

The role of implicit memory in controlling a dynamic
system

Zoltán Dienes and Richard Fahey
University of Sussex, England

Running Head: IMPLICIT MEMORY AND SYSTEM CONTROL

Corresponding author: Zoltán Dienes, Experimental Psychology,
Sussex University, Brighton, Sussex, BN1 9QG. Fax: 01273
678611. E-mail: dienes@epunix.susx.ac.uk

Abstract

The relationship between implicit memory and implicit learning is explored. Dienes and Fahey (1995) showed that learning to control a dynamic system was mediated by a look-up table consisting of previously successful responses to specific situations. The experiment reported in this paper showed that facilitated performance on old situations was independent of the subjects' ability to recognize those situations as old, suggesting that memory was implicit. Further analyses of the Dienes and Fahey data replicated this independence of control performance on recognition. However, unlike the implicit memory revealed on fragment completion tasks, successful performance on the dynamic control tasks was remarkably resilient to modality shifts. The results are discussed in terms of models of implicit learning and the nature of implicit memory.

Acknowledgements

This research was supported by the Medical Research Council. Many thanks to Dianne Berry, Bob Mathews, David Shanks, Michael Stadler, and Daniel Willingham for valuable comments. Richard Fahey is now at Computer Data Systems Inc., 8626 Tesoro Drive, Suite 114, San Antonio, TX 78217, USA.

The role of implicit memory in controlling a dynamic system

Two distinctions in the human learning literature have become very influential: Implicit versus explicit memory (e.g., Roediger & McDermott, 1993; Schacter, 1987); and implicit versus explicit learning (e.g. Berry & Dienes, 1993; Reber, 1989; Shanks & St John, 1994). Implicit rather than explicit memory is shown on tasks that do not require deliberate recollection of a past event although the event influences performance. Implicit rather than explicit learning is shown on tasks for which the subject learns to make the right decision (more often or more quickly) without being able to justify the decision (Berry & Dienes, 1993). The inability to justify is meant to indicate that learning occurs without concurrent awareness of what is being learned (Reber, 1989). The distinctions between implicit and explicit memory and between implicit and explicit learning are logically orthogonal: For example, subjects may remember the episode in which they learned to make a right decision without being able to justify the decision. The purpose of this paper is to explore the relationship between implicit memory and the implicit learning of a complex system (Berry & Broadbent, 1984; Dienes & Fahey, 1995). Initially,

relevant evidence in the implicit memory literature will be described. Then, recent findings and theoretical proposals in the implicit learning literature that indicate a relationship with implicit memory will be overviewed. Next, an experiment will be reported that explores the role of implicit memory in controlling a dynamic system. Finally, the data collected by Dienes and Fahey (1995) will be re-analyzed in the same way to establish the generality of the findings.

Previous research has determined a number of characteristics of performance on implicit rather than explicit memory tasks. A typical implicit memory task is fragment completion: Subjects first study a list of words (e.g., AARDVARK), and are then shown fragments of words (e.g., ARVA) and the subjects are asked to complete the fragment with the first word that comes to mind. The fragments of words that have rather than have not been studied previously are more likely to be completed. This learning, or priming, has two important characteristics that can be used to assess its relevance to implicit learning.

First, the ability to complete the fragment (or perform some other implicit memory task) is stochastically independent of the ability to recall or recognize the word (e.g., Hayman & Tulving, 1989;

Tulving, 1985; Tulving & Schacter, 1990; Tulving, Schacter, & Stark, 1982). The interpretation of this finding is controversial. Hintzman and Hartry (1990) argued that the low level of dependence merely reflects many sources of variance affecting fragment completion in addition to priming; in contrast Tulving (e.g., Tulving & Schacter, 1990) argued that the low level of dependence reflects the operation of distinct memory systems underlying implicit and explicit memory tasks. The dispute is as yet unresolved (see Flexser, 1991; Gardiner, 1991; Hintzman, 1991; Poldrack, 1996). One solution to the problem of many sources of variance is to show that the data are nonsignificantly different from independence, and also that there was sufficient sensitivity to detect a meaningful amount of dependence (Ostergaard, 1992; Tulving & Hayman, 1993; Ostergaard, 1994 see Appendix A). This paper will adopt this solution.

The second characteristic is that amnesics can display unimpaired performance on implicit but not explicit memory tasks (e.g., Warrington & Weiskrantz, 1978). This finding suggests that implicit memory may be based on a different memory system to explicit memory (Roediger, 1990; Tulving, 1985; Tulving & Schacter, 1990; contrast Shanks, 1997). Finding a memory task that is spared in

amnesia would be suggestive evidence that the task relies on the memory system or systems underlying implicit memory tasks generally.

Turning now to the implicit learning literature, one key paradigm is the control of complex systems (e.g., Berry & Broadbent, 1984, 1987, 1988); other key paradigms are artificial grammar learning (e.g., Reber, 1989) and sequential reaction time tasks (e.g., Willingham, Nissen, & Bullemer, 1989). In the "dynamic control tasks" used by Berry and Broadbent, the subject controls the level of one variable (e.g., their friendliness towards a computer personality) in order to reach target values on another variable (e.g., the computer personality's friendliness to the subject). Subjects acquire considerable knowledge about how to control such systems, as indicated by their progressive ability to reach and maintain target values. This knowledge appears to be implicit because subjects find it difficult to describe how to reach and maintain target values (Stanley, Mathews, Buss, & Kotler-Cope, 1989). Stanley et al (1989) asked subjects after every 10 trial block to give complete instructions on how to perform the task. The informativeness of these instructions was assessed by the performance of yoked subjects asked to follow the transcribed instructions. Stanley et

al demonstrated that sudden improvements in performance by the original learners were not associated with simultaneous increases in the informativeness of the instructions. In fact, instructions helped the performance of yoked subjects only if the instructions were taken at least four blocks after the improvement in performance. Note that there are other types of task for which the instructions given by subjects immediately and completely account for their performance (e.g., Mathews, Buss, Chinn, & Stanley, 1988; Schwartz, 1966).

Recent theoretical views on implicit learning suggest a relationship with implicit memory. Both Broadbent, Fitzgerald, and Broadbent (1986) and Stanley et al. (1989) argued for a mechanism for the implicit learning of the control tasks that could rely on implicit memory. They suggested that whereas explicit knowledge is based on a mental model of the system, implicit knowledge is based on memory for specific events related to the control task. This memory-based view of implicit learning is consistent with theoretical proposals in other domains of implicit learning (see, e.g., Brooks, 1978; Cho & Mathews, 1996; Neal & Hesketh, 1997; Perruchet, 1994; Vokey & Brooks, 1992; Whittlesea & Dorken, 1993), but the issue has aroused some

controversy (e.g. Reber, 1967, 1989). At stake is the question of how sophisticated unconscious processing is. Reber argued that implicit learning results in an abstract knowledge base that goes well beyond storage of specific exemplars. Cleeremans (1993) showed how the successful computational models of implicit learning of sequential reaction time tasks could produce knowledge that lay along a continuum of abstractness depending on task conditions: in this case, the mechanism underlying implicit learning tasks can go beyond memories for instances but it falls short of formulating explicit rules. Following Broadbent et al, we have argued (Dienes and Fahey, 1995) that, for the dynamic control tasks, one does not need to postulate a mechanism any more complicated than memory for specific events.

