Experiments in Implicit Learning

Stephen William Kelly, B.Sc. (Hons)

Thesis submitted to the University of Glasgow for the degree of Doctor of Philosophy, December, 1997.

Contents

Abstract

This thesis examines two paradigms from the area of implicit learning in detail. The literature suggests that the invariance detection paradigm of McGeorge and Burton (1990) gives rise to unconscious knowledge held at a conceptual level with the decision process served by a 'nearest-neighbour' similarity mechanism. The experiments in this thesis suggest that several aspects of this task do indeed seem to agree with present conceptions of unconscious knowledge but no evidence could be found that this knowledge is held at a conceptual level or that specific similarity plays any role in this task. Instead the experiments in this thesis suggest that this task may be better understood in terms of an abstraction mechanism which acquires perceptual information.

Using the invariance detection paradigm, this thesis examines the effect of two types of task which measure performance above an 'objective threshold' of awareness. Performance on each task was not the same suggesting that one cannot assume all direct tests measure the same knowledge despite being similar in nature. In addition, the finding that only the more sensitive of the two tasks could elicit information in the invariance detection paradigm suggests that the knowledge is extremely difficult to elicit. This also is a property of implicit learning and points to the digit invariance task being mediated by unconscious mechanisms.

The finding of robust invariance detection in laboratory tasks suggests that one might expect to find similar learning for real world invariance. No evidence for this could be found, which suggests that either implicit learning is a laboratory artefact or that real world invariance learning does not operate in the same way that laboratory experiments suggest. These results suggest that laboratory experiments are required which replicate conditions under which real world learning might occur.

The final part of this thesis examines a different paradigm and investigates whether a rather peculiar finding reported in the complex systems literature can extend to sequence learning. Berry (1991) reports that learning by observation alone can occur in an explicit but not in an implicit version of a complex systems task. Howard, Mutter and Howard (1992) however report that learning by observation alone can occur in a sequence learning task. This thesis indicates

methodological problems with Howard et al's interpretation and presents several experiments which suggest that implicit learning cannot occur by observation alone.

Acknowledgements

I would like to thank Mike Burton for agreeing to supervise this thesis. He suggested many good modifications to my wild schemes, quelled moments of rising panic and proved to be an excellent hillwalking companion. I would also like to thank Peter McGeorge for giving me an interest in implicit learning and for discovering a clever task which I have used extensively in this thesis. Finally I would just like to thank all my good friends at the Universities of Aberdeen and Glasgow who have made my five years working in psychology very enjoyable ones.

Chapter 1. Memory and learning outside of consciousness

1.1 Introduction

The nature of consciousness is presently generating much interest and debate and one aspect of this debate relates to whether or not individuals can acquire or use information without conscious awareness or intention to do so. The acquisition of natural language is put forward as one prime example of this phenomenon (e.g. Reber and Winter, 1994). Information acquisition in the real world is an extremely complex process and researchers attempting to demonstrate unconscious cognition prefer to use much simpler laboratory tasks. This thesis examines some assumptions made about implicit learning and attempts to test these assumptions. The first chapter will describe a brief history of scientific attitudes towards unconscious mentation and the early research into learning which had the possibility of being unconscious. Then, it will review implicit learning and examine evidence from each of the major paradigms which have arisen in this area. Cognitive modelling has also proven useful in elucidating possible mechanisms for implicit learning and provides converging evidence with empirical data for suggesting not only whether implicit learning can take place but also what the underlying representation of any such knowledge might be. This chapter describes several different cognitive models which might account for unconscious knowledge and evaluates their contribution in answering both these questions. The following chapters will each introduce a currently held concept in implicit learning and examine the topic in more detail.

1.2 The historical advance of the unconscious

Schacter (1987) notes that the first reference made to an unconscious memory of any sort was by Descartes in the work from 1649 entitled "The passions of the soul". Descartes observes that a traumatic experience during childhood may remain "imprinted on his [the child's] brain to the end of his life" without "any memory remaining of it afterwards". With the exception of Leibniz and Maine de Biran (see Schacter, 1987 for a review), most of the pre-19th century philosophers restricted their considerations to explicit memory phenomena and it was not until Carpenter coined the term 'unconscious cerebration' in 1874 that interest in unconscious processes became more widespread. Around the same time, the Viennese physiologist Hering criticised earlier theorists for concentrating on explicit phenomena and introduced the term 'unconscious' or 'organic' memory. By the late 19th century, interest in unconscious cognitive processes had ceased to be restricted to the domain of philosophy and empirical evidence from the domains of neurology, psychiatry and experimental psychology were prevalent in the literature. Various now famous cases of amnesia were discussed by Dunn (1845), Korsakoff (1889) and Claparede (1911) who put forward ideas of unconscious memories being very weak memory traces or being dissociated from 'the ego'. It is worth noting that these ideas can be seen as more sophisticated formulations in present day cognitive psychology. The idea that implicit memories are dissociated from the ego can be seen as an early instantiation of the idea of multiple memory systems (e.g. Cohen, 1984; Schacter, 1990; Squire, 1987; Tulving and Schacter, 1990)). Perhaps the most famous early example of unconscious processes affecting behaviour was in the domain of psychoanalysis and popularised by Freud. Here, traumatic events which could not be explicitly remembered were nevertheless expressed indirectly in a person's behaviour.

Work in experimental psychology did not make a distinction between implicit and explicit memory despite the increasing interest in this topic in other fields. One notable exception to this was the work of Ebbinghaus (1885) who describes an experiment in which he himself notices making considerable savings in a test the next day for items which he did not consciously remember studying. One of the first experimental studies in learning, as opposed to memory, was by Thorndike (1898) who rewarded cats with food when they escaped a puzzle box by operating a latch. He concluded that his subjects had no conscious awareness of what they were doing and the behaviour was mediated by automatic strengthening of the link between the stimulus and response. This experiment led to a series of studies on human subjects which seemed to confirm that rule learning could take place without conscious awareness of the rules.

1.3 Modern conceptions of implicit learning

Usage of the term 'implicit'.

Implicit learning has come under tremendous scrutiny in the last 30 years and is now one of the most contentious areas within modern cognitive psychology. The term 'implicit ' itself is used in a variety of ways which only adds confusion to the debate. 'Implicit' can refer to the type of knowledge one holds and usually means such knowledge is not available for verbal report (although this itself is a point of debate). It can also refer to the type of learning carried out by experimental subjects and usually means that there is no conscious effort to learn some aspect of the stimulus which is nevertheless learned. A third use of the term 'implicit' is in differentiating between different types of task. Reber, Walkenfield and Hernstadt (1991) differentiate between the implicit task of grammar learning and the explicit task of a series-completion problem solving task. Of course, none of these terms are absolute. An artificial grammar task can be learned in an explicit fashion and there is evidence that knowledge which is initially implicit can eventually become accessible to verbal report (Stanley et al, 1989).

On the relationship between explicit and implicit.

Reber (1993) warns against what he terms 'the polarity fallacy', namely the idea that implicit and explicit learning are completely separate processes which do not interact with one another. Reber suggests that " implicit and explicit systems should be properly viewed as complementary and cooperative functional systems that act to provide us with information about the world". This is similar to the realisation in the field of implicit memory that explicit processes might contaminate performance on supposedly implicit tasks (Holender, 1986; Reingold and Merikle, 1990) and vice versa (Jacoby, Toth and Yonelinas, 1993). It is now assumed that performance on any task has some conscious and unconscious component and now experimental methods are employed to separate the two (e.g. Jacoby, 1991; Jacoby et al, 1993). Despite the warnings of Reber, however, research in implicit learning still seems to be attempting to show that tasks are either implicit or explicit. As will be discussed, the conclusion that various processes are at work in these tasks is becoming evident but so far little has been done to separate out the explicit and implicit components in implicit learning tasks.

Some definitions of implicit learning.

Mathews et al (1989) suggest that "implicit learning is thought to be an alternate mode of learning that is automatic, nonconscious and more powerful than explicit thinking for discovering nonsalient covariance between task variables". Seger (1994) defines implicit learning as "learning complex information without complete verbalisable knowledge of what is learned". Berry and Dienes (1993) suggest that in implicit learning, a subject "learns about the structure of a fairly complex stimulus, without necessarily intending to do so, and in such a way that the acquired knowledge is difficult to express". These definitions seem to suggest a consensus on what constitutes implicit learning. Some authors however (notably Shanks and St. John, 1994), would suggest that the experimental evidence to date does not show any of these criteria to be true. The next few sections will consider the evidence for implicit learning being:

1. an alternate mode of learning.
2. automatic / dissociable from an intention to learn.
3 unconscious / characterised by non-verbalisable knowledge.
4. more powerful than explicit learning.

<u>Is implicit learning an alternate mode of learning?</u>

There are various dissociations between explicit and implicit learning within the literature which suggest that the two are very different modes of learning. The dissociations peculiar to each task will be discussed later, but there seem to be general properties which separate implicit from explicit learning across tasks. These properties have been summarised by Berry and Dienes (1993) as:

1. Implicit learning shows specificity of transfer.

In general, information acquired implicitly cannot be tapped by free recall tasks. Much more information can be elicited by using more sensitive measures such as forced-choice recognition but even then not all the information known by subjects is always demonstrated. In addition, tasks which are similar to the original learning task can show a transfer of knowledge. For example, Berry and Broadbent (1988) used different versions of a complex systems task, all of which had different surface structures but had the same underlying relationship between the subjects' input and the response which they received. Transfer of knowledge was found between two different person-interaction tasks and between two transport tasks but not between a person task and a transport task or vice-versa. A similar finding was also reported by Squire and Frambach (1990). Bright and Burton (1994) report a lack of transfer in a digit invariance task when the four digit numbers are presented as phone numbers during the learning phase but presented as clock times in the test phase. These findings could be interpreted either as implicit representations being tied to surface characteristics (as is found in implicit memory) or that the specific operations performed on the stimuli during training are necessary during retrieval, as would be suggested by an episodic processing account of implicit learning (e.g. Whittlesea and Dorken, 1993).

2. Implicit learning tends to be associated with incidental learning.

Implicit learning can be shown when subjects are not informed that they have to learn anything during a training and some studies have shown that the amount of learning seems to be reduced if subjects are overtly told to learn certain information (see later).

3. Implicit learning gives rise to a phenomenal sense of intuition.

Reber (1967) first demonstrated that subjects could perform at above chance levels without being able to verbalise why. This finding has been demonstrated in various paradigms including complex systems (Berry and Broadbent, 1984) and invariance learning (McGeorge and Burton, 1990). In fact, subjects tend to offer explanations that cannot account for their performance or believe that they are simply guessing. Chan (1992, cited in Berry and Dienes, 1993) found that subjects who were given a variation of the artificial grammar paradigm, in which they had to rate how confident they were for each decision, rated incorrect responses as having been made with the same level of confidence as correct responses.

4. Implicit learning remains robust in the face of time, psychological disorder and secondary tasks.

Although there is evidence to suggest that implicit knowledge can last for a relatively long time (Allen and Reber, 1980) found that knowledge of a grammar was retained over a period of two years), no study has actually made a direct comparison between explicit and implicit knowledge over time. Thus one cannot say that robustness over time is necessarily a distinguishing characteristic of implicit learning. Intact knowledge from implicit learning studies has been demonstrated in amnesia (e.g. Knowlton, Ramus and Squire, 1992; Nissen and Bullemer, 1987; Squire and Frambach, 1990) despite the fact that such knowledge cannot usually be demonstrated using explicit tasks. Such a finding has also been demonstrated in subjects with psychiatric disorder (Abrams and Reber, 1988) and in normal subjects with scopolamine-induced amnesia (Nissen, Knopman and Schacter, 1987).

Several studies have also found that implicit learning is robust when subjects are given a secondary task to perform. Hayes and Broadbent (1988) argued for two separate modes of learning which they termed s-mode (selective) and u-mode (unselective) and which were consistent with the ideas of explicit and implicit learning respectively. They found that in a complex systems experiment, subjects who were required to perform a secondary verbal task showed a decrement in performance on the explicit task whereas performance on the implicit task was facilitated by the addition of the secondary task. Although this result was not replicated by Green and Shanks (1993), other paradigms show some evidence of robustness of learning with secondary tasks. Dienes, Broadbent and Berry (1991) found that a random number generation task interfered with making a grammaticality judgment. While this suggests that a secondary task does impair performance, a later study (Dienes, Altmann, Kwan and Goode, 1995) suggests that a secondary task only interferes with decisions when the subjects were confident in their decision. When subjects believed they were guessing, the secondary task had no detrimental effect. This is an important point as it demonstrates the danger in believing tasks to be 'process-pure'. It also suggests that two of the characteristics of implicit learning suggested by Berry and Dienes (1993) may tend to occur together.

It seems that there is a reasonable amount of evidence to support most of these suggested characteristics being related to a notion of implicit learning but not to explicit hypothesis testing. These characteristics also show a tendency to co-occur in tasks, which again lends support to the idea that they are representative of different modes of learning as opposed to being simply task-specific variables. A similar argument for distinct learning processes is made by Reber (1993) who views implicit learning as a biological and evolutionary necessity. He argues that implicit learning shows distinctive characteristics which are not shared by explicit learning and, in agreement with Berry and Dienes (1993) suggests that robustness in the face of psychiatric disorder or impairment is one of these characteristics. The basis for such a hypothesis comes originally from other Darwinian treatments of cognition. Rozin (1976) argues that intelligence is derived from hierarchical structures which at one time were distinct from each other and had distinct functions to perform. Over an evolutionary timescale, these structures were able to communicate more and more functions across modules and so eventually were able to function more or less as a whole system capable of performing a wide range of functions and operating with a large degree of adaptability. This idea has also been used by Baars (1988) who argues that the function of consciousness is to provide a means by which individual, self-encapsulated modules can present information to the whole cognitive system. Thus, several modules will provide a solution to a problem facing the organism based on the particular process each module is capable of carrying out. Consciousness allows the different solutions to be presented to and evaluated by the entire system and one solution is then selected either to be put into practice or to be modified by another module.

This idea of modularisation has been used by researchers in implicit memory with Schacter (1984; Sherry and Schacter, 1987) proposing an encapsulated module responsible for unconscious priming effects, which he terms the Perceptual Representation System. Reber (1993) does not advocate a multiple systems view of implicit learning as he argues that there is no strong evidence to date to choose between this and an adaptionist position. He does however, take the view that implicit systems are more primitive and would have evolved earlier than explicit systems. There is some evidence for this being the case with implicit systems. Squire (1986) notes that simple associative processes (which are a form of basic procedural system) take place via neuronal systems which occur in invertebrates whereas declarative or explicit memory systems require systems which are evolutionarily much younger. Shimamura (1986) also suggests that implicit memory is mediated by structures which are evolutionary much older than the hippocampus and Seger (1994) also notes that the hippocampus seems to have little to do with implicit function. Reber (1993) makes the general comment that implicit and explicit function are "dissociable along phylogenetic lines with the evolutionary older, implicit system showing the greatest resistance to insult and injury". This may prove to be too general a comment as Seger (1994) reviews possible anatomical sites of implicit learning and of the three she provides evidence for, only the basal ganglia could be said to be primitive (the other two being the association areas and the frontal lobes). Of course, it may be that these relatively young areas control aspects of implicit learning tasks that are unrelated to actual implicit learning but as yet, the neurology of implicit learning is too undeveloped to make any solid claim about what structures are responsible for learning.

In addition to the claim of robustness, Reber (1993) also makes several other predictions based on the theory of the implicit system being evolutionarily older which point to a dissociation between explicit and implicit learning. These are age independence, IQ independence and low variability. The reasoning behind these propositions is similar to that of robustness in the face of neurological insult. If implicit learning is a 'primitive' process then it should be robust against those aging processes which cause deterioration in more fragile explicit systems. It should show less of a correlation with IQ than explicit learning as IQ is generally seen as a measure of our explicit cognitive abilities. Also, as an older process it should be more stable which would be demonstrable by less individual variability in performance than that seen in explicit learning tasks.

Although each of these points makes clear predictions there has been relatively little direct testing of these predictions in the literature. In implicit memory research there are numerous studies which show that the elderly do not decline as much on implicit memory tasks as they do on explicit (Howard, 1988) and even then the small decline shown on implicit tasks may be an artifact due to explicit contamination of the task (Graf, 1990). In the learning literature, a small number of studies show mixed results. Cherry and Stadler (1995) found that high ability elderly adults showed comparable implicit learning on a sequence learning task to younger subjects but that less learning was found for low ability elderly. Myers and Connor (1992) found that young and older subjects performing on a complex systems task did not differ in terms of performance but the older subjects showed less verbalisable knowledge than the young. McGeorge, Crawford and Kelly (unpublished manuscript) found that performance on the digit invariance task did not decline with age on a census matched sample of the UK population. The data from the learning literature seems to be in agreement with that from the memory literature with respect to aging however there is less of a consensus at the other end of the developmental spectrum. Reber (1993, 1997) cites several studies (Haith, Hazan and Goodman, 1988; Haith and McCarty, 1990) which show that children as young as three months can learn to anticipate the location of a face stimulus. Reber (1993) also reports unpublished studies from his own laboratory using a grammar paradigm modified for children in which age groups from four to fourteen did not differ in terms of performance on the task. Mayberry, Taylor and O' Brien-Malone (1995) however, found that implicit learning improved with age across two groups of children with ages of around six to ages of around eleven.

Their study did lend support to the prediction that implicit learning is independent of the subjects' IQ as the children were divided into low, medium and high ability groups who did not differ on the implicit task but gave the expected pattern on the explicit task. Two other studies (McGeorge, Crawford and Kelly, 1997; Reber, Walkenfield and Hernstadt, 1991) comparing the performance of adults on an implicit and an explicit learning task also reach the conclusion that implicit but not explicit learning is independent of IQ. In addition, these studies also provided support for a related prediction; that of reduced variability. Reber at al (1993) found that the variance of scores on the explicit task were four times as great as those on the implicit task. McGeorge et al also reported much greater variance on the explicit as opposed to the implicit task.

In summary, Berry and Dienes (1993) suggest several criteria which seem to demarcate implicit from explicit learning and the general finding is that these criteria seem to co-occur. The evidence suggests that implicit learning is an alternate mode of learning although whether it is subserved by a different neurological system is debatable. Reber (1993) puts forward an evolutionary argument as to why we should expect implicit learning to exist as an alternate mode of learning and although little research has been done to test the predictions of his theory, the few studies that are relevant show positive results on the whole.

Is implicit learning automatic?

Logan (1990) describes an automatic process as one which requires minimal attentional resources, is effortless and is obligatory as it does not require attention. Eysenck and Keane (1990) also suggest that automatic processes should be fast, unavailable to consciousness and should not interfere with other tasks. These are the general conceptions of what implicit learning "should be like". Despite strong similarities between the two, it is important to realise that implicit learning is not the same as the phenomenon termed automatisation. This is where practice in a given procedure in a given situation changes that procedure from being flexible and characterised by declarative knowledge (as used in Anderson's (1983) ACT* model) to an inflexible procedure which is fast, not available to consciousness, unavoidable and not requiring attention (Berry and Dienes, 1993). Studies using the dynamic systems task have shown that the acquired knowledge is at first not available for verbal report but after extended practice was found to be accessible (Sanderson, 1989; Stanley et al, 1989). Thus it seems that in this respect, automatisation and implicit learning have opposite patterns of effect.

The role of attention in implicit learning tasks has been examined in sequence learning tasks. Both Nissen and Bullemer (1987) and Cohen, Ivry and Keele (1990) required subjects to make a reaction timed keypress corresponding to the position of an asterisk in one of four spatial locations on screen. As is standard in these tasks, the asterisk's position was sequential and the presentation of this sequence was repeated many times. Both studies found that learning of the sequence (as measured by a decrease in the time to respond) was prevented by having the subjects perform a tone counting task during the experiment and having them report how many low-pitched tones had occurred in each block of trials. Cohen et al (1990) extended this finding by showing that it was only sequences without unique associations (one location always following another) that could not be learned under these dual task conditions. These results seem to suggest that attention is necessary for implicit sequence learning. Stadler (1995), however has shown that tone counting prevents learning due to its disrupting effects on the perceived organisation of the sequence rather than any competition for attentional resources. A letter-memory task (subjects are given letters at the beginning of each block and asked to report them at the end of the block) which requires attentional resources but would not affect organisation of the sequence did produce less learning than a control condition where subjects were shown letters but asked to ignore them. Hence it cannot be argued from the results of Stadler (1995) that attention is not necessary for performing implicit sequence learning. In fact, it may well be that the letter-memory task does not impose such a great attentional load in the first place. The results however do lessen the impact of those reported by Nissen and Bullemer (1987) and Cohen et al (1990).

The 'u-mode' learning described by Hayes and Broadbent (1987) was supposed to occur outside of abstract working memory and give rise to knowledge that was unavailable for verbal report. Hayes and Broadbent found that performance on a dynamic systems task which relied on this mode of learning was facilitated by the addition of a secondary task but in a similar version of the task, which required what we would term explicit learning, the secondary task was detrimental to subjects' performance. The fact that the secondary task did not interfere with the 'implicit' task suggests an element of automaticity however as Green and Shanks (1993) could not replicate these results, and in fact found the effects of a secondary load to be equivalent on both tasks the conclusion to be taken from these studies is somewhat doubtful.

Dienes, Broadbent and Berry (1991) examined the effects of additional task load on artificial grammar performance. Dienes et al used a random number generation task during presentation of the study exemplars. They found that performance on a grammar classification task and on a free report measure were equally affected under dual task conditions. This suggests that grammar learning is not 'automatic' according to the definitions given above. One other study into artificial grammar learning suggests a different conclusion with respect to the requirements of attention. McGeorge, Crawford and Kelly (1997) gave a grammar task to 123 census matched subjects and an explicit task used by Reber et al (1991), namely letter sequence completion. The subjects were also tested on a standard IQ test and their sub-scale scores allowed their performance on the two tasks to be indexed in terms of three latent variables derived using structural equation modelling techniques. The three variables were verbal and performance IQ and an attentional factor. McGeorge et al found that performance on the explicit task was linked to attention but no such link was found for the grammar task. Hence it seems that learning on the artificial grammar task could proceed independently of attention.

In relation to automatic processes being fast (Eysenck and Keane, 1990), Turner and Fischler (1993) found evidence that subjects trained explicitly on the rules of an artificial grammar showed impaired performance when required to give a response within a time deadline. Subjects given the normal memorisation instructions showed no such impairment. It could be argued that explicit knowledge could not be retrieved or applied within the short deadline but that knowledge held implicitly was available very quickly. Jacoby and Hay (1997) provide evidence for this idea with a memory study where explicit and implicit processes are placed in opposition to one another. Both young and elderly subjects were required to read a list of words with some words repeated two or three times. Subjects then heard a list of words which they were asked to memorise for a later test. The test was to identify the words that they had heard but not the ones they had read. Hay and Jacoby argue that if familiarity was guiding the choice then the subjects would choose the words that they had read as being on the auditory list however if subjects could consciously recollect having read the word then they would correctly reject it. They found that elderly subjects were more likely than young subjects to misattribute a repeatedly read word to the auditory list and argued that the elderly have less recollective experience with which to make the judgment. However when the experimenters imposed a time deadline on responding, the young subjects gave a similar pattern to that of the elderly suggesting that a time deadline forces the younger subjects to use less recollective experience. This suggests that unconscious processes can act at much shorter durations than conscious, recollective processes and provides additional support from the memory literature for the findings reported by Turner and Fischler (1993).

As can be seen from the previous paragraphs there is no consensus on whether implicit learning can proceed independently of attention. This may be due to the fact that tasks are not process pure and attention may be needed for the conscious aspect but not for the unconscious. There does seem to be evidence for implicit knowledge being accessed much more quickly than explicit knowledge and so in this sense one might claim automaticity for implicit processes. Of course, it may also mean that the type of knowledge which is implicitly acquired may be easier to access.

Is implicit learning mediated by unconscious mechanisms?

Much of the evidence for implicit learning being an unconscious process comes from the apparent dissociation between task performance and verbalisable knowledge about the task. This dissociation is evident from the verbal association paradigm through to artificial grammar learning, dynamic systems tasks and sequence learning. All show that subjects can perform at above chance levels on these tasks without being able to report most of the underlying rules (e.g. Berry and Broadbent, 1984; Nissen and Bullemer, 1987; Reber, 1967) or help yoked subjects perform above the level associated with a novice (Stanley et al, 1989) without extensive practice at the task. Evidence based on subjects' verbal reports has been criticised by several authors, however. First, there is the nature of verbal report itself to consider. Ericsson and Simon (1980) presented a model of verbal report based on retrieval of information from short and long term memory. Inaccurate verbal reports may occur because there was a problem of encoding or retrieval or because all the information held in short term memory may not be reported. Nisbett and Wilson (1977) go further and suggest that, rather than giving inaccurate reports due to a lapse of memory, subjects do not even have direct access to the cognitions responsible for their behaviour. They suggest that the content of verbal reports are determined by the subjects' perceptions of the relation between their responses and the stimuli. Subjects will report a relation which seems to them a plausible account of their actions. They go on to suggest that accurate reports will occur when the critical stimuli are salient to the subject and offer a plausible account of their response. Clearly this is not the case in implicit learning where by definition, the stimuli are non-salient and do not offer an obvious account of why they were chosen. It is possible then, that subjects do not offer a report because they do not consider themselves to have a plausible account of their actions, even if they do know the contingencies of the task. There is also the problem that when giving a verbal report of their actions, subjects may substitute an alternative, more plausible reason for their actions which might mask the fact that they do not know how they performed a particular task. This problem of inaccurate reports may provide false evidence either for or against the unconscious nature of implicit learning.

It is not simply a question of whether subjects can accurately report the true reasons behind their performance on a task, however. Subjects may be perfectly aware of why they performed the way they did on an implicit learning task and give a full verbal account of it. This may not be used as evidence against unconscious implicit learning if what the subject describes is different to that which the experimenter is expecting. This idea forms the first criterion which Shanks and St. John (1994) put forward as one of those which must be satisfied to unequivocally show that implicit learning can occur without conscious knowledge. They term this the 'Information

Criterion'. Put simply, an experimenter must be certain that the knowledge which he or she is trying to elicit in an awareness test is the knowledge which underlies performance on the task. This point is true not just for verbal report but for any type of awareness test. The second criterion which Shanks and St. John regard as essential in showing implicit learning as unconscious is the "Sensitivity Criterion". This has appeared frequently in the literature (e.g. Brody, 1989; Ericsson and Simon, 1980; Reingold and Merikle, 1988) and demands that any test of conscious awareness must be capable of eliciting any conscious knowledge that the subject may hold. This is especially essential where the knowledge may take the form of 'parts of rules' which the subject may not be sure are correct or worth mentioning. It is however, extremely difficult if not impossible to demonstrate that such a test is exhaustive therefore Shanks and St. John suggest that the minimum requirements for an awareness test should be that it is at least as sensitive to conscious knowledge as the supposed unconscious task.

It is due to this "Sensitivity Criterion" that Shanks and St. John regard a verbal awareness test as not sufficient to conclusively show lack of conscious knowledge. Nelson (1978) has shown that recognition is a more sensitive test of memory than free recall, hence it seems that recognition is a more suitable assessment of conscious knowledge in an implicit learning task. An alternative to this which still uses the free recall method is to ask more and more detailed questions about the subjects' strategies on the task (e.g. Berry and Broadbent, 1984, 1988; Bright and Burton, 1994; Lieberman, unpublished manuscript). The problem with this and the recognition task is that both are still very dependent on the experimenter asking the correct questions. Both methods may very easily fall foul of the Information Criterion and subjects may be using very different knowledge than that being tested, thus leading the experimenter to believe that they have no conscious knowledge about the task.

Although Shanks and St. John (1994) suggest that recognition (or prediction) tasks should be the tasks of choice for assessment of conscious knowledge, several authors (e.g. Cleeremans, 1994; Dienes and Perner, 1994; Seger, 1994) point out that such tasks may be contaminated by implicit processes. It is already widely accepted in the implicit memory literature that no task is totally 'process-pure' (Jacoby, Yonelinas and Jennings, 1997) and it seems that such a case may be made for implicit learning. Hence, a recognition task may tap unconscious as well as conscious processes. In certain paradigms, for example the invariance detection of McGeorge and Burton (1990), the task involving recognition of the invariant digit is assumed to be implicit as subjects are not able to report knowledge of an invariant digit or what that digit might be. The alternative view would be that subjects are consciously aware of this knowledge because they choose the exemplar containing the invariant. It is a paradigm such as this which makes clear the flaw in Shanks and St. John's argument. They would have to argue that subjects were conscious of the invariant because they could choose the exemplar which contained it, whereas awareness interviews which are structured to include increasingly detailed questions on the task knowledge suggests that subjects do not have any such knowledge.

Dienes and Berry (1997) argue that a subjective threshold is a useful way of distinguishing conscious processes from unconscious. This terminology was first used in the literature on subliminal perception (Cheesman and Merikle, 1984). An objective threshold is a point below which subjects score at chance on a response task and would be analogous to a recognition or cued recall test of awareness. Below a subjective threshold, subjects may be performing at above chance levels but believe that they are simply guessing. Hence, Dienes and Berry would argue that an implicit task is unconscious if the subjects believed they were guessing (termed the guessing criterion by Dienes et al, 1995) or if their confidence was unrelated to their accuracy (termed the zero-correlation criterion by Dienes et al, 1995). Dienes and Perner (1996) describe a second category of metaknowledge which is where subjects know they are making correct responses but do not know why they are correct. This second type of category is used by Reber (1993) to argue against findings such as those reported by Dulany et al (1984) and Perruchet and Pacteau (1990) who suggest that because subjects can accurately say what part of, for example, a letter string in a grammar task is the part that makes it grammatical or ungrammatical then the knowledge held must be explicit. Reber distinguishes between these two different types of knowledge by arguing that "it is very different from *knowing that* to *knowing why* ". His argument is that because subjects can make a correct decision overtly then this does not necessarily mean that the knowledge used to make that decision cannot be held implicitly.

Possibly the strongest evidence for this type of learning being at an unconscious level is the data gathered from special populations in several different paradigms. Knowlton et al (1992) found that amnesics could perform a grammar classification task at the same levels of performance as control subjects but performed more poorly than controls on a subsequent recognition task. Squire and Frambach (1990) found that amnesics could control a complex systems task as well as controls but were unable to complete an post-task awareness measure. Nissen and Bullemer (1987) and Nissen et al (1989) demonstrated that amnesics show faster RTs in the same way as control subjects do but with no awareness of the inherent sequential structure of the display. In addition, Nissen et al (1989) showed that amnesic subjects retain implicit knowledge of the sequence for one week. Shanks and St. John (1994) offer several criticisms against this sort of evidence being indicative of implicit learning. Firstly, they point out that in all these examples the amnesic subjects actually performed worse than the control subjects on the implicit tasks. They suggest that as long as amnesics show this impaired performance on an implicit task whenever they show impairment on the explicit task, this may mean that the effect is due to differential sensitivity of the two types of task (with the 'implicit' one being more sensitive to small amounts of fragmentary knowledge in the system). Whether this claim is true or not will probably only be discovered with a study using sufficient power to test whether the differences between normals and amnesics really are statistically non-significant or a case where amnesics outperform normals on the implicit task but underperform on the explicit. One finding from the study by Knowlton et al (1992) however, suggests that the information held by amnesics is different from that which they would hold explicitly. Although the patient group displayed normal grammar learning under the usual paradigm, when they were instructed to try to recall the specific strings from the study phase to perform the task their performance fell well below that of the control subjects. While it could be argued that this change of instructions caused the subjects to neglect the use of fragmentary information this does not seem very likely. Given the same learning phase and the same type of test stimuli one might imagine that more or less the same knowledge would contribute to the decision. Rather, the large decrement in performance suggests that they were instructed to use information which they simply did not have. This finding has also been reported in the memory literature by Graf, Squire and Mandler (1984). The second criticism they raise against amnesic data is that although amnesics have no explicit knowledge of the learned information at test, this does not mean that they did not have explicit knowledge of the contingencies between stimuli at study. Thus the amnesic data, they claim does show 'implicit retrieval' where there is no conscious recollection of the information at test but not implicit learning. They argue that for implicit learning to have occurred it must be shown that conscious awareness was absent during the learning phase. Reber and Winter (1994) however argue that this criticism "reveals a basic misunderstanding about amnesic functioning" and the finding that amnesics show normal implicit learning "indicates that the information necessary for successful performance can be acquired without the participation of explicit learning mechanisms".

It is clear then, that there is no simple answer to the question of whether implicit learning is mediated by unconscious mechanisms or not. The answer will very much depend on what the individual will accept as evidence for lack of awareness. There seems to be problems associated with both verbal report and recognition type awareness tasks. To attempt an answer at the question of whether implicit learning is mediated by unconscious processes, it is probably necessary to examine each paradigm separately as the criterion for one may not be the same as the criterion for another. Additionally, some paradigms may lend themselves to conscious processing more than others. A later section of this chapter will examine each of the main paradigms in detail and review the evidence for unconscious processes being involved.

Is implicit learning more powerful than explicit learning?

The definition by Mathews et al (1989) suggests that implicit learning is more powerful than explicit learning for learning covariations between nonsalient task variables. This is certainly true as the early studies on grammar learning showed that subjects could learn information about the deep structure of the grammar under incidental conditions but that instructions to look for underlying rules was detrimental to performance (Reber, 1976). Work on complex systems has shown similar results

(Berry and Broadbent, 1988) although in both these paradigms, explicit searching for rules can have a facilitatory effect under certain conditions (e.g. Berry and Broadbent, 1988; Cantor, 1980; Millward, 1981). These situations seem to be exclusively where the rule structures were either manipulated to be or were inadvertently more salient. Although this seems to suggest a clear distinction between explicit and implicit processes, there have been studies which have reported a failure to replicate this effect (Abrams, 1987; Dulany et al, 1984; Rathus et al, 1990). In sum, implicit processes do seem to be more powerful than explicit when the rules are nonsalient although the converse is not necessarily true; explicit processes may not be better when the relationships are salient as implicit processes may be learning these rules just as well. It is simply more difficult to show the effect of implicit learning when there is accompanying explicit learning.

1.4 Early implicit learning paradigms

Verbal conditioning procedure

One of the first experiments reported in this paradigm was by Thorndike (1932). Subjects were presented with between five and six hundred cards over several days with each card having four lines of equal length printed on one side. Subjects had to say which line they thought was the longest and they were provided with feedback. Most of the cards had other distinguishing features on them such as an ink smudge or a certain number and Thorndike reports that subjects seemed to learn the relationships between the correct response and the particular feature on the cards but when questioned, could not always report the correct association. Thorndike and Rock (1934) presented a verbal conditioning procedure as a word association task. The experimenter would say a word to which the subject would have to respond with the first word which came to mind. The experimenter would then tell the subject whether the response was right or wrong and the subject would get a monetary bonus or penalty depending on this outcome. A 'correct' response was one which could follow the probe word in a normal sentence (e.g. if the experimenter said 'across', the subject would be correct in saying 'the street' or 'the bridge'). An 'incorrect' response was one which was semantically related to the probe word (e.g. if the experimenter said 'over' and the subject's response was 'under'). Subjects did increase the number of correct responses between the first and last block of trials and as with the earlier experiments on cats, Thorndike concluded that the gradual increase in responding correctly meant that the subjects were not explicitly aware of the relationship between their answer and the experimenter's response. This assumption was questioned by Irwin, Kauffman Prior and Weaver (1934) and the series of experiments shown to have a number of methodological flaws (Postman, 1962).

