTRRM

XMVTTRMTM

XXRRM

XMVTRXRRM

XMVRMTM

XMVTRMTRM

XMVRMTRM

VVRMVTRXM

VVRXRRRM

XMVRXRRM

XMVRMVRXM

VVTTRMTRM

VTRRRM

XMVRMTRRM

VVTRMTM

VVTTRXRRM

VTRRM

VVTTTRMTM

XMVRXM

VVRXRM

XMTRRRM

XMVTRXM

VVRMTRM

VVTRXRRM

XMVTTRXM

VVRMVRXRM