Broadbent et al. (1986) suggested that in learning a dynamic control task a subject could construct a "look-up table" which would determine the appropriate action by matching the current situation to the most similar of the entries already in the table. Using methodology introduced by Marescaux, Luc, and Karnas (1989), Dienes and Fahey (1995) tested and confirmed several predictions of the look-up table approach in two experiments. In Experiment 1, subjects were trained on a simulated

sugar production factory, in which they manipulated the level of work force in order to reach target values of sugar production. After training, subjects were given a specific situation task in which the subject was presented with hypothetical situations (e.g., "if you had just employed 400 workers and if the sugar production was then 8000 tons, what should you do next to bring the sugar production to target?"). Some of the situations were new (i.e. the subject had not experienced them in the training phase). Some of the situations were old (i.e. the subject had experienced the situation in training) and the subject had given a correct response to them (i.e. a response that lead to target or closeby) or an incorrect response; these old situations were called 'correct situations' and 'incorrect situations', respectively. (Note that calling a situation correct refers to whether the response given in the context of the situation was correct in the training phase, not to whether the situation contained the target as a feature.) Crucially, subjects' experience in controlling a dynamic system did not allow subjects to perform above chance on new situations; their knowledge seemed only to apply to old correct situations. Further, subjects tended to give the same response to old correct situations as they had given

originally; the tendency to give the same response to old incorrect situations was smaller. These results were replicated in Experiment 2 with a person interaction task in which subjects manipulated the friendliness of a computer personality. Further, detailed aspects of subjects' performance in both experiments could be fitted with a one-parameter computational model of a look-up table (Logan, 1988).

Given that control of these tasks appears to be mediated by memory of specific instances, one can ask whether the memory is the sort that typically underlies performance on explicit memory tasks or if it is of the sort that underlies performance on implicit memory tasks. Evidence consistent with the latter alternative was provided by Squire and Zola-Morgan (1991): They found that amnesics were not initially impaired on the dynamic control tasks, just as Warrington and Weiskrantz (1978), for example, found that amnesics were not impaired on implicit memory tasks. More direct evidence for normal subjects can be provided by the data collected by Dienes and Fahey (1995). They presented subjects with a recognition test on the same situations that had been used on the specific situations task, allowing a test of stochastic independence between recognition and performance on

the specific situations task. Dienes and Fahey did not test the relation between recognition and control performance: This paper will do so. First a new experiment will be described that tests the relation between recognition and control performance, and then a reanalysis of the Dienes and Fahey data will be reported. If there is a look-up table based on explicit conscious recollection of previous situations, there should be a dependence between performance on the specific situations task and recognition of the situation as old; if it's based on implicit memory there should be no dependence.

The dependency of the subjects' tendency to respond correctly to a given situation on explicit memory for that situation is important in assessing whether the control tasks are learned unconsciously. If correct responding is dependent on the subjects' ability to recall the previously successful episode, and subjects could justify their responding in terms of the episodic memory, the claim for unconscious learning (Hayes & Broadbent, 1988; Stanley et al., 1989; contrast Sanderson, 1989; Shanks & St John, 1994) would need reconsidering.

Experiment 1

In Experiment 1 subjects were trained on the sugar production factory task and then presented

with specific situations from the training phase. Each subject was asked to both give a control response to the situation (specific situations task) and to indicate whether they recognized the situation as old or not (recognition task). The primary aim of the experiment is to assess whether the look-up table used by subjects in controlling the sugar production task was based on implicit memory. This was achieved by assessing (1) the degree by which the tasks differed from stochastic independence; and (2) the degree to which the tasks differed from a plausible model of dependence (the Ostergaard, 1992, model, described in the results section below and in more detail in Appendix A, and also endorsed by Poldrack, 1996).

In order to conclude that subjects used an implicit look-up table, the following three points need to be established:

- A. Subjects learnt the task by acquiring simply a look-up table. Dienes and Fahey (1995) provided the evidence for this claim.
- B. The tasks do not differ significantly from stochastic independence.
- C. The tasks do differ significantly from the model of dependence.

This paper will attempt to establish the validity of points B and C. If B is established but

not C (i.e. the tasks do not differ significantly from either stochastic independence or from the model of dependence) then the data are consistent with subjects using either an implicit or an explicit look-up table, or some combination of the two. If the data are significantly different from independence but not significantly different from the model of dependence then the data suggest subjects are using an explicit look-up table.

Method

Subjects. The subjects were 20 paid volunteers aged between 18 and 35 from the Sussex University subject panel.

The task. Subjects in Experiment 1 were trained on the sugar production task introduced by Berry and Broadbent (1984). Subjects were asked to imagine they were in charge of a sugar production factory. They were told they could change the amount of sugar produced by changing the size of the work force. Their goal was to achieve and maintain a target sugar output of 9,000 tons. The starting work force was 600 workers and starting level of sugar output was 6,000 tons. On each trial, subjects entered a number between 1 and 12 on the computer keyboard to represent the number of hundreds of workers they wished to employ on that trial. The level of sugar

production on trial n was determined by the equation $P_n = 2*W - P_{n-1} + N$, where P_n is the number of thousands of tons of sugar output on trial n , W is the number of hundreds of workers employed by the subject, and N is noise (N could be -1 , 0 , or $+1$ with equal probability). If the equation resulted in a sugar output of less than 1,000 tons, the output was simply set at 1,000 tons; similarly, if the equation resulted in an output of greater than 12,000 tons, it was set at 12,000 tons. Subjects were aware of these lower and upper limits. Because of the noise, N , in the equation, subjects' responses were counted as correct if they resulted in an output on target or one level off.

Optimal performance on this task requires the subject to take account of the current level of sugar production. In fact, W_{n+1} should be exactly halfway between P_n and target in order to reach the target on the next trial.

Procedure. All subjects were trained for two blocks of 40 trials of the sugar production factory. Each block started with a sugar production of 6,000 tons and a work force of 600 workers. On each trial, subjects saw a graph indicating the level of performance on all previous trials of that set. A horizontal line indicated the target performance. In addition, written information about the level of

work force and the level of sugar production for the last trial was presented above the graph (see Figure 1).

Insert figure 1 about here

The computer kept a record of all situations the subject came across. A situation was defined as the current level of sugar production. The situations were tabulated into those for which the subject entered a workforce that resulted in the target sugar output or only one level off (loosely correct situations), and those that were followed by an output more than one level from target (incorrect situations). Note that calling a situation correct refers to whether the response given in the context of the situation was correct, not to whether the situation contained the target as a feature.

Next subjects were exposed to specific situations on the computer and performed either the specific situations task introduced by Marescaux et al (1989) or a recognition task. In both tasks subjects were presented with a level of sugar production: The statement 'The current level of sugar production is X tons' appeared in the centre of the screen in an identical way for both tasks. All 12 levels of sugar production were presented in

a different random order for each task. In the specific situations task, subjects were told to enter the level of work force they thought would achieve or maintain the target level of output, based on their previous experience of controlling the factory. The same target level was used as in the training phase. Subjects were told that after each situation, the next situation to be shown would be unrelated to the work force they had just entered; it would simply be another possible situation, and thus, they would get no feedback on how successful they were being. In the recognition task, subjects were told that some of the situations will have been ones they came across in their first task. Subjects responded "old" or "new" to each situation depending on whether they could remember seeing it in the training phase. The specific situations and recognition tasks were presented in counterbalanced order.