Interest in this paradigm was rekindled by an experiment reported by Greenspoon (1955) in which subjects were required to say as many words as they could think of in 25 mins. Plural words were reinforced by the experimenter saying "mmm-hmm" and Greenspoon reported that the frequency of plural words would increase over time given this reinforcement despite the fact that subjects displayed no apparent awareness of the reinforcing contingency. This latter conclusion has been criticised by Levin (1961) as the test of awareness comprised only one or two questions which were quite vague (e.g. were you aware of the purpose of this experiment ?). Subjects may have been aware of the contingency but were unaware that this was the information that the experimenter required. Dulany (1961) replicated Greenspoon's study and found that subjects who did show learning were also aware that the experimenter said "mmm-hmm" after a correct response. Further, subjects did not realise the reinforcement was for plural nouns but rather members of the same category as the original reinforced word. Thus subjects who realised that "diamonds" was a correct response might go on to say "rubies, emeralds, pearls, etc." all of which are plural and so would seem to support Greenspoon's hypothesis as they would not have reported responding with plurals.

In a review of the literature on the verbal conditioning procedure, Brewer (1974) notes that of 19 experiments which have used rigorous methods for testing awareness, only four have shown learning without awareness and even then they are open to methodological criticism. One problem with this type of research which Brewer (1974) recognises is that subjects will actively attempt to discover what the experimenter requires of them and so will respond with many different hypotheses and strategies some of which will give similar results to those supporting the experimental hypotheses. Lieberman (1995) reports an ingenious method to address this problem which is to make subjects believe that they are participating in an ESP experiment. Subjects were put in isolation and could communicate with the experimenter through a set of lights only. They were given two words on cards and told the experimenter was thinking of one of the words and that they were to indicate which word they thought it was. Each pair of words had one which contained a double letter and it was this one which was reinforced. Thus subjects might see 'apple' or 'pear' on one trial and 'sword' and 'battle' on the next. Lieberman found that subjects chose the word with the double letter more often across trials and did not show any awareness of the contingency when tested using a rigorous 12 item questionnaire.

It seems that evidence for learning without awareness is not particularly forthcoming in the verbal conditioning paradigm although in a recent formulation, Lieberman does seem to provide evidence of such learning without knowledge of the contingency. If his results are replicable then it seems that in addition to a rigorous test of awareness, a crucial aspect of finding positive results is to give subjects a believable cover story so that other explicit hypotheses cannot account for the results. This theme will be discussed further in later sections of the thesis.

Probability learning

Reber and Millward (1968) required subjects to observe a display on which one of two lights would be activated. These events would occur with different probabilities and would be shown to the subject at the rate of two events per second for several minutes. Subjects were then required to predict which event will happen next. Reber and Millward (1971) found that subjects' predictions matched the probabilities of the events occurring even when the probabilities changed over trials according to some pattern. The evidence for this type of learning being independent from explicit knowledge of the pattern comes from a manipulation reported by Millward and Reber (1972) in which subjects saw sequences where Trial *n* was stochastically dependent on the outcome of a previous trial. This previous trial could have taken place 1, 3, 5 or 7 outcomes before the test trial and Millward and Reber found that the subjects seemed able to exploit dependencies occurring even at the longest lag of seven trials. Earlier experiments (Millward and Reber, 1968; Reber and Millward, 1965) however had found that on a recall task, beyond a lag of five trials subjects performed at no better than chance levels. Thus it seems that performance on this task can use information beyond that which subjects can consciously recall.

1.5 The paradigm which re-launched implicit learning.

The discovery of unconscious artificial grammar learning

Reber (1967) used a finite-state, Markovian system to generate examples of an 'artificial grammar' (see Figure 1). By simply moving through the diagram from start to finish and obeying the pathways, various letter strings can be formed. Letter strings which are formed by obeying the structure of the system are termed 'grammatical' (from Figure 1, examples of grammatical strings would be VXTTTV or MTVRXR). 'Ungrammatical' strings could also be formed simply by violating the order of working through the system (e.g. VXMRXV or MVRTR). Reber found that subjects became sensitive to the constraints of the rule system from exposure to a subset of grammatical items and could typically classify novel strings correctly as grammatical or ungrammatical with 60 - 70 % accuracy. This learning was termed 'implicit' as subjects were not aware of the rule-governed nature of the exemplars until the test phase and were in fact led to believe that they were participating in a memory experiment at study. Subjects were also unable to report the rules that led to a correct grammaticality decision.

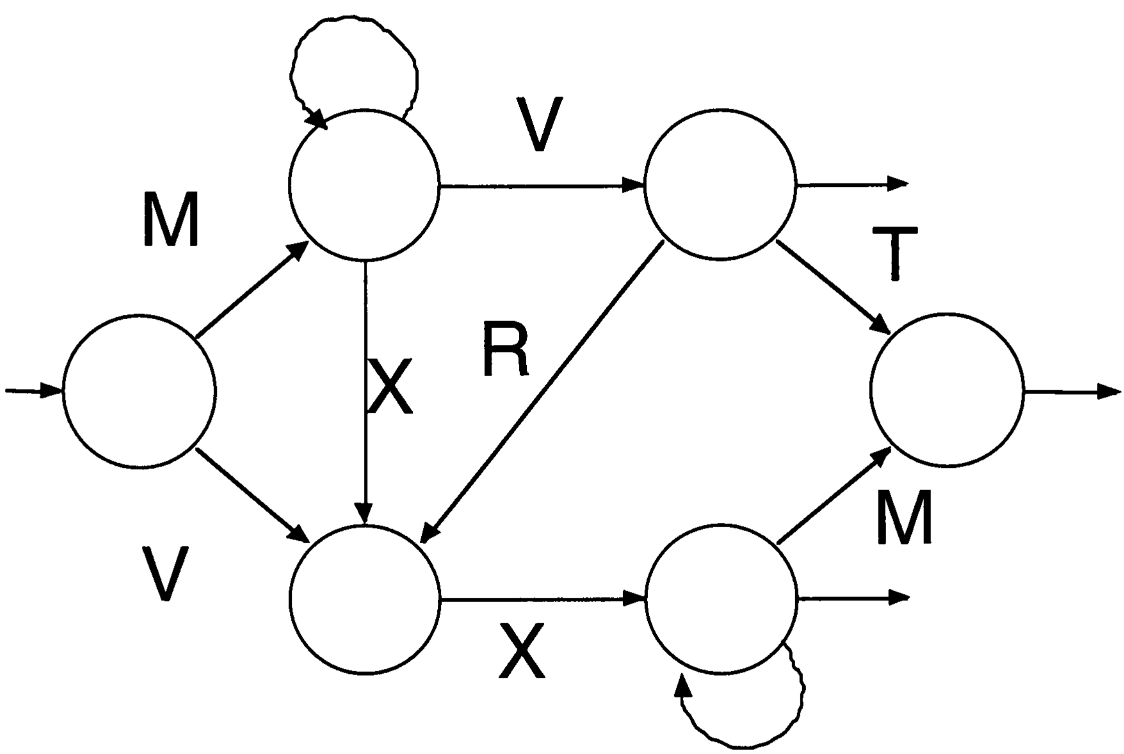

Figure 1: Example of a finite state grammar taken from Dienes (1992)

Artificial grammar has been the predominant paradigm with which to investigate implicit learning and has generated intense debate over two important questions. Firstly, what is it that subjects learn from studying grammatical exemplars and secondly, is what they learn in any sense unconscious? These two questions are sometimes seen as related although there is no strong evidence to suggest one is necessarily linked to the other. Of course these two questions are pertinent to other implicit learning paradigms however most of the fundamental research has taken place in grammar learning.

What is learned?

Initially, Reber (1969) argued that subjects had unconsciously abstracted a partial representation of the rules underlying the Markovian system from the subset of exemplars presented at test. This, Reber admits, is a rather naive view of the process underlying performance and is shown to be incorrect in a study by Reber, Kassin, Lewis and Cantor (1980). Subjects were given an extended learning phase and were shown the Markovian system used to generate the strings either at the beginning, the end or in the middle of the learning phase. Reber et al (1980) predicted that the best performing group would be those who had been shown the rule structure at the end of the learning phase as this would be explicitly imposing the correct formal structure on a mental representation which had been already formed. In fact this group gave the poorest performance suggesting that the subjects were using a mental representation that was very different from the structure used to generate the exemplars. Reber (1989) has updated his argument on the nature of the abstracted information to "a partial but representative subset of the patterns among the various elements of the stimulus display that are reflected in the environment". This view is similar to that put forward by Dulany, Carlson and Dewey (1984) who suggest that subjects learn 'correlated grammars' which although do not exactly match the underlying rule structure, do allow the subjects to respond with reasonable accuracy to that rule structure.

The abstractionist position.

Although Reber (1993) recognises that subjects may be using a rule structure that is not identical to that underlying the Markovian system used to generate a particular grammar, he is still committed to the idea of the abstraction process at learning. The evidence for such a process comes primarily from experiments showing transfer of knowledge across letter sets. Reber (1969) trained subjects on exemplars from one set of letters. Subjects were then informed about the rule governed nature of the exemplars and asked to classify new exemplars. The new exemplars however were formed from a new set of letters although they were governed by the same underlying rule structure. For example, the exemplar MRTVX would have become WSPNZ in the changed set. Of course, subjects did not even see exemplars from study in the changed letter set but exemplars which were new to them. Hence, because subjects show transfer of learning from one letter set to another, Reber argues that it is the deep structure of the grammar which is abstracted. This finding of transfer of learning between letter sets has been replicated many times in the literature (Brooks and Vokey, 1991; Gomez and Schvaneveldt, 1994; Mathews et al, 1989; Shanks et al, 1997; Whittlesea and Dorken, 1993). Altmann, Dienes and Goode (1995) report a more impressive display of transfer of structural knowledge. Subjects were trained on a letter set formed from a Markovian system and tested on a set of musical tones. Each tone corresponded to one of the letters used in the original training set. Although performance levels dropped in the crossed modality condition relative to same modality, subjects nevertheless showed some savings in performance. This suggests that at least some part of the structural information underlying the letter strings was abstracted and was subsequently available to aid decisions about the tones.

Reber and Lewis (1977) also report evidence for abstraction of the underlying rule structure. Subjects were trained on a subset of grammatical strings and were then given a test phase which required them to make grammatical strings from anagrams of a set of grammatical test items. Reber and Lewis found that the frequency of bigrams produced by the subjects had a higher correlation to the bigrams in the whole grammar than to the bigrams in the training set alone, suggesting that subjects were in fact learning more than simple bigram frequencies from the training set. Perruchet, Gallego and Pacteau (1992) argue however that such a result is artifactual in nature and a consequence of bigrams in the training set being underrepresented in the test items. Thus the frequencies of the bigrams that the subjects must produce will necessarily have a correlation with the frequency of the bigrams in the whole grammar. Perruchet et al (1992) provided experimental evidence for their explanation by training subjects only on bigrams at study (so no rule learning could take place). As with Reber and Lewis's study, the frequencies of bigrams produced by the subjects correlated more with the frequencies of bigrams in the full grammar than with those in the training set. With any sort of rule abstraction ruled out in this study, this result clearly demonstrates that Reber and Lewis's (1977) results do not provide support for an abstractionist account of learning.

<u>The non-abstractionist position.</u>

In contrast to the view that some aspects of the rule structure are abstracted from exemplars at study, the non-abstractionist position claims that subjects are learning something about the surface features of the exemplars. Reber (1993) divides non-abstractionist accounts into two separate views: the exemplar based view and the fragmentary view. The former holds that the exemplars are stored in memory as they appeared in the study task and a store of all the exemplars is formed. In the test phase of an artificial grammar task, subjects match the novel test exemplar to either one or more instances and decide to accept or reject the exemplar on some criterion of similarity. The fragmentary view rejects the notion of whole exemplars being stored in memory and proposes that subjects code fragments of the exemplars. These fragments could take the form of bigrams or trigrams (Dulany et al, 1984; Perruchet and Pacteau, 1990, 1991) or of 'chunks' of varying sizes (Servan-Schreiber and Anderson, 1990). According to this view, subjects compare the fragments of letters in the test stimuli with the fragments stored in memory from the study phase and as with the exemplar based view, calculate an index of similarity between the two to accept or reject the test string.

Subjects' reports on how they made the grammaticality judgments seem to favour a non-abstractionist account, particularly the fragmentary view. Reber and Allen (1978) found that the commonest type of information used by subjects was violation or non-violation of expected bigrams, especially those at the extremes of the string. Expectations about single letters (usually at the extremes of the string again), trigrams and longer sequences were also reported. Perruchet and Pacteau (1990) demonstrated that subjects do have considerable knowledge about the occurrence of letter pairs in a grammar. After subjects received a normal training phase on an artificial grammar, they were given a recognition test on old and new letter pairs. They judged 88% of the old stimuli as more familiar than the new stimuli. Perruchet and Pacteau (1990) went on to demonstrate that subjects are far more likely to reject an illegal bigram than a legal bigram in an illegal position within the string. They also found that subjects' performance in the grammaticality task is the same whether they were trained on strings generated from the grammar or trained on the bigrams which made up those strings (as long as no test items contained an illegal initial letter). Taken together, these findings provide strong evidence for knowledge of fragments as being solely sufficient to account for performance on a standard artificial grammar task (that is, when both study and test strings are instantiated in the same letter set).

Gomez and Schvaneveldt (1994) also found that subjects trained on bigrams were more likely to reject an illegal bigram than a legal bigram in an illegal position, however they also found that subjects trained on whole strings could reject either equally well and subjects trained on trigrams performed at an intermediate level between the other two groups. While this suggests that learning of intact exemplars at study affords better learning of some aspect of the grammar, it leaves open the question of whether it is abstraction of information from whole exemplars or if more statistical knowledge of the letter relations are the basis for this finding of better performance from studying whole exemplars. This study does demonstrate that performance on an artificial grammar task may well be determined by the study strategy adopted by subjects.

This last point makes the study of knowledge representation in artificial grammar extremely difficult. Many of the studies which claim to have discovered how subjects represent the knowledge acquired in an artificial grammar paradigm have used methods during the study phase which are very different from the 'normal' memorisation instructions. For example, Brooks (1978) gave subjects exemplars which had been generated by two separate grammars. Each string was paired with either the name of an animal or a city. This did not differentiate between the two grammars but was in fact orthogonal to them and used as a cover story for a much more subtle manipulation. The grammars could be separated by Old World animals and cities and New World animals and cities. At test, Brooks reports that subjects could classify unseen exemplars as either Old World or New World or in the case of ungrammatical items, neither. This result suggests that subjects were comparing each test string with exemplars held in memory from the study phase as it is extremely unlikely that subjects could abstract information about two categories as obscure as Old World and New World. While it is very likely that Brooks is correct in arguing the task was performed using exemplar knowledge, it is worth bearing in mind that subjects were not performing an artificial grammar task as such. Reber and Allen (1978) were able to show that it was the paired-associate technique used in the study phase that accounted for this type of representation and were able to replicate Brooks's (1978) results without the use of two grammars and an elaborate cover story.

Vokey and Brooks (n.d., cited in McAndrews and Moscovitch, 1985) noted that grammaticality and similarity are confounded in the usual artificial grammar paradigm introduced by Reber. They manipulated the test exemplars of a grammar so that grammaticality and similarity to study exemplars were orthogonal dimensions. They found that by separating the influence of grammaticality and similarity in this manner, the factor which had the greater influence on classification was similarity. This seems to point to exemplar encoding as the mode by which classification judgments are made. McAndrews and Moscovitch (1985) however criticise the study by Vokey and Brooks on several grounds, all of which McAndrews and Moscovitch claim would have the effect of biasing results against an abstraction explanation. Using more study exemplars, specific study instructions and a forced-choice task, McAndrews and Moscovitch found that the contribution of grammaticality and similarity to classification performance was approximately equal. Results from additional recognition and completion tasks suggested that both rule and exemplar based information was available, depending on the task requirements. One interesting aspect of the results was that McAndrews and Moscovitch report differences in the types of strategies used by subjects, a finding that ties well with the conclusions reached from the Gomez and Schvaneveldt (1994) study (discussed above).

Vokey and Brooks (1992) examined what they term the 'dual knowledge' hypothesis which suggests that performance is dependent on both knowledge of individual items and information abstracted across those items (Mathews et al, 1989; McAndrews and Moscovitch, 1985; Reber, 1989; Reber and Allen, 1978). They found that grammaticality and similarity were independent but additive dimensions in the classification process and that the differential influence of each depended upon the particular encoding operations at study and task requirements at test (which explains why different contributions of similarity to classification performance were found in McAndrews and Moscovitch (1985) and Vokey and Brooks (1992). They also found that when experimental conditions encouraged better learning of instances, the effect of similarity decreased on both a classification and a recognition task. The effect of grammaticality remained unchanged however. This goes against the idea proposed by Reber and Allen (1978) that encouraging better learning of individual items renders any abstraction process less efficacious. Rather, these results suggest that any difference in classification performance is more likely to be due to differential retrieval of specific items from memory.

As noted above, the non-abstractionist position has difficulty in accommodating results which show that knowledge can be transferred across surface form as long as the deep structure of the grammar remains unchanged. This can be across different letter sets (Brooks and Vokey, 1991; Gomez and Schvaneveldt, 1994; Mathews et al, 1989; Shanks et al, 1997; Whittlesea and Dorken, 1993) or across modality and stimulus form (Altmann, Dienes and Goode, 1995; Manza and Reber, in press). Brooks and Vokey (1991) however argue that this evidence for abstraction is confusing two different meanings of 'abstraction'. They suggest that subjects can show knowledge across surface forms by use of an 'abstract analogy'. In this sense, what is abstracted is a representation of each item beyond that of its surface form (e.g. the string 'MVXXT' has an underlying '12334' pattern). This is not abstraction in the same sense that Reber (1967,1969, 1976) means. He used the term abstraction in the sense of information which is collapsed across items (e.g. items must begin with an M or a V). Brooks and Vokey (1991) argued that abstraction in terms of the former sense is sufficient to account for transfer of knowledge across surface forms. To support this hypothesis, Brooks and Vokey (1991) report a similar study to that of Vokey and Brooks (1992) where the dimensions of similarity and grammaticality in the test exemplars were orthogonal to one another. Brooks and Vokey however used test exemplars which were instantiated in a different letter set to that of the study exemplars. They found that similarity to a specific study exemplar had a large effect on classification judgments even under these changed letter set conditions. This suggests that some knowledge of specific items is used in this task. Similarly, Vokey and Higham (manuscript submitted for publication) report a study where they manipulate similarity and grammaticality within a set of test exemplars instantiated with a changed letter set. The difference between this and Brooks and Vokey (1991) is that in the 1991 study, a consistent letter mapping was used for each exemplar. In this more recent study Vokey and Higham use mappings that are unique to each exemplar. This is a more stringent test of whether knowledge for each exemplar can be held at a truly abstract level as it rules out the possibility that subjects simply calculate the mapping rules and translate each exemplar back to the original letter set. Vokey and Higham found that the effect of grammaticality was attenuated in switching from the same letter set to a changed letter set but the effect of similarity was not. Subjects' performance at the same level for decisions based on similarity, whether the test is on the same letter set, a changed letter set generic to all items or a changed letter set unique to each item suggests that the level of abstraction is far greater than surface characteristics or simple remappings of such surface characteristics. As the effect is independent of the grammaticality dimension, the abstraction does not seem to be that of general rules of the system.

Taken together these results diminish the force of the argument that transfer across form is evidence for abstraction across a set of items. Note that grammaticality still has an independent role to play in the classification performance and it may be that this is evidence of this sort of abstraction. As Vokey and Higham point out however, this effect of grammaticality may be due to aspects of similarity that are uncontrolled for in present experiments.

The weight of experimental evidence then seems to favour some form of abstraction but not abstraction of information across instances, rather abstraction of regularities within each exemplar. As shown by Gomez and Schvaneveldt (1994), Perruchet and Pacteau (1990), Reber and Allen (1978) and Vokey and Brooks (1992) differences in the study and test exemplars and in task instructions to subjects may vary the way in which subjects encode the information from the grammar. This viewpoint is emphasised by Whittlesea and Dorken (1993) and is termed the 'episodic-processing' account of implicit learning. This account does not necessarily support either the abstraction or the exemplar based viewpoints but instead states that what is learned is that aspect of the stimulus structure that is necessary for the required task demands. Whittlesea and Dorken demonstrated this point by requiring subjects to make grammaticality decisions to items such as ONRIGOB, having seen study items such as ENROLID which were generated from a finite-state grammar. During the study phase however, subjects were required to either say the item or spell it out loud. At test, subjects were required to do the same before they made a grammaticality judgment. Whittlesea and Dorken (1993) found that where there was an overlap in processing conditions, subjects could perform the classification task reasonably well. This was not the case when the processing conditions were mismatched with subjects being unable to correctly classify the exemplars.

1.6 Models of implicit learning in artificial grammar

Recent interest in cognitive modelling has led to various implicit learning paradigms being implemented on different models of cognitive architecture. Use of such modelling techniques allows for testing of assumptions underlying implicit learning and greater specification of the processes that are necessary and so introduces much more theoretical constraint than is usually achieved through experimental methods. The models can address both main questions of interest in implicit learning: in what form is the knowledge that subjects acquire and to what extent is that knowledge acquired without conscious awareness? Most attention has been paid to modelling the artificial grammar task as this seems to be the predominant paradigm in implicit learning. The following section will introduce the models which have been used in relation to artificial grammar and summarise how each model performs in accounting for data from real subjects. Models relating to other paradigms will be introduced later in the relevant sections.

The ACT* architecture

Anderson (1983) developed a model of human cognition based upon three types of memory store: working, declarative and production. (see figure 2). Working memory contains the knowledge which is currently available to the system and is stored in declarative format. Information can enter working memory from three sources. It can be retrieved from the long term declarative memory store, it can be encoded from the outside world or it can be deposited as a result of the actions of productions. Productions are rules which follow an 'IF x THEN y 'structure where x is a particular state in working memory and y is a mental and / or physical action.

Information in declarative memory can be translated into a procedural format and in this way, the ACT* architecture can account for skill learning where knowledge which is originally declarative in nature, with practice becomes more automatic and less available for verbal description. This is done by a process which is termed 'knowledge compilation'. This is a two stage process which consists of 'proceduralisation' and 'composition'. Proceduralisation is where general productions act on declarative knowledge to produce specific productions. This formation of specific productions means it is no longer necessary to retrieve declarative information from the declarative store but rather initiate action as soon as the relevant 'IF' information appears in working memory. The second process in proceduralisation is composition which is where a series of productions are collapsed into a single production which has the same effect as the series.

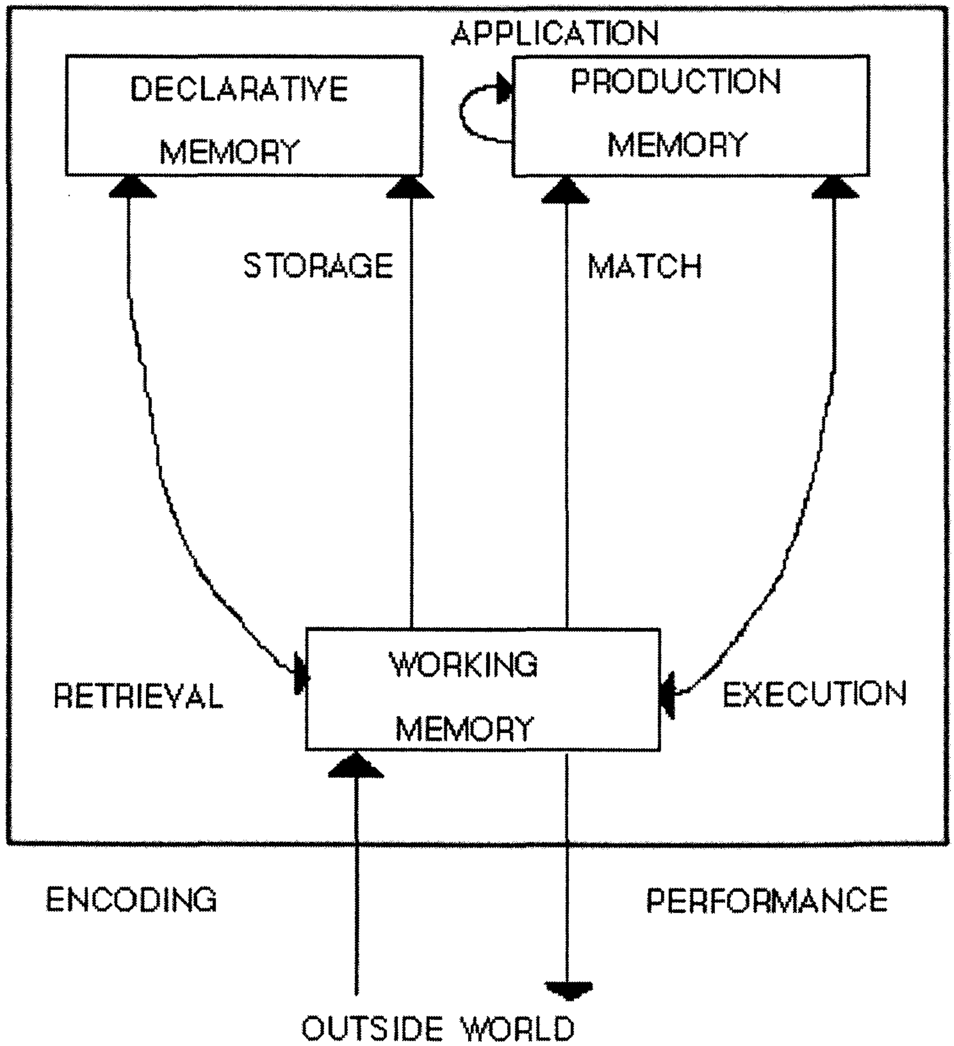

Figure 2: Diagram of ACT* architecture.

These processes can account for speed-up in responses and specific stages within a skilled action being less open to introspection. Skill acquisition continues even after this however, with learning of when to apply a particular production taking place. ACT* has three learning mechanisms for achieving this: generalisation, discrimination and strengthening. Generalisation creates a new production which captures what a pair of other productions have in common. Discrimination, as its name implies, seeks to ensure that a particular production is used only in the correct circumstances and is achieved by adding a test to the production to identify situations in which it has previously failed. The strengthening function changes the likelihood of a given production being used by examining the number of correct applications that the production is involved in.

There are several ways in which information could be inaccessible to conscious awareness in the ACT* architecture. Firstly, proceduralisation constructs rules which can act on information which appears in working memory without needing to retrieve information from the long-term declarative store. If only the information which is currently held in working memory is accessible to consciousness then much of the declarative information which has been built into a production would thus not be available. A second way is through the process of composition where several productions are merged to form a single new production. This has the effect of increasing the 'THEN' components which, when placed in working memory to be executed, may overflow the limited capacity available (Neves and Anderson, 1981). This would have the effect of some of the actions which are performed being lost from memory.

The architecture of ACT* is inherently rule-based. While some effects in implicit learning seem rule based there is also evidence (e.g. Vokey and Brooks, 1992) that learning can be based on similarity to individual items and also to fragments of items (e.g. Perruchet and Pacteau, 1990). This would suggest that ACT* would have to make a separate rule for each instance or fragment which is not a particularly parsimonious explanation. In addition, the ACT* architecture requires that conscious information is initially available to the system as encoding takes place via working memory. As discussed earlier, implicit learning differs from automatisation in that the information required for the task is acquired unconsciously right from the start but may become accessible with extended practice (Sanderson, 1989; Stanley et al, 1989). With both these considerations in mind it seems unlikely that the ACT* architecture can provide an adequate mechanism for implicit learning.

The Competitive Chunking Model

Servan-Schreiber and Anderson (1990) carried out a study which demonstrated that subjects tended to recall strings learned in a grammar task as 'chunks'. Additionally, some subjects learned strings that had already been split into chunks. Subjects could learn these strings just as easily as strings presented in the normal fashion as long as the strings were chunked in a consistent manner across the study set. Subjects had much more difficulty in memorising strings which had been chunked at random. Subjects were also better at classifying nongrammatical strings which contained invalid chunks than those which contained valid strings in an invalid order. These findings suggested that subjects engaged in a chunking procedure in order to perform the task. The Competitive Chunking Model is based on these results and posits that during the test phase each string is analysed in terms of the single letters which make up the string. Then these base units are formed into higher level chunks by a competitive process whereby the two or three item chunks that are possible are given a certain level of support from their sub-chunks. The amount of support is determined by an average of a parameter known as 'chunk strength' and is a measure of the frequency and recency with which chunks have been used. The amount of support a higher level chunk is given will determine the likelihood of that chunk being retrieved and used by a subject as a representation of the stimulus. This process carries on until the string is represented by a single superordinate chunk or until no higher order chunks receive enough support to be retrieved. A string which is represented with only one superordinate chunk is taken as being more familiar than a string which is represented by several higher order chunks.

Servan-Schreiber and Anderson modelled several aspects of the data gathered by Reber (1967) and found that like real subjects, the Competitive Chunking Model recalled more and more grammatical strings as opposed to random strings on successive trials. The model could also classify strings at the same levels of performance as subjects and was more likely to correctly classify a string as nongrammatical if the invalidity occurred at the end rather than in the middle of the string. The model can also account for the finding by Allen and Reber (1980) that knowledge of the grammar is robust over a two year period. It does this by assuming that the base elements (i.e. the individual letters) are not subject to decay as are the higher level chunks and so a subject can reconstruct the grammar from these elements.

The model also fits in with the idea of implicit learning being an automatic and unconscious process. Servan-Schreiber and Anderson suggest that the processes of creating new chunks and strengthening existing ones are basic processes which happen all the time and so may be considered automatic. While they suggest that knowledge of the actual contents of the various chunks may be accessible to conscious awareness, the related strength and support parameters may not. While this model seems an impressive candidate for the mechanism by which artificial grammar tasks are performed there is one piece of data which is problematic. The chunking procedure is tied to the surface forms of the stimuli and hence cannot account for the data from real subjects which show that they can transfer their knowledge across surface form (e.g. Altmann, Dienes and Goode, 1995; Brooks and Vokey, 1991; Gomez and Schvaneveldt, 1994; Mathews et al, 1989; Shanks et al, 1997; Whittlesea and Dorken, 1993).

THIYOS - THe Ideal YOked Subject

Mathews et al (1989) found that when naive subjects were given periodic instructions from other subjects who were at the time learning to classify grammar strings, the naive subjects' subsequent classification performance fell between chance levels and the level of the original subjects, suggesting that the original subjects could not or did not verbalise all of the information which they were learning during the study phase. They also found that the instructions given to the naive subjects varied across individuals and did not converge even after extensive practice. Druhan and Mathews (1989; cited in Berry and Dienes, 1993) proposed a model which could account for both of these effects. The model was a classifier system (c.f. Holland, Holyoak, Nisbett and Thagard, 1986) which was given the rules suggested by the original subjects. Classifier systems work on the basis of a message list (which is similar to a working memory) which can have messages posted to it by classifiers or rules. Classifiers are composed of 'conditions' which tell the classifiers what message to look out for on the list and an 'action' which tells the classifier what message to post should a critical message appear on the list. For a classifier to post its message, it must have the necessary condition satisfied and it must also reach a criterion via a probabilistic function of certain values within the system. One of these values is 'strength' which is a measure of past performance of the particular classifier. Every time the classifier posts a message its strength is reduced and every time it participates in a successful decision its strength is increased. This process in known as 'tuning'. Druhan and Mathews (1989) found that with no tuning of the system, THIYOS performed between chance and the original subjects; i.e. similarly to the yoked subjects. When the system was maximally tuned, THIYOS performed as well as the original subjects.

The conclusion that Druhan and Mathews reached was that the rules were accessible to consciousness and were reported. What was not accessible to consciousness was the relative weightings which each rule had and it was this information which the yoked subjects lacked. The explanation for the rules which subjects offered not converging was that seeking new rules should be driven by a failure of the current rules used (Holland et al, 1986) and so the rules formulated by different subjects should not become more similar over time if they operate with a moderate degree of success.

Rousel and Mathews (n.d. cited in Berry and Dienes, 1993) extended THIYOS to learn rules rather than have them as original input by the use of a forgetting algorithm. Thus THIYOS was able to learn fragmentary knowledge from exemplars and use these to classify novel strings. The interesting thing about this model is that it could also encode abstract information by noting whether the previous letter in the string was the same as or different from the previous letter. With this function, THIYOS could classify test strings instantiated in a changed letter set which is a major stumbling block for other types of model.

Exemplar models

Several theorists (e.g. Brooks, 1978; Vokey and Brooks, 1992) have suggested that implicit learning involves the memorisation and later use of whole exemplars. As discussed previously, several studies (McAndrews and Moscovitch, 1985; Vokey and Brooks, 1992) show evidence that similarity is important in the grammar classification task but there is another important orthogonal factor of grammaticality to be considered. This factor of grammaticality is possibly the effect of Reber's (1989; Reber and Allen, 1978) idea that implicit knowledge is abstracted across a set of learning items. This finding of some information seeming to be item-specific and some information seeming to have been abstracted across a set of items is not only found in implicit learning but has also been documented in learning of random dot patterns (Posner and Keele, 1968, 1970). Exemplar models may offer an explanation for these results in that they encode item specific information and so would account for the similarity effect but they also can act as though they were abstracting information across study exemplars. Several authors have proposed exemplar models of memory (e.g Estes, 1986; Hintzman, 1983; Medin and Schaffer, 1978; Nosofsky, 1986), most of which differ in terms of how they code similarity.

To take one of the exemplar models in the literature as an example, Hintzman (1986) developed a multiple trace model called MINERVA 2 which encodes all exemplars as memory traces. Each exposure to an exemplar lays down a trace (the process of laying down multiple traces of the same exemplar is known as multiplexing) and these traces are subject to decay over time. This decay operates on parts of the individual exemplar and means that partial recollection of an exemplar can take place (Hintzman encodes the various parts of the exemplar using the values of 1 and -1 for feature present and feature absent, and a zero for feature not encoded or not relevant). As well as being subject to decay, MINERVA 2 assumes that encoding is not a perfect process and incorporates a learning parameter which can be adjusted to reflect variations in encoding. Retrieval is assumed to operate by means of a probe which contains information similar to that of a trace item. When a probe enters the 'short-term memory' store of the system, it activates all traces stored in the 'long-term memory' store in parallel. Each trace is activated as a function of the probe's similarity to that particular trace. It is this parallel activation that can account for abstraction-like effects as any commonalities across traces are activated together. When a probe is presented it causes an 'echo' to be returned to short-term memory.

This echo has two properties: intensity and content. Intensity is the sum of the activation across all traces and it produces a single number corresponding to the overall activation of the trace and is thought to be an index of the familiarity of the item. Content is the sum of the activation for each individual element in the traces and is thought to index the explicit memory of a particular item. In terms of specific tests of memory, the echo intensity would be used for recognition and content for recall.

These processes are both assumed to be available to consciousness as the echo intensity and content are both returned to the short-term memory store. Dienes (1993) notes that what may remain inaccessible to consciousness is each individual stored exemplar, the similarity between each of these exemplars and the probe, and how each of these combines with the others to form the overall similarity. This would account for subjects being able to make a classification decision without being able to say how they reached that decision. It is less clear however that if subjects are consciously aware of the average similarity of the test string to the traces in memory that their confidence would be unrelated to their classification accuracy as Chan (1992, cited in Berry and Dienes, 1993) has observed. In terms of what is stored, MINERVA 2 has shown that the encoding of individual exemplars can account for abstraction like effects (and also prototype abstraction) without the abstraction actually taking place. Although this accounts for the grammar task findings in a very satisfactory way, there are problems when one tries to assume a mechanism such as this applies to real life situations. McClelland and Rumelhart (1985) give their word perception model as an example. The model assumes linguistic rules emerge from a "conspiracy of partial activations of detectors for particular words" without the need to posit rule abstraction. They suggest that this is inconsistent as there has to be an abstraction of every single occurrence of a word to put it into the same single word detection unit. The response to this is to claim that each representation of a word is itself a conspiracy of individual exemplars (which McClelland and Rumelhart term an 'enumeration of specific experiences') but this leads to the acceptance of a mechanism with an unlimited storage capacity and tools capable of searching such a store. These problems led McClelland and Rumelhart to suggest an alternative model by which individual exemplar storage might account for abstraction effects.