Results

Order had no effect on any measure (p s > .20), including the measures of dependence and independence, so this factor will not be discussed further. Not all subjects produced data for which all dependent variables could be calculated (for example, if the subject had experienced all levels of sugar production in the learning phase there

would be no new situations in the later phases, and so the Ostergaard (1992) measure could not be calculated for that subject), and so the degrees of freedom of the subsequent analyses vary slightly as a consequence.

Learning. The number of trials that were loosely correct improved from the first block of trials (8.2, SD = 4.9) to the second block of trials (13.0, SD = 7.5), $t(19) = 2.85$, $p = .01$.

Specific situations task. Reliable priming was obtained: Subjects performed better on old levels of sugar production (.26, SD = .20) rather than on new levels (.07, SD = .27), Wilcoxon $p = .03$.

Recognition. Mean recognition (64%, SD = 26%) was significantly above chance, $t(19) = 2.46$, $p = .02$.

Stochastic dependence. The contingency table for the specific situations and recognition tasks is shown in table 1. The entries were summed over all subjects and correct situations. The proportion of situations expected to be both correctly recognized and correctly responded to on the specific situations task based on assumed independence between the tasks was separately calculated for each subject. The difference between the actual proportion and that expected on the basis of independence was, averaged over subjects, .001 (SD = .062), which is non-significantly different from

zero, $\underline{t} < 1$, analysing over subjects. That is, there was no evidence that giving a correct response to a situation in the specific situations task depended on recognizing the situation.

A null result is only compelling to the extent that the data were sensitive enough to detect a reasonable effect size. Appendix A describes a model based on Ostergaard (1992) that can be used to estimate the amount of dependence expected if learning consisted of a look-up table based on explicit memory. The logic of the model is as follows. Subjects can to some extent give correct answers to each task (specific situations task and recognition task) not because they have learnt to respond to specific situations but because of the pre-existing strategies they bring to the task or because they simply guess. These additional sources of variance increase the expected amount of independence between the tasks; it is only to the extent that learning has occurred that dependence can emerge between the tasks. To take into account these additional sources of variance, the proportion of correct responses on new situations was subtracted from the proportion of correct responses on old situations to estimate the 'study effect' of each task; i.e. the effect that exposure to situations has on ability to perform the task above

and beyond what subjects can do without learning. It was assumed that subjects could recognize all the situations that contributed to the study effect of the specific situations task (this is the assumption of an explicit look-up table); thus the proportion of situations that subjects performed correctly on both tasks will be the study effect for the specific situations task, plus an additional component due to correct performance on both tasks for independent reasons. Thus, if the study effect for the specific situations task were zero, the Ostergaard (1992) model would make the same prediction as the model of independence. As more learning occurs on the specific situations task, the more the Ostergaard (1992) model departs from the model of independence. For more explanation see Appendix A. The difference between the actual proportion correct on both tasks and that expected on the basis of Ostergaard's model of dependence was, averaged over subjects, $-.11$ ($SD = .15$), $t(13) = 2.74$, $p = .017$. That is, the dependence between the tasks was less than would be expected if correct performance on the specific situations task was based on recognition of the same situations (see Appendix B for a further analysis showing the robustness of this result). One measure of the effect size of the difference of subjects' data from the model of dependence is Cohen's (1988)

d , defined as the mean difference divided by the standard deviation, i.e. $.11/.15 = 0.73$, close to what Cohen calls a large effect ($d = 0.80$).

The key result of Experiment 1 was that performance on the specific situations task did not depend on the subjects' ability to recognize old situations as old. The analysis of Experiment 2 of Dienes and Fahey (1995), presented next, tests the generality of this finding.²

Experiment 2 of Dienes and Fahey (1995)

Experiment 1 used a simulated sugar production factory as the control task. Another commonly used task is the person interaction task. Berry and Broadbent (1988) and Hayes and Broadbent (1988) employed two versions of a person interaction task that they claimed were learned in distinctive implicit and explicit modes. Experiment 2 reported by Dienes and Fahey (1995) employed both versions of the person interaction task. The major purpose of this analysis of Experiment 2 was to determine the relationship between recognition and performance on the specific situations task for the two versions of the person interaction task.

A secondary aim of this analysis of Experiment 2 was to determine if there was any effect of modality shift on control task performance. One of the characteristics of implicit rather than explicit

memory is its frequent sensitivity to changes in surface features, and in particular to shifts between the auditory and visual modalities (e.g. Bassili, Smith, & MacLeod, 1989; Berry, Banbury, & Henry, 1997; Roediger & Blaxton, 1987). Dienes and Berry (1997) also argued that the knowledge acquired in a number of implicit learning paradigms is inflexible and perceptually bound. For example, in the artificial grammar learning paradigm, subjects trained on colour patches show a decrement in performance when tested on colour names, and vice versa (Dienes & Altmann, 1997). Berry and Broadbent (1988) found there was no transfer between two control tasks with the same underlying equation when the cover story changed (e.g. from being a transport task to a person interaction task), even when subjects were informed of the critical relationship between the tasks.

Schacter (1990) presented one interpretation of the finding of modality specificity in implicit memory tasks. According to Schacter, performance on implicit memory tasks is mediated by a set of modular perceptual representation systems (PRSs) that process information about the form and structure of objects but do not represent associative information about them. Different types of objects have their own PRS; for example, there is

a word form system underlying performance on fragment completion, and a structural description system underlying the form processing of common visual objects. According to Schacter, priming on many implicit memory tasks reflects the establishment of highly specific representations within an appropriate PRS; it is these representations that underlie the sensitivity of priming to modality shifts. However, not all implicit tasks rely on a PRS. For example, in conceptual priming, after studying the word 'limerick', subjects may be given a conceptual cue ('What name is given to a lighthearted five line poem?'). The priming obtained in this paradigm is stochastically independent of recognition but also insensitive to modality shifts (Challis, Chiu, Kerr, Law, Schneider, and Yonelinas, 1993). Priming on such tasks may be due to associative links being formed in a semantic memory system, and not in a PRS. Variability in performance on the dynamic control tasks is not a measure of variability in the subject's perception of the structure of any type of object. Perceiving the situation is trivial in the dynamic control tasks; the problem is purely in learning the right response to a given situation, that is, in learning the associative links. Thus, Schacter's (1990) model does not predict any effect

of modality shift for the dynamic control tasks, in contrast to some other implicit tasks (e.g. Servan-Schreiber & Anderson, 1990, regarded learning in the artificial grammar learning paradigm as an example of perceptual learning). The current reanalysis of Experiment 2 tested whether knowledge acquired about the person interaction tasks behaved like stem completion and artificial grammar learning (perceptually bound) or like conceptual priming (not perceptually bound).

Method

In Experiment 2 of Dienes and Fahey (1995), each of 48 subjects was randomly allocated to one of the cells of a 2 X 2 X 2 (Person [Person S vs Person U] by learning modality [visual vs auditory] by testing modality [visual vs auditory]) between subjects design. Equal numbers of subjects were allocated to each of the cells. Subjects were told that they would be meeting a computer person called Ellis, and that they could communicate their level of friendliness to Ellis through the keyboard. Ellis would respond with how friendly he is to the subject. The aim of subjects was to move Ellis to a target value of friendliness and keep him there. Friendliness varied along a 12 point scale: Very Rude, Rude, Very Cool, Cool, Indifferent, Polite,

Very Polite, Friendly, Very Friendly, Affectionate, Very Affectionate, and Loving.

Berry and Broadbent (1988) and Hayes and Broadbent (1988) employed two equations for controlling Ellis' behaviour, a "salient" equation and a "nonsalient" equation. The salient equation was $E_n = S_n - 2 + N$, where E_n is a number between 1 and 12 representing Ellis' behaviour on the 12 point scale on trial n , S_n is the subject's behaviour on the 12 point scale on that trial, and N is noise (-1, 0, or +1 with equiprobability). The nonsalient equation was $E_n = S_{n-1} - 2 + N$. Because of the noise, a response of Ellis up to one off target was counted as correct. Hayes and Broadbent (1988) called the personality controlled by the salient equation "Person S" and that controlled by nonsalient equation "Person U". Note that for the nonsalient equation Ellis' behaviour depends not on how the subject just responded on that trial, but on how the subject responded one trial back. The target in Experiment 2 was Polite for both Person S and Person U, and so the optimal strategy when interacting with either would be to always enter Friendly.

In the learning phase, subjects interacted with Ellis for one block of 30 trials for Person S and one block of 50 trials for Person U. Subjects in

the visual learning condition saw two bar charts, one representing Ellis' behaviour on the previous trial, and the other their behaviour on the previous trial. A horizontal line in Ellis' chart indicated the target behaviour (Polite). The bars moved up and down according to Ellis' and the subject's respective behaviours (see Figure 1). Subjects in the auditory learning condition did not see the computer screen. The experimenter simply said "Your behaviour was X; Ellis' behaviour was X". The scale of possible behaviours was placed in front of the subjects to remind them. All subjects entered their response by typing in the corresponding initials (e.g., VA for Very Affectionate). The computer kept a record of all situations the subject came across. A situation was defined as what the subject could see, i.e. the current level of Ellis' behaviour and the subject's behaviour on the last trial.

Next subjects were exposed to specific situations and performed either the specific situations task or the recognition task on the computer. In the specific situations task, subjects were shown possible situations consisting of Ellis' and the subject's behaviours on the preceding trial. Subjects in the visual testing condition saw the information displayed as bar charts; subjects in the auditory testing condition heard the experimenter

read the information out. Subjects were told to enter the behaviour they thought would achieve or maintain the target level of Ellis' behaviour, based on their previous experience interacting with Ellis. The target level was the same as used in the training phase. Subjects were told that after each situation, the next situation to be shown would be unrelated to the behaviour they had just entered; it would simply be another possible situation, and thus, they would get no feedback on how successful they were being. In the recognition task, subjects were told that half the situations will have been ones they came across in their first task, and half would be new situations. Subjects responded "old" or "new" to each situation. The old situations for both tasks were all situations experienced in the training trials.

The specific situations and recognition tasks were run in counterbalanced order using the design employed by Tulving, Schacter, and Stark (1982). Specifically, for each subject, the computer randomly generated a number of new situations equal to the number of old (i.e. half the situations were not "study primed" - the new situations - and half were study primed - the old situations). Half of the new situations and half of the old situations were randomly assigned to one stimuli set (set A),

the other half to set B. Half the subjects were run, first, on the specific situations task with set A stimuli; then the recognition task with sets A and B; and finally, the situations task with set B. The other half of the subjects were run on, first, the recognition task with set A stimuli; then the specific situations task with sets A and B; and finally, the recognition task with set B. (Sets A and B constituted different situations for different subjects.)

Results

Learning. Subjects were scored for the number of trials correct in the first set of 10 trials and in the last set of 10 trials. A 2 X 2 X 2 X 2 (Block [first set of 10 trials vs last set of 10 trials] by Person [Person S vs Person U] by learning modality [visual vs auditory] by testing modality [visual vs auditory]) analysis of variance on the number of trials correct indicated a significant effect of person, $F(1,40) = 10.47$, $p < .005$, and of block, $F(1,38) = 14.29$, $p = .0005$. That is, subjects scored more trials correct for Person S (6.5) than for Person U (4.9). Also, subjects performed better on the second block (6.5) rather than the first (5.0). The interaction of block with person was not significant, $p > .10$. The improvement from the first to the second block was 1.3 for Person U

($t(23) = 2.65$, $p < .05$) and 1.7 ($t(23) = 2.92$, $p < .01$) for Person S. Specific Situations Task. A response was scored as correct on the specific situations task if it would lead to target, or at most one level off, at least two thirds of the time. A 2 X 2 X 2 X 2 (person [Person S vs Person U] by learning modality [visual vs auditory] by testing modality [visual vs auditory] by priming [old versus new]) mixed model analysis of variance on proportion of correct responses indicated a significant effect for person, $F(1,40) = 13.62$, $p < .001$, and priming, $F(1,40) = 50.62$, $p < .0001$. That is, subjects performed better with Person S (62%) rather than Person U (43%). Also, subjects performed better on situations which they had (63%) rather than had not seen before (41%) in the person interaction task.

An important result is that a change in modality between learning and testing had no effect on performance. Table 2 shows the means for Persons S and U for subjects in the same and different modalities. In fact, subjects performed numerically better in the different rather than the same modality for both Persons S and U. Could the experiment have suffered from a lack of power to detect the effect? In the case of implicit memory as revealed by stem completion, Bassili et al. (1989) found that priming reduced from 32% in the

same modality groups to 20% in the different modality groups (averaged across their experiments one and two); i.e., a proportional reduction of 37%. Brown, Neblett, Jones, & Mitchell (1991) reviewed the effect of modality shift on implicit memory in 10 other studies, and in all cases the proportional reduction was greater than 37% (their own study was an exception). The power of the current experiment to detect a proportional reduction in performance of 37% in both Persons S and U was greater than .99; the power to detect such a reduction in Person U alone was .89; and the power to detect an interaction in which the drop occurs for Person U but not Person S was .65.

Recognition. Overall recognition performance was 74%, significantly greater than chance, $t(40) = 16.32$, $p < .0001$. A 2 X 2 X 2 (person [Person S vs Person U] by learning modality [visual vs auditory] by testing modality [visual vs auditory]) analysis of variance indicated a significant effect of person, $F(1,40) = 6.33$, $p < .05$. That is, subjects had a greater recognition performance with Person S (78%) rather than Person U (70%).

Stochastic Dependence. The contingency table for the specific situations and recognition tasks is shown in table 3 for Persons S and U. The entries were summed over all subjects and correct

situations. The proportion of situations expected to be both correctly recognized and correctly responded to on the specific situations task based on assumed independence between the tasks was separately calculated for each subject. The difference between the actual proportion and that expected on the basis of independence was, averaged over subjects, $-.00$, for Person U, and $+.01$ for Person S, both non-significantly different from zero, $t_s < 1$, analysing over subjects. That is, there was no evidence that giving a correct response to a situation in the specific situations task depended on recognizing the situation.

The difference between the actual proportion correct on both tasks and that expected on the basis of Ostergaard's model of dependence was, averaged over subjects, $-.10$, $t(20) = 3.48$, $p < .01$, $d = 0.76$, for Person U, and $-.09$, $t(23) = 4.65$, $p < .001$, $d = 0.95$, for Person S. That is, the dependence between the tasks was less than would be expected if correct performance on the specific situations task was based on recognition of the situations (see Appendix B for further analyses showing the robustness of this result for Person U).

Discussion

The aims of this analysis of Experiment 2 of Dienes and Fahey (1995) were, first, to determine

the relationship between recognition and performance on the specific situations task, and, second, to determine whether the dynamic control tasks are sensitive to modality shifts. In terms of the first aim, the results indicated that subjects' ability to respond to a situation in the specific situations task did not depend on their ability to recognize that situation as old, for both Persons S and U. In the case of Person U, Dienes and Fahey provided evidence that subjects learned by storing appropriate responses to specific situations: subjects' did not perform at above baseline levels on new situations and further details of subjects' performance could be fit by a one-parameter look-up table model.³ The amount of dependence was less than that predicted by Ostergaards' model, suggesting that performance on the specific situations task was based on implicit memory.