McClelland and Rumelhart's (1985) distributed model of memory

With their distributed model of memory, McClelland and Rumelhart (1985) posit a system which like exemplar models views memory as a storage of traces from specific instances. Also, abstractive effects can be simulated by the activation (or 'superposition' as McClelland and Rumelhart define it) of these memory traces. Where this model differs is that the superposition happens at the time of encoding so that each trace is not stored individually but as a composite of traces. The models consists of a number of simple processing units which have connections to many other units. This has led these types of models to be compared with the neural substrate and so considered more neurologically plausible than other cognitive models. A particular mental state arises from a particular pattern of activation across these simple units. Hence, viewing a particular word or grammar string would give a particular pattern of activation which would be replaced by another pattern when another word or string was encoded. The memory of the event takes the form of a change in the connection weights between the relevant units. Presenting part of a pattern should lead to the whole instantiation of the pattern being reproduced. The model deals with abstraction by being able to capture any structure which is inherent in a set of patterns. The model allows abstracted prototypes and the individual exemplars to co-exist. McClelland and Rumelhart make this explicit by asserting that "given part of a visual pattern, the model will complete it; if the part corresponds to the prototype, then that is what is completed, but if it corresponds to one of the repeated exemplars, that exemplar is completed."

It seems then that a distributed model can account for the specific effects of individual exemplars and the abstraction effect that takes place across exemplars. The model also assumes that the patterns which are formed are accessible to consciousness but the weights between the units would not be as they are embedded in the processing mechanism (McClelland, 1988; cited in Dienes, 1993). While this model addresses the problem of postulating unlimited storage capacity which would be required by exemplar models, like the exemplar models the distributed model cannot account for transfer effects in implicit learning of grammars although Dienes (1993) suggests that using hidden units to form abstract structures may resolve this problem.

It seems than that of all the models proposed to account for grammar learning, the distributed model is the one which is the most theoretically plausible, especially if the modifications suggested by Dienes can allow the model to acquire abstract information which can be used across form. An experimental demonstration of the superiority of the distributed model was provided by Dienes (1992) who tested several different types of exemplar and distributed models against the performance of real subjects on measures such as average classification performance, proportion classified correctly twice, once or not at all and prediction of the rank order of exemplar difficulty. Dienes found that the only model among those he tested was a distributed model which used a delta rule for updating the link weights could account for the pattern of results shown by real subjects. Hence this gives convergent evidence with that of the experimental studies to show that subjects may be capable of a form of abstraction across study exemplars.

1.7 Other major paradigms in implicit learning

Hidden Covariation Detection

The hidden covariation detection paradigm described by Lewicki (1985, 1986a, 1986b) involves the presentation of complex stimuli to a subject. Each stimulus has two nonsalient features which covary and a number of other irrelevant features. These features do not covary in real life and so can only be learned from the study phase of the experiment. During the test phase the subjects see a stimulus with only one of the critical features present and are instructed to make a judgment about the other critical feature. The results suggest that subjects learn the covariation between features and respond accordingly in the test phase. Examples of the types of covariations used are: complexity of brain scans and intelligence (Lewicki, Hill and Sasaki, 1989; Lewicki, Hill and Czyzewska, 1994), hair length and personality (Lewicki, 1986a, 1986b) and facial proportion and personality (Hill, Lewicki, Czyzewska and Schuller, 1990). Most of Lewicki's experiments used a semi-structured post task interview to assess awareness of the covariation. Some experiments used pilot studies where subjects were explicitly told to look for a covariation between features. Negative results for both the above methods suggest that learning in a hidden covariation paradigm is implicit although this conclusion has recently been questioned by Hendrickx, De Houwer, Baeyens, Eelen and Van Avermaet (1997). They point out that while subjects do not report the correct covariation in an awareness test, this may be because subjects are basing their decisions on another partially correct covariation between the critical features and a supposedly irrelevant feature (these were not controlled but simply assumed to vary randomly). Thus it is possible that they did report a covariation that would allow success on the task but not a covariation which the experimenter would classify as correct. Hendrickx et al (1997) attempted to control for this and other problems associated with the assessment of awareness in this paradigm but found it very difficult to replicate Lewicki and colleagues' results. In fact, of three exact replications Hendrickx et al found evidence for learning in only one. In addition, other experiments similar in nature to those reported by Lewicki yielded little evidence for the learning of the covariation and so Hendrickx et al conclude that their results "cast doubt on the generality and robustness of hidden covariation detection". Lewicki, Hill and Czyzewska (1997) suggest several reasons why Hendrickx et al (1997) had difficulty in finding learning of the covariation. They suggest that the stimuli used by Hendrickx et al were not as rich and complex as those used in their original studies and also that the covariations between elements in the stimuli were extremely subtle (e.g. 21 vs. 23 individual hairs on schematic faces as opposed to long vs. short hairstyle in Lewicki (1986)). As noted by Hendrickx et al (1997) this is the first time that a suggestion of boundary conditions for hidden covariation detection has been proposed in the literature although one might reasonably have expected there to be some sort of constraints in learning.

There are precedents in the literature, however for such boundary conditions in other tasks. Lewicki et al's (1997) suggestions that the task cannot be too subtle, must have stimuli which are subject to some sort of holistic analysis and must be believable to the subject seem to have already acquired the status of unspoken assumption. Reber (1989) suggests that the failure to replicate the effect of rule search instructions having a detrimental effect on learning by Abrams (1987) was due to the use of a computer to present the instructions. Reber suggests that "simply presenting the explicit instructions on a computer screen may not have the compelling quality that a real experimenter reading them has". In several studies on perceptual processing Smith and Kemler Nelson (1984) found that certain conditions (e.g. speeded deadline, concurrent tasks and incidental learning instructions) designed to force subjects to use a "fallback mode of cognition" induced more holistic (i.e. a comparison of overall similarity) processing of stimuli. Holistic categorisation was also found by Smith and Shapiro (1989) using these types of condition. Hence there seems to be a link between holistic types of perception and categorisation with task conditions such as those which have been used to demonstrate effects of implicit learning. This suggests that the boundary conditions proposed by Lewicki et al (1997) may not be specific to a particular implicit learning task but may be more general conditions under which the perceptual or categorisation processes possibly associated with implicit learning might work.

Control of complex systems

Complex (or dynamic) systems usually involve some sort of complex task which has input and output parameters governed by a complex equation. The surface structure of the task involves a task based on some simplified real life scenario such as controlling a factory or a city transport system with subjects entering input values and trying to maintain the output values to within a given range. The interesting thing about this task was that although subjects showed an improvement in controlling the task with practice, their verbal descriptions of how they were controlling the task did not match their performance (e.g. Broadbent, 1977; Broadbent and Aston, 1978, Broadbent, Fitzgerald and Broadbent, 1982). Berry and Broadbent (1984) note that although these data seem to suggest a dissociation between conscious knowledge and task control, this difference can always be accounted for by methodological inadequacies. Berry and Broadbent therefore examined three different factors which might have had differential effects on task performance and verbalisable knowledge.

The tasks used were the simulation of a sugar production factory and a interaction with a 'computer person'. In the sugar production task, subjects were required to control the production rate by varying the number of workers from 100 to 1200 in steps of 100. The output was tons of sugar ranging from 1000 to 12000 and was given by doubling the subjects' input subtracting the previous sugar output (all divided by 100 or 1000 to give integers from 1 to 12). The computer randomly added 1, -1 or 0 to this number on each trial before converting the value into thousands of tons. The use of this equation and the random factor ensured that several different outputs could be associated with a particular input. The subjects were required to reach and maintain an output of 9000 tons with a starting workforce of 600 workers. The person interaction task used exactly the same relationship between input and output variables although the surface features of the task were very different. Here, subjects were told that they would be introduced to a computer person (called Clegg) who could display twelve different types of behaviour ranging from 'very rude' up to 'loving' (and these corresponded to the tonnage of sugar output in the other task). The subject was required to show one of these particular behaviours towards Clegg in order to make Clegg 'very friendly' and keep him that way. In order to access verbalisable knowledge about the task, Berry and Broadbent gave subjects three types of question. The first was a description of a change of state in the system and a multiple choice answer about final output. The second was three pairs of input and output values and subjects had to make a prediction about another output given the input. The third question simply asked subjects how they went about controlling the task.

Berry and Broadbent (1984) found that previous results were replicated using these two tasks in that subjects performance did not bear any relation to their ability to answer questions about the system. Additionally, practice at the task showed an improvement in performance but not in answering the questions. They found however that giving subjects detailed verbal instructions on how to control the system had no effect on performance but did increase subjects' ability to answer questions. Concurrent verbalisation however had a mixed effect by improving task performance when used with verbal instruction on how to perform the task but no effect on its own. They argue that these results do not support a general dissociation of task performance and verbal knowledge but rather that under certain circumstances these tasks may be performed in an implicit manner. Berry and Broadbent (1988) went on to examine some of the circumstances in which implicit performance might occur and what sort of mechanism might account for the observed dissociation. They suggested that two modes of learning might be successful in a task such as the complex systems they devised. One way would be to learn all the contingencies involved and to store a large number of condition-action links. Performance would improve as the correct links became stored but verbalisation might be difficult due to the sheer number of links and the low levels of confidence associated with them. This would be a slow process as there is no selection of links. The other mode is a selective one where the contingencies learned are few but provided they are the correct ones, learning will be fast and accurate. The drawback of this mode is that if the wrong contingencies are learned then performance will be bad. The knowledge associated with this type of learning however is assumed to be verbalisable due to the small number of contingencies that would have been learned. Berry and Broadbent suggested that explicit search instructions would be likely to induce the selective mode whereas incidental instructions in a complex task would be more likely to initiate the hypothesised unselective mode. They also manipulated the likelihood of the correct contingencies being selected in two ways; firstly by making one version of a complex systems task 'salient' (by having a computer person always respond lower on the output scale by a fixed amount) and another version 'non-salient' (by using the same relationship as the salient task but having the output lagged by one response). Secondly, they informed one group of subjects about the exact relationship between input and output variables. The salient and non-salient tasks were validated as 'explicit' and 'implicit' by control subjects who showed an improvement in performance for both tasks but could only answer questions relating to the task for the salient condition. The explicit rule search instruction had the proposed effect of facilitating performance on the salient task but inhibiting performance on the non-salient. Explicitly telling subjects what the relationship was had the effect of facilitating performance on the task and on the post-task questions (the subjects controlling the salient task learned this anyway). This suggests that saliency of the relevant variables involved in a task is critical in determining whether that task will be carried out using a selective or an unselective mode of learning. Berry and Broadbent (1988) also found that for a non-salient version of a complex systems task, the underlying knowledge can be transferred to a task which is similar (e.g. bus transport system to a train transport system) but not to a task which has dissimilar surface characteristics but the same underlying relationship between variables (e.g. bus transport system to a computer person task). It has been suggested that implicit learning in this task takes the form of a 'look-up table' (Broadbent et al, 1986) with specific instances being stored and used again in situations in which they were previously successful. This would account for the findings of lack of transfer to a task which was dissimilar at a surface level as previous instances would not be recognised and a new table would probably be required. Marescaux, Luc and Karnas (1989) tested several predictions of a look-up table model on a test involving questions about specific situations from the sugar production task and found that as predicted, subjects performed better on situations which they had already experienced and subjects showed a consistent response to the same situation. These results were confirmed and extended by Dienes and Fahey (1995) who also used the salient and non-salient versions of the computer person task to train subjects. They found that subjects trained on the non-salient version could not use the acquired knowledge beyond previously encountered situations but subjects trained on the salient version could.

Dienes and Fahey (1995) also reported the outcomes of modelling the data with an instance based model and a rule based model. The instance based model was based on Logan's (1988, 1990) instance theory in which each representation of a stimulus is encoded separately and 'race' against each other and explicit strategies to provide a solution to a task. Dienes and Fahey assumed that any input-output pairing that was what they termed 'loosely correct' would be encoded as an instance and explicit knowledge was represented by a certain number of instances that were activated by any stimulus. When presented with a particular situation, the model would 'race' all the instances assumed to represent explicit knowledge and those instances which were specific to the situation to generate a solution. The rule based model was similar to this in that the rules the model had were 'raced' against one another to determine a solution. Each time a rule produced the correct response the model included another instantiation of that rule thus increasing the probability of it being used again. Dienes and Fahey found that the two different models accounted for performance on different types of task. A non-salient version of the person interaction task could be performed by the instance based model to the same level as subjects but the same was not true for a salient version. Conversely, the rule based model could account for the data from a salient task but not from a non-salient version. These findings suggest that a look-up table can indeed account for performance on a complex systems task and fit with the data reported by Berry and Broadbent (1988) where explicit instruction on the underlying rules did not improve task performance on the non-salient task.

Sequence learning

Nissen and Bullemer (1987) developed a task (the sequence reaction task or SRT) which seemed to suggest that sequentially organised information can be acquired without conscious knowledge of the sequence. The task required subjects to press a particular key as quickly as possible when a light associated with that key lit up. There were four lights arranged horizontally and these would light either randomly or according to a repeating ten-item sequence, although subjects were led to believe all displays were random. Nissen and Bullemer found that subjects' reaction times (RTs) decreased over exposure to the sequential condition but not to the random sequence and that subjects who had been exposed to the sequential condition showed a large increase in RT when switching to a random stimulus display. Most (but not all) subjects could also report parts of the sequence but an amnesic (Korsakoffs' psychosis) group displayed the same facilitation in RTs to the sequential task as the normal subjects but without noticing any pattern. Similar findings have been reported in other special population groups such as Alzheimer' disease patients (Knopman and Nissen, 1987) and also normal subjects with scopolamine-induced amnesia (Nissen, Knopman and Schacter, 1987). This suggests that a dissociation between the performance measure and declarative knowledge of performance should also be observed in normal subjects under certain conditions. Willingham, Nissen and Bullemer (1989) gave subjects a task in which they were presented with a ten-item sequence consisting of an asterisk which could appear on a screen at one of four horizontal positions. Subjects had to press a key corresponding to the location of the asterisk as quickly as possible. Subjects were questioned about the presence of a pattern after the task and also required to perform an additional measure of their explicit knowledge by predicting the next element of the sequence rather than reacting to it. Although the level of explicit knowledge concerning the sequence was different for different subjects, Willingham et al demonstrated that subjects who were classified as having no awareness of the sequence still showed a facilitation in the performance measure. Shanks and St. John (1994) however suggest that subjects may simply have been learning the relative frequencies of each part of the sequence as in a ten item sequence each item could not appear an equal number of times. They argue that subjects were not asked about this information and so this weakens the claim that they were in fact learning unconsciously. Shanks and St. John also argue that the 'no-awareness' subjects in Willingham et al's study did have some level of awareness as they had prediction scores which were higher than control subjects although not significantly so. In a replication, Shanks, Green and Kolodny (1994) found that this category of subjects did have significantly greater prediction scores than control subjects. Although this does seem to suggest that subjects may have had some explicit knowledge of the sequence which can be tapped by a more senaitive test, it does not rule out unconscious learning. To do this one would have to assume that tasks are process pure and that any level of conscious awareness excludes the possibility of unconscious learning. Perruchet and Amorin (1992) criticise the use of the generation task by noting that Willingham et al did not inform the subjects that the same sequence would be used in the generation task as was used in the RT phase. The subjects were informed however that "the procedure would change in such a way that instead of pressing the key below the stimulus, they were to press the key corresponding to where they thought the next stimulus would appear". It seems that implied within these instructions, nothing else would change and one might suspect that subjects would realise the stimulus display would be the same (especially if they had noticed a sequence, been asked about a sequence and then asked to predict where the next item would appear). Thus this criticism raised by Perruchet and Amorin does not seem to be a particularly strong one. They also argued that because Willingham et al only caused the stimulus in the generate task to change once the subject had made the correct response, this made recalling the sequence more difficult. Perruchet and Amorin changed the generate task to account for these two criticisms and found that subjects could give partial information about the sequence. They went on to use a recognition task for measuring explicit knowledge where subjects saw four-items chunks and had to say whether they were part of the sequence or not. Subjects were able to do this and they argue that this shows subjects have explicit knowledge of the sequence although as several authors (e.g. Cleeremans, 1994; Dienes and Perner, 1994) point out, such a task may tap implicit knowledge as well as explicit.

Lewicki, Czyzewska and Hoffman (1987) used a similar type of task where subjects were required to press a key corresponding to a target appearing in one of the four screen quadrants. The first six trials showed only the target but the seventh had the target embedded in a field of 35 distractors. The measure of learning was subjects' RT to the seventh target. The position of the seventh target was actually determined by the position of the target on the first, third, fourth and sixth trial and 24 such combinations were used as predictors. Lewicki et al found that subjects' RTs to the seventh target gradually decreased with practice but when they changed the position of the seventh target to the diagonally opposite quadrant from the expected one, the subjects showed a large increase in RT. Despite this apparent acquisition of knowledge on a performance measure, Lewicki et al found that subjects were unable to report anything about the rule-governed nature of the display. Again, Shanks and St. John (1994) have criticised this task as knowledge of only part of the sequence proved informative as to the location of the seventh target but Lewicki et al only accepted knowledge of all four predictive trials plus the location of the seventh trial as evidence of conscious awareness. They also again mention the lack of sensitivity of a free recall task to tap awareness. Stadler (1989) replicated Lewicki et al's results and included a prediction task as an additional test of awareness with no feedback for subjects (thus avoiding Perruchet and Amorin's (1992) criticisms). Stadler found that neither the test of verbal awareness nor the prediction task showed that subjects had any awareness of the rules underlying the task yet they still showed clear evidence of learning as measured by RT. Dienes and Berry (1993) point out however that the prediction task tested each rule only twice and the performance task used eight repetitions of each rule. They suggest that increasing the power of the prediction task might have led to some awareness being shown.

Other evidence for unconscious processing.

Some evidence has suggested that attention may be implicated in these implicit learning studies. This would of course, question such learning being truly unconscious. Nissen and Bullemer (1987) found that introducing a secondary tone-counting task caused subjects responding to a sequence to improve in RT by no more than a control group responding to a random sequence. Cohen, Ivry and Keele (1990) suggest that this may be due to the type of sequence used and went on to show that learning second order conditional sequences (these are where there are no unique relationships between items in the sequence) required attention but that learning first order conditionals (where each item is uniquely associated with the next item) can occur without attention (a finding also replicated by Frensch, Buchner and

Lin, 1994). McDowall, Lustig and Parkin (1995) suggest that Nissen and Bullemer's failure to find learning may be due to the secondary task affecting not the learning but the expression of the knowledge and show that with this taken into account, subjects do show evidence of sequence learning under dual-task conditions. Stadler (1993, 1995) provides evidence that the apparent disruption of learning with a concurrent tone-counting task is not in fact due to increased attentional load but rather to a disruption of the sequence organisation. Hence, there seems to be good evidence for learning without awareness of sequences being able to occur without requiring attentional resources and contradictory findings being due to factors other than attention.

Eimer, Goschke, Schlaghecken and Sturmer (1996) provide electrophysiological evidence of a dissociation of awareness in the sequence learning task. This is especially interesting as Eimer et al argue that ERPs provide an online measure of awareness and therefore are unaffected by memory-retrieval processes or effects of informing subjects about the structured nature of a task prior to a direct test. Eimer et al gave subjects a repeating sequence which had 'deviant' stimuli replacing the standard items at unpredictable points in the sequence. In this way, ERPs could be collected separately for standard and deviant items. They found that a negative component of the ERP occuring at around 200msec after stimulus presentation was associated with those subjects who displayed explicit awareness of the sequence on verbal and recognition tests. This component was either smaller or absent from those subjects classified as unaware. Eimer et al view this as "at least suggestive evidence that the acquisition of sequential knowledge in the serial RT task in explicit and implicit groups was indeed mediated by different underlying brain processes".

Buchner, Steffens, Erdfelder and Rothkegel (1997) have developed a model based on Jacoby's (1991) process dissociation procedure. Using this model they have attempted to separate the individual components of conscious recollective experience and unconscious perceptual fluency. They found that subjects given intentional learning instructions had a parameter related to conscious processing which was significantly higher than that of subjects given incidental instructions. The parameter related to unconscious processing did not differ significantly between the two groups however. This suggests that there are two components present in sequence learning and that a process dissociation procedure may be a useful tool in separating out their respective influence on the task although at present the separate influence of these parameters is not well defined.

In sum, there seems to be evidence that some subjects learn the sequence without awareness of the sequence as measured by verbal report. Whether awareness can be shown by more sensitive measures such as prediction and recognition measures is still arguable as is the underlying processes which these type of test might actually tap into. The study by Eimer et al (1996) lends support to the ideas that this process is linked with brain processes that are not necessarily associated with conscious awareness although obviously replication of these ERP results is required.

As with other implicit learning paradigms, the orthoganal problem to the implicit / explicit debate is that of underlying representation. Subjects may be acquiring knowledge about the sequence itself and / or knowledge about motor responses. Cohen et al (1990) provide evidence that motor learning is not exclusively involved in sequence learning when they found almost perfect transfer of knowledge between different effector muscles. This finding has been replicated by Keele, Jennings, Jones, Caulton and Cohen (1995) who also demonstrated transfer between modality.

Howard, Mutter and Howard (1992) have shown that subjects simply watching a display learn as much as subjects making overt RT responses to it. One possibility suggested by Stadler (1989) involves learning to anticipate a spatial quality in the stimulus display. Stadler trained subjects to respond to targets appearing at the corners of a large square. When Stadler gave subjects the same sequence but using a small square which did not require eye movements to view different targets, no transfer of learning was found. Although this suggests that the mechanism of learning may involve spatial knowledge or even memory for eye movements, several studies (Willingham et al, 1989; Mayr, 1996) have shown that it is possible to learn sequence tasks which do not include a spatial component. At present, it is unclear which of these positions are correct or, as Mayr (1996) suggests, if two separate systems (one spatial and one non-spatial) need to be postulated.

Cleeremans (1994b) has modelled a type of sequence learning task introduced by Kushner, Cleeremans and Reber (1991). The sequence prediction task involved target which could appear at the vertices of a triangular template. The subjects saw five successive stimuli each of which appeared randomly at one of the three target locations. The subjects were asked to predict the location of the sixth target which unbeknown to them was determined by the locations of the second and fourth stimuli. The results show that subjects could make accurate predictions at above chance levels without awareness of the rules as measured by free report or forced choice test. Cleeremans (1994b) used a simple recurrent network (which can learn to predict the next item in a temporal sequence) to model subjects' performance. If the model was allowed to store all five stimulus presentations at once then it showed similar performance to real subjects. Cleeremans also used this model to investigate claims made by Perruchet (1994) that subjects perform this task by an instance based mechanism that is not necessarily unconscious in nature. Perruchet noted that as

Kushner etal used all possible combinations of sequences during the training phase, it was not possible to say whether subjects seemed to be using rule abstraction or memory for instances to predict the target location. Perruchet trained subjects with only two-thirds of the possible rules and tested subjects with the remaining third. Perruchet argues that the results favour an instance based model rather than one which posits an abstractive mechanism. Cleeremans (1994b) trained the SRN model on two-thirds of the rules and found that it also behaved as though it were using specific instances however the model did not actually learn in this way. The model learned by recognising the relationship between the two critical predictive stimuli and the target. The non-critical stimuli played virtually no role in selecting the location of the target stimulus. This demonstrates that despite the empirical results of Perruchet, the underlying representation which subjects use in this task is not yet clear.

Acquisition of invariant knowledge

McGeorge and Burton (1990) introduced a task which seemed to suggest subjects could learn information from a relatively simple display and use this information during a test phase with no apparent overt knowledge of what had been learned. Subjects were exposed to a series of four digit numbers one at a time, and were given some orienting task to perform on them. Each number contained one particular digit (the invariant) although this was not made known to subjects. Subjects were then given a surprise recognition test where they were presented with two four digit numbers and asked to choose the one which they saw in the previous phase. Subjects had seen neither number before although one contained the invariant digit (the positive) and the other did not (the negative). McGeorge and Burton (1990) found that subjects chose the positive over the negative at above chance levels. When asked what strategy was used to choose between the two numbers, McGeorge and Burton found that subjects' responses could not account for their performance. Bright and Burton (1994) found a similar effect training subjects on a series of clock faces which obeyed a rule (the hour hand always appeared on the left hand side of the clock face) and testing them on two new clock faces , one of which obeyed the rule and one of which violated the rule. Subjects again chose the positive over the negative but could not verbalise the underlying rule or give correct answers to a series of questions about the rule. This task assumes that a recognition task is an index of implicit knowledge so perhaps authors such as Shanks and St. John or Perruchet and Amorin would argue that this task does not demonstrate implicit learning. One thing that seems clear about the task is that the rule is a simple one (unlike the rule structure underlying artificial grammar) so subjects would surely be able to verbalise it if they had consciously been aware of it. Thus this task seems very good evidence of implicit learning at a subjective level of awareness.

One problem with this conclusion is that it assumes no violation of Shanks and St.John's (1994) Information Criterion. It may be however, that subjects are using information other than that of the invariant to choose the positive over the negative. Wright and Burton (1995) suggest that as the test items are constructed randomly in the McGeorge and Burton study, there is a possibility that some numbers may contain repetitions of digits (e.g. 3558). As the positive is constrained to have one occurrence of the invariant digit, there is less chance of the positive containing a repetition than the negative. Wright and Burton argue that as a repetition is very salient, subjects are likely to remember not seeing such a combination of digits at study and so reject it (and hence reject a greater proportion of negatives). They demonstrated this empirically although there was a small effect of the positive even though repetitions were removed from the test exemplars. This demonstrates that there are other rules which subjects can use to choose between the positive and negative exemplars and these rules are not included in a subsequent test of awareness. Although the evidence is only that of verbal report and there is a question of correlated hypotheses which might give the positive an advantage, the evidence seems to indicate that invariance learning can occur without conscious awareness.

The question of how this knowledge is represented is similar to that in the artificial grammar literature. It is possible that subjects are abstracting some 'rule' across the exemplars or alternatively it is possible that subjects are comparing similarity to a previously seen exemplar or exemplars. McGeorge and Burton (1990) provide evidence that subjects are not performing the task using simple perceptual similarity. They showed that transfer of knowledge can occur between numbers in digit format at study (e.g. 1234) and word format at test (e.g. one thousand two hundred thirty four). Bright and Burton (1994) show a similar result with clock stimuli with transfer of knowledge from analogue clocks at study to digital clocks at test. As Bright and Burton did not use a particular invariant but rather a more abstract rule, the transfer suggests that knowledge is held in a form more abstract that surface form. They also compared performance for actual items seen during the study phase with novel positive items and found no difference in performance between the two. They argue this favours a non-similarity based mechanism for performance. Cock, Berry and Gaffan (1994) however, show that performance on the McGeorge and Burton paradigm can be explained in terms of simple similarity. They manipulated the test exemplars to be similar or dissimilar to a particular study exemplar and showed that subjects chose a negative which was similar to the previous instance over a positive which was dissimilar. This clearly seems to indicate that it is a similarity mechanism which governs performance rather than a more general rule abstraction process. This type of mechanism may have difficulty in accounting for the data presented by Bright and Burton (1994) although it is unclear whether the same mechanism would necessarily underlie both tasks.

There are no published studies which have attempted to model performance on this task however McGeorge (1990) attempted simulations of the task in his doctoral thesis. McGeorge (1990) implemented a model based on Hintzman's (1986) MINERVA 2 which assumed perfect encoding of thirty four digit numbers as separate traces in memory. Test exemplars were also constructed and used to probe the memory base. McGeorge assumed that the exemplar which generated the strongest echo would be the one selected as having 'previously occurred'. The simulation was run 500 times for all ten test pairs and McGeorge found that the positive exemplar was far more likely to be chosen than the negative. McGeorge found a similar result when he used stimuli which had a dissimilar positive and a similar negative test exemplar. McGeorge also modelled the effect of an increased pool of study exemplars and found that the model performed to a higher degree when the exemplars were increased. This was not the case for subjects however who showed the same level of performance no matter how many exemplars they were shown. This suggests that the mechanism which might underlie performance on this task does not depend on the encoding and retrieval of every presentation of the stimuli. The mechanism may be much simpler and rely on matching to a specific instance (as the findings of Cock et al (1994) suggest) or the mechanism may rely on a subset of exemplars beyond which no advantage is conferred. Finally there is the possibility that a prototype might be extracted from the exemplars and this protoype will only be very negligibly changed once a certain number of exemplars contribute to the abstraction.

1.8 Summary and future considerations

Researchers such as Shanks and St. John (1994) view the conscious nature of cognition as the default position for information acquisition and utilisation. Reber and Winter (1994) argue that if this were the case then how is the acquisition of natural language or categories explained. They make the point that implicit processes seem to govern vast areas of knowledge acquisition necessary for everyday life, such as those cited above. Dienes and Berry (1997) make the point that in order to distinguish whether the mechanisms responsible for performance on explicit and implicit tasks really are dissociable in terms of a conscious / unconscious divide, there must be appropriate criteria to decide what is conscious and what is unconscious. It is clear from the above discussions that in each paradigm, whether the knowledge is viewed as explicit or implicit depends very much on what criteria the individual researcher is willing to accept. Shanks and St.John (1994) are not willing to accept verbal report as an adequate measure of conscious knowledge as they may not be sensitive enough to elicit low confidence or fragmentary information which recognition or prediction tests would elicit. Perruchet and Amorin (1992) echo this view and further suggest that even recognition and prediction tests might be insensitive to the information which the subject used at test due to forgetting and interference from feedback. Using these tests of awareness as the benchmark for the demarcation between conscious and unconscious has lead these researchers to conclude that there is very poor evidence for any learning which is truly unconscious.

Dienes and Berry (1997) however argue that this conclusion is not correct if one takes what they term a "subjective threshold" as the most appropriate criterion. As has been shown for all the major paradigms, subjects' verbal knowledge lags behind their performance on the various tasks. Shanks and St. John (1994) criticise the use of verbal report for assessing awareness on the grounds that it does not adequately meet what they term an Information and a Sensitivity Criterion. Dienes and Berry however suggest two alternative methods which may be useful in assessing whether knowledge is above or below a subjective threshold. The first is subjects' metaknowledge about their responses, i.e. knowledge is below the subjective threshold if subjects believe they are guessing. The second is that subjects' confidence should be unrelated to their accuracy if this knowledge is below a subjective threshold. They argue that these reflect only one type of metaknowledge and propose another which is that subjects may know that they know something but be unaware of what it is that they know. This is an important additional criterion because, as pointed out by Perruchet, Vinter and Gallego (1997), in the case of natural language we are clearly confident about making grammatical sentences and are usually confident in judging a sentence to be ungrammatical if it is indeed so. Thus, the use of subjects' metaknowledge about their responses would not allow the classification of language acquisition as implicit when it is in fact, thought to be one of the fundamental implicit procedures (e.g. Reber, 1997).

This point by Perruchet et al may hold the key as to the future direction for research. The criterion for deciding between conscious or unconscious processes may be better defined if we can be sure what it is in the real world that we can learn unconsciously (if anything) and if that learning really is different from explicit learning. Dienes and Berry (1997) hint at this by suggesting that a subjective threshold may be in keeping with what a 'layperson' might mean by implicit learning. Mathews (1997) states explicitly that implicit knowledge in real life such as in face recognition or natural language does not seem to have the properties such as inflexibility which has been ascribed to it by the current empirical research paradigms. It seems that the research into implicit learning sometimes gets 'bogged down' by the details of a particular task and researchers lose sight of the real questions (a sentiment shared by Reber, 1993). In short, converging evidence from many different types of investigation will probably be better evidence for implicit learning rather than the one 'true' task and evidence from tasks less tied to the experimental rigour of the laboratory setting seems to be required if a thorough understanding of implicit learning is to be gained.

Chapter 2

Implicit acquisition of an invariant characteristic

2.1 Introduction

One definition of implicit learning is that the subject learns about the structure of a fairly complex stimulus, without necessarily intending to do so, and in such a way that the acquired knowledge is difficult to express (Berry and Dienes, 1993). However, complexity in these tasks can be replaced by simplicity if the underlying rule(s) remains difficult to notice explicitly. McGeorge and Burton (1990) used an extremely simple type of 'rule'; merely an invariant in each study exemplar. In their task, subjects saw 30 four digit numbers and were required to perform an arithmetical operation on them. Subjects did not notice that each number contained the digit '3'. Immediately following the test phase, subjects received an unexpected forced choice recognition test. Subjects were presented with two numbers and told that they had seen one of them in the previous task and were to choose which one they thought it was. In fact, subjects had seen neither number before but one of the numbers contained a '3' (the positive) and the other did not (the negative). McGeorge and Burton found an above chance preference for the positives over the negatives without subjects reporting any awareness for using knowledge of the invariant. They interpreted this as evidence for implicit learning. Furthermore, they found transfer of this implicit knowledge between surface forms of the stimuli ('1234' at study; 'one two three four' at test) suggesting that the knowledge that subjects acquire is at a conceptual level, and that this effect remains when subjects are discouraged from sub-vocal reading. This result is in agreement with other findings of implicit knowledge transfer when the surface structure is different between study and test but the deep structure remains unaltered (Reber, 1969; Mathews et al, 1989; Altmann et al, 1995).

There is evidence in the literature for this apparent effect of choosing the item containing the invariant being due to subjects' use of correlated hypotheses. Wright and Burton (1995) found that above chance performance on this task can be attributed largely to rejection of negatives which were more salient due to a repeated digit (e.g. 2997). These repetitions occur more frequently in the negatives than in the positives and therefore suggest a way in which subjects can explicitly achieve above chance performance. However, because subjects can use a strategy, it does not therefore follow that they will use such strategies in every circumstance. Wright and Burton found that when items were equally distinctive subjects still displayed a small preference for the positive item over the negative.

Other less specific strategies are more difficult to evaluate experimentally however. Churchill et al (1995) found that subjects reject negatives on the basis of unfamiliarity and Churchill et al (1996) and Gilmore and Churchill (1996) report the use of a type of 'transfer-appropriate procedure' strategy by subjects who seemed to re-use the study orienting task as an aide-memoire. Cock, Berry and Gaffan (1994) suggest that subjects choose positives on the basis of familiarity to previously seen exemplars and not on the basis of an invariance rule. These two mechanisms were confounded in the original McGeorge and Burton experiments. As Cock et al show from various different calculations of similarity, positive strings will be strongly favoured if the test strings are matched to study strings on the basis of summed similarity. The positive strings are almost always more similar due to the presence of the invariant. Manipulating the test items so that the positive was quite dissimilar to a specific study item, and the negative was quite similar, Cock et al found that subjects chose the more similar negative item over the positive. This clearly suggests that the apparent effect of invariance learning is an artifact of a more general similarity matching process. In finding evidence for a similarity matching process, Cock et al fall firmly on the side of instance based accounts such as those proposed by Medin and Florian (1992) and Shin and Nosofsky (1992). While Cock et al remain neutral on the question of whether this similarity matching process is an explicit or an implicit process, others (Perruchet, 1994; Shanks and St. John, 1994) would suggest that such instance-based learning is explicit. It is an impossible debate to resolve without another index of what is meant by explicit and implicit.

One possible index which may be used to separate explicit from implicit is that of a time deadline at test. In a typical artificial grammar paradigm, subjects are not constrained to make a response within a certain time limit. Presumably they can use a large number of complicated explicit heuristics to make a response decision. These heuristics may involve any and all of the strategies identified by Churchill et al (1995a, 1995b), Gilmore and Churchill (1996), Cock et al (1994) and Wright and Burton (1995) plus other strategies idiosyncratic to the individual subject. Consequently, it is difficult to claim that subjects solve the grammar task in any one given way and demonstrating that the task can be solved using one specific mechanism does not exclude the possibility of implicit learning taking place. Hay and Jacoby (1996) provide experimental evidence to support the idea that a time deadline reduces the amount of explicit control in a memory task. Using paired associates and a process dissociation technique, they found that conscious recollection was lower for a short deadline condition than for a long deadline condition. The deadline condition did not affect estimates of automatic processing.