Person S also showed independence between task performance and recognition. Dienes and Fahey (1995) provided evidence that subjects learnt Person S (unlike Person U) by inducing rules that generalize beyond specific situations; specifically, Dienes and Fahey showed that subjects could perform well on new situations and that subjects' performance could be fit well by a rule-based model. Once a subject has learnt a rule, there is no reason

to suppose that its successful application depends on recognition of old situations. Thus, the lack of dependence in the case of Person S does not indicate that subjects acquired an implicit look-up table.

The results showed that a modality shift (from visual to auditory presentation or vice versa) had no influence on either Person U or Person S⁴. In terms of Schacter's (1990) model, this result suggests that variability in implicit learning on the person interaction task does not reflect variability in the perceptual representation systems underlying performance on standard implicit memory tasks. Rather learning the person interaction tasks involves forming associative links between easily perceived situations and appropriate responses. Berry (1991) did find a remarkable inflexibility in transfer of subjects' knowledge of this link from situations to responses: Subjects could not transfer their knowledge gained in the context of making a verbal response on a dynamic control task to make a typing response. Further, Roediger (1990) and Tulving and Schacter (1990) argued that implicit tasks are often sensitive to modality shifts not by virtue of their implicitness but because they predominantly involve perceptual processes. According to Tulving and Schacter, implicit priming effects can occur on conceptual tasks that do not

involve changes in a perceptual representation system but modifications to semantic memory. For example, both normals and amnesics show priming in a task in which subjects are given the name of a category and are asked to produce the first instance that comes to mind. Although the knowledge acquired in learning to control a complex system may be inflexible over some surface changes (e.g. the change in cover story used by Berry and Broadbent, 1988; the change in response requirements used by Berry, 1991), it may be less perceptually bound than the knowledge acquired in artificial grammar learning tasks (Dienes & Altmann, 1997).

Discussion

This paper has reported data from two experiments showing that, when controlling a dynamic system (the sugar production factory and person U), being able to respond successfully to specific situations does not depend on recollective experience. Rather, the subject appears to know the correct response because of implicitly formed associations: implicit learning in these cases is based on implicit memory.

The conclusion that implicit learning is based on implicit memory, depends on the claim that subjects used almost exclusively a look-up table in controlling the sugar production factory and person

U. If subjects used abstract rules then there is no reason why performance on the specific situations task should be dependent on performance on the recognition task. Dienes and Fahey (1995) argued that control of the sugar production task and of person U was indeed based almost exclusively on a look-up table for the amount of training used in the experiments reported in this paper. Specifically, Dienes and Fahey found that subjects performed at baseline levels on new situations; and they also showed that a one-parameter look-up table model could simulate levels of performance, subjects consistency of responding in different situations, and rate of learning.

Nosofsky (1988, 1991) showed that previous apparent dissociations between recognition and classification performance were perfectly consistent with one memory system underlying both tasks. Specifically, he showed that if classification is based on the relative similarity of the test exemplar to stored exemplars in the different categories, and recognition is based on the total similarity of the test exemplar to all stored exemplars, there can be a zero correlation between recognition and classification performance. This dissociation resulted from the different decision rules for recognition and classification: that is,

total compared to relative similarity. Could there be a similar explanation for the dissociation between recognition and dynamic control task performance found in this paper? Dienes and Fahey (1995) argued that learning to respond appropriately in a specific situation involves remembering that specific situation, or a very similar one. Memory for a specific situation is what the recognition test was supposed to directly measure. But it may be that subjects, when making a recognition judgement, relied on some overall similarity of the test situation to all training situations. To test this possibility in Experiment 1, the similarity was computed between each to-be-recognized situation and each training situation. Of course, any test depends on a model of how similarity is computed and used by subjects. Initially, for simplicity, similarity was measured on a linear scale that varied between 1 and 12:

$$\text{Similarity} = 12 - |\text{sugar production in test situation} - \text{sugar production in training situation}|.$$

For each test situation, similarity was summed over all training situations. For each subject, test situations were grouped into those that contributed to hits and those that contributed to misses on the recognition test; similarity was averaged separately for hits and misses. If

similarity determined recognition judgements, then it should be higher for hits than for misses. Indeed, the mean similarity was 675 (SD = 40) for hits and 623 (SD=79) for misses, $t(17) = 3.25$, $p = .005$, analysing over subjects.

These data are consistent with subjects using a summed similarity strategy, but also could arise for other reasons. For example, similarity probably covaried with how often a test situation occurred in the training phase (we will call this the frequency of a situation). On almost any theory, frequency would affect probability of recollecting a specific situation⁵. Performance on the specific situations task was based not only on memory for the exact situation, but also on situations numerically adjacent (Dienes & Fahey, 1995). It may be that subjects' recollection of situations confused numerically adjacent situations⁶. Frequency was computed by counting the number of times either the exact test situation occurred in the training phase, or another situation that was one only one level of sugar production different. The frequency for hits (21.1, SD = 6.2) exceeded that for misses (16.6, SD = 3.4), $t(17) = 2.92$, $p = .010$. When similarity difference was regressed against this frequency difference, the intercept (21) was not significantly different from zero, $t(16) = 1.34$, $p = 0.20$. That

is, as far as we can tell, the similarity of a test situation to training situations more than one level of sugar production different did not influence subjects' tendency to say 'old' on the recognition test.

In summary, performance on both the specific situations task and the recognition task were insensitive to situations that were more than one level of sugar production different. Thus, our response to the Nosofsky (1988, 1991) type challenge is that both the specific situations task and the recognition tasks seemed to test memory for specific situations in the same way; there did not appear to be different decision criteria. The analyses of dependence between the two tasks showed they were independent even when recognition of a situation only one level of sugar production different was treated as a correct recognition of the original situation (Appendix B). We suggest that this independence arose because the two tasks tapped different memory systems⁷. If both tasks depended on familiarity, it is unclear why they should be stochastically independent. Thus, we suggest that the recognition task relied on recollective experience (Gardiner & Parkin, 1990; Tulving, 1985). Subjects were told on the recognition test to say 'old' if they remembered seeing the situation

before; that is, if they had a recollective experience and not just a sense of familiarity. On the other hand, performance on the specific situations task may have relied on implicit mechanisms that give rise to a sense of familiarity. Future research could test whether encouraging subjects to respond on the recognition task according to familiarity would produce correlations between the two tasks.

One point of debate in the implicit learning literature has been whether subjects acquire unconscious or conscious knowledge (Berry & Dienes, 1993; Hayes & Broadbent, 1988; Reber, 1989; Lewicki, 1986; Shanks & St John, 1994). Our results indicate that in the case of the dynamic control tasks the knowledge is conscious in the sense used by Shanks & St John: Given a specific situation, the subject can say what the appropriate response should be. That is, knowledge of the link between situation and response is conscious. On the other hand, the knowledge is unconscious in the sense that subjects do not know why they know the correct response (this is what Shanks & St John call implicit retrieval): They do not know that it is because they have been in that specific situation before. There are also, of course, other senses in which knowledge can be unconscious or conscious (see Chan, 1992; Dienes,

Altmann, Kwan, & Goode, 1995; Dienes & Berry, 1997; Dienes & Perner, 1996) which these results do not address: future research could usefully do so.

Finally, we note that the results reported in this paper may be bounded by the amount of learning subjects receive. With more extended learning periods, subjects acquire greater amounts of explicit knowledge (Stanley et. al., 1989; Squire & Frombach, 1990). Subjects learn general rules, but they may also come to explicitly remember specific situations. Further, the implicit learning itself may become less well approximated by a look-up table as learning proceeds, and subjects may progressively extrapolate and interpolate to greater degrees around the situations trained on. This is also a matter for future research to address.

Appendix A

Ostergaard's (1992) model of dependence

Ostergaard argued that in the implicit memory literature previous studies had not considered a model of what the greatest degree of dependence could be between two tasks. He provided a model of dependence between memory tests that takes into account the variance not related to the study episode, as a way of assessing the greatest degree of dependence that could be expected between two tasks. Showing that the dependence between two tasks is nonsignificant is only informative if the degree of dependence is less than the amount that could be expected. We can use the Ostergaard model to determine the degree of dependence that would be expected if correct responding on the specific situations task was based on explicit recognition. The effect of study on each task is calculated as (proportion of positive responses to studied items) - (proportion of positive responses to nonstudied items); or, in symbols, $P(S+) - P(N+)$. 'Positive response' refers to a correct response on the specific situations and to a 'yes' response on the recognition task. Call the task with the smallest study effect task 1; the study effect for this task can be represented as $P(S1+ \text{ from study}) = P(S1+) - P(N1+)$. If all the traces of items accessible on

task 1 are also accessible on task 2, then the $P(S1+$ from study) of the study items given a positive response on task 1 will also be given a positive response on task 2. That is, in our case, it is assumed that the subject can recognize the situations that contributed to the study effect of the specific situations task; this is the assumption of an explicit look up table. There will be other situations that subjects can give a correct response to in the specific situations task, however, because of chance, or the use of a pre-existing explicit strategy that is independent of learning. The proportion of such situations can be estimated by $P(N1+)$. Some of these situations will be recognized but could not form the basis of an explicit look-up table; for example, they may be recognized as old, but the subject could not remember the response to them. The proportion of situations jointly correct on both tasks for independent reasons can be estimated by $P(N1+) \times P(S2+)$. In total, the proportion of positive responses given to both tasks is $P(S1+) - P(N1+) + P(N1+) \times P(S2+)$. Thus, the proportion of situations responded to correctly on both tasks can be predicted by the model to be the smaller study effect plus an additional chance component.

This calculation assumes that the study effect for task 1 is constant across the possibly different contexts and instructions of the two tasks. This is a reasonable assumption for our application: Situations were presented to subjects in exactly the same detail on the specific situations and recognition tasks. Thus, we could reasonably expect the degree of dependence between the specific situations and recognition tasks predicted by the Ostergaard model if performance on the specific situations task relied on explicit recognition.

In this paper, degree of dependence was calculated separately for each subject and analyzed over subjects. This is so that a Simpson's paradox could not be present in the data due to collapsing over subjects (but it could of course arise due to other reasons). We cannot discount all possible artifactual reasons for why two tasks may fail to correlate - there can always be unknown covariates, but we can discount various plausible reasons, and therefore make a plausible case that the demonstrated independence is theoretically informative. Making plausible cases is all we can do as scientists in any situation.

The Ostergaard model takes into account the fact that the effect of study could be smaller for one task rather than the other, and the level of

dependence is calculated accordingly. For example, If subjects use an explicit look up table, they would have to (1) recognize the current situation as an old one or not; (2) if it is old, recall whether the target followed the situation or not; and (3), if it did, recall the response given to the situation. The recognition test used in this paper only assessed the first component of this process, and thus only a subset of the situations recognized as old would be given a correct response in the specific situations task. This just means that the effect size on the recognition task should be larger than for the specific situations task. Consistent with the assumption of an explicit look-up table, we assumed that the population effect of study would be greater for the recognition task than for the specific situations task for each subject. The calculations were conducted accordingly: In the expression used to predict the proportion of positive responses to both tasks [i.e. $P(S1+) - P(N1+) + P(N1+) \times P(S2+)$], task 1 (i.e. the task meant to have the smallest effect of study) was always the specific situations task.

Because the Ostergaard analysis was calculated separately for each subject, the expression was often determined for cells where there were small amounts of data. The effect of small numbers of

observations is always to reduce power. In terms of Type I errors, note that the expected value of the Ostergaard estimate of proportion of positive responses on both tasks can be obtained by substituting the expected value of each of its terms in the expression. (This statement assumes that the proportion of positive responses to nonstudied items on task 1 is largely independent of the proportion of positive responses to studied items on task 2 for each subject, because these two proportions are involved in a product term. This is a reasonable assumption and in fact is an assumption made by the Ostergaard analysis in general.) The expected value of each term is not changed by small numbers. Thus, small numbers of observations do not distort the expected value of the Ostergaard prediction, and an analysis based on small numbers will not involve any inflation of Type I error rate. That is, finding a significant difference between the Ostergaard predicted proportion and the obtained proportion could not be explained by the small numbers of observations.

Note that independence between the recognition and specific situation tasks could arise simply by responding "old" to everything. This type of problem is addressed by Ostergaard's analysis. If subjects responded "old" to everything, the

difference between independence and the degree of dependence predicted by Ostergaard's model would be zero. If data are significantly less than the prediction of Ostergaard's model, then trivial explanations of this sort are ruled out.

Appendix B

Effect of generalization between situations on the
Ostergaard model

The analyses of the degree of dependence between the tasks reported in the results sections of Experiments 1 and 2 assumed that there was no generalization between situations. In fact, Dienes and Fahey (1995) showed that there was generalization between situations. For the sugar production task, Dienes and Fahey showed that subjects treated situations that differed by only one level of sugar production as if they were the same (there was little generalization between situations that differed by two or more levels of sugar production). This generalization could increase the apparent amount of independence between performance on the recognition and specific situations tasks. For example, subjects may fail to recognize a situation as old, but still produce the correct response on the specific situations task because they remember a similar situation. To deal with this potential problem, the analyses on the data from Experiment 1 were repeated but subjects were regarded as having correctly recognized an old loosely correct situation either if (1) they correctly recognized that particular situation; or (2) they correctly recognized another loosely

correct situation that was only one level of sugar production different and to which they gave a correct response in the specific situations task. The difference between the actual proportion of situations in which subjects performed correctly on both tasks and the proportion expected on the basis of independence was, averaged over subjects, .004 (SD = .043), which is non-significantly different from zero, $t(19) < 1$. That is, there was no evidence that giving a correct response to a situation in the specific situations task depended on recognizing that situation or a similar situation. The difference between the actual proportion correct on both tasks and that expected on the basis of Ostergaard's model of dependence was, averaged over subjects, -.05 (SD = .08), $t(13) = 2.16$, $p < .05$, $d = 0.63$. That is, the dependence between the tasks was less than would be expected if correct performance on the specific situations task was based on recognition of the situation or a similar one¹.