Manipulations involving a time deadline can be seen as a subset of procedures which may reduce the amount of explicit processing utilised in performing a task. Smith and Shapiro (1989) suggest a link between implicit learning and holistic categorisation of stimuli and that certain "primitivising" conditions such as incidental exposure or concurrent load will induce holistic processing. Speed may be seen as such a "primitivising" condition. Smith and Kemler Nelson (1984) found speed linked to holistic categorisation and as mentioned above, Hay and Jacoby (1996) and Turner and Fischler (1993) also suggest that a speeded decision may promote more implicit processing.

In a standard artificial grammar paradigm, Turner and Fischler gave subjects either 2 seconds or 6 seconds to make a grammaticality decision to novel grammar strings. Another group of subjects were given the same task but were instructed at study to search for rules. This form of explicit rule searching leads to a decrement in performance relative to that of subjects given incidental instructions at study (Reber, 1976; Howard and Ballas, 1980). Turner and Fischler found a significant decrease in response accuracy for the short (2 sec) deadline over the long (6 sec) deadline, but only when subjects had been given rule-discovery instructions at study. For the subjects given memory instructions, which are not thought to inhibit implicit learning as are the rule-discovery instructions, subjects exhibited no such decrement in performance. So under conditions thought to promote implicit processing, a task constraint forcing subjects into a 'primitive' mode of cognition had no effect. In contrast, when subjects had been using explicit rule search strategies during the study phase, forcing a 'primitive' mode of cognition on them had a detrimental effect. Although Turner and Fischler attributed these results to the possibility that "subjects who acquired an implicit base of knowledge were conferred an advantage when forced to make fast decisions", there is another way to view the results. Rule-discovery instructions may not prevent implicit knowledge acquisition, but may make subjects more likely to respond to the test phase with explicit hypotheses. A short time deadline would make it more difficult to respond via explicit heuristics and so would enable any implicit knowledge that subjects might have, to be accessed with relatively little interference from competing explicit processes. This alternative explanation of Turner and Fischler's results would seem to be in accord with the results found by Hay and Jacoby (1996) where a deadline manipulation did not affect the amount of implicit processing but did affect explicit processing.

The same alternative reasoning can apply to other results in Turner and Fischler's study. They found that subjects in a condition where the rules were presented in a salient fashion had almost identical accuracy under both short and long deadlines, whereas accuracy for a nonsalient group fell under a short deadline. Turner and Fischler suggest that these results support facilitation of abstraction of implicit knowledge by nonsalient displays. The alternative explanation is that salience helps subjects to attend to the crucial aspects of the display, enabling explicit decisions to be made even under short deadline conditions. What Turner and Fischler are suggesting is that instructions at study or display salience constrains the information which is learned at study. If subjects engage in more explicit processing of the stimuli, they suggest that virtually no implicit processing of the stimuli can take place. An alternative hypothesis and one which cannot be separated by the paradigm they use, is that subjects can learn explicit hypotheses about the stimuli at study as well as implicitly abstracting information about those stimuli and it is the constraints placed upon the subjects at test which determine which of these modes of cognition will be of most value under those constraints. The argument is one of hypothesis testing preventing implicit abstraction of information at study or of inhibiting retrieval of implicit information at test.

The following series of experiments prevents the former by using incidental study instructions for subjects across all testing conditions. Thus the possibility that hypothesis testing at study prevents implicit abstraction cannot be addressed, but if there exist significant differences across conditions due to response deadline, then the hypothesis of explicit heuristics inhibiting implicit processes will find some support. While this explanation may account for Turner and Fischler's finding, it does not rule out their explanation as the tasks are different and there is not a manipulation of possible knowledge abstracted at study.

2.2 Experiment 1: An attempt to replicate Cock, Berry and Gaffan (1994)

The paradigm used in the following experiments is that of McGeorge and Burton (1990), with the modifications to the stimuli suggested by the work of Wright and Burton (1995) and Cock et al (1994). Therefore, care was taken to ensure that none of the study or test exemplars contained any repeated digits (e.g. 2775 or 3763) as Wright and Burton demonstrated these combinations to be more salient and likely to be rejected at test, irrespective of the presence of an invariant. Cock et al found that similarity of the test exemplars to previously seen study exemplars was a major determinant of subjects' performance in this task. To control for this factor, each test exemplar was constructed to be either quite similar to a previously seen study exemplar or quite dissimilar. Hence, the test pairs could be matched for similarity quite independently of the presence or absence of the invariant. The calculation of similarity between study and test exemplars was taken directly from the Cock et al study and refers to similarity to a specific exemplar which was previously encountered. A similar item is one which has only one digit changed from the original study exemplar. A dissimilar item has two digits which differ from the original exemplar and has those digits repositioned within the number. These manipulations apply to both positive and negative exemplars and form a similarity dimension which is orthogonal to the presence or absence of the invariant. A more detailed description of the construction of the test exemplars is given in the Materials section. Experiment 1 uses three conditions from the study by Cock et al (1994): a condition where the two exemplars in the 2-AFC are equally similar to the study string from which they were modified but differed from each other with respect to presence of the invariant, Positive Similar exemplar versus Negative Similar exemplar (PSNS); a condition where the positive exemplar was quite dissimilar to the original study exemplar whereas the negative exemplar was quite dissimilar, Positive Dissimilar versus Negative Similar (PDNS). Note that this is the crucial condition which separates the effect of similarity from the presence of the invariant. The final condition is one where the exemplars are equally dissimilar to the study exemplar but differ from each other with respect to presence of the invariant, Positive Dissimilar versus Negative Dissimilar (PDND). Those conditions that could not distinguish between rule and instance-based accounts (those which included no invariant or had both similarity and invariance explanations favouring the same exemplar) were not included.

The results of Cock et al (1994) suggest that subjects would choose the exemplar which has the greater similarity to a previously seen instance with the positive-negative manipulation having no effect. However, McGeorge (1990) performed a manipulation designed to test between specific similarity and the invariance rule and found the opposite result: that subjects would choose the exemplar conforming to the rule over the exemplar which was the more similar. While this study was subject to some methodological criticisms (the exemplars contained repetitions and the similarity dimension was not as stringently controlled as in later studies) it nevertheless casts doubt on the replicability of the Cock et al result. Each of the conditions also included a time deadline manipulation whereby subjects had to make a response within 2s (fast condition) or 10s (slow condition). In accordance with the data reported by Turner and Fischler (1993) and Hay and Jacoby (1996) which suggest that speeded tasks access implicit mechanisms to a greater extent than self-paced tasks, subjects in the 'fast' conditions should choose the positive exemplar over the negative more often than subjects in the 'slow' conditions. It is possible that the two factors will interact, with speeded deadline and therefore more implicit processing giving an advantage to the invariance rule, and the slow (explicit) deadline giving a similarity advantage.

Method

Subjects

90 University of Glasgow undergraduates participated in return for a small payment.

Materials.

A computer program was used to generate random study sets of 30 four digit numbers with the constraints that each digit could not appear more than once in a number and one digit was to appear in all thirty numbers. 5 different sets of stimuli were constructed with each set having a different invariant digit. For each set, 10 numbers were chosen at random and used to generate the four different types of test stimulus. Care was taken to ensure that the test strings were constructed in the same way as that described by Cock et al (1994). For the negative string, it was the invariant itself that was changed. For the positive string, a digit other than the invariant was changed; e.g. if the study string that the subjects saw was 5382 with 5 being the invariant, the positive similar string could be 5381. The negative similar string might be 6382. Care was taken to ensure that the new digit pairing had not actually occurred in any other study string. Taking one of the examples above, when 5382 was transformed into 5381, the digit pair 81 could not have occurred in any other study string, either as the 1st or the 2nd pair of digits. The same criterion was used to construct the dissimilar strings. For a positive dissimilar string, the digit pair not containing the invariant was changed to another digit pair, not encountered in the study phase, and the digit pairs switched sides (e.g. 5382 at study becomes 4753 at test, with 47 not occurring as a pair in the study set). For negative dissimilar exemplars, the procedure was the same except that the digit pair containing the invariant was replaced. The frequency of occurrence of the individual digits at test was matched as best as possible to their frequency in the study set and never differed by more than 25% and usually by a lot less. The digits were all in courier font, size 24, and constructed using graphics software.

Design

Subjects were randomly allocated to one of five sets of stimuli and one of six test conditions. These conditions were determined by 2 variables; firstly, the types of test strings subjects would be required to choose between. These were positive similar vs. negative similar (PSNS condition), positive dissimilar vs. negative similar (PDNS) and positive dissimilar vs. negative dissimilar (PDND). The second variable was time deadline at test, with half of the conditions requiring a response within 2 secs (Fast deadline) and half with a deadline of 10 secs (Slow deadline).

Procedure

Subjects were tested individually and all study strings and test pairs were presented in a randomised order to each subject. The study phase was the same for all subjects; 30 four digit numbers were presented one at a time on a computer screen and subjects were required to add the two digits on the left hand side and the two on the right hand side and make a key press as to which sum gave the higher total. This was immediately followed by an unexpected test phase in which subjects were told that they would be presented with two numbers, one of which they had seen in the earlier phase and were required to make an appropriate key press as to which one they thought it was. Subjects were told to guess if they were unsure. They were also told that they would have to answer within a specified time limit (2 or 10 secs), that being measured by how long the stimuli remained on the screen. In order to orient subjects to using a time deadline, they were given four arithmetic problems such as '16 + 8 = 24' and asked to respond by key press to a true/false decision for each one within the appropriate time limit. They then received the 10 pairs of numbers, each one remaining on screen for the appropriate time or until a key press was made. The pairs of numbers were presented in the centre of the screen with seven character spaces between them and any response made after the deadline was not recorded.

Finally, each subject's awareness of the invariant rule was examined using a short semi-structured questionnaire. This questionnaire was comprised of four questions: a) In the earlier part of the task, did you notice anything unusual about the 30 four digit numbers as you did the task?

b) Did you notice if they had anything similar or dissimilar in relation to all the other numbers?

c) In fact, all the four digit numbers had one digit in common. It could occupy any position in the number but it appeared in each number. Did you notice this as you were doing the task?

d) Can you tell me what that digit was or guess what it might have been?

Subjects were classified as aware if they mentioned the rule at any point or if they gave a positive response to question c and could report the invariant digit.

Results and Discussion

11 out of the 90 subjects missed some of the response deadlines. Any subject not responding to at least 6 out of the ten test pairs would have been discarded and replaced, however no subject scored below this minimum and only one subject at the minimum. Six of the remaining ten subjects only missed one trial, so overall the response totals were very high. Those subjects with missing trials were marked *pro rata* out of ten. The results are shown in Table 1. The main effects of Similarity/Invariance and Response Deadline were not significant but the interaction between them was ($F(2,84) = 3.41$, $p < 0.05$). Simple main effects showed that the only condition which was significantly different from chance was fast PSNS ($F(1,84) = 5.75$, $p < 0.05$). As can be seen from Table 1, the standard deviation for this condition as well as the mean is noticeably different from the other conditions. In fact comparing the variance of the fast PSNS condition to the next highest reveals that the standard deviation of 0.79 is significantly different from those of the other 5 conditions. Given that this breaks the assumption of homogeneity of variance for

using the F statistic, a correction procedure suggested by Keppel et al (1992) was used to re-examine the data. This involves calculating an F ratio by dividing the largest variance by the smallest variance. If this ratio is greater than 3.0, Keppel suggests using a more stringent alpha level. The adjusted alpha level was set at .025 and reveals no change to the original interpretation of the data. The result of Cock et al (1994) was not replicated in these data with no significant effect of similarity being evident. In fact, although the mean in the crucial slow PDNS condition was at chance, the trend was in the opposite direction to that predicted by a similarity effect.

Table 1: Mean no. of positives chosen (out of 10)

	PSNS	*PDNS*	*PDND*
FAST	6.37 (0.79)*	5.03 (1.44)	5.55 (1.64)
SLOW	5.07 (1.79)	5.73 (1.44)	5.40 (1.64)

*p<.05, corrected for violation of assumption of homogeneity of variance.

From the responses to the awareness questionnaire, six subjects were classified as aware and were discarded from the analysis and replaced with new subjects. Five subjects reported noticing an invariant digit but when asked what that digit was they gave an incorrect answer. These subjects were classified as unaware. Of the remaining 85 subjects, who all reported a failure to notice an invariant digit, 12 gave the correct invariant digit when asked to guess what it might have been. It can be concluded that subject's guessing was at chance for choosing the invariant digit, providing compelling evidence that subjects were unaware of the invariance rule.

In the original formulation of this paradigm, McGeorge (1990) provided evidence that the mechanism underlying performance in this incidental learning task was not one of similarity matching. However, Cock et al (1994), using more carefully constructed test stimuli across a number of manipulations of similarity, could find no evidence for an invariant rule underlying performance but a strong effect of similarity. The present study could find no effect of similarity on either the speeded or non-speeded deadline task, using the same criteria for test stimuli construction as Cock et al and three of the same similarity manipulations. In the crucial condition which directly opposes the invariance rule against similarity (PDNS), the mean number of positives chosen were in the direction supporting the rule, although not significantly so. This contrasts sharply with Cock et al's result for PDNS where subjects' scores are significantly below chance for choosing the positive over the negative, indicating that previous confounding of the positive as the more similar exemplar led to the false conclusion of a rule-based account.

Are there any differences between the present study and that of Cock et al which might explain this failure to replicate? Two possibilities arise; the first is that the stimuli were written on cards in the Cock et al study and in this study they were presented in identical font on a computer screen. Without seeing the actual stimuli used by Cock et al it is impossible to be sure but perhaps using written stimuli emphasised some type of similarity between the strings. A second and more probable explanation is that Cock et al only used one set of learning strings whereas the present study used five different sets across subjects. The similarity effect could be due to as little as one item in that learning set being memorable to subjects and so 'forcing' them to choose a negative over a positive. Using 5 different sets of materials significantly decreases the impact of a salient item and may well balance such an effect out. While at first these methodological points may seem like trivial considerations, it is possible that similarly small considerations may underlie the inability to replicate many effects reported in the literature. The inability to replicate

Cock et al's finding suggests that it would be premature to deny the existence of invariance-based implicit learning.

The results do seem to provide scant evidence in support of invariance-based implicit learning, however. When previous methodological criticisms are taken into account, five out of the six manipulations fail to find above chance performance for this reportedly robust finding of choosing the positive over the negative. Only in the fast PSNS condition did subjects show any evidence of learning. While it was predicted that a fast deadline should allow for greater expression of implicit processes and so give a higher mean number of positives chosen over negatives, the fact that only one of the three fast deadline conditions was significantly different from the slow conditions and from chance makes for a difficult interpretation of the results. While replication is obviously required, if the results are robust it suggests a violation of one of the assumptions underlying this experiment i.e. that the point at which explicit processing is sufficiently reduced to allow the robust expression of implicit processes would be fixed at the same time deadline across all conditions. The time deadline was designed to remove explicit processes from an implicit test. It is interesting that it seemed to succeed only when the two test exemplars were very similar to one another and were both 75% similar to a previously studied instance. It is possible that such a high degree of similarity making the choice a difficult one combined with the two second deadline successfully attenuated the degree of explicit weighting given to one exemplar. However, when one or both items is dissimilar to a previously seen exemplar, the choice becomes far easier, even within two seconds, and an explicit heuristic can be utilised.

Turner and Fischler (1993) argued that their finding of a differential effect of deadline for subjects in a rule-discovery condition but not for memorisation instructions suggested that it was the instructions at study which determined whether subjects learned information which could be accessed via implicit processes or information which could be accessed by explicit processes. The results of Hay and Jacoby (1996) suggest an alternative possibility which is that a short time deadline reduces the amount of explicit processing at test but does not affect the amount of implicit processing. The results of Experiment 1 agree with the latter view in that there was no experimental manipulation at study and yet there seemed to be an effect of time deadline at test. This suggests that as in the Hay and Jacoby study, subjects learned information which could be accessed by implicit processing and also information which could be accessed by explicit processing and the time deadline at test constrains which information can be utilised.

What is very evident from Experiment 1 is that subjects do not choose the more similar test exemplar. If the mechanism does not involve such specific similarity, then subjects cannot be comparing the test exemplars with individual stored instances. This effect seems robust as across the six experimental conditions there was not even a trend to suggest that subjects might choose a more similar exemplar over one which was dissimilar but contained the invariant. Note that the inclusion of the invariant is essential as it would be counter-intuitive to suggest that subjects would choose a less similar item over a more similar item in a memory test. The argument is that an invariant property is processed subconsciously and will have an effect on performance when implicit processes are utilised. The fact that only one condition out of the six was significantly different from chance raises another question, however. When similarity and salient items are controlled for, does the effect of an invariant disappear? Replication of the fast PSNS condition is necessary to ensure that the effect obtained in Experiment 1 is not simply spurious.

2.3 Experiment 2: Replication of the deadline effect

Experiment 2 attempts to replicate the fast PSNS effect using a fast and a slow time deadline. The hypothesis from Experiment 1 was that subjects will use an explicit hypothesis in the test phase whenever they can and that implicit processing will only be evident when subjects are prevented from using such explicit heuristics. The orienting task in Experiment 2 was designed to allow subjects to re-use this task with 100% accuracy in the test phase but only when they had sufficient time to do so. When the subjects had a two second deadline to meet, it was hypothesised that they would have insufficient time to re-use the orienting task and so have to rely on implicit knowledge of the invariant to make their choice.

A third group of subjects was included in this experiment. These subjects did not receive the same test phase as the other subjects but one in which neither number contained an invariant. Subjects in this condition had to choose between a Negative Similar and a Negative Dissimilar (NSND) test exemplar. This condition was included to demonstrate that the presence of the invariant is an important factor in this task. It is possible that for some reason, a fast deadline might bias subjects against the Negative Similar exemplar as opposed to actively choosing the exemplar containing the invariant.. If this were the case then an NSND condition should show an advantage for the ND exemplar. If subjects favour the Negative Similar exemplar (as might be expected from a memory account based on similarity matching) then this would give additional support to the idea that implicit learning of the invariant really is something other that a simple similarity matching process as the Positive Similar exemplar in Experiment 1 was favoured over the Negative Similar.

Method

Subjects

30 University of Glasgow students who were naïve to this paradigm were used. All were paid a small amount for their participation.

Materials

Five sets of study and test materials were constructed in the same manner as in Experiment 1 with the following differences. All the study numbers were chosen so that on addition of all the individual digits in the number, subjects would arrive at a two digit number which when in turn had the two digits added would produce an even number. For example, 5249 adds to give the number 20 which when added again gives 2 which is even. From the sets of study exemplars, ten numbers were chosen to be modified into the test exemplars. Positive Similar, Negative Similar and Negative Dissimilar test exemplars were constructed according to the method described in Experiment 1. The exemplars were constrained so that a negative would always add up in the manner described above to give an even total and the positives to give an odd total. The materials were presented via a computer in exactly the same manner as in Experiment 1.

Design and Procedure

Subjects were randomly assigned to one of three groups. Two of these groups received the same test materials which had both four digit numbers equally similar to a previously seen instance form the study phase. One of these numbers contained the invariant digit and the other did not (the PSNS condition from Experiment 1). One group of subjects were required to perform the pseudo-recognition task under a two second deadline while the other group of subjects had ten seconds to make their decision. A third group of subjects were required to choose between two numbers neither of which contained the invariant. One number was similar (as defined in the Method section of Experiment 1) to a previously seen instance and the other number was dissimilar (in other words this condition was a manipulation of the similarity dimension but not of the invariant and will be abbreviated as NSND).

Task instructions for the study phase were that subjects had to add up all the digits in the four digit number to produce a two digit number. These two digits were then to be added together to produce a single digit number and subjects were required to make an odd / even keypress to this digit. All responses in the study phase were the same; the numbers added up in this fashion to always give an even number. Subjects were asked after the study phase if they had noticed anything that the items had in common. Most responded that they always gave an even response and where the few subjects that had made addition mistakes did not give this response, they were explicitly told that they should have responded with the even key for every item.

Subjects were given an immediate surprise test phase with the same instructions as in Experiment 1. One group of subjects had ten seconds in which to make a response to the PSNS condition and it was assumed that this was long enough for subjects to perform the addition task from the study phase on the test exemplars. The other group of subjects who received the PSNS test exemplars were required to respond within a two second deadline and it was assume that this was not enough time in which to perform the addition task nor even to make a reasonable guess at the answer. A third group of subjects received the NSND condition under a time deadline of two seconds. This condition does not manipulate the presence of the invariant but examines the effect of similarity under a fast deadline.

On completion of the test phase, all subjects were given the same awareness questionnaire as used in Experiment 1. More emphasis was placed on uncovering what strategies subjects actually used in this experiment and this question was specifically asked to each subject.

Results and Discussion

From the awareness questionnaires, no subject was found to be aware of the presence of the invariant and again when asked to guess at which digit it might have been, subjects performed at no better than chance levels. Table 2 shows the mean scores of subjects in the three test conditions.

Table 2: Mean Scores and Standard Deviations for Subjects in fast and slow PSNS and fast NSND conditions.

Condition	fast PSNS (Subjects choosing PS)	slow PSNS (Subjects choosing PS)	fast NSND (Subjects choosing NS)
Score (out of 10)	6.2 (1.3)	4.4 (2.4)	6.7 (1.7)

The comparison of most interest is between the slow and fast PSNS conditions. From the results of Experiment 1, it was predicted that subjects in the fast condition should choose the positives over the negatives and subjects in the slow condition should use the explicit rule from the training phase and choose negatives over positives. It can be seen from Table 2 that this prediction is supported and the two groups are significantly different on a one-way ANOVA ($F(1,19) = 4.63$, $p < 0.05$, one-tailed). It is interesting to note that as in Experiment 1, the variances for both groups are statistically different with the fast deadline group showing the smaller variance. Using the same correction procedure as in Experiment 1, re-analysis shows no change in the interpretation of the results. What is also apparent is that when there is no invariant in either of the two test exemplars, subjects do seem to choose the item which is most similar to a previously seen instance. Thus it cannot be argued that introducing a speeded deadline at test biases subjects against using similarity as in the fast NSND condition subjects chose the more similar exemplar at above chance levels ($t(9) = 3.12$, $p < 0.05$).

This second experiment provides strong support for the existence of an effect of the invariant. In Experiment 1, the only condition which showed an above chance effect for choosing the invariant was fast PSNS and it was hypothesised that choosing between two items which were equally similar to a previously seen instance within a time deadline of two seconds was too difficult a choice to resolve with explicit heuristics. This experiment provides a replication of the fast PSNS effect and also shows that when subjects are given an obvious explicit heuristic to use, that they will in fact do so if time permits. Subjects in the slow PSNS condition could re-use the orienting task from the study phase and the results suggest that some of the subjects did choose the number which gave an even response. Subjects in the fast condition did not have time to utilise this or presumably any other explicit strategy and chose the positive over the negative at above chance levels, despite the positive adding to an odd number which therefore could not have been seen during the training phase.

During the de-briefing, subjects were asked what sort of strategies they used to choose one exemplar over the other. In the slow PSNS condition four out of the ten subjects said they specifically re-used the orienting task to help them decide. Most others said they used this strategy for some of the test items. In contrast, subjects in the fast PSNS condition said they used a feeling of familiarity or just guessed. Most said they tried to use the task from the training phase but just did not have enough time. This may explain why the variance in the slow PSNS condition was much greater than that in the fast condition. Subjects in the slow condition used a variety of strategies, some of which would lead to very high scores and some leading to low scores and so lead to much variability in performance. Subjects in the fast condition did not have time to use a wide range of strategies and so performance might be expected to be less variable. Reduced variability for implicit tasks as compared to performance on explicit tasks was predicted by Reber (1993) and has been experimentally found in studies using artificial grammars (McGeorge, Crawford and Kelly, 1997; Reber, 1991). The finding of reduced variability in the speeded conditions in Experiments 1 and 2 provides further evidence that these manipulations may involve more implicit processing than the non-speeded counterparts.

The fast PSNS effect seems to be robust and again no evidence is found which supports the idea of subjects implicitly choosing the more similar item over one which is dissimilar but which contains the invariant, as long as the assumption of a speed deadline allowing greater implicit processing holds. With the knowledge that there does seem to be an effect of the invariant under certain circumstances, the next question is what sort of mechanism might be operating to account for this effect? The results of Experiments 1 and 2 clearly rule out a mechanism based on specific similarity to a previously seen instance as the determinant of the subjects' choice. Another possibility is one at the opposite end of the spectrum with regard to categorisation theories. Is choosing the exemplar corresponding to the invariance rule perhaps mediated by a prototype generated by instances? Experiment 3 performs a manipulation on the McGeorge and Burton paradigm which is taken from the categorisation literature on prototype effects and also provides a replication of the fast PDNS effect from Experiment 1.

2.4 Experiment 3: Effect of a lag between study and test

The results of Experiments 1 and 2 suggest that performance on the McGeorge and Burton task is not explainable in terms of a mechanism which compares similarity of test exemplars to specific study strings. An alternative mechanism which might account for the effect of the invariant is one where the test exemplar is compared to some sort of summed representation of the study exemplars. Mechanisms which act on this basis have already been put forward to account for processes such as memory and categorisation, both of which seem related in some sense to the invariance effect. The first experimental evidence for this effect comes from the work of Posner and Keele (1968, 1970) who gave subjects a categorisation task and showed that subjects categorised the prototype of a category more accurately than the individual exemplars of that category. These results can be implemented in a computer model such as MINERVA 2 (Hintzman, 1986) or a PDP model such as that described by McClelland and Rumelhart (1985). Posner and Keele (1968) found that a significant delay between study and test in the categorisation task led to greater forgetting of the old training stimuli than for the extracted prototype. This finding is consistent with the idea that category judgments rely on a mixture of specific information from particular items and more general information abstracted across all items and that the specific item information is forgotten more rapidly. If performance on the digit invariance task is due to such an abstraction mechanism, then a lag between study and test might show increased performance based on choosing the invariant. From the speeded task in Experiments 1 and 2 there was the suggestion that decreasing the probability of subjects using explicit strategies led to robust above chance performance in choosing the exemplar containing the invariant. It was suggested that this effect was due to decreased interference from the explicitly remembered individual traces. Such an effect may also be shown by adopting a lag procedure by assuming that over time the explicit individual representations are forgotten and allow more influence from an implicit prototypical exemplar. Experiment 3 tests this prediction using the PDNS strings from Experiment 1. The PDNS stimuli were chosen because the subjects did not score significantly above chance for the speeded version of the task. By including a fast PDNS condition with immediate test phase as a baseline condition, Experiment 3 also provides a replication of this result from Experiment 1. We would predict that subjects would be performing at chance for the immediate condition as in Experiment 1, but if the abstraction mechanism is plausible then subjects should show above chance performance when there is a lag between study and test.

Method

Subjects

24 subjects were recruited from the student population at the University of Glasgow. All were naive to the experiment and they received a small payment in return for their participation.

Materials

The materials were the same as those used for the PDNS manipulation in Experiment 1 except that the set with 9 as the invariant digit was not used. This was due to a concern that subjects adding 9 to everything may be particularly salient as 9 is the highest digit in the set.

Design and Procedure

The same general procedure used in Experiment 1 was followed in Experiment 3. Subjects were randomly allocated to one of two groups, either the digit invariance task with the test phase immediately following the study phase or with the test phase following a lag of 30 mins which was filled with face recognition experiments. Both groups received the fast PDNS manipulation.

Results and Discussion

Subjects in the immediate condition chose the positive item over the negative item with a mean of 5.05 out of 10 (standard deviation = 0.79) and subjects in the 30 min lag condition chose the positive item with a mean of 6.35 (standard deviation = 1.37). Thus the immediate condition replicates the finding of Experiment 1 that subjects do not perform at above chance levels on fast PDNS ($t = 0.2$, $p > 0.1$). Subjects who did not receive the test phase until 30 mins after the study items did perform with above chance accuracy ($t = 3.27$, $p < 0.05$), and the scores for the two groups are significantly different from one another on a one-way ANOVA ($F(1,22) = 8.14$, $p < 0.05$). As in Experiment 1, the variances of the two groups are significantly different from one another, but using the same correction procedure as in Experiment 1, the two groups remain statistically different. As in Experiment 1, the number of subjects who correctly guessed the invariant was not any greater than that which might be expected by chance.

The results of Experiments 1 and 2 did not support an argument for performance on the digit invariance task being mediated by similarity to specific exemplars. Another possible mechanism for this task was suggested to be extraction of a prototypical exemplar which is then compared to the test exemplars. If this is to be considered a viable mechanism for performance on this task then we would expect such performance to be affected by variables known to influence prototype extraction. One such variable was shown to be a delay between study and test (Posner and Keele, 1968). In Experiment 3 a delay of 30 mins gave robust above chance performance for choosing the positive item whereas an immediate test condition showed no such effect. The mechanism by which Posner and Keele explained the delay effect was the decreasing influence of specific exemplars over time. If the assumption is made that explicit knowledge is essentially knowledge of the specific exemplars and implicit knowledge is the extracted prototype (if the prototype was explicit then we would surely expect subjects to be able to identify the invariant) then these data may be considered as converging evidence for the conclusions reached in Experiments 1 and 2. In resume, the action of the invariant is implicit as subjects cannot state which item was common to all study exemplars, but this knowledge is secondary to any explicit hypotheses the subjects may have. Thus in Experiment 1, the invariant only seemed to have an effect when the subject was required to make an extremely difficult response (both exemplars 75% similar to the specific study exemplar) under a speeded deadline condition and in Experiment 3 when the surprise test phase occurred 30 mins after the study phase. In the former, it is suggested that subjects would not have time to make such a difficult decision based on explicit information and in the latter, that subjects had forgotten most of the explicit information, both of which will lead to unconscious use of the prototype information.

As mentioned in the introduction, one possible mechanism for above chance performance on this task was not that of choosing positives but of rejecting negatives (Wright and Burton, 1995; Churchill et al, 1995). This type of explanation would find it very difficult to account for above chance performance in this experiment as rejecting negatives on the basis of unfamiliarity or idiosyncrasy might account for such performance in the lagged condition but one would also expect such strategies to give above chance performance in the immediate test condition and this was not the case. This experiment provides further evidence that the presence of the invariant is the crucial aspect which determines performance on this task and knowledge of the invariant cannot be utilised unless decreased weighting is given to information from specific exemplars.

2.5 Experiment 4: Effect of a week long lag

Experiment 4 follows on from Experiment 3 in an attempt to discover if this apparent formation of a prototype is a relatively long lasting effect. It is possible that the effect of the invariant will be greater if more knowledge of specific exemplars is forgotten and such an effect may explain why, under usual testing conditions , the magnitude of effect in any implicit learning paradigm never seems to reach above 60 - 70 % accuracy. It may be the case however, that the prototype does decay but at a slower rate than the specific exemplars. Experiment 4 increases the length of the delay between study and test to one week.

Method

Subjects

12 undergraduates from the University of Glasgow were recruited and paid a small amount to participate in this experiment. All were naive to the method.

Materials

The same materials as in Experiments 3 were used.

Design and Procedure

The procedure was identical to that of Experiment 3 except that the delay between study and test was one week and only one group of subjects (fast PDNS) was tested. It should be noted however, that the study phase was embedded in a much larger battery of tasks so subjects would not be able to guess why they were coming back one week later.

Results and Discussion

The mean number of positives chosen out of 10 was 4.7 with a standard deviation of 1.5, which is not significantly different from chance on a one sample t-test ($t(11) = 0.66$, $p > 0.1$). After a 30 min delay between study and test subjects chose the positive over the negative exemplar. When this delay was increased to one week the effect of the invariant disappeared. This suggests that if a prototype is extracted, it is more resistant to decay than the individual exemplars, however it does decay at some time period after 30 mins. This last fact suggests that the mechanism responsible for learning of the invariant is different from that which mediates priming effects. Seger

(1994) suggests that the main difference between implicit memory and implicit learning is that the former involves memory for specific stimuli and the latter involves memory for patterns. It may be suggested that learning an invariant property is simply perceptual priming but Sloman, Hayman, Ohta and Law (1988) show that priming effects are robust across several months and maybe even years (Maylor, 1995) for different types of stimuli. Finding no such effects for invariance learning after a period of one week strongly suggests separate mechanisms responsible for priming and for invariance detection.

An alternative suggestion is that there is no prototype formation at all and the effects are mediated by explicit memory. Thus after one week, enough is forgotten about the exemplars to cause the effect to vanish. This does not account for the null effect using an immediate test phase however. Nevertheless, an argument for explicit memory may be forced by claiming that there is some explicit function which may be time dependent and show such a pattern as the previous results show. For example, one suggestion by Gilmore et al (1996) is that subjects use the orienting task to a great extent to help them decide which test exemplar to pick. It may be that with an immediate test phase, subjects can remember all the additions which they made in the study phase and use this to help them in the test phase. With such a large number to choose from, this strategy may not be much help at all. After 30 mins, it is easy to imagine that subjects can remember only the more common results of additions that they made and this may lead to increased choosing of the positive exemplar. At one week, one might expect subjects to have no memory of any additions which they might have made and so performance on the task drops to chance levels. This simple correlated strategy, which Gilmore et al have shown subjects do use, would seem a plausible way to account for the data without recourse to a mysterious implicit prototype. In addition, the effect of a lag on a fast PDNS condition has only been shown in one experiment and may simply be spurious. Experiment 5 was designed to replicate the lag effect of Experiment 3 but using what would seem to be a totally different task during the encoding phase.

2.6 Experiment 5: The invariant nose experiment

This experiment is designed to replicate Experiment 3 but using a task which is totally dissimilar to that of the digits task at a surface level but which retains all the characteristics of an invariance detection task. Identikit type faces were constructed using a computer generation programme and one feature (the nose) was kept invariant. Examples of the training stimuli are given in Figure 3. Other aspects of the faces were made up randomly. Using these type of stimuli had several major advantages over the digits. Firstly, the task was to locate a possible suspect from a written description. All subjects thought that this was an experiment on eyewitness testimony and none suspected there would be another phase to the experiment. Even though the digits task has a surprise test phase, it is possible that subjects are suspicious of a task at study which only has them doing simple arithmetic. Asking subjects to choose a likely suspect also has none of the task demands that addition has and so there is much less for them to remember and use as an explicit heuristic from the orienting task. One major problem with digits is that the stimuli cannot be counterbalanced very easily. This means that idiosyncratic exemplars may have disproportionate effects on the results. Using these 'mugshots' as stimuli means the invariant feature can be swapped across test exemplars. Finally, due to the cover story and the complexity of stimuli, the invariant would possibly be much more difficult to notice.

These various properties of the stimuli should make it much less likely that any explicit strategies will be used by the subjects or if any are used, they will have less of an impact on the results. However the effect of the invariant is likely to be very weak. Speeded responses were not required (the manipulation of most interest in this experiment is that of a lag) and so each subject may apply various explicit hypotheses to the stimuli when choosing which one had already appeared. This may have the effect of increasing the amount of variance in comparison to a task with a time deadline. The number of subjects used in this experiment was increased in order to compensate for this problem.

Method

Subjects

80 subjects were recruited from the undergraduate population of the University of Glasgow. All were new to the task and all were paid a small amount for their participation.

Materials

100 study faces were constructed at random using a programme called Mac-a-Mug (see Figure 4 for examples). The exception was that all 100 faces had the same nose, which was chosen to be somewhat nondescript. These faces were bound into 5 booklets each containing 20 faces. Five descriptions of 'suspects' were also constructed (see Figure 3) with vague references to general facial features. They were not constructed to represent any face in particular and the nose was specifically not mentioned in any of the descriptions. A further 20 faces were constructed at

random, 10 of which contained the invariant nose and 10 which each contained a different shape of nose. A counterbalanced set of test faces were constructed from these simply by swapping the noses over. These were kept loose on A4 sheets of paper to enable each test exemplar to be paired with every other test exemplar. Thus, using both sets of test materials, the stimuli could be fully counterbalanced over 20 subjects. The size of each image printed on the A4 paper was 8 cm in height and 6 cm in width and all images were in black and white.