In terms of the person interaction task, Dienes and Fahey (1995) showed that when interacting with Person U, subjects treated situations as similar if they differed by no more than four levels of Ellis' behaviour or the subjects' behaviour. As for Experiment 1, this generalization could increase the

apparent amount of independence between performance on the recognition and specific situations tasks. To deal with this problem, the analyses conducted on the data from Experiment 2 were repeated but the subject was regarded as having correctly recognized an old loosely correct situation either if (1) they correctly recognized that particular situation; or (2) they correctly recognized another loosely correct situation that was no more than four levels of Ellis' behaviour or the subjects' behaviour different, and to which they gave a correct response in the specific situations task. The difference between the actual proportion and that expected on the basis of independence was, averaged over subjects, $-.01$ for Person U, and $.00$ for Person S, both non-significantly different from zero, $t_s < 1$. That is, there was no evidence that giving a correct response to a situation in the specific situations task depended on recognizing the situation or a similar situation. The difference between the actual proportion correct on both tasks and that expected on the basis of Ostergaard's model of dependence was, averaged over subjects, $-.06$, $t(20) = 2.70$, $p < .02$, $d = 0.59$, for Person U, and $.00$, $t < 1$, $d < 0.20$, for Person S. That is, for Person U, the dependence between the tasks was less than would be expected if correct performance on the specific

situations task was based on recognition of the situation or a similar one. For Person S, the data did not distinguish between Ostergaard's model and the model of independence.

References

- 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. (1991). The role of action in implicit learning. Quarterly Journal of Experimental Psychology, 43A, 881-906.
- Berry, D. C., Banbury, S., and Henry, L. (1997). Transfer across form and modality in implicit and explicit memory. The Quarterly Journal of Experimental Psychology, 50A, 1-24.
- Berry, D. C., & Dienes, Z. (1993). Implicit learning: theoretical and empirical issues. Hove: Erlbaum.
- Berry, D.C. & Broadbent, D.E. (1984). On the relationship between task performance and associated verbalizable knowledge. Quarterly Journal of Experimental Psychology, 36, 209-231.

Berry, D. C., & Broadbent, D. E. (1987). The combination of explicit and implicit learning processes. Quarterly Journal of Experimental Psychology, 39, 585-609.

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

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

Broadbent, D. E., Fitzgerald, P., Broadbent, M. H. P. (1986). Implicit and explicit knowledge in the control of complex systems. British Journal of Psychology, 77, 33-50.

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

Brown, A. S., Neblett, D. R., Jones, T. C., & Mitchell, D. B. (1991). Transfer of processing in repetition priming: Some inappropriate findings.

Journal of Experimental Psychology: Learning, Memory, and Cognition, 17, 514-525.

Challis, B. H., Chiu, C.-Y., Kerr, S. A., Law, J., Schneider, L., Yonelinas, A., & Tulving, E. (1993). Perceptual and conceptual cueing in implicit and explicit retrieval. Memory, 7, 127-151.

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

Cho, J.-R., & Mathews, R. C. (1996). Interactions between mental models used in categorization and experiential knowledge of specific cases. Quarterly Journal of Experimental Psychology, 49A, 572-595.

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

Cohen, J. (1988). Statistical power analysis for the behavioural sciences (2nd ed.). Hillsdale, NJ: Lawrence Erlbaum Associates.

Dienes, Z., & Altmann, G. T. M. (1997). Transfer of implicit knowledge across domains: How implicit and how abstract? In D. C. Berry (Ed.), How implicit is implicit learning? (pp. 107-123). Oxford: Oxford University Press.

Dienes, Z., Altmann, G., Kwan, L, Goode, A. (1995) Unconscious knowledge of artificial grammars is applied strategically. Journal of Experimental Psychology: Learning, Memory, & Cognition, 21, 1322-1338.

Dienes, Z., & Berry, D. C. (1997). Implicit learning: Below the subjective threshold. Psychonomic Bulletin and Review, 4, 3-23.

Dienes, Z., & Fahey, R. (1995). The role of specific instances in controlling a dynamic system. Journal of Experimental Psychology: Learning, Memory, & Cognition, 21, 848-862.

Dienes, Z., & Perner, J. (1996). Implicit knowledge in people and connectionist networks. In G. Underwood (Ed), Implicit cognition (pp 227-256). Oxford: Oxford University Press.

Flexser, A. J. (1991). The implications of item differences: Commentary on Hintman and Hartry.

Journal of Experimental Psychology: Learning, Memory, and Cognition, 17, 338-340.

Gardiner, J. M. (1991). Contingency relations in successive tests: Accidents do not happen. Journal of Experimental Psychology: Learning, Memory, and Cognition, 17, 334-337.

Gardiner, J. M., & Parkin, A. J. (1990). Attention and recollective experience in recognition memory. Memory and Cognition, 18, 23-30.

Hayes, N.A., & Broadbent, D. E. (1988). Two modes of learning for interactive tasks. Cognition, 28, 249-276.

Hayman, C., & Tulving, E. (1989). Contingent dissociation between recognition and fragment completion: The method of triangulation. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 228-240.

Hintzman, D. L. (1991). Contingency analyses, hypotheses, and artifacts: Reply to Flexser and to Gardiner. Journal of Experimental Psychology: Learning, Memory, and Cognition, 17, 341-345.

Hintzman, D. L., & Hartry, A. L. (1990). Item effects in recognition and fragment completion: Contingency relations vary for different subsets of words. Journal of Experimental Psychology: Learning, Memory, and Cognition, 16, 955-969.

Lewicki, P. (1986). Nonconscious social information processing. New York: Academic Press.

Logan, G. (1988). Towards an instance theory of automatization. Psychological Review, 95, 492-527.

Marescaux, P-J., Luc, F., & Karnas, G. (1989). Modes d'apprentissage selectif et nonselectif et connaissances acquies au control d'un processus: Evaluation d'un modele simule. Cahiers de Psychologie Cognitive, 9, 239-264.

Mathews, R. C., Buss, R. R., Chinn, R., & Stanley, W. B. (1988). The role of explicit and implicit learning processes in concept discovery. Quarterly Journal of Experimental Psychology, 40, 135-165.

Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. Journal of Experimental Psychology: Learning, Memory, and Cognition, 14, 700-708.

Nosofsky, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. Journal of Experimental Psychology: Human Perception and Performance, 17, 3-27.

Ostergaard, A. L. (1992). A method for judging measures of stochastic dependence: Further comments on the current controversy. Journal of Experimental Psychology: Learning, Memory, and Cognition, 18, 413-420.

Ostergaard, A. L. (1994). Who is mistaken about priming in "recognition/identification" experiments? A reply to Tulving and Hayman. European Journal of Cognitive Psychology, 7, 1-11.

Perruchet, P. (1994). Learning from complex rule-governed environments: On the proper function of conscious and unconscious processes. In C. Umiltà & M. Moscovitch (Eds), Attention and Performance XV: conscious and nonconscious information processing pp (811-835). Cambridge, MA: MIT Press.

Poldrack, R. A. (1996). On testing for stochastic dissociations. Psychonomic Bulletin & Review, 3, 434-448.

Reber, A. S. (1967). Implicit learning of artificial grammars. Journal of Verbal Learning and Verbal Behaviour, 6, 855-863.

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

Roediger, H. L., III (1990). Implicit memory: Retention without remembering. American Psychologist, 45, 1043-1056.

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

Roediger, H. L., III, & McDermott, K. B. (1993). Implicit memory in normal human subjects. In F. Boller J. Grafman (Eds.), Handbook of Neuropsychology, Vol. 8. Elsevier Science Publishers.

Sanderson, P. M. (1989). Verbalizable knowledge and skilled task performance: Association, dissociation, and mental models. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 729-747.

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

Schacter, D. L. (1990). Perceptual representation systems and implicit memory: Towards a resolution of the multiple memory systems debate. Annals of the New York Academy of Sciences, 608, 543-571.

Schacter, D. L., & Graf, P. (1989). Modality specificity of implicit memory for new associations. Journal of Experimental Psychology: Learning, Memory, & Cognition, 15, 3-12.

Schwartz, S. H. (1966). Trial-by-trial analysis of processes in simple and disjunctive concept attainment tasks. Journal of Experimental Psychology, 72, 456-465.

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. (1997). Dissociating long-term memory systems: Comment on Nyberg and Tulving (1996). European Journal of Cognitive Psychology, 9, 11-120.

Shanks, D. R., & Gluck, M. A. (1994). Tests of an adaptive network model for the identification and categorization of continuous dimension stimuli. Connection Science, 6, 59 - 89.

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

Squire, L. R., & Frombach, M. (1990). Cognitive skill learning in amnesia. Psychobiology, 18, 109-117.

Stanley, W.B., Mathews, R.C., Buss, R.R. & Kotler-Cope, S. (1989). Insight without awareness: On the interaction of verbalization, instruction, and practice in a simulated process control task.

Quarterly Journal of Experimental Psychology, 41,
553-577.

Tulving, E. (1985). How many memory systems are there? American Psychologist, 40, 385-398.

Tulving, E., & Hayman, C. G. (1993). Stochastic independence in the recognition/identification paradigm. European Journal of Cognitive Psychology, 5, 353-373.

Tulving, E., & Schacter, D. L. (1990). Priming and human memory systems. Science, 247, 301-306.

Tulving, E., Schacter, D. L., & Stark, H. A. (1982). Priming effects in word fragment completion are independent of recognition memory. Journal of Experimental Psychology: Learning, Memory, and Cognition, 8, 336-342.

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

Warrington, E. K., & Weiskrantz, L. (1978). Further analysis of the prior learning effect in amnesic patients. Neuropsychologia, 16, 169-176.

Weldon, M. S., & Roediger, H. L. (1987). Altering retrieval demands reverses the picture superiority effect. Memory & Cognition, 15, 269-280.

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.

Willingham, D.B., Nissen, M.J. & Bullemer, P. (1989). On the development of procedural knowledge. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 1047-1060.

Footnotes

¹Generalization between instances may increase performance on new situations in the specific situations task. However, performance on new situations is used in the Ostergaard analysis as a measure of baseline performance in the absence of learning. Performance on new situations which are more than one level of sugar production different from any old correct situation was .10 (SD = .32), no less than the performance on all new situations (.07, SD = .27). That is, generalization did not detectably increase performance on new situations, and so the Ostergaard analysis would not be affected.

²An analysis of the data from Experiment 1 of Dienes and Fahey (1995) also revealed that the relation between performance on the specific situations task and recognition was nonsignificantly different from independence and significantly lower than that expected by Ostergaard's model. In the Dienes and Fahey experiment, unlike the experiment reported in the current paper, situations were presented in the specific situations and recognition tasks as combinations of sugar production and work force on the previous trial. This creates a potential problem with drawing strong conclusions from the

apparent stochastic independence in the Dienes and Fahey experiment because they provided evidence that subjects were responding to the sugar production task on the basis of single features. That is, subjects responded to a given level of sugar production in a similar way regardless of the level of work force with which it was combined, and they also responded to a given level of work force in a similar way regardless of the level of sugar production. Thus, even if subjects did not recognize a particular combination of sugar production and work force as being old (and incorrectly responded 'new') in the recognition task, they may still have recognized the level of sugar production by itself as being old and responded appropriately in the specific situations task. This would reduce the expected degree of statistical dependence between the tasks below that predicted by the Ostergard model.

³Further, Dienes and Fahey (1995) presented evidence that subjects responded to the specific situations task on the basis of the combination of Ellis' behaviour and the subjects' behaviour on the last trial. Because in this case, unlike Experiment 1 of Dienes and Fahey (see footnote 2), subjects were responding to both the specific situations task and

the recognition task on the basis of the same features, Ostergaard's model provides a plausible estimate of the amount of dependence expected if subjects' responses on the specific situations task were based on explicit recollection.

⁴Brown et al. (1991) argued that modality shifts in implicit memory tasks may be more likely to appear when study modalities are manipulated within rather than between subjects so as to focus subjects' attention on perceptual features of the stimuli. However, there have been a number of studies finding substantial modality shifts on implicit memory tasks using between subjects designs (Bassili et al., 1989; Weldon & Roediger, 1987).

⁵Considering new test situations, the similarity for false alarms (493, SD = 128) was higher than for correct rejections (470, SD = 92), though not significantly so, $t(13) = 0.4$ (95% CI on the difference went from -151 to 105, indicating that the data were not sensitive enough to distinguish the interesting alternatives).

⁶Or that familiarity depended only on exact matches and immediately adjacent situations. Similarity is

known to drop off at a faster than linear rate with distance in other paradigms (Shanks & Gluck, 1994).

⁷A similar analyses was conducted on the data from Experiment 2. Similarity was defined by a city block metric (as the dimensions were presumably perceptually separable; Shanks & Gluck, 1994):

Similarity = 23 - |difference in Ellis' behaviour| - |difference in subject's behaviour|. Thus, the minimum similarity is 1 and the maximum is 23. The difference between hits (961, SD = 56) and misses (885, SD = 76) was significant, $t(19) = 3.97$, $p = .001$. For frequency, computed as the number of times that exact situation occurred in training, the difference between hits (2.6, SD = 0.9) and misses (1.3, SD = 0.3), was also significant, $t(19) = 5.55$, $p < .001$. When similarity difference was regressed against frequency difference, the intercept (20) was not significantly different from zero, $t(18) = 0.76$ (the intercept lay within the body of the data). That is, on this model of similarity, the data provided no evidence that subjects' recognition decisions were determined by any situation in the training phase other than the exact situation tested for.

Table 1

Contingency table for situations to which subjects gave the correct response in the training phase of Experiment 1

	Situations Task	
	Correct	Incorrect
Recognition Task		
Subjects response:		
'Old'	29	53
'New'	12	16

Note. On the recognition task, 'old' is the correct response, 'new' is the incorrect response. On the situations task, 'correct' refers to the subject giving the correct response when tested on this task, and 'incorrect' refers to the subject giving the incorrect response when tested on this task: All the situations had been given the correct response in the training phase.

Table 2

Specific situations task/proportion correct

Person:	U		S	
Modality:	Same	Diff.	Same	Diff.
Study priming:				
New	.30 (.13)	.33 (.18)	.47 (.19)	.54 (.15)
Old	.53 (.25)	.54 (.23)	.70 (.20)	.74 (.23)

Note. Standard deviations appear in parentheses.

Table 3

Contingency table for Person S

	Situations Task	
	Correct	Incorrect
Recognition Task		
Subjects response		
'Old'	105	24
'New'	39	10

Contingency table for Person U

	Situations Task	
	Correct	Incorrect
Recognition Task		
Subjects response		
'Old'	56	21
'New'	50	18

Note. On the recognition task, 'old' is the correct response, 'new' is the incorrect response. On the situations task, 'correct' refers to the subject giving the correct response when tested on this task, and 'incorrect' refers to the subject giving the incorrect response when tested on this task: All the situations had been given the correct response in the training phase (i.e. a response that led to target or just one level off).

Figure 1.

The displays seen by subjects during the learning phases of Experiment 1 (top panel) and Experiment 2 (bottom panel).