Figure 3: Written descriptions of 'suspects' in Experiment 5.

Youngish man. Dark straight hair, not very long. Thin eyebrows which were close together. Thin face with prominent cheekbones. Small cold eyes.

Chubby man with a moustache and beard. Dark hair, fairly long, which may have been a bit curly. Staring eyes. Early 30's

Man of medium build probably in his late 30's. Bushy eyebrows
and close together eyes. He might have had a small moustache but maybe not. Short hair but scruffy looking.

Young man probably around 24. Clean-shaven but with dark menacing eyes.
Girlish looks, especially the lips. Short light hair.

Small man with receding hairline. Narrow lips and small chin.
Looks intelligent but dishevelled appearance and had not shaven for a few days.

Figure 4: Examples of training stimuli used in Experiment 5.

Design and Procedure

Subjects were divided into two groups: those who received an immediate test phase and those who received the same test phase but after a delay of 30 mins, which was filled with other experiments. Subjects in both groups received the same study phase where they were given each of the five booklets in turn and a suspect description for each booklet. The booklets and the descriptions were given out to subjects at random. Subjects were told that this was an experiment on eyewitness testimony and they were to read the suspect description and then try to find a match for that description from the booklet. They were told they could take as long as they wanted and could go forwards and backwards through the booklet. They were informed that they could only choose a maximum of one suspect but if they thought no-one matched the description well enough then they did not have to choose anyone. Subjects did this for all five booklets. Then the subjects in the immediate test group were told that they would see pairs of faces and that one of each pair was had been in one of the booklets. They were to choose which one they thought it was and if they weren't sure then they were simply to guess. Subjects had not seen either face before but were led to believe that they had seen one of each pair. Each face pair was placed in front of the subject with the face containing the invariant nose (the positive) randomly assigned to one side or the other and the subject was asked to choose which one they thought they had seen before. Then these items were removed and the next pair was presented. For the next subject the face pairs were changed, so that across 10 subjects every positive item had been paired with every negative item. The same was done for the next 10 subjects but with the set of faces with the invariant nose swapped over. After this subjects were asked if they had noticed anything similar about the faces in the booklets. Subjects in the delay condition received exactly the same test procedure but after a delay of 30 mins which contained other cognitive experiments.

Results and Discussion

The measured variable was the number of times out of a maximum of ten that subjects chose the positive exemplar (the one containing the invariant nose) over the negative exemplar. Table 3 shows the means for both groups.

Table 3: Mean no of times positive item chosen over negative (out of 10).

Test Delay	Immediate	30 min lag
mean positive score	5.15 (1.66)	5.65 (1.80)

One sample t-tests show that the there is no effect of the invariant with a group of subjects tested immediately after the study phase ($t(39) = 0.56$, $p > 0.1$) but the group which had a test delay of 30 mins showed scores which were significantly above chance for choosing the positive item over the negative ($t(39) = 2.26$, $p < 0.05$). A one-way ANOVA between these groups was not significant however ($F(1,78) = 1.67$, $p > 0.1$). This complicates the interpretation of the results. However, remembering that the counterbalancing procedure was designed to remove between subject variance and not within subject, we might expect both groups to show a large range of scores which will give disproportionately large standard deviations in comparison to the difference between means. Taken together with the results from previous experiments (which were designed to control for within subject variance) it may be argued that while interpreting these results as showing a difference between the groups may not be the most conservative interpretation, it can be supported from the methodological approach taken. With respect to awareness of the invariant feature, only one subject noticed the invariant feature when questioned. This subject was discarded and replaced. Most other subjects reported the eyes as being invariant (which was not the case) and even when given a booklet to look through, most subjects found it very difficult to report the nose as being invariant.

This experiment was designed to test the results found from the digit experiments using a completely different method. This method was designed to remove between subject variance by fully counterbalancing the items whereas the digit experiments were designed to remove within subject variance by limiting the time they had to make a decision. Thus a smaller effect than was found in the digits experiment was expected for the Identikit test. The results of Experiment 3 were replicated. Subjects given an immediate test phase did not choose the positive test exemplar at above chance levels but subjects given a test phase with a delay of 30 mins did choose the positive over the negative. This seems to rule out the suggestion that subjects are using some explicit feature of the stimuli to make the decision as one would expect this to be equally if not more effective immediately after the study phase as opposed to 30 mins afterwards. The fact that the stimuli were fully counterbalanced also would seem to suggest against this alternative hypothesis. In addition, the orienting task was extremely believable and the invariant seemed extremely difficult to detect so these should limit the amount of explicit strategy available to the subject. These results and those of Experiment 3 would seem to fit in with the sort of findings by Posner and Keele (1968) and Homa et al (1981) which suggest that a prototype is extracted from a set of specific items and over time this prototype has more influence on memory and category decisions as compared to the individual instances.

If a prototype is being extracted from a set of instances the next question becomes that of determining what form the prototype might take. The most interesting finding reported in McGeorge and Burton (1990) is that knowledge of the invariant seems to be encoded in a conceptual manner. They report transfer of knowledge form study strings of the form “5927” to test strings of the form “five nine two seven”. They also attempted to remove any perceptual overlap between study and test by asking subjects to perform a perceptual task on the study strings as opposed to the usual addition task. Thus, if prototype formation was taking place it could not have been at a perceptual level.

It is difficult however, to imagine how the invariant nose in Experiment 5 could have been differentially coded from other noses at a conceptual level. There are few studies which show transfer of implicit knowledge across different surface forms (e.g. Altmann et al, 1995; Manza and Reber, 1992; Mathews et al. 1989; Reber, 1969). In dynamic systems tasks, Berry and Broadbent (1988) found that subjects were unable to transfer knowledge learned on one system to a system which had identical underlying equations but was based on a different cover story. In fact, it seems plausible that effects such as those found in the McGeorge and Burton paradigm should be mediated by perceptual mechanisms. Jacoby and Dallas (1981) term such perceptual learning "perceptual fluency" and suggest that there is no reason for subjects to hold any verbalisable knowledge about such learning. This idea is closely related to that of Schacter (1991) who suggests a modular mechanism which he calls the Perceptual Representation System, and suggests that such a system holds perceptual information which is independent of conscious awareness. Clearly the result obtained by McGeorge and Burton suggest that an explanation of such perceptual learning is not sufficient to account for the invariance effect as one would not therefore expect to find knowledge transfer across surface forms.

However, given that the original McGeorge and Burton study confounded invariance and similarity and also did not control for extremely salient items (Cock et al, 1994; Wright and Burton, 1994) it is unclear whether these will have had an effect on such transfer. In addition, the method used by McGeorge and Burton to remove perceptual overlap did not actually prevent subjects from using such means. Subjects were required to count the number of horizontal lines in each stimulus which while not requiring phonological encoding of the stimulus certainly does not prevent it. Some subjects might develop a strategy to aid them in this orienting task which specifically involves phonological encoding of the stimuli. Instead of individually counting the lines on each digit for all 30 training exemplars subjects may have found it easier to have subvocalised "eight has three lines, seven has one line" and so in fact attain phonological overlap in encoding between study and test. It is clear that the transfer of invariance knowledge between different surface forms requires replication using the more stringent methods of stimulus construction already described and also using an orienting task which prevents subjects from phonologically encoding the training stimuli. Experiment 6 examines whether transfer across surface forms is possible when these considerations are implemented.

2.7 Experiment 6: Is knowledge of the invariant held at a conceptual level

One of the most important claims made about implicit invariance detection is that the 'rule' is held at a conceptual level. Evidence for this comes from McGeorge and Burton (1990) who reported transfer from digits to words. McGeorge also (1990) reports transfer for the converse. It has been established that performance on this task may be mediated by a number of possibly explicit strategies and so it is unclear whether this transfer across surface forms is truly a phenomenon of implicit learning or whether it may be due to explicit strategy use. Furthermore the procedure used by McGeorge (1990) and McGeorge and Burton (1990) does not totally rule out an explanation of perceptual fluency as phonological encoding was discouraged but not prevented. This experiment remedies these problems by requiring subjects to perform an articulatory suppression task at study in addition to a number encoding task. The encoding task is to look for the position of the lowest digit in the number. This task must access semantic information about the digits and so any null effect cannot be due to a lack of semantic activation. The stimuli used were those constructed for Experiment 1 with both test exemplars equally similar to a previously seen instance (PSNS). Thus any effect cannot be to due a similarity effect rather than an effect of the invariant. Experiment 6 examines three experimental groups. One is a replication of the fast PSNS effect of Experiment 1 with no change in surface form (digit-digit). The other two conditions both examine whether learning will transfer from numbers in word form at study to test exemplars composed of digits. One of these conditions used the fast deadline manipulation and the other gives subjects unlimited amount of time to make a response. Above chance performance is expected in the digit-digit group as a replication of the results found in Experiment 1. No learning is expected in the slow word-digit condition as there is no learning found in a slow digit-digit from Experiment 1. One would not expect knowledge to transfer across surface form if it is not present when the test exemplars are in the same form as the study exemplars. The condition of most interest is that of fast word-digit. If the information abstracted from the study exemplars is held at a conceptual level then we would expect above chance performance in this condition. However if the results of McGeorge (1990) and McGeorge and Burton (1990) are due to either phonological fluency or explicit contamination of the task, then we would expect no learning in the fast word-digit condition.

Method

Subjects

75 University of Glasgow undergraduates participated in return for a small payment.

Materials.

The materials were the PSNS exemplars as used in Experiment 1 except that only three of the digit sets described in Experiment 1 were used in this experiment. The invariant digits in these sets were 5,6 and 7.

Design and Procedure

Subjects were randomly allocated to one of three conditions. One group saw the study numbers in digit format and the other two groups saw them in word format (e.g. the number "1234" would be displayed in word format as "one thousand two hundred thirty four". The subjects were told that they simply had to make a keypress corresponding to the position of the lowest digit in the number. Subjects were also required to perform this under conditions of articulatory suppression. This involved the subjects repeating the phrase "alpha beta gamma" continuously until the task ended. After a 5 min intervening task involving face recognition experiments, subjects were given the surprise test phase. As with the preceding experiments, subjects were presented with two four digit numbers and falsely informed that they had seen one of the numbers before. Both test numbers were in digit format for all three groups of subjects. The group who received study numbers in digit format and one of the groups who received word format study exemplars were required to perform the pseudo-recognition task under speeded conditions using a 2 sec deadline. The remaining group had unlimited time to make a response. Following the test phase, all subjects were tested for awareness of the invariant rule.

Results and Discussion

No subject reported awareness of the invariant digit. One subject in the slow word-digit condition reported noticing the invariant in many of the numbers and using this information in the test phase. This subject's results were discarded and replaced. Table 4 shows the mean number of positives chosen over negatives for each condition.

Table 4: Mean number of positives chosen over negatives.

Number format during study.	Deadline at test.	Mean no. of positives chosen.	Standard Deviation
Digit	2 sec	6.02	1.58
Word	2 sec	5.06	1.50
Word	Unlimited	4.94	1.77

A one-way ANOVA shows that there is a significant difference between the three groups ($F(2,74) = 3.34$, $p < 0.05$). Planned comparisons show that the fast digit-digit condition is significantly different from the fast word-digit ($t(72) = 2.35$, $p < 0.05$) and from the slow word-digit ($t(72) = 2.10$, $p < 0.05$) groups but that the two word-digit groups were not statistically different from one another ($t(72) = 0.26$, $p > 0.1$). A one-sample t-test confirms that performance in the fast digit-digit group is above chance ($t(24) = 3.23$, $p < 0.05$).

As predicted, above chance performance was found in the fast digit-digit condition. This replicates the finding of an effect of the invariant under conditions which make the use of explicit heuristics extremely difficult. No learning was found for the slow word-digit condition but this is unsurprising as with materials which were specially constructed to avoid the problems of similarity and idiosyncratic items, no learning was found in a slow PSNS condition where the study and test items were in the same format. Learning was not found in the fast word-digit condition. This suggests that the apparent transfer of knowledge across surface forms reported in McGeorge

(1990) and McGeorge and Burton (1990) was a consequence of methodological inadequacies in the task and that any rule abstraction does not take place at a conceptual level.

2.8 General Discussion

McGeorge and Burton (1990) reported learning of an invariant feature of a stimulus display when subjects had no apparent explicit knowledge of the invariant feature. This finding provided strong evidence for the existence of implicit learning as it could be argued that subjects would be able to report such a simple rule as an invariant digit if they had any explicit knowledge of the rule. Furthermore, they reported that the acquired knowledge could transfer across surface features and suggested that such knowledge was therefore held at a conceptual level. This finding supported the view of Reber (1969) who suggested that implicitly acquired information is held at an abstract level of representation. Although McGeorge and Burton's (1990) task seems clearly to demonstrate implicit learning of an invariance rule, various objections to this conclusion have been raised. Wright and Burton (1995) provide evidence that the task can be performed using a strategy of rejecting salient negatives. Churchill and colleagues (1995, 1996) report various other strategies for performing the task such as unfamiliarity of negatives and re-using the study task as a memory cue. Cock et al (1994) manipulated the test exemplars in terms of their similarity to a previously seen instance and also whether they conformed to the invariance rule. Cock et al found that when these two factors were tested in opposition to one another, subjects chose the more similar exemplar as opposed to the exemplar obeying the rule. This finding is strong evidence against the invariance learning hypothesis and suggests that the apparent rule learning is actually an artifact of a similarity matching process.

Experiment 1 of this chapter examined three of the conditions used by Cock et al. In two of these conditions the test exemplars were equally similar to a previously seen instance and the results of Cock et al would predict no preference for either the positive or negative exemplar. Under similar testing conditions to Cock et al this was what was found. However, in the third condition which opposed the invariance rule with item similarity (PDNS) subject performance was again at chance. If the underlying mechanism was one of similarity matching then it would be expected that the Negative Similar exemplar would be chosen over the Positive Dissimilar. The same three conditions were also presented under time deadline conditions and again the PDNS manipulation did not show that subjects displayed a preference for the more similar item. In fact the trend for both PDNS manipulations was for a preference of the Positive Dissimilar exemplar. This inability to confirm Cock et al's results suggests that a similarity matching process does not underlie performance on this task but the mostly chance performance on the task does suggest that the robust effects shown in McGeorge (1990) and McGeorge and Burton (1990) may be due to the construction of items which are salient or similar in addition to obeying the invariance rule.

One of the conditions (PSNS) in Experiment 1 did show above chance preference for the positive exemplar but only in the time deadline condition. Although significant, the fact that the other five conditions were at chance in this experiment suggested that little weight should be given to this possibly spurious result without replication. Experiment 2 provided that replication and as in Experiment 1, a PSNS condition under time deadline showed a preference for the positive exemplar. One possible explanation is that when the two items are equally similar to a previously seen instance, it becomes much more difficult to use an explicit heuristic to choose between them and implicit processes become the deciding factor in the choice. What is clear from Experiments 1 and 2 is that memory for specific study exemplars cannot account for performance on this task.

Experiment 3 examined the hypothesis that if memory for specific exemplars was not involved in performance then prototype abstraction may play a role. Posner and Keele (1968, 1970) found that a delay between study and test in a categorisation task led to increased weighting being given to a prototype abstracted across all exemplars and less weighting being given to individual exemplars. If an invariance rule is abstracted across the 30 study numbers in the McGeorge and Burton (1990) paradigm, then it might be expected that a lag between study and test will show a robust preference for the positive exemplar. Using the PDNS condition (which was not significant for a slow or a fast deadline in Experiment 1) with a 30 min lag between study and test, Experiment 3 showed that subjects choose the positive over the negative exemplar at above chance levels. Subjects do not show this performance with an immediate test, however, and the chance result for a fast PDNS condition in Experiment 1 was replicated. This provides support for the idea that a prototype is abstracted from the instances learned at study.

Experiment 4 shows that this lag effect does not last for an interval of one week. Hence, one can discount priming mechanisms as being the basis for learning of the invariant digit, as priming has been shown to last months or even years (Sloman et al, 1988; Maylor, 1995). Experiment 5 tests this lag effect again but without a time deadline and using pictorial stimuli. Subjects saw 100 computer generated faces and were asked to find five 'suspects' from written descriptions. Hence, subjects were not performing an orienting task that could be used at test as was the case with the digits. Also, subjects believed they were performing a task related to face identification for forensic purposes and did not suspect a later test phase. This prevents subjects re-using the study task at test as a strategy which Churchill et al (1996) and Gilmore and Churchill (1996) suggest may be one way of performing on this task. Despite no explicit awareness of the nose being invariant during the study phase, subjects chose an exemplar with the same nose over one with a new type of nose at above chance levels after a 30 min delay. Subjects given an immediate test were at chance for choosing the positive exemplar. Unfortunately, the two conditions were not statistically different from one another. Nevertheless, this finding does lend some support to the conclusions of Experiment 3 in that it seems some abstraction takes place and the prototype becomes more influential over time. The findings of Experiments 3 and 5 also argue against the sort of explanation suggested by Wright and Burton (1994). If subjects are simply rejecting salient negatives (or accepting salient positives) then one would expect these salient numbers to be best remembered immediately after the study phase and so performance on the task to be at least equal if not better than after a delay of 30 mins. As the opposite result seems to be true it is reasonable to reject the account of performance by Wright and Burton (1994) save to say that saliency no doubt would play a role as an explicit heuristic if allowed to.

Experiment 6 examines whether the conclusion reached by McGeorge and Burton (1990) that conceptual abstraction takes place in this task is valid given that their exemplars did not take saliency or familiarity into account and they did not adequately control for phonological overlap between study and test. Experiment 6 found no evidence for conceptual abstraction but did replicate the results found for a PSNS condition under a time deadline with articulatory suppression conditions added.

Given the results of these experiments, it seems that the explanations put forward by Cock et al (1994), Churchill and colleagues (1995, 1996) and Wright and Burton (1994) for above chance performance on the digit invariance task cannot account for the performance of subjects as shown in this chapter. While their explanations do describe ways in which this task can be performed, they do not resolve whether implicit learning can still take place when the factors they describe are taken into account. The experiments in this chapter have attempted to test this hypothesis and the results suggest that robust implicit learning effects can still be found under certain circumstances. These circumstances seem to be whenever access to specific instances of the training stimuli is reduced or is of little benefit. Experiment 3 would suggest that a prototypical exemplar is abstracted. While it may be that this prototype is extracted over time, the results of Experiments 1, 2 and 6 suggest that this is not the case. In the PSNS condition, the positive is still favoured immediately after test but it will be argued that because the two test stimuli are both very similar to a previously seen instance, 'fuzzy' recollection of that instance will be of little benefit. Hence, the prototype seems to be abstracted immediately but does not have much effect unless memory for a specific instance does not help differentiate between the two test exemplars or the individual exemplars are given much less weighting over time. This suggestion could be formulated as a distributed model as suggested by McClelland and Rumelhart (1985). The model they propose conceptualises memory as a collection of individual traces which are superimposed. They assume that this superposition occurs at time of storage and that this superposition of traces "automatically results in abstraction though it can still preserve to some extent the idiosyncrasies of specific events", which would fit with the data presented in this chapter.

As Seger (1994) notes, the problem of the implicit / explicit distinction is more difficult to resolve with regards to connectionist models. One would surely have to argue that the abstracted prototype is implicit knowledge, or else subjects would be able to give a verbal report of the invariant. The status of the individual instances is more difficult to define. One might argue that as subjects can probably recall some specific instances that they observed then the instances must have the capability of being held as explicit knowledge.

McClelland and Rumelhart (1985) emphasise the importance of the traces being context dependent. It is possible that as each of the individual traces have an associated context (e.g. they are four digit numbers, they occurred in a particular learning set etc.) it is this which can allow them to be explicitly recalled. An abstracted prototype would have no such context (if the mechanism is some sort of invariance learning then a notion of 'threeness' or 'fiveness' bears little contextual relation to a specific four digit number) and so would not be open to explicit processing. Context may be a way to distinguish implicit from explicit in this particular type of modelling.

Another possible model with which to account for the data presented in this chapter has been put forward as a theory of perceptual priming. The perceptual specificity (as demonstrated in Experiment 6) in conjunction with lack of conscious awareness is reminiscent of Schacter's (1987) Perceptual Representation System (PRS). The PRS is supposedly presemantic and modular in that the information that it holds is not available to conscious awareness. This idea of system modularity which encapsulates information has also been suggested by Fodor (1983) and Rozin (1976) and seems to have parallels with what is known about implicit knowledge (Berry and Dienes, 1993). The results of Experiment 4 suggest that the invariance effect is not the same as that of perceptual priming in that the effect was not evident after a delay of one week whereas priming would still be expected (Sloman et al, 1988; Maylor, 1995). Presumably the difference is that priming refers to the study of one particular instance whereas the digit invariance task relies on abstraction across a set of instances. Nevertheless, the fact that the digit invariance task is perceptually mediated without explicit awareness of knowledge (as is perceptual priming) suggests that a system which can mediate the effects found in this chapter may be viewed as a modular system which encodes invariant information. One would then have to account for the possible explicit status of the individual instances. This could be simply achieved by postulating that the individual instances then undergo semantic processing which is associated with explicit awareness and that the semantic processing system does not have the capability of automatic abstraction. This could be viewed in the same way as allowing contextual cues to be necessary in the distributed model discussed above. It could be said that both models are describing the same process but at different levels. Hence, the data described in this chapter seem compatible with currently existing models of memory, abstraction of prototypes and other unconscious processes.

Finally, if abstraction does occur at a perceptual level then this result may hold an important suggestion for research on implicit learning. From the other experiments reported in this chapter there is evidence that using a time deadline and a lagged test phase constrains the amount of explicit contamination of the task. Where these procedures were implemented in Experiment 6, no conceptual transfer of information was found. This suggests that other implicit learning paradigms which claim to show conceptual transfer between study and test should be re-examined using procedures such as time deadlines to ensure that the effect which is claimed really is a manifestation of implicit as opposed to explicit processes.

Chapter 3

An examination of differential test demands

3.1 Introduction

One of the problems in implicit memory or learning is assessing the contribution of explicit awareness to subjects' performance. While dissociations between task performances may be reflecting actually dissociable cognitive mechanisms there is an alternative interpretation that there is just one mechanism for learning but that the different tasks are sensitive to different levels of information. This idea is the basis for Shanks and St. John's (1994) Sensitivity Criterion. This argument is illustrated most clearly in normal memory research where it is known that recognition has access to more information than recall. It is not suggested however, that there is a cognitive system to learn the information that will be recognised and another separate system for the information to be recalled. It is merely that one test of memory is more sensitive to the stored information than the other test.

Early subliminal perception experiments (e.g. Balota, 1983; Marcel, 1980, 1983) were criticised on methodological grounds in that subjects were simply asked whether they had seen a word or not. Cheesman and Merikle (1984) defined this type of awareness task as testing whether the information held by the subject is below a "subjective threshold". This threshold is where subjects claim not to have noticed a word. A more powerful test of awareness would examine what Cheesman and Merikle (1984) define as the "objective threshold" which is where subjects actually respond to a forced-choice task at chance levels. Any priming found from subjects who were classified as having knowledge below the objective threshold would be much stronger evidence for semantic activation without conscious identification than priming from those with knowledge between the subjective and objective thresholds (Holender, 1986).

Similar criticisms have been applied to artificial grammar learning. and Dienes, Broadbent and Berry (1991); Dulany, Carlson and Dewey (1984) and Perruchet and Pacteau (1990) all used recognition tasks which showed that the information which subjects can consciously retrieve can explain classification performance without recourse to an additional implicit knowledge system (Shanks and St. John, 1994). This may also be seen as a criticism of the claims of the digit invariance task, as the test phase is a forced choice recognition task and it thus could be argued that this taps explicit knowledge at a different level of sensitivity than a verbal awareness test. While it is still a contentious issue as to whether or not recognition tasks tap implicit or explicit knowledge, it may be informative to know if the invariance learning effect is present under a task less sensitive than a forced choice procedure but more sensitive than verbal awareness. In artificial grammar tasks, the test phase is similar to a yes / no recognition task with only one test exemplar at a time being presented. This chapter will examine performance on the digit invariance task when the test items are presented individually.

3.2 Experiment 7: Comparison of 2-AFC and Y / N recognition tests

This experiment examines the effect of giving subjects a yes / no recognition task during the test phase of the digit invariance task (McGeorge and Burton, 1990) instead of the more usual 2-AFC task. If the objective threshold is indeed a criterion which demarcates one type of knowledge from another then subjects should show above chance performance on both types of task.

Method

Subjects

40 undergraduates from the University of Glasgow participated in the experiment in return for a small cash payment.

Design and Procedure

The materials were those used in Experiment 3 and the four sets of materials were counterbalanced across all subjects. One group of subjects was assigned to a test condition which involved a binary recognition judgement (as in the experiments reported in the previous chapter). The second group of subjects performed a yes / no decision on individually presented items. In the training phase all subjects were presented with 30 four digit numbers one at a time on a computer screen. They were required to add the two digits on the left hand side and the two on the right hand side and make a key press as to which sum gave the higher total. This was followed by a 10 min lag period during which subjects performed face recognition experiments and then an unexpected test phase in which subjects in the binary judgement condition were told that they would be presented with two numbers, one of which they had seen in the earlier phase. Subjects were required to make an appropriate key press as to which one they thought it was. Subjects in the other condition saw the same test exemplars except they were presented individually and not in pairs. They were asked to say whether the item had appeared in the previous study phase or not. Subjects were told to guess if they were unsure. They were instructed that they had a maximum of 10 seconds to make a response for each pair. This was included to ensure that subjects did not use overly elaborate recall strategies. After 10 secs the numbers would disappear and a response would not be allowed. After subjects had responded to all 10 test pairs they were asked a series of questions about how they had made their choice. These questions were the same as those used in Experiments 1-6.

Results

No subject was aware of the invariant digit and most claimed to have used general familiarity or guessing as a strategy. Three subjects in the two-alternative forced choice condition (2-AFC) did miss the 10 second deadline: two missed one deadline and one missed two. Their data was scored *pro rata* out of 10.

Subjects in the 2-AFC test phase chose the positive exemplar over the negative with a mean of 5.86 out of 10 and standard deviation of 1.27. A one-sample t-test shows this to be significantly different from chance ($t\ (19) = 2.94$, $p < 0.05$).

The subjects in the single item test phase did not perform significantly above chance with a mean score for items correctly chosen or rejected of 10.6 out of 20 and standard deviation of 2.5 ($t\ (19) = 1.05$, NS). This null effect was not due to bad performance on just choosing the positives or rejecting the negatives. Subjects chose a mean of 5.35 positives out of 10 (std dev = 1.67) and rejected a mean of 5.25 positives out of 10 (std dev = 1.60), neither of which is significantly above chance ($t\ (19) = 0.91$; $t\ (19) = 0.68$, respectively).

Despite subjects' knowledge of the invariant not being accessible according to a subjective threshold (i.e. they could not report the invariant digit), this knowledge is accessible above the level of an objective threshold (subjects could make the correct response in a forced choice recognition task). However it is not accessible on a yes / no recognition task. This suggests a further division between the subjective and objective thresholds. The next experiment tests whether this result is reliable in a related paradigm.

3.3 Experiment 8: Y / N recognition in the 'Clocks' invariance task

Bright and Burton (1994) used a similar paradigm to that of McGeorge and Burton (1990) but instead of the invariant being that of a particular digit appearing in every number, it took the form of an abstract rule. The rule was that on a series of clock faces, the hour hand always appears between the 6 and the 12 on the left hand side of the clock. In a test phase Bright and Burton (1994) found that subjects would choose a clock with a new time when the hour hand was between 6 and 12 over a similar clock which had six hours added to the time. They propose that in order to accomodate these results, one must either posit a rule-based mechanism or an instance based mechanism with much deeper levels of abstraction than is usually considered. While this task has not been investigated as fully as that used in McGeorge and Burton (1990) it will provide a test as to whether this null learning effect found in a yes / no recognition version of the digit invariance task is robust and can be applied generally. This experiment follows the procedure described in the previous experiment but used the clocks as stimuli instead of four digit numbers.

Method

Subjects

24 undergraduates from the University of Glasgow participated in the experiment in return for a small cash payment.

Materials

For purposes of replication, stimuli were constructed to match closely the description given in Bright and Burton (1994) from which the paradigm was taken. A database was constructed which contained all possible times in 5 min intervals from 6.00 to 12.00. 40 items were chosen at random from this database to form one set of task materials. 30 of these times were used as the study items and the remaining 10 were used as the positive test items. The negative test items were made by adding 6 hours to the times of the positive items. Four such sets of items were constructed.

The sets of times were presented on analogue clock faces (see Figure 5 for examples). Three different shapes of clock face were used: octagonal, circular and diamond. Four different legend designs were used: Arabic numerals, Roman numerals, straight lines at the 12,3,6 and 9 positions and no legend at all. The clock hands were simply a short hour hand and a long minute hand on all clocks and both of rectangular design. These variations allowed a plausible incidental learning task to be used at study and were spread evenly across all four sets of times. The clock faces were constructed using a drawing application on an Apple Macintosh computer and all clocks were printed in black and white onto A4 paper. Each clock was approximately 15 x 15 cms. For the study phase, a five point Likert scale was constructed containing the scale, 'Very Poor, Poor, Average, Good and Very Good'.

Design and Procedure

The four sets of materials were counterbalanced across all subjects and all subjects received the same training phase. Each subjects was given a sheet of paper which had a five-point Likert scale printed on it for 30 items. The subject was informed that they would see 30 clock faces one at a time and had to rate each clock on the scale as to how good they thought the clock face was, either aesthetically or practically. The orienting task was self-paced and subjects were not informed that there would be any follow-up task. On completion of the study task, the materials were removed and subjects were informed that they would now get a recognition test for some of the clocks.

Half of the subjects received the test phase described in Bright and Burton (1994). They were presented with a pair of test items, one positive and one negative, and falsely told that they had seen one of them before. The left/right position of the positive and negative was varied randomly. Subjects were asked to choose which one they remembered seeing and were told that they could guess if they weren't sure. All 10 test pairs were presented in this way. The other half of the subjects were given exactly the same test stimuli but were shown them one at a time, not in pairs as is usual to this paradigm. They were told that they had seen some of the clocks before and were to respond yes if they thought they had seen each particular clock in the earlier phase and no if they thought the clock had not appeared in the earlier phase. They were allowed to guess if unsure and were not constrained by a time deadline.

Figure 5: Examples of the clock stimuli

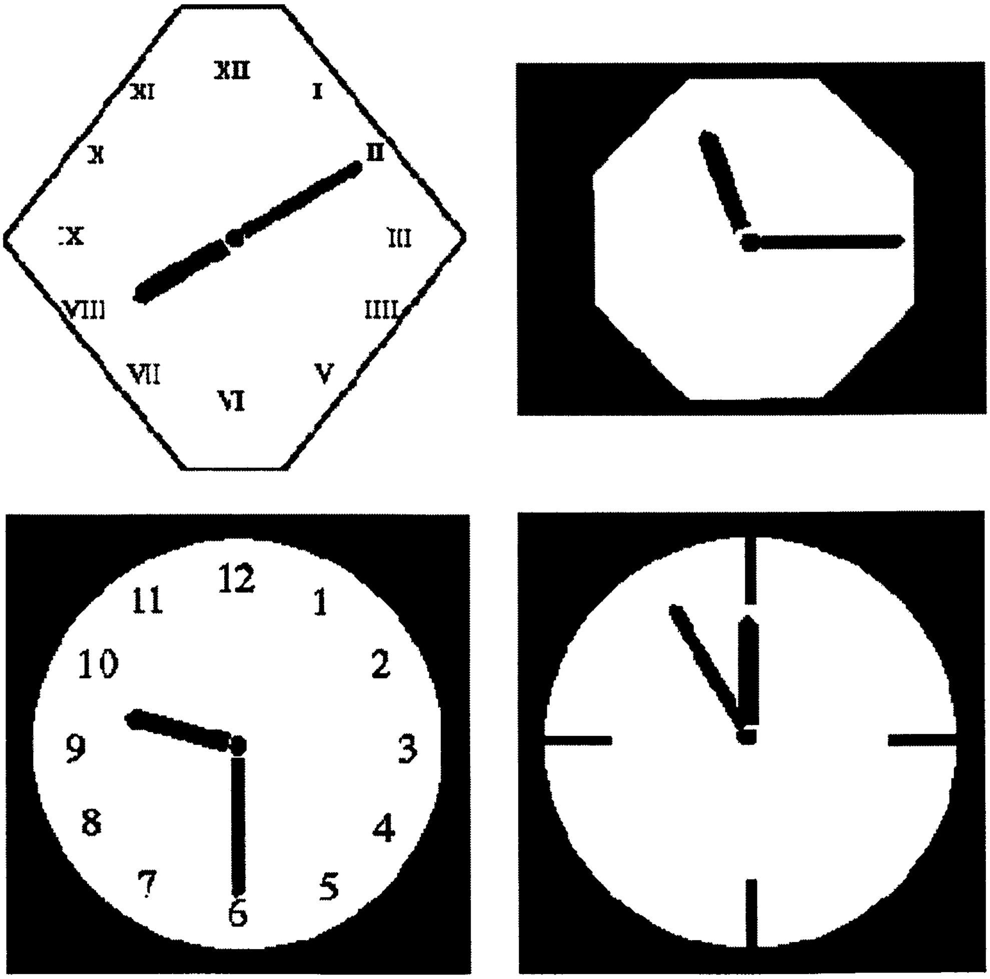

To ensure that subjects' responses were not based on explicit knowledge of the rule, some of the same awareness evaluation questions used by Bright and Burton were given to the subjects immediately after the test phase. These were:

Do you have any comments to make about the experiment?

Did you notice anything peculiar about the clocks I showed you earlier?

Did you notice if any of the items had anything unrelated to each other or in common with each other?

Can you think of anything the times had in common?

Results

One subject was classified as aware from the post-task questions. The subject was in the group which had the two-alternative forced choice task at test and scored nine out of ten for choosing the positives. This subject was discarded and replaced. The remaining subjects gave no indication that they were aware of the rule and most seemed surprised when they learned of it.

Subjects in the 2-AFC test group selected a mean of 6.5 positives out of a possible 10 with a standard deviation of 1.2. A one-sample t-test comparing this to a chance score of 5 out of 10 reveals a significant effect, $t(11) = 4.1$, $p < .01$. This replicates the effect of the invariant rule found by Bright and Burton (1994).

Scores for the subjects in the single item test group are out of 20 and are composed of positives to which the subject said yes and negatives to which the subject said no. The mean score for subjects in this group was 10.92 out of 20 with standard deviation 2.57 which was not significantly different from chance on a one-sample t-test ($t(11) = 1.19$). This null effect was not due to bad performance on just choosing the positives or rejecting the negatives. When examined separately, it seems that subjects could do neither (mean no. of positives accepted out of 10 = 5.67, std dev = 1.72, $t(11) = 1.29$, NS; mean no of negatives rejected = 5.33, std dev = 1.72, $t(11) = 0.64$, NS).

As in the previous experiment, a forced choice task seems to elicit information which is not demonstrated by verbal report or by a yes / no recognition task. There may be other testing conditions which will elicit such evidence, however. Robust evidence of implicit learning was demonstrated in the PDNS condition of the digit invariance task only with a two second time deadline and a 30 min lag between study and test. If a further yes / no recognition procedure does not elicit evidence of implicit learning under these conditions then this would represent stronger evidence for a difference between single item and forced choice recognition in an implicit learning task. Experiment 9 uses the McGeorge and Burton (1990) paradigm with a 30 min interval between study and test and presentation of single items under a response deadline.

3.4 Experiment 9: Yes / no recognition test after lag

Method

Subjects

16 University of Glasgow undergraduates were paid for their participation in this experiment.

Materials

The materials were the same as those used in Experiment 3.

Design and Procedure

Subjects were given one of the four sets of exemplars and the test exemplars for a PDNS condition. Subjects performed the addition orienting task at study (as in Experiment 7) and were given the pseudo-recognition task 30 mins afterwards. The 30 min interval was filled with face recognition tasks. As robust learning effects have already been demonstrated for forced choice recognition under these conditions, only the single item presentation condition will be used. All other aspects were the same as in Experiment 7.

Results and Discussion

Subjects correctly classified a mean of 10.59 exemplars out of 20 (std dev = 1.65) which was not significantly different from chance (t (15) = 1.33). This supports the finding of Experiment 7 in that even under conditions where robust learning was demonstrated for a forced choice task, evidence of similar learning is not elicited under yes / no recognition conditions. This suggests that the finding is indeed reliable across different invariance learning paradigms and under conditions known to promote implicit processing. As in previous experiments there was no advantage in choosing positives or rejecting negatives (mean = 5.66, std dev = 1.92, t (15) = 1.29; mean =4.93, std dev = 0.97, t (15) = 0.27, respectively).

One possible mechanism for this discrepancy between the two types of recognition task can be illustrated by referring to signal detection theory. In signal detection theory, one can conceptualise a yes / no recognition task as choosing between two populations; one of signal plus noise and one of noise alone. This would correspond to the digits tasks described as choosing between a PD (signal plus noise) or an NS (noise alone) exemplar (see Figure 6). This type of decision process requires the observer to set a threshold at some point between these two distributions. For each stimulus presentation the observer assigns a likelihood ratio to the exemplar and accepts it as signal plus noise if this ratio exceeds the threshold value.

Figure 6: Probability density functions of signal plus noise and noise.

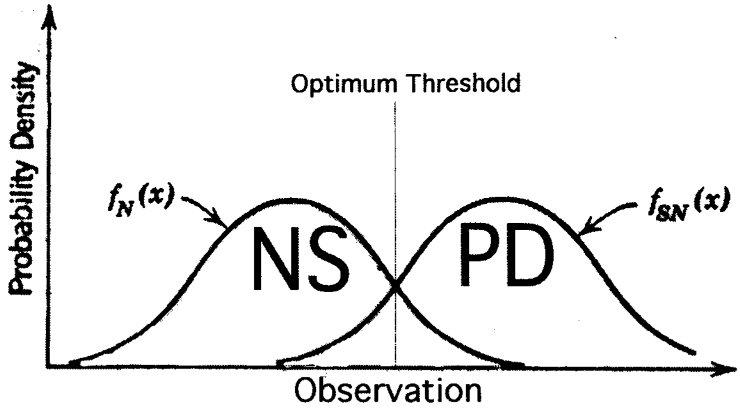

In a forced choice procedure, the observer does not set a criterion. The optimal procedure for the observer is to assign a likelihood ratio to each test exemplar and choose whichever one has the greater value (Green and Swets, 1966). Thus in the forced choice procedure the presence or absence of the invariant would have a dynamic effect in each trial whereas in a yes-no recognition task the criterion for accepting or rejecting the exemplar is set at some point along the continuum of observations (in this case the study set) and will not alter during the task unless the task parameters are altered. Thus one might expect the threshold in a yes / no recognition task to be determined by how similar a test item is to the continuum of instances and not by the presence or absence of the invariant, as it is presumably with reference to the study instances that a threshold is chosen. The subjects in Experiments 7 and 9 did not choose the more similar exemplar over the less similar one, however. Although this result does not support the explanation of the results by signal detection theory, one might expect that subjects' performance would be poor when based on 'similarity'. This is because the notion of similarity used in this thesis is a narrow definition; that of similarity to a previously seen instance as was used by Cock, Berry and Gaffan (1994) and may not be an appropriate usage in terms of signal detection theory. It is also likely that the subjects' memory for the individual exemplars is poor due to the lag between study and test and also to the incidental learning instructions. Hence it might be expected that the criterion based on memory for the exemplars would be less than optimal if it can be assumed that the optimal criterion is one based on perfect encoding and retrieval of all study exemplars.

It may be possible to address these problems and test whether the invariant can have an effect in yes / no recognition. The test must have incidental instructions (or else the rule becomes transparent) therefore this is still something which may cause the criterion that subjects adopt to be less than optimal. The other methodological procedure of concern was that of introducing a lag between study and test. This may have lead to some exemplars being affected by decay of the exemplar traces and so causing the threshold to be less than optimal. It has already been demonstrated that there seems to be a robust effect of the invariant under immediate testing conditions if a speeded PSNS condition is use.

Another benefit of using a speeded PSNS condition is that it reduces the discrepancy in similarity between the positives and negatives. Although one cannot be sure of the accuracy in measuring 'similarity', if both populations are constrained to be 75% similar to an individual instance then this must have the effect of reducing the noise due to subjective similarity.

3.5 Experiment 10: Yes / no test with deadline and PSNS exemplars

Experiment 10 is identical to Experiment 9 except that the speeded PSNS exemplars are used and the test phase follows immediately after the study phase. As with

Experiment 9, the speeded PSNS condition has shown to give robust results under forced choice conditions and so will not be included in the present experiment.

Method

Subjects

12 undergraduates from the University of Glasgow were paid for their participation. All were naive to the task.

Materials

The materials were four sets of those used in the PSNS condition of Experiment 1.

Design and Procedure

The design and procedure were the same as those used in Experiment 9 except that the study and test exemplars for the PSNS condition were used and the test phase followed immediately after the study phase.

Results and Discussion

Subjects were exactly at chance for choosing positives over negatives (mean = 10.00, std dev = 2.15). This time however, subjects were better at selecting positives (mean = 5.95, std dev = 1.46, $t(11) = 2.15$, $p < .05$) than rejecting negatives (mean = 4.06, std dev = 2.01, $t(11) = 1.55$, ns). This result however seems more a reflection of the larger standard deviation in rejecting the negatives as the two means are the same distance away from chance. The main result of 10 items correctly accepted or rejected out of 20 is the most important indicator that as in the previous experiments, no learning of the invariant is evident when subjects are tested with a yes / no recognition task.

It was hypothesised that when the two populations of test exemplar are designed to be 'equally' similar to a previously seen instance, the yes / no recognition task would not be influenced by subjective similarity and the influence of the invariant may become more apparent. There was no evidence of any invariance effect in the categorisation of exemplars as 'old' or 'new' with the results indicating exactly chance performance. This replicates the results of the previous three experiments.

3.6 General Discussion

McGeorge and Burton (1990) demonstrated that subjects seemed to acquire some knowledge of an invariant property of a stimulus set and apply that knowledge in a test phase without having any verbal awareness of the property itself. This has been proposed as evidence for implicit learning although it could be argued that a verbal awareness test is not as sensitive to explicit information as the forced choice (pseudo)recognition used at test in the McGeorge and Burton (1990) paradigm. Berry and Dienes (1993) suggest that information not elicited by verbal report may become accessible at some later point in time (as demonstrated by Erdelyi and Becker (1974)) or alternatively because the subject can opt not to respond, they might fail to report low confidence information. Some authors (e.g. Dulany et al, 1984; Perruchet and Pacteau, 1990; Shanks and St. John, 1994) would argue that a recognition test taps into explicitly held information.

Another analogous way to demarcate implicit from explicit would be to use the definitions of subjective and objective thresholds (Berry and Dienes, 1993; Cheesman and Merikle, 1984; Dienes and Berry, 1997). The subjective threshold is where subjects believe they are responding at chance and the corollary of this in the McGeorge and Burton (1990) paradigm would be the inability to report verbally the invariant. The objective criteria is defined as the point where subjects are actually responding at chance on forced choice recognition test. Clearly awareness in the McGeorge and Burton (1990) paradigm lies below the subjective threshold but above the objective threshold.

Those who claim recognition tests are a measure of explicit knowledge would use the objective threshold as a criterion for implicit learning (and so to be classified as implicit, knowledge would need to be shown in some way other than a recognition test, e.g. as a facilitation in reaction time). The artificial grammar and digit invariance paradigms use the subjective threshold as the criterion to show implicit learning as the test phase is essentially a recognition task. Although it is difficult to reconcile these two definitions of a criterion for implicit processes, Berry and Dienes (1993) suggest that "if the knowledge is accessible by one task (e.g. classification) all relevantly similar tasks will also access it (e.g. perhaps rating parts of exemplars)". They suggest that the difficulty is in specifying what 'relevantly similar' actually means. One possibility which they suggest is that it means a test having the right context to elicit the knowledge and that " if the right context was not there, the knowledge would be inaccessible on forced choice tasks".

The results from the experiments in this chapter suggest that for the digit invariance task of McGeorge and Burton (1990) and the related clocks task of Bright and Burton (1994), the evidence for implicit learning cannot be interpreted as easily as that. In all four experiments, information which can be elicited under a two-alternative forced choice task is not demonstrated in a yes-no recognition task. Invariance detection thus seems to have the paradoxical nature of being both above and below an objective threshold as it is defined above. With reference to the suggestion of Berry and Dienes (1993) it is difficult to imagine how a yes / no recognition task could not be 'relevantly similar' to a two-alternative forced choice task. The fact that one type of recognition task will elicit information in a digit invariance task whereas another type of recognition task will not can only strengthen the claim of this task as being mediated by implicit processes. As Berry and Dienes (1993) note, one of the main characteristics of implicit learning is that the knowledge is difficult to access and as these experiments have shown, knowledge of the invariant cannot be accessed by free report or yes / no recognition and only by a two alternative forced choice task.

While these results suggest important considerations for the definition of demarcation criteria for implicit and explicit processes, it is interesting to speculate why the two recognition tasks differ in the results they give. As was noted earlier, signal detection theory draws a distinction between a yes / no recognition task and a two-alternative forced choice task. In the former the observer is required to set a threshold for accepting or rejecting a probe whereas in the latter, the observer decides which probe has the higher likelihood ratio. Presumably, it is memory for the studied instances which detemine the threshold for accepting or rejecting the test exemplar. One could speculate that as the subject has no conscious awareness of the invariant this information will not be used in determining the threshold. This would be in agreement with the yes / no recognition tasks described earlier where the presence or absence of the invariant seemed to make absolutely no difference to the subjects' decision. In a two-alternative forced choice task, no threshold is set by the subject. The subject makes an observation for each exemplar and chooses the one with the higher likelihood ratio. The hypothesis is that under implicit processing conditions the information on which that choice is made is different from that which determines the threshold in a yes / no recognition task.

While this may or may not be the mechanism by which one task is sensitive to implicit information while the other is not, it is clear that a two-alternative forced choice task and a yes / no recognition task should not be assumed to have the same status with respect to the information which they can elicit. In the case of invariance detection in the two paradigms discussed, one of these measures of recognition may tap explicit (and as such useless for performing the task) information while the other can elicit the knowledge held implicitly. These results also suggest caution when interpreting evidence drawn from paradigms which differ in the type of task demand (such as single item vs. 2-AFC test) as what is found for one type of task may not generalise to others unless the task demands can be reinstated.

Leaving aside the possible explanations of the underlying mechanisms by which the two tasks differ, these results do not conflict with the hypothesis put forward in Chapter 2 to explain invariance detection. If one imagines a perceptual prototype which captures knowledge of the invariant, then in a two-alternative forced choice task, the comparison is between the two psychological distances between the prototype and each exemplar. The one which is the closer is chosen. With only one exemplar as in a yes / no recognition task, there is no metric to decide 'closeness' so the prototype information is redundant. The hypothesis involving signal detection theory described above, is a more formal instantiation of this argument. Hence the data are not in opposition to a prototype theory as the underlying mechanism for invariance detection.

The data however, do not rule out alternative hypotheses such as similarity as one could replace the prototype concept with a notion of overall similarity and again say that it is the comparison of the two distances which is taking place and not a setting of the threshold. What these data do rule is an explanation of the sort suggested by Wright and Burton (1994). They reasoned that a salient negative would allow the subject to consciously reject that item and so choose the positive item by default. One might imagine a complementary strategy where a salient positive might be consciously accepted. If subjects are using this sort of strategy which would account for above chance performance on the invariance task then there is no reason not to accept a salient positive or reject a salient negative when they are presented individually. The hypothesis proposed by Wright and Burton (1994) would predict above chance performance on the yes / no recognition task as well as the two-alternative forced choice task and this is clearly not the case.

Chapter 4

Searching for real world invariance learning.

4.1 Introduction

Reber (1993) has proposed the idea that implicit learning depends on brain systems which are phylogenetically primitive as compared to systems responsible for explicit learning and supports this view with experimental evidence which agrees with various predictions made from evolutionary theory. He also regards phenomena such as language acquisition to be fundamental aspects of implicit learning in real life (e.g. Reber, 1997). Given that there is this idea that everyone has such a mechanism which can unconsciously learn and that this mechanism is supposed to be able to encode the complex regularities which we encounter in everyday language then it might be expected that various other regularities in the real world should be learned without any intention on the part of the learner. The experiments in the previous chapters suggest that there is good evidence for invariance detection without any intention to learn such invariances being evident. Experiment 6 suggests that the mechanism underlying such learning relies on perceptual information. Is there any evidence that knowledge of real world perceptual regularities are acquired and can exhibit a causal effect on behaviour?

Morton (1967) reports a complete lack of evidence to suggest subjects can learn real world regularities without intention. The task was to report the position of the numbers and digits on an old style of British telephone dial. Of 151 subjects tested formally, not one could successfully complete the task. Even those who were experienced users (e.g. telephonists) performed extremely poorly on this task. This study was conducted before much research on the present day concept of implicit learning had been done, so little was known about the circumstances under which evidence for unconscious processes could be elicited. It is possible that the recall task used by Morton was simply not sensitive enough to elicit low confidence information that subjects might have acquired.

Related evidence has been found in more recent experiments on memory for the features on coins (Nickerson and Adams, 1978; Rubin and Kontis, 1983; Jones, 1990; Jones and Martin, 1992); Richardson, 1992, 1993). These authors all found subjects' performance to be extremely poor on a variety of memory tasks including drawing coins from memory, drawing the coins with information on what the criterial features are, choosing the criterial features from a list and choosing the correct depiction of a coin from various alternatives. One interesting feature from both the British and American studies is that subjects draw the head so that it is facing in the wrong direction. Jones (1990) found that only 19% of subjects drew coins with the monarch's head facing in the correct direction. This is particularly surprising considering that every coin which bears the head of Queen Elizabeth II has a depiction of a right facing profile. As we come into contact with coins almost every day of our lives, the proposition that we have a mechanism to detect perceptual regularity without intention to do so should mean that we would know which direction the monarch's head faces on a coin. There is no lack of positive exemplars in the real world and there are no negative exemplars; the Queen's head never faces to the left on a coin. Possible explanations for this apparent lack of knowledge are that incidental encoding of head orientation on coins does not happen or that a bias towards left-facing stimuli is masking any knowledge which may be held implicitly. Jones (1990) suggests that a leftward bias may be as a result of interference from the orientation of the monarch's head on postage stamps. As evidence of this Jones and

Martin (1992) found that this leftward bias does not exist for Canadian subjects and suggest that it is the lack of the Queen's left facing portrait on Canadian stamps which accounts for this.

Bright (1994) presents evidence that suggests subjects do know the orientation of the Queen's head but such information is not accessible when tested with a recall task. Bright notes that the invariance detection studies require a binary recognition judgement whereas the experiments on coin memory used recall tasks. The previous chapter shows that information which can be elicited by a binary recognition judgement is not necessarily demonstrable in a yes / no recognition task. This same criticism could be applied to Morton's (1967) task where subjects were asked to recall features on a telephone dial. Bright (1994) conducted an experiment where subjects were required either to recall the orientation of the Queen's head on a coin or choose between two images of a coin. One of the images was a coin with the head in the correct orientation and the other image had been manipulated so that the Queen's head was facing in the opposite orientation. Bright also performed the same experiment using stamps as stimuli. The results reported by Bright show that the leftward bias is only observed under recall conditions. Using the recognition task, subjects could correctly identify that the head was left facing on a stamp and right facing on a coin. Bright interprets these results as being consistent with episodic processing accounts (Tulving, 1983; Whittlesea and Dorken, 1993) in that the 'trivial' aspects of the stimuli are less richly encoded and so require a more sensitive test to elicit information about them.

4.2 Experiment 11: Replication of Bright (1994)

The previous chapter reported findings of robust implicit learning of a perceptual invariant when tested via a binary recognition judgement. That Bright (1994) found such an effect for real life invariance detection when using such a task a test fits well with these results. However it is surprising that Bright reports such robust results for recognition when the findings of the previous chapter suggest the effect of a perceptual invariant is robust only under conditions where explicit heuristics can have a minimal effect and / or when the individual instances have less of an (explicit) influence and a prototype has an increased (implicit) influence. While exposure to the learning items in the laboratory based experiments and those in Bright's (1994) real life learning experiments are obviously not comparable, if Bright's results are indeed robust then they provide an interesting extension to the results reported earlier. Specifically, robust results were found when testing was not immediate but followed a 30 min lag. The retention interval for this information however did not seem to last the duration of one week. The results reported by Bright suggest that continued spaced exposure to learning stimuli provides a reliable knowledge base under binary recognition conditions. Experiment 11 is an attempt to replicate the type of findings reported by Bright (1994) for coins and stamps in a recognition task.

Method

Subjects

60 subjects were recruited from the undergraduate population of Glasgow University and from final year high school pupils visiting the university at an open day.

Materials

A digital image of a British 10p coin, a 50p coin and a first class stamp was captured using a video camera. The 50p coin was reproduced actual size and the 10p coin and the stamp enlarged using graphics software. Versions of each image which had the Queen's head facing in the wrong direction were also produced simply by flipping that part of the image about its vertical axis. Using this technique, the information on the edge of the coin remained unchanged between the correct and the incorrect version and so the only discriminating information was the orientation of the head. Both versions of each image were printed on an A4 sheet of paper, one above the other, with a light grey background. Two copies of each sheet were made with the correct version of the coin or stamp being topmost on one and the incorrect version being topmost on the other. The size of each image was 7.5 cm diameter for the 10p coin, 4 cm diameter for the 50p coin and 4 x 3.5 cm for the stamp (see Figure 7 for examples of coin and stamp stimuli).

<u>Figure 7: Correct and incorrect orientation of the Queen's head on 10p coin and stamp.</u>

Design and Procedure

Each subject received only one of the images during the test. Half of the subjects were shown the stimuli with the topmost image being the correct one and half with the topmost image being the incorrect one. Subjects were presented with the stimulus and told that only one out of the pair of coins or stamps was correct and asked to choose which one they thought it was. Subjects were given as much time as they wanted to reach a decision.

Results and Discussion

14 out of the 20 subjects correctly chose the left facing stamp as the correct one and 14 chose the right facing 50p coin as correct. 13 chose the right facing 10p coin. Although most subjects in each condition were correct, none of the conditions reach statistical significance on a sign test. One might argue, however that the data shows a trend for recognition of the invariant feature. It may be that the experiment lacked sufficient power to demonstrate significant learning and increased numbers of subjects

may have allowed the data to reach significance. The data will be discussed as non-significant although one must bear in mind that lack of power may be sufficient to explain the null result.

The seeming lack of knowledge in the stamp condition is most surprising as Bright (1994) reports significant results for stamps even under recall conditions. As there is no obvious discepancy between the two methodologies which might account for the difference in results it must be concluded that the results concerning real world knowledge are not as robust as one would believe from the data reported by Bright (1994). The results of Experiments 1-6 suggest that under certain testing conditions the influence of invariant knowledge held unconsciously may become more apparent. The presentation of two coins which differ only in the orientation of the Queen's head might be said to be analogous to a PSNS condition in a digit experiment (both coins would be 'equally similar' to the previously seen instance, differing only in the presence of the invariant quality). In the digits task, subjects did not show any preference for the positive exemplar under normal forced-choice procedures but introduction of a time deadline did elicit robust preference for the exemplar containing the invariant. If the coin experiment could be performed under a time pressure then subjects might show robust knowledge of the invariant. A problem in the design of such an experiment is that 2 seconds proved to be a suitable deadline for excluding explicit heuristics when it came to subjects having to choose between two four digit numbers. It it is not at all clear what a suitable deadline would be when comparing two coins. A possible solution to this problem would be to look at qualitative differences in responses to correct and incorrect items. Experiment 12 still has the requirement of a fast response to items but without imposing a time deadline. Subjects will be required to make reaction time responses to coins and stamps.

Another possible reason for lack of robust results in the previous experiment is that subjects were asked directly which way the head faced. The results of the digit invariance experiments suggest that when subjects can use explicit heuristics to provide an answer the effect of the invariant is lost. While it is reasonably obvious that in a digit invariance task there are many heuristics which the subject might believe to be an aid in solving the task but which in fact do not differentiate between positive and negative, it is difficult to imagine many similar confouding heuristics in the coin task. Presumably subjects simply try to remember the last time they looked at a coin. However, there is evidence to suggest that simply attempting to remember an explicit episode can have a detrimental effect on unconscious processes. Graf, Squire and Mandler (1984) report data from an implicit memory task given to amnesic patients and normal controls. Both groups of subjects saw a list of words during the study phase and were given a test phase after a short delay. The task was to complete the remaining parts of words when given the first three letters. The indirect version of the task was to complete the stems with the first word that came to mind. The direct version was to complete the stems with a word seen during the study phase. Graf et al report that under indirect task instructions the amnesic group performed just as well as the controls. However, when given direct test instructions the performance of the amnesic group fell dramatically.

The implications of these results seem to be that unconscious processes are optimal when tested indirectly. The digit invariance task has an indirect test phase in that subjects are asked to remember which item they have seen before when in fact they have seen neither. They are not told a simple rule differentiates between the two test exemplars. If subjects do have information about the orientation of the head on stamps and coins then it may be that such knowledge will be more evident under indirect testing. Furthermore, it is impossible to say how much explicit memory is involved in deciding which way the head faces on a coin or stamp and so showing above chance performance on this task does not necessarily suggest the existence of real world implicit learning. In fact some (Perruchet and Amorin, 1992; Shanks and St. John, 1994) would argue that such a recognition test is indicative of explicit knowledge. The procedure used by Bright (1994) violates Graf and Schacter's (1985) definition of implicit processing which stipulates that an implicit task should make no reference to previously encountered data yet performance on such a task should be facilitated by the previous episode.

4.3 Experiment 12: Indirect test of head orientation knowledge on coins and stamps

Experiment 12 does not ask subjects directly which way the head faces but embeds use of that knowledge in a semantic decision task. Subjects must make a British / foreign decision to a range of items as quickly as they can. It is hypothesised that if subjects do hold knowledge as to the orientation of the Queen's head on coins and stamps, they should be slower at identifying a stimulus as British when the head orientation is incorrect.

Method

Subjects

32 undergraduates were recruited from the University of Glasgow. Due to the nature of the task all subjects were British.

Materials

30 digital images of objects were collected with half of the objects classifiable as typically British or very commonly seen in Britain (e.g. a red post box, a carrot) and the other half classifiable as foreign to Britain (e.g. a French car registration plate, a chillie pepper). In addition to these stimuli were images of a 10p coin and a first class stamp (see Figure 5) with the correct orientation of the Queen's head and of the same coin and stamp with a reversed orientation of the head. All these images were presented in the centre of the screen on a black background. The dimensions of the coin and stamp were 7 cm diameter and 4 x 3.5 cm respectively.

Design and Procedure

It was decided that presenting correct / incorrect as a within subjects factor could lead to a problem of explicit episodic influences on the decision to the image which would be appearing second. A between subjects design would eliminate this possibilty and so subjects were assigned to one of four conditions with each combination of seeing a coin and a stamp alternated between subjects (i.e. both correct, both incorrect or one correct and one not). The task was presented to subjects as a reaction time task involving a binary decision. For each image subjects were instructed to press a key corresponding to 'British' or 'foreign' as quickly as they could. The first ten images presented did not contain either of the two critical items and were comprised of equal numbers of British and foreign images. This was to ensure subjects were familiar with the correct keys for each response. The two critical items were embedded within the other 20 items and all of these were shown in a random order.

Results and Discussion

Although the experimental design comprised four conditions the data will not be analysed in that form. As an orthogonal design was not used, the responses for correct and incorrect coins and correct and incorrect stamps will be analysed separately. Table 5 shows the mean RTs for responses to each type of critical item.

Table 5: Mean RT (msec) and std dev to correct and incorrect coins and stamps.

	Coin	Stamp
Correct	769 (205)	687 (130)
Incorrect	820 (194)	691 (142)
Difference	51	4

While subjects seemed to show a decrement in RT when making a decision on the incorrect versions of both the stamp and the coin, neither of these differences were statistically reliable. For the coins, a one-way ANOVA reveals no difference in reaction times between correct and incorrect items ($F(1,27) < 1$), and a similar finding for the stamps ($F(1,29) < 1$). These results suggest that subjects can recognise a British coin or stamp equally well when the Queen's head is in the correct or opposite orientation. Hence there does not seem to be any evidence that knowledge of the correct orientation of the Queen's head can be elicited by an indirect test such as that used here. Of course, it is possible that the British / foreign decision was mediated by other, more overpowering cues such as the dimensions of the object. The dimensions on a British stamp are fairly distictive and so subjects may possibly have been ignoring the inner features. A circle is not in any way distinctive and so the coin should have had its inner features inspected. In addition, amongst the foreign objects were three images of currency, two banknotes and one coin. This should also have encouraged subjects to inspect the inner features of the images or at least of the coins. While these results suggest that implicit learning of the orientation of the Queen's head does not happen through exposure to coins and stamps in real life, it cannot be concluded that such learning cannot take place for other types of information.

4.4 Experiment 13: Is there acquisition of real world colour invariance?

Bahrick (1958) reports incidental learning for colour information. If knowledge of colour can be acquired indirectly in a laboratory then colour may be a type of knowledge which can be implicitly acquired in real life. An experiment investigating this hypothesis would require an invariant stimulus which is composed of colour and prevalent in everyday life but the colour information be incidental to the function of the stimulus (i.e. subjects would not have attempted to explicitly learn the colours as would be the case with traffic lights, for example). One such stimulus which would be ideal for these purposes is the logo for the British television broadcasting company, Channel 4. The logo is a figure 4 made up of five lines, each of which is a different colour. This logo is invariant and has been displayed for 15 years before each Channel 4 programme starts. Experiment 13 examines whether subjects have any knowledge of the Channel 4 logo colours using again a direct and an indirect test. The procedure for the direct test is simply asking subjects to choose the correct version of the logo out of 6 choices. The indirect test is presented to subjects as a tachistoscopic identification task. This was to attempt to reduce the number of explicit strategies used by subjects. Again, it would have been impossible to choose a correct threshold for the test phase with six test exemplars presented simultaneously. The way in which to control for explicit strategy is to tell subjects that they would see a stimulus very quickly and would then have to choose which one they thought it was. Due to the difficulties of ensuring presentations really are subliminal when presented via computer, subjects do not see anything but a backward mask on the critical presentation. If the colours on the Channel 4 logo really are held as a mental representation in memory then it might be expected that subjects would misattribute familiarity of this stimulus to the study part of the task. Thus, if implicit learning of a real-life invariant can take place, it is expected that this will be demonstrated by subjects choosing the real Channel 4 logo in the indirect task but not in the direct task.

Method

Subjects

30 undergraduates participated in this experiment in return for a small payment.

Materials

A digital image of the Channel 4 logo was captured from the Channel 4 website. Using graphics software it was resized to 17 x 10 cm and five other versions of the logo were constructed simply by swapping around the various colours which comprised the logo. In addition, similar types of images were constructed for the digits 6 and 8 and these were similar in size to the images of the logo. For the study stimuli, one version of the '6' and one of the '8' was chosen at random and presented in the centre of the screen. For the test stimuli all six versions of each digit were copied onto the same picture file and presented in two columns of three (see Figure 8 for an example). For the test stimulus which was comprised of the Channel 4 logos, two versions were constructed. One had the correct version in the middle right position and the other had the correct version in the bottom left. These two test stimuli were counterbalanced across subjects between the implicit and explicit tests. In addition a backward mask was constructed by using the airbrush function on a graphics software package to colour a square just larger than the logo with the various colours used in the logo. The stimuli were presented on the screen of a Power Macintosh computer. All were displayed with a black background.

Design and Procedure

All subjects were given both implicit and explicit testing conditions. The implicit condition was designed to make subjects believe they were participating in a subliminal perception experiment. They were informed that they would first see a digit flash up on the computer screen very quickly and that this digit would be made up of various colours. Then they would see a square of colour (the backward mask) appearing where the digit was. After this, subjects were told they would be given six versions of the digit with only the positions of the colours differentiating between them. The task for subjects was to choose the digit with the correct pattern of colours. Subjects were given as long as they wanted for the test phase. This was done with the digit 6 on the first trial and the digit 8 on the second trial. Subjects were informed that the stimuli would be presented for shorter durations each time although this was not actually the case as both were presented for 20 msec. On the third trial subjects were told they would see the digit 4 at a very short duration. Subjects did not actually see a stimulus on this trial before the backward mask appeared. They were then presented with six versions of the Channel 4 logo with only one having the colours in the correct position. Subjects were asked to choose which one they thought had been presented. Once that choice was made subjects were then shown the test stimulus with the correct logo in a different position and asked which one was the real Channel 4 logo. This was the explicit version of the task.

Figure 8: Display in forced choice recognition of Channel 4 logo.

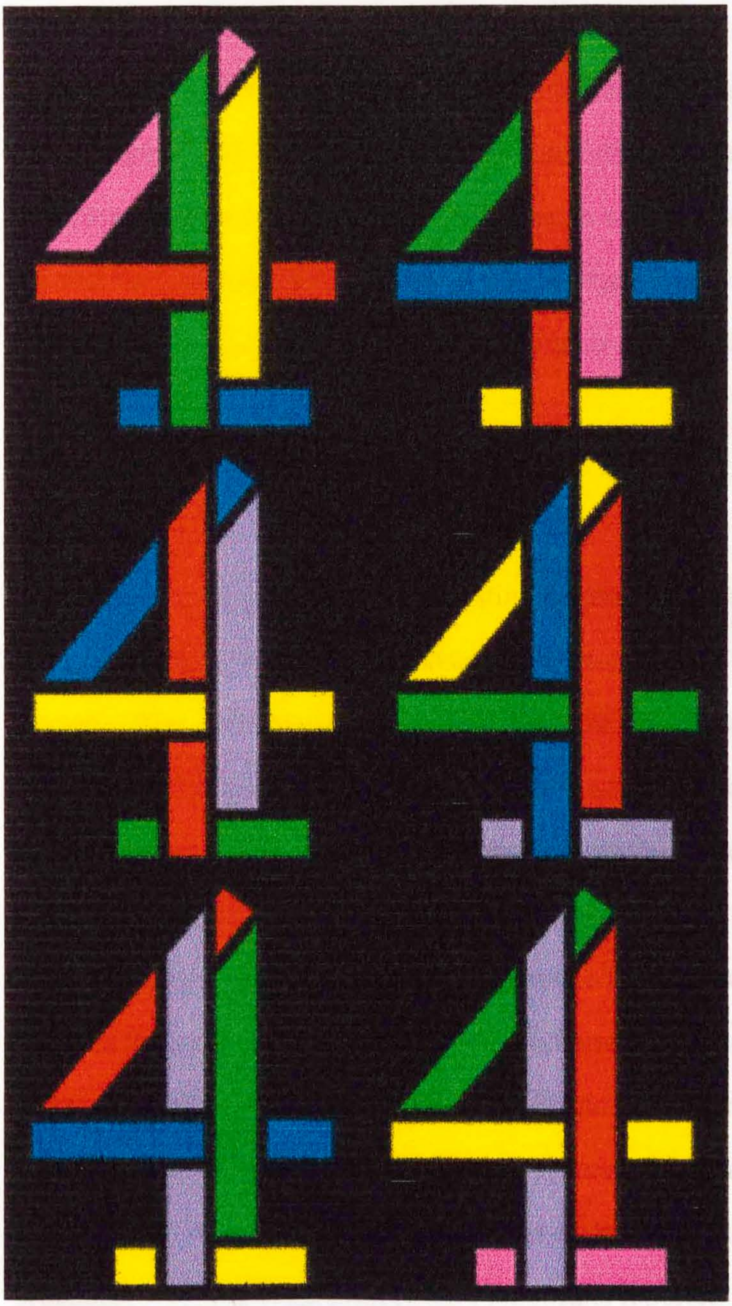

Results and Discussion

Eight out of 30 subjects in the implicit condition chose the real Channel 4 logo. If chance is assumed to be 0.167 then this value is not significantly different from chance according to the binomial test (z approximation, $p > .1$). Out of 30 subjects in the explicit condition however, 19 could correctly identify the real Channel 4 logo. This is significantly different from chance according to the binomial test (z approximation, $p < .05$). Thus the hypotheses were not supported. Subjects in the indirect condition did not misattribute knowledge of the logo to the study phase whereas when asked directly they could identify the real logo.

4.5 General Discussion

Detection of an invariant quality in a stimulus display without any verbal awareness of that knowledge has been demonstrated several times in the literature (Bright and Burton, 1994; McGeorge and Burton, 1990). If such implicit learning is indeed a consequence of a ubiquitous primitive system which can abstract rule-like information from the environment without conscious awareness (Reber, 1993, 1997) then one might expect invariant qualities in the real world to be learned without the learner having conscious access to that knowledge. Morton (1967) could find no evidence for such an effect although the task used to assess any knowledge of the regularity was not particularly sensitive. Bright (1994) reports that subjects could correctly report the direction in which the Queen's head faces on coins and stamps although this information was only evident under a two-alternative forced choice procedure. This is in agreement with the results reported in the previous chapter where implicit knowledge of an invariant was not detectable in a yes / no recognition task but was elicited in a two-alternative forced choice procedure.

The results reported by Bright (1994) could not be replicated in Experiment 11, however. This suggests that if the effect does exist, it is not robust. Another problem with the interpretation of Bright's (1994) results is that there is no way to be sure that the information used by subjects is not episodic in nature. It was hypothesised that if the information about head orientation was held at some level then it should have an effect in the categorisation of whether a coin or stamp is British or foreign. A stronger effect might also be evident if the time deadline procedures used in Experiments 1-6 were included in the design. This hypothesis was not supported and no difference in reaction times to a coin or stamp with the correct orientation and to one with the incorrect orientation could be found. If orientation does not seem to be encoded implicitly then perhaps some other type of invariance might. Bahrick (1958) reports incidental learning of colour information. One source of invariant colour information available to the British public for several years is the logo used by Channel 4 which is broadcast during each interval in the programming schedule.

The final experiment in this chapter examined whether subjects would misattribute knowledge of the colours in the logo to a pseudo tachistoscopic recognition task. In a six-alternative forced choice design, subjects did not choose the Channel 4 logo at above chance levels. When given the same task but asked to choose which stimulus was the real Channel 4 logo subjects performed better than would be expected by chance. One could argue that this demonstrates incidental if not implicit learning as subjects would not have attempted to memorise the colours in the logo before the test. However the fact that subjects are asked a question to which they can give an answer makes it impossible to rule out episodic memory as the mechanism by which subjects responded.

None of the experiments reported in this chapter could find convincing evidence of implicit learning of real world regularities. No knowledge was elicited even when circumstances which increase the likelihood of implicit processing taking place (e.g. incidental test instructions, speeded task instructions) were used. This suggests a discrepancy between the knowledge of regularities that was hypothesised to exist in these experiments and the knowledge of regularities that is elicited by laboratory implicit learning tasks. Of course it may be that laboratory findings are mediated purely by episodic mechanisms (Whittlesea and Dorken, 1993) and subjects do not have a clear memory for episodes where the Queen's head faces to the right on a coin. This view would also explain why subjects could choose the correct Channel 4 logo when asked to do so.

These experiments do not provide any support for an abstractionist position, however. It may be that for abstraction to take place the learning stimuli must have a temporal grouping as is the case for laboratory experiments. In order to qualify what the necessary conditions are for any possible implicit abstraction of information, more laboratory work is needed using spaced presentation of learning items. In addition, Experiment 4 suggests that implicit knowledge of an invariant does not last for intervals as long as one week. This was after one presentation of the study items. It may be that given many presentations of an invariant over time a more durable representation might be formed. If these types of results are not found in subsequent laboratory experiments then Reber's (1993) position supporting an implicit learning system which we are all supposed to have and which can learn complex real life tasks such as language acquisition would seem less likely to be correct. The question would be what sort of real world knowledge could such a system abstract if not the type described above. There cannot be many examples of information in the real world which is presented to an organism in the way that a laboratory study phase presents information and so if one advocates an evolutionary based system, it is crucial that one can find an evolutionary adaptive behaviour which fits the parameters of the proposed system.

Chapter 5

Learning by observation in a sequence learning task: a re-examination

5.1 Introduction

Implicit learning can be defined as the acquisition of knowledge without concomitant verbal awareness of that knowledge. This definition follows from the early implicit learning studies by Reber (1967) in which he gave subjects letter strings which were formed by following the rules of a finite-state grammar. In a surprise test phase, Reber found that subjects could classify new grammatical and ungrammatical strings with above chance performance but could not explain how they had made their decisions and could not articulate any of the rules of the grammar. This finding has been replicated and extended by many other studies into artificial grammar learning (Reber, 1976; Reber and Allen, 1978; Reber and Lewis, 1977) and this dissociation between task performance and verbalisable knowledge found in other types of implicit learning task (e.g. dynamic systems, Berry and Broadbent, 1984; sequence learning, Nissen and Bullemer, 1987; invariance detection, McGeorge and Burton, 1990).

The available evidence based on verbal report, it is argued, does not provide adequate support for the existence of learning without conscious awareness. Shanks and St. John (1994) suggest that verbal reports are likely to violate both of their criteria for assessing learning without awareness: namely the "Information Criterion" and the "Sensitivity Criterion". They suggest that a task used for assessing conscious awareness must be as similar as possible to the task which is claimed to be performed without such conscious awareness. For example, in the artificial grammar paradigm, subjects may be required to complete missing letters from a letter string in order to test the extent of conscious knowledge about string formation. Many tasks reported in the literature do not use such a task when assessing explicit awareness and rely only on verbal report. As these studies which show a dissociation between performance and verbal awareness do not meet either of Shanks and St. John's criteria, they provide only very weak support for a divide between conscious and unconscious learning.

Such a dissociation may provide converging evidence when used in conjunction with other types of dissociation between supposedly conscious and unconscious knowledge. Experiments which examine these other dissociations typically manipulate experimental conditions to favour either explicit or implicit processing, showing different performance by subjects on the same task. For example, Reber (1976) contrasted performance on artificial grammar learning in two groups of subjects, one given incidental learning instructions and one being explicitly told that the strings were rule-based and that they should search for the rules. Subjects in the rule-search condition showed markedly worse performance than those in the condition requiring incidental learning. Hayes and Broadbent (1988) found that a secondary task designed to place additional load on working memory had a detrimental effect on a task thought to use an explicit mode of processing, but no effect at all on a task thought to tap implicit processes. Used in conjunction with tests of explicit awareness, these dissociations provide stronger support for distinct modes of conscious and unconscious processing.

One of the more peculiar dissociations in the literature on the explicit/implicit distinction concerns the role of the subject in interacting with the task. Berry (1991) used a task which involved the subject controlling a computer simulation of a complex scenario. Examples of such scenarios used by Berry include controlling a sugar production factory, a city transport system and an interaction with a 'computer person'. Subjects could enter a range of values into the system and were required to keep the system's output within certain parameters. The relationship between the subject's input and the system's output was determined by a complex equation. One version of each task used a salient relationship between input and output variables and another version used a non-salient relationship. It has been argued from previous work (Berry and Broadbent, 1984) that subjects use explicit processing for the salient task to achieve optimal performance but such performance is reached on the non-salient task only via implicit learning. Berry (1991) found that simply observing someone interact with the salient task resulted in the observer learning the relationship, whereas observing the non-salient task had no facilitative effect on subsequent control of the task. This suggests that asking subjects simply to act as observers removes some component which is necessary for implicit learning but not for explicit learning. Requiring subjects to observe the task being performed removes the need for the subject to make decisions about the system, and physically to interact with the system. The fact that subjects can learn the salient but not the non-salient version of the task by observation suggests the involvement of a decision-making component or a physical interaction component in implicit learning. As Berry demonstrates, neither of these on their own are sufficient. Subjects making decisions but not actually physically interacting with the system do not learn, and nor do subjects typing in someone else's decisions. It seems that both action and observation are essential for learning a non-salient relationship. This result replicates and extends results of a study reported by Funke and Muller (1988) in which subjects were either controlling or observing a complex system. They found that observers were worse at controlling the system but were better at explaining how the system worked. It is unclear whether this phenomenon reported by Berry (1991) is specific to some aspect of controlling a complex system with an underlying non-salient relationship between variables or whether it is a general feature of implicit learning.

The role of action and observation in implicit learning has been investigated in another task which has some similarities to a complex systems task. A sequence learning task such as that used by Nissen and Bullemer (1987) requires that subjects observe a stimulus, which is displayed according to a predetermined sequence, and respond to it. Typically, this task is portrayed to subjects as a reaction time task with multiple responses. No mention is made of the fact that items appear in a fixed sequence, and hence any learning shown by faster reaction times is deemed implicit. The sequence can either be conceptual (the subjects' responses will vary depending on the stimulus presented) or spatial (the subjects' responses correspond to where the stimulus appears). This task has some parallels with the complex system task in that subjects must physically respond to the display and they must also make some decision about which key to press, although this is a much simpler choice than that required in the complex systems task. If the result reported by Berry (1991) is a general phenomenon of implicit learning then it should be impossible to learn a sequential pattern by observation alone.

Howard, Mutter and Howard (1992) report a study in which subjects were presented with four boxes at the bottom of a computer screen and four keys which corresponded to those boxes. An asterisk appeared in one of the boxes and the subject was required to press the corresponding key as quickly as possible. The asterisk appeared according to either a short 10 item or a long 16 item sequence although the subjects were simply told that the position of the asterisk was randomly determined on each trial. The usual finding in these type of tasks is that the RTs to the asterisk show a steady decrease per block of trials and then exhibit a marked increase once subjects encounter a block where the asterisk appears according to a random sequence.

Subjects who were required to respond to just the first 10% of trials in each block and then simply observe the remaining trials showed just as much learning as subjects responding to 100% of the trials in each block. In a second experiment, Howard et al removed the requirement for subjects to respond to the first 10% of trials and had them simply observing the sequence until the random block and found the same results. However, the learning by observation effect is only apparent when RT is collapsed across the groups with the long and short sequence. The data are not conclusive in determining whether learning by observation occurs within each group. Table 6 reproduces the data reported in Howard et al (1992) for both these experiments. The numbers represent the increase in msec for subjects' RTs when they move from the sequenced blocks to the random block, and hence gives an index of learning.

It is clear from Table 6 that while learning is observed for the short pattern in both observation and response conditions, there appears to be no learning for those given the long pattern. Neither the 18 msec increase for the response group nor the 61 msec increase for the observation group is significantly different from chance on a one-sample t-test ($t(7) = 1.22$, $p > 0.1$; $t(7) = 0.97$, $p > 0.1$, respectively). For the 2nd experiment in which significant learning was found, only the short sequence was used.

Table 6: Data reported in Howard et al (1992)

	Short Pattern		Long Pattern		
Experiment	MRT	SD	MRT	SD	M
1.					
Observation	88	45	61	132	75
Response	128	100	18	49	73
2.					
Observation	170	108	----		
Response	108	69	----		

Note :Dash indicates that the long pattern was not used in Howard et al's 2nd experiment.

Howard et al argue that due to particular aspects of their methodology, the data reported for the long sequence is an underestimate of the amount of learning shown by subjects. There are however, a number of other reasons why one might expect learning to be found for the short sequence but not for the long sequence. Firstly, Shanks and St. John (1994) would argue that the short sequence did not meet their requirements to demonstrate implicit learning as the four base items did not appear with equivalent frequency. Subjects may simply have been able to predict which items would most probably occur which in turn would have allowed fast responses to develop. Secondly, both sequences used by Howard et al contained reversals (i.e. if the four spatial positions in which the stimulus could appear were identified as A, B, C and D then an example of a reversal would be ABA) which Reed and Johnson (1994) suggest may be a particularly salient feature of a sequence. As Berry (1991) found that salient rules could be learned by observation, this limits the Howard et al study in making claims about implicit learning in general. Finally, there is evidence from a prediction task given by Howard et al that subjects in the observation condition had a high degree of explicit knowledge about the short sequence but not for the long sequence. It would not be surprising to find explicit learning by observation alone. The experiments in this chapter are designed to rectify these possible confounds in the Howard et al study and test whether learning without awareness through simply observing a stimulus structure is possible.

5.2 Experiment 14: Observational learning of a non-salient sequence

Experiment 14 applies Berry's (1991) observational method to a sequence learning task. One subject is required to respond to a sequence of asterisks appearing on a computer screen. The sequence is carefully constructed to contain equal presentations of base items and no reversals. Another subject is seated next to the first and is told to observe the display as they would be asked to perform the task later on. Both subjects are then required to respond to the sequence using a keypress. Both Howard et al and Berry would predict that subjects in the response condition should show learning. The results of Berry (1991) however would suggest that an observer would not learn the sequence.

Method

Subjects

12 pairs of subjects were recruited from the student population of Glasgow University. All were naive to the task and all were paid a small amount for their participation.

Materials

The stimuli were presented to the subjects on a Power Macintosh via the "Superlab" experimental package. Subjects used the numerical keypad to make responses.

Design and Procedure

The sequence was presented using an asterisk appearing on the computer screen. The sequence in which the asterisk appeared was DBCDACBADCBA where each letter represents one of the screen quadrants. This sequence was repeated without pause for 30 cycles per learning block. Each asterisk appeared for 2000 msec or until a response was made and the sequence used an RSI of 200 msec. Subjects were tested in pairs with one of each pair assigned to the action condition and one to the observation condition. The subject in the action condition was seated 50 cm in front of the computer screen and was told that they would see an asterisk appearing in one of the four screen quadrants. They were told that this asterisk would appear randomly and that they had to press a key corresponding to the location of the asterisk as quickly as they could. The response keys were 1, 3, 7 and 9 on the numerical keypad. Subjects were instructed to use the middle and index fingers of each hand for responding with each finger touching the appropriate key for the entire duration of the experiment. They were told that they would receive 10 blocks of trials and that each block would contain 360 individual presentations of the asterisk but there would be an opportunity to rest between blocks. Subjects in the observation condition were seated to the right of the action subject and were told that they had to watch what was happening carefully as later on in the experiment they would be performing the task. No mention was made of the sequence to either subject. Instead, subjects were told the asterisk would appear at random. Subjects were told that the aim of the experiment was to find out if watching someone perform a reaction time task would have a beneficial effect on actually doing the task. After subjects had completed 9 blocks, a message warned that the final block was next. Both subjects were now required to respond manually to the sequence. Half of the subjects in the action condition performed the test phase on the same computer with subjects in the observation condition being moved to another room to complete the task, and for the remaining half the converse was true.

The subject in the action condition was told to carry on as before and complete the final block. The sequence for the final block changed to DBACBDCABDAC but the subject was not informed of this. The dependent measure was the change in RT from the ninth block to the final block with the changed sequence. The subject in the observation condition was instructed that they would now have to respond to the display using the keyboard. They were informed that they would only have to complete three blocks of trials with time to rest between blocks. The first block (which used the same sequence as had been observed) was given simply to familiarise the observation subjects with the correct keys and to allow for any increase in simple motor speed. The second block provided the baseline RT and corresponded to the ninth block given to the action subject. The third block presented the asterisk in the new sequence and as in the action condition, the index of learning was the change in RT from the ninth block to the final block with the changed sequence. The subjects in the observation condition thus received two blocks of responding via keypress to the sequence whereas the subjects in the action condition received nine such blocks.

Finally, a measure of explicit awareness of the sequence was taken for both action and observation subjects. On completion of the final block, both subjects were informed of the sequential nature of the stimulus display. They were asked to think back to the learning blocks and to use any pattern knowledge they might have gained to predict the position of the asterisk through one cycle of the sequence. They were not given any information about the sequence but were told that it had changed during the final block and to ignore this block when performing the prediction task. Subjects were instructed to use the same fingers and same keys as before but to concentrate on accuracy rather than speed. Whenever subjects made a keypress, the asterisk appeared in the position which they should have chosen. Subjects were instructed to pay attention to this feedback and to orient the knowledge they had of the sequence to the correct version being displayed. Thus if subjects had full knowledge of the sequence but were unsure where the starting point was, they should still be able to achieve a high score on the prediction task by responding to the feedback. The score for the prediction task was simply the number of correct responses made, from 0 to 12. High scores indicate complete knowledge of the sequence and low scores indicate a lack of any sequential knowledge.

Results and Discussion

The RT measure is the median value of the 360 recorded RTs excluding any errors, which is then averaged across all 12 subjects. The RTs for the first nine blocks in the learning phase for subjects in the action condition were calculated and show a decrease in RT across all nine blocks. These data are displayed as a graph in Figure 9. Errors were calculated for the learning blocks and any pair of subjects scoring above 20% errors in any one learning block were excluded from the study and replaced. One subject pair's data was replaced for this reason.

Figure 9: Mean RTs for each block of trials in the action and observation conditions.

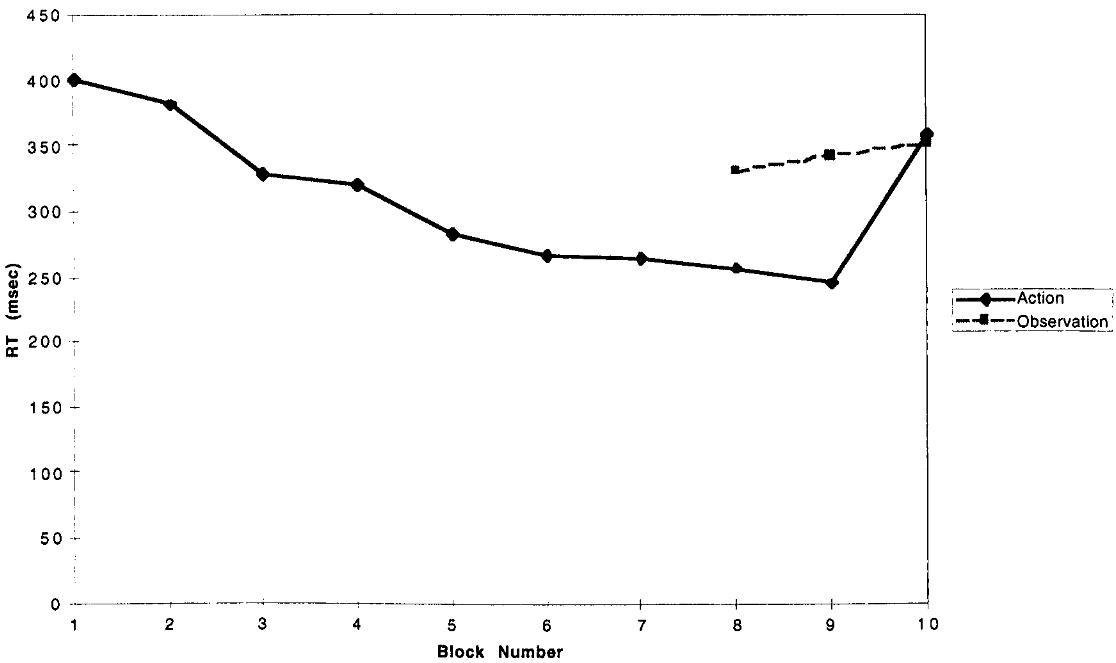

Note: The new sequence was introduced to subjects in Block 10. Subjects in the observation condition had observed the sequence for the first nine blocks before performing the keypress task so the observation condition plot corresponds to that of the action subjects in terms of scores for baseline RT from the learning phase and RT to the new sequence.

There was no effect of performing the test on the same computer used in the learning phase or on a different computer ($F(1,20) < 1$) so the results were collapsed across this condition. The dependent measure was the mean RT for the final block (which used the new sequence) minus the mean RT from the preceding block (which used the sequence subjects had encountered in the learning phase). Any decrement in RT across these two blocks can only be due to subjects holding some information about the sequence encountered in the learning phase. The mean change in RT between these two blocks is presented in Table 7 for both experimental conditions. A positive value indicates a decrement in RT when responding to the new sequence and hence suggests some knowledge of the old sequence. There was no significant difference between the errors made between the two blocks for either condition and the overall error rate was low (approx. 5% for all blocks and conditions).

Table 7: Mean RT (msecs) and Standard Deviations for Action and Observation Conditions.

Task	RT to old sequence.	RT to new sequence.	Difference
Action	246 (112)	359 (127)	113
Observation	342 (108)	352 (59)	10

A mixed design two-way ANOVA on task (action or observation) x learning (same or different sequence) reveals no main effect of task ($F(1,22) = 1.24$, $p > 0.1$) but a significant main effect of learning ($F(1,22) = 14.67$, $p < 0.05$) and a significant interaction ($F(1,22) = 10.52$, $p < 0.05$). Simple main effects reveal that subjects in the action group show evidence of learning ($F(1,22) = 25.2$, $p < 0.05$) but subjects in the observation group show no such learning ($F(1,22) < 1$) as indexed by a decrement in RT.

The mean predictability score (from 0 to 12) as measured by the prediction task for subjects in the action condition was 5.9, std dev = 3.0, and for the observation condition was 6.8, std dev = 3.1. These two levels of awareness are not statistically different from one another ($F(1,22) < 1$) and will not be considered further.

This result very clearly indicates that while subjects performing a sequence learning task can learn something about the sequence, subjects merely observing someone else responding to the sequence show no evidence of learning. This is in agreement with results reported in a very different implicit learning paradigm where subjects interacting with a complex system show evidence of implicit learning but subjects merely observing the display do not (Berry, 1991). However, this result contrasts with that reported by Howard et al (1992), where equivalent learning was found for subjects who simply observed the sequence as for those who made RT responses. It is possible that this discrepancy can be explained by differences in the salience of the sequence presented to subjects. The sequence used by Howard et al had the 4 base items presented an unequal number of times in one cycle of the sequence, which may have allowed subjects to respond more quickly to the sequence on that basis alone. It could have been this responding to probability that can be learned by observation alone. Howard et al's sequences also contained reversals which may have had the effect of increasing the salience of the sequence to subjects. As shown by Berry (1991), learning by observation is possible where the underlying rule is salient. The sequence used in this experiment corrected both of these possible problems. In addition, the sequence used in the changed sequence block was a different 12 item sequence. Howard et al used a random sequence, which although widely used in this area can only maximise the possibility of differences being found when the sequence is changed. As Stadler (1995) points out, RT will decrease faster in a repeating sequence condition than in a random condition because in the former there are fewer unique episodes and each one is practiced more often. The two sequences used in this experiment also had the same items on the first two trials of the sequence to ensure that any decrease in RTs was not simply due to a discrepancy between the positions that were likely to be the best learned. Thus the sequence in the test phase was constructed so that any change in RT was not likely to be due to any differences between an ordered sequence and a random one.

Although the observers did not show any evidence for learning the sequence, it is possible that learning did take place but that such learning is modality specific.

Jordan (1995) has developed a connectionist model of sequencing which embeds the representation of the sequence within the effector system carrying out the movement. Thus learning may be taking place but such learning cannot transfer to a different modality. Keele, Jennings, Jones, Caulton and Cohen (1995) provide evidence that this is not the case and support an idea of modularity where one module is responsible for learning the locations and encodes them in an effector-free spatial description, whereas another module is responsible for computing articulatory activities to carry out the movement. Keele et al gave subjects a sequence learning task and changed the effector system which carried out the movement. In changing from using three fingers on three keys to using one finger on three keys (effectively changing the effector from finger muscles to arm muscles) they found total transfer of knowledge. In a follow-up experiment, they found significant transfer of knowledge from a key-press to a vocal response. Although not all the sequential learning transferred between these two modalities, it does suggest that the acquired knowledge is not modality specific.

5.3 Experiment 15: Cross modality transfer of knowledge

Although Keele et al demonstrate that knowledge can be transfered across modality, their results do not totally rule out the possibility that subjects in the previous experiment were learning information that could not transfer from one modality to another as Keele et al used a keypress transferring to a vocal response. If subjects in the previous experiment were learning information about the sequence by observation alone, the transfer would be to a keypress. In order to rule out an explanation of modality specificity, Experiment 15 trained subjects using vocal responses to the sequence, and tested whether any knowledge of the sequence could transfer to a response via keypress.

This experiment may also provide a clearer understanding of the components necessary for learning a non-salient relationship. Berry (1991) concludes that 'action' in conjunction with observation is necessary for learning in a non-salient complex systems task. Even when making decisions about the system, subjects could not learn the relationship in the absence of 'action'. What this seems to suggest is that some sort of physical interaction with the rule-governed system is necessary for learning to take place. With respect to sequence learning, the results of Experiment 14 might also suggest that physical interaction is the necessary component for learning. It is unlikely that this could simply be due to learning taking place within the motor system alone, as Keele et al (1995) have reported transfer of learning between modalities. In using a vocal response to the learning blocks, Experiment 15 will also examine what could be meant by the term 'action'. If subjects exhibit no learning of the sequence then it may be that a vocal response does not constitute 'action' in terms of what is necessary for implicit learning. If subjects can learn the sequence then it suggests that a wider concept of action may be necessary than is suggested by Berry's (1991) results for complex systems tasks.

Method

Subjects

12 undergraduates were recruited from the student population at the University of Glasgow. Subjects were paid a small amount for their participation and all were naive to the experiment.

Materials

The materials were the same as in Experiment 14.

Design and Procedure

The general procedure used in the action condition for Experiment 1 was used here. The only difference was that subjects were tested individually and were required to respond vocally to the position of the asterisk (A, B, C or D) during the learning blocks rather than with a keypress. The asterisk remained displayed on the screen for 2000 msec or until the voice key was triggered but the RSI was increased to 600 msec during the vocal phase. This was because saying a letter out loud could take longer than 200 msec for some subjects and so would inadvertently trigger a response to the following item. Following 9 blocks of 30 presentations of the sequence, the subject was informed that they would now have their RTs measured via a keypress. They were given one block which had the same sequence as the vocal blocks and then a second block which obeyed the new sequence. A practice block to familiarise the subjects to the keys was not given as from Experiment 1 the increase in RT from the practice block to the first test block was found to be negligible. On the keypress part of the experiment the RSI was kept at 200 msec. No mention was made of the stimuli following a sequence. As in Experiment 14 subjects were told the asterisks would appear in a random order.

Results and Discussion

The data for the learning indicated that all subjects showed a decrease in RT across blocks and no subjects made more than 20% errors on any learning block. The test phase required the subject to respond using the keyboard rather than the voice key. Figure 10 shows the mean RT across all blocks.

Figure 10: Mean RT for vocal response and keypress blocks.

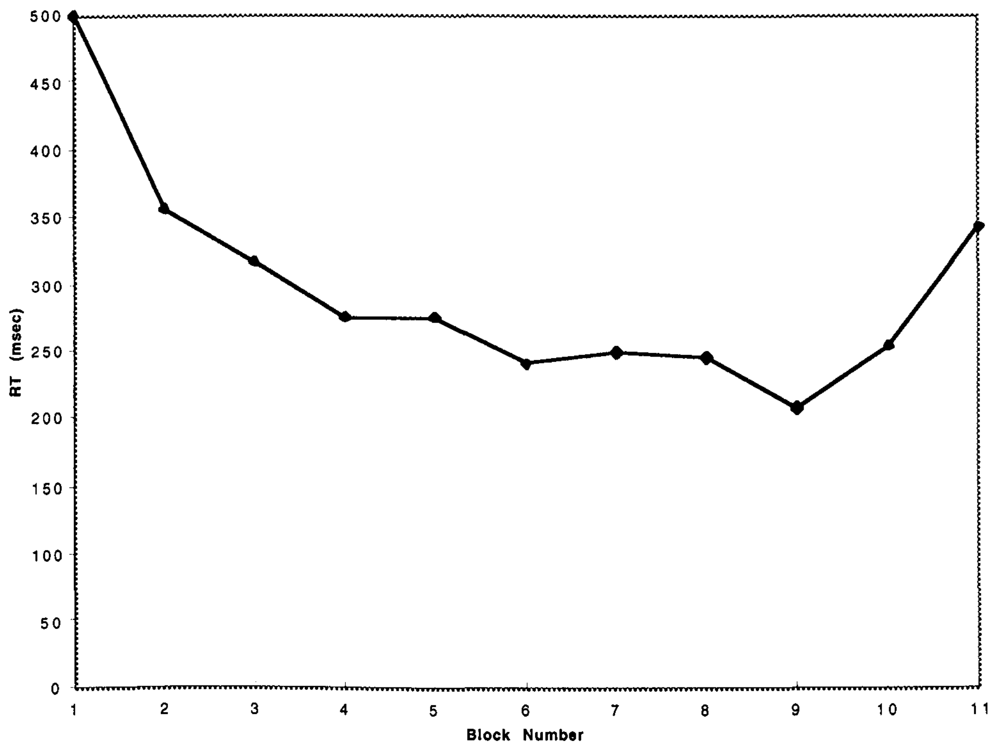

Note: Blocks 1 - 9 show the mean vocal RT for the learning phase. Block 10 had subjects switch to using a keypress response but still to the same sequence as in Blocks 1 - 9. Block 11 had subjects respond via keypress to a new sequence.

Table 8 shows the mean RTs for the block where the subject responds to the old sequence using the keys, the block where the sequence changes and the difference between these blocks.

Table 8: Mean RTs (msecs), Errors and Standard Deviations to Old and New Sequences.

	Old sequence	New sequence	Difference
Mean RT	254 (110)	344 (65)	90
Errors	11.7 (8.3)	15.1 (6.3)	3.4

A one-way ANOVA between the two blocks shows there is a significant difference between the RTs for subjects responding to a sequence that they have learned vocally and to a different sequence ($F(1,11) = 11.92$, $p < 0.05$). This result replicates and extends that of Keele et al (1995) in finding that knowledge of a sequence can transfer across modality. This result also means that the finding of no learning from observation of the sequence in Experiment 14 cannot be due to such learning being modality specific. The mean awareness score for the 12 subjects was 6.75 with a std dev of 4.0. This is almost identical to the awareness score for the subjects in the observation group in Experiment 14 and suggests that learning in Experiment 15 cannot be due to an increase in explicit knowledge of the sequence over that observed in Experiment 14. There was a small but reliable difference between errors in the two test blocks ($F(1,11) = 2.34$, $p < 0.05$). This difference is consistent with the effect in RT and provides further evidence of learning.

The finding that subjects can learn the sequence using a vocal response suggests that if action is indeed necessary for learning in these non-salient tasks, the definition of action must be extended to include various overt physical responses to the stimuli. Berry (1991) found that subjects must physically respond via keypress to the complex system task in order to learn about the relationship. Experiment 14 in this thesis found that subjects making overt physical responses to the stimuli could learn the sequence but in Experiment 14, subjects instructed simply to watch exhibited no such learning. It seems parallels can be drawn between these two types of task and suggests that their learning might share some common process.

However, there is one other possibility that might account for the difference in learning between Experiments 14 and 15. In Experiment 15, subjects were responding to the sequence but in Experiment 14, subjects were watching someone else respond to the sequence. This latter condition might have led to subjects watching a sequence which was occurring too fast for them to learn or of a sequence which seemed stilted or disjointed. Stadler (1993) demonstrated that insertion of random pauses in a sequence severely reduced the amount of learning shown by subjects and argues that the organisational component of an implicit sequence learning task is vital for any such learning. It may be that subjects in the observation condition of Experiment 14 were watching a sequence which seemed to them to have pauses or items occurring too quickly and this disruption of organisation is the factor which accounts for no learning, rather than lack of physical overt responding.

5.4 Experiment 16: Are organisational factors responsible for disrupted learning

Experiment 16 tests whether failure of learning in Experiment 14 was due to observationally disrupted presentation of the sequence, or the fact that no overt responses were made to the sequence. The methodology of this experiment returns to the two subject design used in Experiment 14. One subject (the observer) views the screen on which the sequence appears. The second subject (the actor) controls the keyboard, but cannot see the screen. In this arrangement, the observer calls out appropriate responses to the actor, who then implements them (i.e. presses the appropriate key). At test, both subjects make keypress responses to the visual sequence.

The observer in this experiment is, to some extent, watching a sequence which is controlled by another person, and hence may be disrupted. If this is the explanation for absence of learning in Experiment 14, one might expect no learning by observers. However, if learning is mediated by the need to make any response during the first phase, as suggested by Experiment 15, then one might expect the observers to show evidence of learning at test. The actor, on the other hand, is implementing the responses to the sequence. Learning by the actor could arise either as a result of acquiring the appropriate physical responses, or as a result of learning an abstract representation of the sequence which is unchanged by different presentation types (i.e. from vocal to spatial presentation).

Method

Subjects

25 pairs of subjects were recruited from the student population of Glasgow University. All were naive to the task. Subjects were paid a small amount for their participation.

Materials

The materials were the same as Experiment 14.

Design and Procedure

The general procedure was the same as in Experiment 14 except for the following changes. The subjects in the observation and action conditions were seated at right angles to one another. The subject in the observation condition could see the screen but not the keyboard and was required to call out the position of the asterisk to the other subject as fast as possible. The subject in the action condition could not see the display and had to rely on the instructions called out by their partner to make the appropriate key-press on the keyboard. The action subject was told this had to be done as fast as possible as their RTs were being measured. Subjects were told that this was an experiment on correct and rapid communication of a random stimulus display and no mention was made of the sequence. The asterisk would appear for 2000 msec or until the action subject pressed a key. There was an RSI of 600 msec between asterisk presentations during the learning phase and 200 msec during the test phase. The learning phase was reduced from 9 blocks to 7 blocks as this manipulation took far longer than that in Experiments 14 or 15 and there was not much gain in speed of response for blocks 8 and 9 in the first two experiments. The test phase occurred immediately after the learning phase and had both subjects responding via a keyboard to the display on a computer screen. The subjects were in different rooms during the test phase. After that, they were given the standard prediction task (as in Experiments 14 and 15) in an attempt to measure their explicit awareness of the sequence.

Results and Discussion

All pairs of subjects exhibited a decrease in RT over the learning blocks and no subjects were excluded from the analysis due to errors. The means for the learning blocks will not be shown as they are determined by the response of both subjects and so are uninformative as to which subject is learning. The test phase had subjects from both conditions respond to the sequence using the keyboard. Again, the dependent measure was the difference in RTs between the two blocks, the first of which retained the sequence from the training phase and the second presenting a different sequence. Table 9 shows the mean RTs for these blocks and the mean difference in RT for both the action and observe conditions.

Table 9: Mean RT (msecs) and Standard Deviations for Action and Observation Conditions.

Task	RT to old sequence.	RT to new sequence.	Difference
Action	293 (91)	315 (58)	22
Observation	268 (83)	302 (43)	34

A two-way ANOVA between task (action vs. observation) and learning (same or different sequence) reveals no main effect of task ($F(1,48) = 1.15$, $p > 0.1$), a significant main effect of learning ($F(1,48) = 9.15$, $p < 0.05$) and no interaction ($F(1,48) < 1$). Hence, subjects show learning of the sequence irrespective of which task condition they were in.

If lack of learning for subjects observing the sequence in Experiment 1 was due to the disruption of organisational principles, then one might have expected subjects observing the sequence in this experiment to show no learning. This was not the case, although the RT decrement when moving to a new sequence was much less than that observed in Experiment 15, suggesting that organisational principles may play a role in implicit acquisition of the sequence. The difference between observation conditions in this experiment and those in Experiment 14 is that subjects in this experiment are making overt physical responses to the stimulus presentation.

As was suggested by Berry (1991), it is this difference that is crucial if learning is to be found.

Subjects in the action condition of Experiment 16 could not see the screen on which the sequence was being displayed but had to respond to the location of the asterisk via vocal instructions from their partner. These subjects displayed small but significant amounts of learning during a test phase where they were required to respond to the asterisk appearing on a computer screen. This result would not have been predicted by the results of Berry (1991) and suggests that implicit learning of a sequence is possible by action alone. It could be argued however, that just as the definition of 'action' must be extended to incorporate overt verbal responses, so must the definition of observation be extended to include auditory presentations. Nevertheless, this finding illustrates that learning of a conceptual sequence (hearing A, B, C or D called out) can transfer to responding to the same sequence but given spatially.

These results, however, rest on the assumption that any learning is occurring without conscious awareness. The findings above would not be at all surprising if subjects were consciously aware of the sequence. The results of the prediction task show that subjects in the observation group displayed significantly higher levels of conscious awareness of the sequence than subjects in the action group. Subjects in the observation group had a mean awareness score of 7.48 (std dev = 3.95) whereas subjects in the action group had a score of 5.12 (std dev = 3.32) and a one-way ANOVA shows these scores to be different ($F(1,48) = 5.23$, $p < 0.05$). This finding mirrors that of Howard et al (1992) where the prediction scores for observing subjects were initially high and did not rise over cycles of the prediction task. Thus it may be that the effects described above are mediated by conscious knowledge of the sequence and do not apply to learning of a non-salient sequence.

There is a suggestion that learning can occur independently of awareness as Willingham, Nissen and Bullemer (1989) report that although some subjects do become aware of the sequence, both aware and unaware subjects show a pattern of RT decrease. In order to test whether awareness could be a confounding variable in Experiment 16, the RTs for both groups were correlated with their prediction scores. Subjects who were required to respond vocally to the sequence during the learning phase (the observation group) showed a weak but non-significant relationship between implicit knowledge gained as measured by the RT task and explicit knowledge shown in the prediction task ($r = 0.36$, $p = .08$). The correlation between RT difference and awareness for the action group however, shows a strong relationship between the two ($r = 0.51$, $p < .05$). The greater the score on the prediction task then the greater the RT decrement in switching to a new sequence.

Examination of scatterplots (see Figures 11 and 12) of the data suggests that these correlations may be affected by outliers in the distribution and so nonparametric correlations were also calculated. These revealed a significant relationship between awareness and RT decrement for both action and observation subjects. (Spearman's r for action and observation groups was 0.43 and 0.42 respectively, both of which are significant at below an alpha level of .05)

Figure 11: Scatterplot of Awareness and RT difference for observation group.

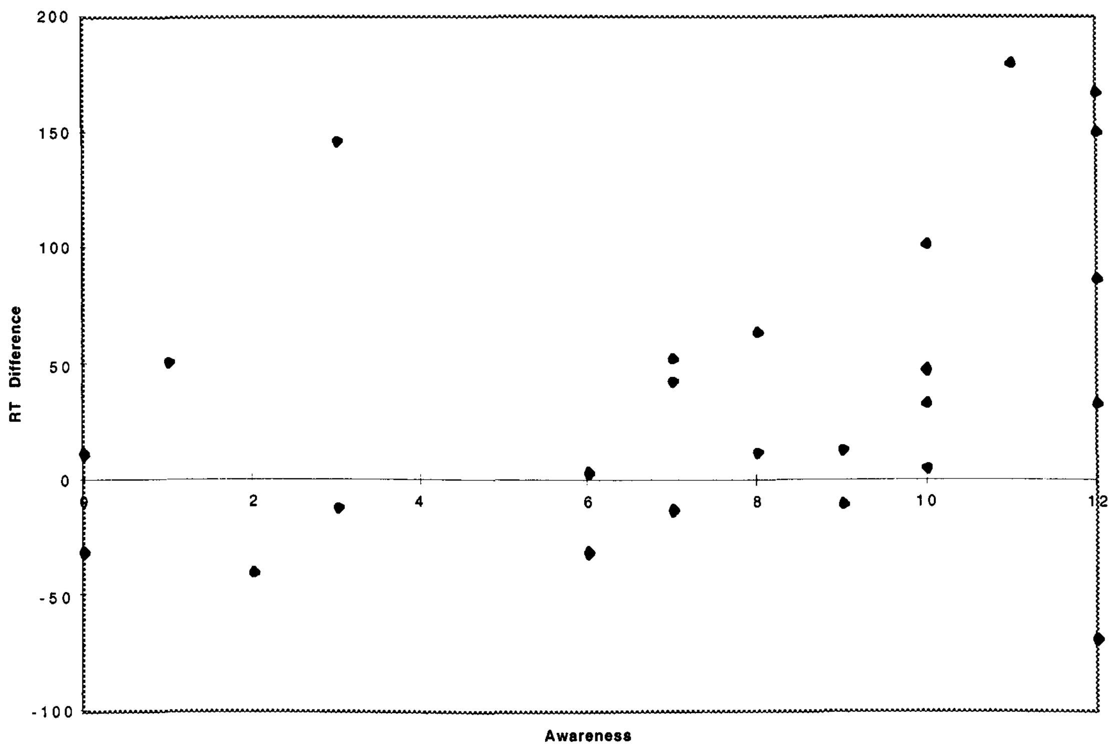

Figure 12: Scatterplot of Awareness and RT difference for action group.

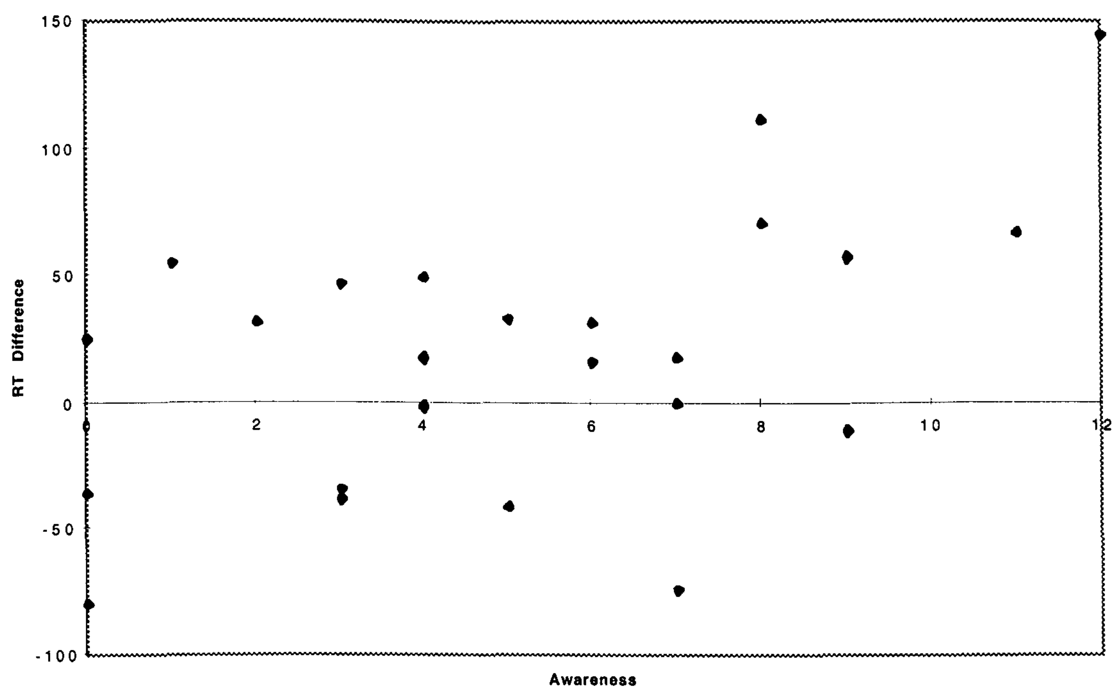

While learning is shown for both action and observation groups it is unclear whether this learning is mediated by explicit or implicit processes. The sample size is too small to allow reliable post-hoc tests on a subset of the data and the subjects do not display levels of awareness which are normally distributed making for additional problems in interpreting these correlations. Clearly this is a matter for concern as both these data and that of Howard et al (1992) have shown a possible relationship between awareness of the sequence and learning by observation. This problem suggests that while we may conclude that learning was shown for both action and observation subjects, these results must be treated with caution.

5.5 Experiment 17: The effect of rule search instructions

Experiments 14, 15 and 16 demonstrated learning of a non-salient sequence when subjects made overt responses during the learning phase. Experiment 14 found no learning when subjects were asked simply to observe the sequence. Experiment 15 shows that this lack of learning is not due to modality specificity and Experiment 16 shows that it is not due to an observed sequence seeming discontinuous, which Stadler (1993) has shown to be able to disrupt learning. The result found by Howard et al (1992) of learning by observation might be explained by the sequence in their experiments being salient to subjects. There are however other differences between the two experiments which might account for the discrepancy and these must first be discounted before a strong claim can be made about the existence of implicit learning simply by observation.

Both Howard et al (1992) and Experiment 14 used similar instructions to subjects. In Experiment 14 subjects were told that they were to watch what was happening carefully as they would do the task later on with the aim of the experiment being to discover whether watching another perform on the task would have a beneficial effect. Howard et al told subjects to "watch the asterisks carefully". It is difficult to know how subjects will interpret experimental instructions and even slight differences in wording may lead to different strategies or goals for groups of subjects. There is a possibility that in Experiment 14 subjects did not believe the cover story and would invent their own hypotheses as to what the experimenter really required of them. It is known from a number of implicit learning paradigms that subjects who actively try to discover underlying rules do not show as much learning as those subjects who are passive towards the task (e.g. Berry and Broadbent, 1988; Brooks, 1978; Reber, 1976; Reber, Kassin, Lewis and Cantor, 1980). Reber (1993) suggests that the decrement in performance is due to subjects looking for rules which are simply too difficult to find and so subjects have to use non-representative rules. There are a number of studies which report facilitation of performance when subjects are instructed to specifically look for the underlying rules but this only seems to occur when the rules are salient (e.g. Berry and Broadbent, 1988; Howard and Ballas, 1980, Reber et al, 1980). With regards to these findings, it is possible that if subjects suspected the items would appear according to a sequence in Experiment 14 and if they actively tried to discover that sequence, then this might have a detrimental effect in actually learning the non-salient sequence. Such a possibility would not affect (or might even facilitate) the learning performance of the subjects reported in Howard et al (1992) as it is argued the sequence that they used was in several possible ways more salient than the one used in this set of experiments.

Experiment 17 examines whether subjects given rule-discovery instructions show a decrement in performance as compared to subjects given incidental instructions. This effect has not been shown in the sequence learning paradigm and would be powerful evidence against the conclusion that "sequence learning in the absence of awareness has not been shown" (Shanks and St. John, 1994). In addition, if subjects do not seem to learn when given rule discovery instructions, then it is possible that subjects were using a similar rule-discovery strategy in the observation condition of Experiment 14. This might account for the lack of learning in that experiment as opposed to the hypothesis that such learning is impossible by observation alone. Experiment 17 follows the same procedure as Experiment 15 with subjects giving vocal responses for the learning phase and responding via the keyboard for the test phase. This keeps the learning phase similar to that of a learning by observation alone condition and ensures that any peripheral motor learning cannot account for the results.

Method

Subjects

20 subjects were recruited from the undergraduate population of the University of Glasgow. All were naive to the experimental aims and all received a small payment in return for their participation.

Materials

The materials were the same as those used in Experiment 15.

Design and Procedure

Subjects were randomly assigned to one of two groups. One group received the same instructions as the subjects in Experiment 15; namely that the asterisk would appear at a random location and the experiment was to measure their reaction time. This group will be called the incidental learning group. The other group (called the rule-discovery group) were informed that the asterisk would appear in a particular spatial location according to a sequence and that this sequence would be repeated many times in the learning phase. One feature of this experiment which differed from the previous ones is that subjects were only exposed to one learning block which contained 30 cycles of the sequence. In the previous experiments where subjects were given nine learning blocks, many subjects reported boredom and lack of attention at some point. Nine experimental blocks may have caused subjects to cease searching for rules and so act as incidental learners.

All subjects received the same learning and test blocks and the RSI was the same as that used in Experiment 15. The two test blocks were identical to those used in the experiments already reported. The only difference between the two groups was in the instructions they received at study. Vocal responses were required during the learning phase and RTs were measured via keypress during the test phase (as in Experiment 15). After the test blocks were completed, all subjects were given the prediction task and subjects in the incidental learning condition were made aware of the structured nature of the display prior to completing the prediction task.

Results and Discussion

It can be clearly seen from Table 10 that subjects displayed some learning when given incidental learning instructions but no such learning when told that there was a sequence which they could find and exploit.

Table 10: Mean RTs and Std Devs for the Incidental and Rule Discovery Conditions

Condition	RT to old sequence.	RT to new sequence	Difference
Incidental	301 (73)	361 (69)	60
Rule Discovery	340 (73)	334 (66)	-6

A mixed design two-way ANOVA between condition (incidental vs. rule discovery) and learning (old or new sequence) reveals no main effect of condition ($F(1,18) < 1$), a trend for main effect of learning ($F(1,18) = 4.0$, $p = .06$) and a significant interaction ($F(1,18) = 5.86$, $p < .05$). Simple main effects show that subjects in the incidental learning condition exhibit a significant decrease in RT when responding to the new sequence ($F(1,18) = 9.80$, $p < .05$) but subjects in the rule discovery condition show no such evidence of learning ($F(1,18) < 1$). Awareness as measured by the prediction task was not different between the two groups of subjects ($F(1,18) = 1.60$, $p > .1$)with those in the rule discovery condition having a mean of 7.6 (std dev = 2.95) and those in the incidental condition having a mean of 6.2 (std dev = 1.87).

This experiment provides a replication of the effect found in Experiment 15, namely that knowledge of the sequence transfers between response modality. The effect was smaller in this experiment (60 msec as opposed to 90 msec in Experiment 15) however subjects in this experiment received only one learning block whereas those in Experiment 15 received nine.

In line with other implicit learning paradigms, Experiment 17 demonstrates that subjects in a sequence learning task who were informed of the rule governed nature of the presentation did not show any learning. Firstly, this argues against the possibility that performance in this task is due to subjects actively searching for a pattern and using this knowledge explicitly. Further, the rule discovery group did not show a differential pattern of awareness from the incidental learning group suggesting that the mechanisms responsible for the two forms of learning may indeed be dissociable. Secondly, this suggests a possible reason why no learning by observation was apparent in Experiment 14. If the task instructions were not believable and subjects actively began to search for a structured sequence, then learning would not be observed, as demonstrated by the results shown in this experiment. This issue will be addressed in Experiment 18.

5.6 Experiment 18: Increasing observers' attention to the display

There is a final possibility as to why subjects in an observation condition would not show learning of a sequence compared to subjects in an action condition. As Berry (1991) notes, "one problem with studying learning while observing is that it is difficult to have control over what people actually do while observing". Subjects may simply not be attending to the stimulus display if they do not have to overtly respond to it. The problem may not even be as obvious as simple lack of attention. Willingham et al (1989) suggest that one possible mechanism for facilitation in a sequence learning paradigm is that subjects learn which spatial position to next focus attention on. Results from cuing experiments suggest that signals which are presented at an unexpected location are not processed to the same extent (and so produce a decrement in speed and / or accuracy measures) as signals which are not associated with any spatial expectancy. Conversely, a facilitation in RT is usually observed when subjects attend to the position where a signal is expected (e.g. Eriksen and Yeh, 1985; Muller and Findlay, 1987). It is possible that subjects in an observation condition either do not pay as much attention to the display as subjects in an action condition or they do not shift their attention in the same way. Experiment

18 removes the spatial component by coding the sequence with coloured circles which do not change location. In doing so, this experiment should circumvent the problems associated with spatial expectancy and test the more general hypotheses about implicit learning of sequences. In addition, Experiment 18 will contribute to the small number of studies which examine learning of a non-spatial sequence (Mayr, 1996; Willingham et al, 1989).

The methodology of Experiment 18 attempts to address the problem of diminished attention in the observation condition by giving subjects a plausible reason for watching colours appearing on a computer screen. This serves a dual purpose in that firstly, it will give subjects a reason to pay full attention to the display. Secondly it will address the possibility that subjects are actively looking for the experimental manipulation when being told simply to watch the display. In using a plausible orienting task subjects will be intent on following the given instructions and so will not be as likely to engage in rule searching strategies. As Experiment 17 shows, such strategies inhibit learning of the sequence as compared to incidental study instructions.

Experiment 18 appears to the subjects as an experiment investigating mood alteration by colour. Subjects are asked to complete a mood adjective checklist. They are then told that they will be shown four colours appearing on the computer screen for 10 mins and will subsequently be required to complete an identical checklist. Subjects are led to believe the experimenter will be testing whether any of the mood dimensions on the checklist will have differed between the pre- and post-test sessions. In an additional attempt to maximise subjects' attention to the display, subjects were also asked later to report anything that they might have seen appearing within the coloured circle. This was to make subjects suspect the inclusion of subliminal suggestions within the display and would hopefully act as an additional incentive for subjects to attend to the colours. This individual testing of subjects in the observation condition differs from that in Experiment 11 where an action and an observation subject were tested together. This makes Experiment 18 more similar in methodology to that of Howard et al (1992) than Experiment 14 and so should allow for a stronger comparison of results in the observation condition.

In addition to the observation condition, a second group of subjects performed the sequence learning task in an action condition, i.e. they were required to make overt responses to the stimuli during the learning phase. The purpose of this group is to verify that implicit sequence learning can occur when the spatial component of the display is removed. Subjects were required to say the first letter of the colour which appeared on the screen as quickly as they could. This was the same for both the learning and the test phase.

Subjects in both the action and the observation conditions of the experiment received the same test phase. This was comprised of a vocal response to a colour sequence which had previously been encountered and then a vocal response to a new colour sequence. Subjects in the action condition were required to make vocal responses to the learning set of stimuli and in accord with results from Experiment 14 it is predicted that these subjects will show slower RTs to the new sequence as compared to the old. Subjects in the observation condition were simply required to watch the sequence during the learning phase. However, unlike the observation group in Experiment 14, subjects in this experiment were given a cover task for watching the colours. This cover task should have the effect of discouraging the subjects from attempting any explicit rule search strategies but should also have the effect of encouraging them to pay close attention to the display. From the hypothesis supported by Experiment 14, namely that subjects cannot learn implicitly by observation alone, we would expect subjects in this observation condition to show no difference in RT between the old and the new sequence.

Method

Subjects_

40 subjects were recruited from the undergraduate population of the University of Glasgow. All were naive to the experimental aims and all received a small payment in return for their participation.

Materials

Subjects' vocal responses were timed with a voice key connected to an Apple Power Macintosh computer. The stimuli were generated using graphics software and comprised four circles which were 4.5 cm in diameter. The circles were red, blue, green and yellow in colour. The mood adjective checklist was constructed specially for the experiment with the only requirement being that it show face validity. 29 adjectives describing various mood states (e.g. happy, sad, angry, thoughtful) were printed on an A4 sheet of paper with a five point Likert scale for each adjective.

Design and Procedure

The sequence was presented as a succession of coloured circles on a white background. Each circle was centred on the computer screen. The sequence used was exactly the same as that used in the preceding experiments except that each stimulus position now referred to a particular colour of stimulus. Using this mapping the sequence subjects saw during the learning phase was YGBYRBGRYBGR (where each letter is the first letter of the colour). Using the same mapping the new sequence during the test phase is YGRBGYBRGYRB.

Subjects were assigned to one of two groups; an action or an observation condition. Those in the action condition were told that the experiment was investigating differences in RT to respond to certain colours. They were instructed to respond to each stimulus as quickly as they could by saying the first letter of the colour into the voice key. The stimulus disappeared as soon as the subject responded and the next stimulus was presented 500 msec later. Subjects were told that the colours would appear at random. The learning phase comprised four blocks each containing ten presentations of the sequence. Subjects received a 30 sec break between each of the four blocks. Subjects in the observation condition were told that the experiment was investigating the effect of colour on mood. They were asked to complete the mood adjective checklist describing how they felt at that exact moment in time. They were then told that after a computer display they would be required to complete another identical checklist. Subjects in the observation condition were instructed simply to watch coloured circles appearing on the computer. The display was exactly the same as that shown to the subjects in the action condition except that the stimulus duration in the observation condition was not controlled by the subject. To equate the two conditions for stimulus duration, the action condition was run first and the mean stimulus duration for all four blocks across all subjects was calculated. This duration (525 msec, std dev = 110) was used for each stimulus presentation in the observation condition. The testing phase followed immediately from the learning phase. All subjects were told that they would now receive a further two blocks in which they would be required to respond vocally to each stimulus with the first letter of the stimulus colour. All subjects were told that this was a reaction time measure and encouraged to go as quickly as they could. The first block was identical to those in the learning phase and the second block had the colours following the new sequence. Immediately after the test phase subjects in both conditions were told that the colours had appeared in a particular sequence which was repeated many times and that this was true for all blocks except the final one which they were to ignore. They were informed that they would now be required to predict one cycle of the sequence and accuracy rather than speed was stressed. Again, this task required subjects to press a key corresponding to a particular colour and feedback as to the correct colour was given on each key press.

Results and Discussion

Subjects in the Action condition showed a decrease in RTs across the four learning blocks. The mean RT for each block is shown in Figure 13.

Table 11 shows the mean keypress RTs to the old and new sequences for subjects in the action and the observation conditions.

Figure 13: Mean RT for each block for both the action and observation conditions.

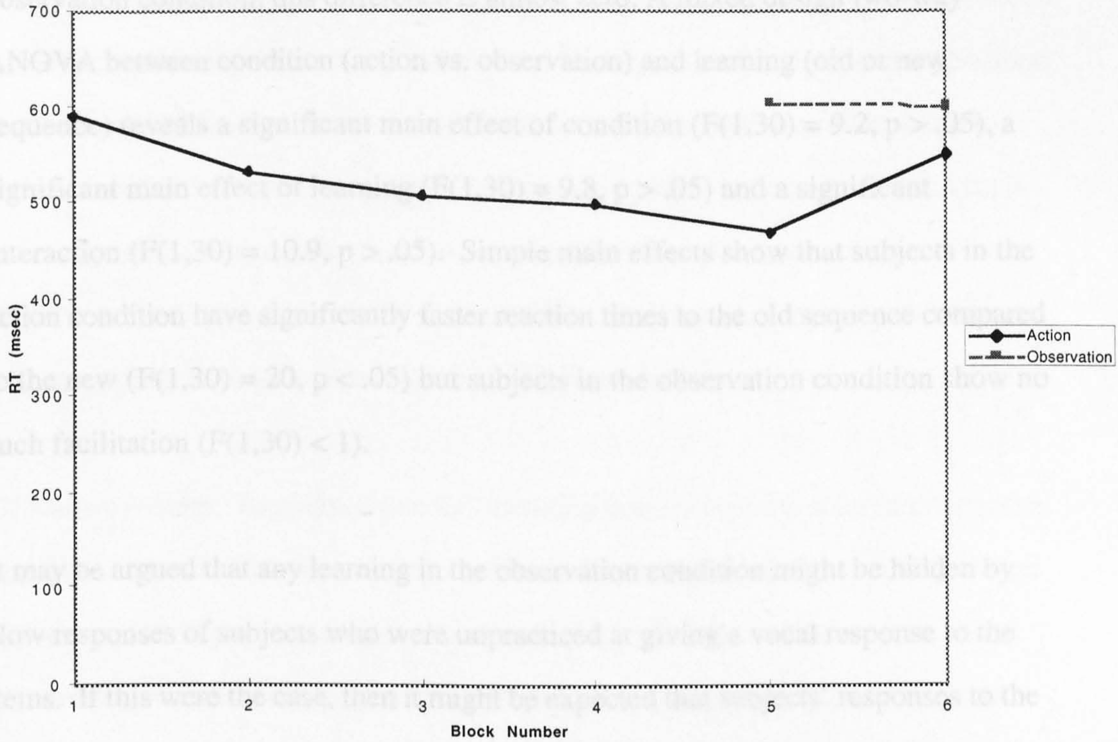

Note: Subjects in the Action condition made a vocal response to the stimuli during Blocks 1 - 4. They then responded via keypress to the same sequence in Block 5 and to a new sequence in Block 6. This was the same for subjects in the Observation condition except that they only observed the stimuli in Blocks 1 - 4 and made no overt response.

Table 11: Mean RTs (msec) and Std Devs for the Action and Observation Conditions

Condition	RT to old sequence.	RT to new sequence	Difference
Action	471 (112)	551 (93)	80
Observation	604 (90)	602 (72)	-2

It can be seen from Table 11 that subjects in the action condition show a large difference in RT between the old and new sequence whereas for subjects in the observation condition, this difference is almost zero. A mixed design two-way ANOVA between condition (action vs. observation) and learning (old or new sequence) reveals a significant main effect of condition ($F(1,30) = 9.2, p > .05$), a significant main effect of learning ($F(1,30) = 9.8, p > .05$) and a significant interaction ($F(1,30) = 10.9, p > .05$). Simple main effects show that subjects in the action condition have significantly faster reaction times to the old sequence compared to the new ($F(1,30) = 20, p < .05$) but subjects in the observation condition show no such facilitation ($F(1,30) < 1$).

It may be argued that any learning in the observation condition might be hidden by slow responses of subjects who were unpracticed at giving a vocal response to the items. If this were the case, then it might be expected that subjects' responses to the new sequence would be significantly faster in the action condition as opposed to those in the observation condition and this was not found ($F(1,30) = 1.41, p > .1$).

Scores on the prediction task were 6.44 for the action subjects and 6.56 for the observation subjects and were not statistically different from one another ($t(30) = .17, p > .1$). Awareness as measured by the prediction task will thus not be considered further in this experiment.

The results clearly show subjects can learn a non-spatial sequence. This is in agreement with results reported by Mayr (1996) which suggest that subjects can learn a spatial or a non-spatial sequence. Subjects in Experiment 14 could not learn a spatial sequence by observation alone. It was hypothesised that this may have been due to subjects in the observation condition having different patterns of spatial attention to those in a response condition. Experiment 18 shows that even when spatial attention is not necessary for the task, subjects still cannot learn by observation alone. Both action and observation conditions required subjects to attend to the same part of the screen and instructions to subjects in the observation condition emphasised paying attention to the coloured circles. Subjects were even asked to report black lines or words that they might think they have seen within the circles in an attempt to maximise their attention to the display. Despite these attempts to force subjects in the observation condition to attend to the display no learning was found, which supports the conclusions reached in Experiment 14.

Of course, it cannot be claimed that this methodology equated the amount of attention between the action and observation conditions. The criticism that responding forced action subjects to pay more attention to the display than observation subjects might still account for the data reported in this experiment. However, equating subjects' attention while one responds and the other does not seems to be an impossible task. Although this criticism is a valid one, it does seem extremely unlikely that absolutely no learning would take place in the observation condition if the only difference were simply a matter of attenuated attention. Hence it seems that lack of learning in the observation condition cannot be attributed to a variation in attentional factors.

In addition to encouraging subjects in the observation condition to attend to the display, Experiment 18 also gave subjects what seemed to be a plausible reason for observing such a display. In Experiment 17, subjects who were told to look for a sequence displayed no learning at all on a RT measure. It is possible that subjects in an observation condition become suspicious at simply being asked to observe a display and so begin to actively attempt to work out what is going on. This may have been why subjects in an observation condition seem to show no learning of the sequence. The results of Experiment 18 appear to rule out this possible explanation. Subjects were led to believe that the different coloured circles were part of a mood induction experiment and all subjects said they had believed the cover story when asked during the debriefing. It is therefore unlikely that they would have attempted to discover any sequence. These results show that subjects under definite incidental learning conditions do not learn the sequence by observation alone.

5.7 General Discussion

Berry (1991) found that for maximally effective transfer of learning across blocks of a dynamic systems task, action had to be tied to decision. Neither on their own were as effective during the task which had a non-salient underlying rule but either was sufficient when the rule was salient. This finding suggests an important dissociation between implicit and explicit processes. In a totally different implicit learning paradigm, Howard et al (1992) found that subjects who observed a sequence of asterisks presented many times showed just as much learning as subjects who responded to the sequence via keypress. This result runs somewhat contrary to that reported by Berry (1991), however there were a number of problems associated with the Howard et al study which were corrected in the present study.

Experiment 14 required one subject to respond to a sequence learning task using a sequence of 12 items. Another subject was simply instructed to watch. Learning was measured by RT decrement when the sequence was changed to another 12 item sequence. The results from Experiment 14 do not replicate the finding of Howard et al (1992) in that the subjects who simply observed showed no evidence of learning. This would agree with the results found by Berry (1991) in the dynamic systems task in that both observation and action are necessary to learn without awareness. The lack of learning for the observation group may have been due to the knowledge being modality-specific, although previous studies in sequence learning (Keele et al, 1995) suggest that this should not be the case.

In Experiment 15, subjects were trained on the sequence using vocal responses and then tested via keypress. These subjects showed significant learning of the sequence using this method. Such learning cannot then be modality specific. This also suggests that if as Berry concludes, action is necessary for learning, then the definition of action must be extended from interacting via the keyboard to giving a vocal response. It seems that the crucial component missing from learning a non-salient rule structure is then any overt physical response to that sequence.

It is possible that the reason mediating learning in Experiment 15 but not in Experiment 14 was that subjects were imposing their own response pattern onto the sequence in Experiment 15. In Experiment 14, subjects were watching a sequence which was responding to another subjects' RT pattern and as Stadler (1993) suggests, this may have had the effect of disrupting the organisation of the sequence for the observing subject. Experiment 16 utilised a design which split action and observation in the learning phase. The subjects who were observing the sequence had to say the sequence to their partner but were not having their RTs recorded. The next position of the asterisk would appear after a keypress by the subject in the action condition. Thus the subject in the observation condition of Experiment 16 was overtly responding to the sequence but the sequence responded to the other subject's RT pattern. Both subjects showed learning of the sequence and there was no statistically significant difference between performance of the two groups. This suggests that for subjects in the observation condition, the crucial factor missing from Experiment 14 was some sort of overt physical response while interacting with the task.

Experiment 17 investigates the possibility that subjects in an observation condition may be using rule-search strategies to solve the task. It is known that such strategies can have a detrimental effect on learning in other implicit learning paradigms and Experiment 17 shows that this is also the case for sequence learning. This problem is addressed in Experiment 18 where subjects in the observation condition are led to believe that observing the stimuli is part of a mood induction process. Experiment 18 also removed the spatial component from the task as another possible difference between action and observation conditions was the pattern of subjects' spatial attention. Experiment 18 replicated the result of Experiment 14 in that subjects in the observation condition showed no evidence of learning but subjects giving an overt response to a non-spatial sequence showed robust learning.

The experiments reported in this chapter investigate a number of possibilities as to why learning was not found for subjects observing a sequence unless they were overtly responding in some way to the stimuli. None of the alternative explanations could adequately account for this finding and the conclusion must be that some form of overt action is necessary for learning in an implicit sequence learning paradigm. The reason for an overt physical response as a requirement for learning a non-salient rule-governed task is unclear. Berry's (1991) discussion also clearly indicates that any mechanism for learning with such a requirement is not an obvious one. There are several studies however, on the relationship between perception and action which may provide a speculative account of how this mechanism may be mediated.

Christovich, Fant, Serpa-Leitao and Tjernlund (1966) presented subjects with a synthesised speech continuum between adjacent vowel categories and the subjects were required to repeat each individual vowel. Christovich et al found that there was no one to one relationship between the formants of the stimulus presentation and the formants in the subjects' response set. In fact, subjects achieved a good imitation for only three of the canonical vowels. Repp and Williams (1985) found similar results for subjects imitating synthetic vowels. Massaro (1990) suggests that these results show that perception and action are only indirectly related and suggests that the step which links them is the process of categorisation. It may be that in the absence of action, a non-salient sequence is not categorised. For a salient sequence one might suggest action is not necessary for categorisation but this argument is not convincing. A more plausible argument might be that it is possible that subjects will use inner speech to code the responses for a salient display and this inner speech will also serve the broad definition of 'action' that Experiment 15 suggests must be used.

There is another theory which suggests a dichotomy between perception with action and perception without action. Neumann (1990) has posited a theory which claims that the higher vertebrates have developed a particular type of action which Neumann terms 'exploration'. Neumann suggests that 'exploration' serves the purpose of establishing and updating an internal representation of the environment. Although it is defined as a particular type of action, the end result of exploration need not be action but simply the acquisition of information. Neumann argues that all actions include central components and these can be activated without actually performing the overt action, however the intended goal can be attained only if the physical action finally takes place. By contrast, he argues that exploration can be performed without carrying out the peripheral movements. In Neumann's theory, there is thus a clear distinction between perception with action and this concept of 'exploration' whereby an internal representation of the environment is either established or updated without the need for overt action.

Can the proposed dichotomy between these two types of action be useful in suggesting a reason why we should not expect implicit learning by observation? Neumann (1990) links exploration with conscious awareness and so one might expect that in the absence of conscious awareness, an internal representation may not be established in the absence of action. Although this does not lead to the conclusion that without awareness action will establish the internal representation, there is a final part of Neumann's theory which links implicit processes with action. Neumann suggests that the "functional basis of visual attention combines phylogenetically old mechanisms that select information for the immediate control of action with more recent mechanisms that subserve the updating of the internal representation". In other words, it is primitive mechanisms that are used in the control of perception leading directly to action. It is also suggested by Reber (1993) that it is phylogenetically primitive mechanisms which subserve implicit processes. Hence, if there is a primitive unconscious learning system, we might expect it to use information from other primitive systems and so use information from the system which ties perception to action. While these ideas are nothing more than speculation, they do demonstrate that there may be a difference between perception and action and perception without action. In light of this, the dissociation described in this chapter and by Berry (1991) between action and observation in implicit learning tasks does not seem as peculiar and as unlikely as would at first be thought.

The possibility of such a dissociation existing is not interesting simply in the academic sense. It can clearly be seen that knowledge of such a dissociation would be very useful for research into applied aspects of implicit learning. In so far as implicit learning has been implicated in the use of knowledge elicitation systems (Berry and Dienes, 1993), it is easy to imagine that simply being shown how to use such a system may not be as useful as actually participating in scenarios for oneself. On more theoretical terms, there has been a recent tendency to subdivide implicit learning tasks into various groupings (e.g. Frick and Lee, 1995; Seger, 1994) which draws into question the existence of this general notion of implicit learning. While the results described in this chapter obviously do not apply to tasks such as grammar learning, it does suggest that there may be general properties of implicit learning which are common to tasks requiring similar processing and responses for subjects. This rather peculiar dissociation may then provide another index by which implicit learning tasks may be grouped together.

The separation of action and observation led to an interesting result regarding awareness. For both the observation and action groups, greater levels of sequence knowledge as indexed by RT was associated with greater levels of explicit knowledge as measured by the prediction task. While it is clear that further work must be done, it does suggest that awareness is a possible confound to any results. It is clear that to refine our knowledge about the mechanism by which implicit learning might operate, it is necessary to break tasks down into their components and to ensure that each of those components are operating with an absence of conscious awareness.

Chapter 6 - Summary and Conclusions

This thesis examined several assumptions made about certain implicit learning paradigms and aspects of implicit learning in general. Chapter 1 provided an overview of the main properties associated with implicit learning and evaluated evidence relating to these properties. One main question relating to implicit learning is whether such learning really is unconscious in nature. From the discussion of this in Chapter 1 the conclusion seems to be a rather unsatisfactory "it depends on how you define 'unconscious'". The second main question relating to implicit learning asks how the knowledge is actually represented. There is evidence for both instance storage and abstraction. This issue of representation may even be different for different paradigms and different learning conditions within a specific paradigm. Chapters 2-5 presented evidence relating to these questions and the conclusions which this evidence suggests will now be summarised.

Chapter 1 investigated several aspects of the digit invariance task of McGeorge and Burton (1990). They discovered that subjects could seemingly acquire and use information about an invariant quality of a stimulus display without any apparent knowledge of that quality. Furthemore, McGeorge and Burton found that the information acquired showed transfer across form suggesting that a conceptual representation of the invariant quality was abstracted. Wright and Burton (1995) suggested that subjects were simply rejecting salient items and Cock et al (1994) provided evidence that 'nearest-neighbour' similarity could account for the results of McGeorge and Burton. Experiment 1 attempted a replication of this finding of specific similarity and could find no evidence of subjects being able to use such similarity to give above chance performance on this task. The data however did not show any evidence for learning the invariant quality of the stimuli under normal testing conditions. The only evidence for such learning was found under conditions of a 2 sec test deadline and when the items in the 2-AFC were both 75% similar to a previous instance. This effect was replicated in Experiment 2. This suggests that, in accordance with results such as those reported by Turner and Fischler (1993), a time deadline in some way affords the use of implicit knowledge. One possibility is that the use of a time deadline prevents explicit heuristics from being used and so the only information the subject can use is that which can be used without such mechanisms. It is interesting to note that this effect did not occur when the test exemplars were anything other than very similar to a previous instance. One would imagine a time deadline would only be able to exclude heuristics of a certain complexity and perhaps this complexity would be necessary to choose between two such exemplars but a much simpler one could be used with other types of relationship between the two exemplars.

Experiment 3 provides evidence for this notion of decreasing influence of heuristics and / or episodic memories having a facilitative effect on performance due to the invariant digit. Under conditions where an immediate test phase shows chance performance, subjects given the test phase after a delay of 30 mins showed robust above chance performance. This fits with theories of prototype abstraction where individual items are given less weight in decision making over time and qualities abstracted over the set of these items are given more weighting. This type of effect can be instantiated as a neural network of the type described by McClelland and Rumelhart (1985). This knowledge was shown to extend to parts of pictorial stimuli (Experiment 5) and that this knowledge had decayed fully at an interval of one week (Experiment 4). These experiments also suggest that the Wright and Burton (1995) account of performance is unlikely as one would expect salient items to be rejected immediately as well as after 30 mins. Experiment 6 examined the claim by

McGeorge and Burton (1990) that this type of knowledge could show transfer across form. Only subjects who received the test exemplars in the same surface format as the study exemplars demonstrated above chance performance. The original study of McGeorge and Burton did not have exemplars constructed to avoid use of simple similarity or saliency in decision making plus they did not use a procedure which prevented phonological overlap between study and test. Any of these factors could account for their earlier finding of transfer across form.

In sum, this chapter suggests a type of mechanism for learning which is similar to the original idea proposed by McGeorge and Burton, i.e. abstraction across a set of instances. The information which is abstracted seems tied to surface properties of the stimuli and the knowledge is not robust over intervals of one week. The information is not specific to individual instances (as shown by the PDNS manipulation) but is reliant on a pool of exemplars. These points all seem at odds with recent suggestions as to the general properties of implicit learning (e.g. Dienes and Berry, 1997). According to these suggested properties, implicit learning is usually transferable to items with relatively similar surface structure (although certain studies, e.g. Altmann, Dienes and Goode, 1995) suggest transfer is possible across grossly different surface structure as long as the deep structure is preserved). It is also supposed to be associated with specific items rather than underlying rules and should be robust over time. These data open up two possibilities. The first is that the suggested properties for implicit learning in general are wrong. Dienes and Berry (1997) cite experimental evidence supporting each of these suggestions yet it could be argued that none of the experiments have examined these properties under conditions of equivalent similarity and a time deadline. Indeed it would be a mistake to argue that if these conditions were adopted by a grammar task for example, then it would lead to the same influence of implicit processes upon task performance. As yet, it is probable that virtually all of the studies reported have some explicit contamination of results. Until better methods of controlling for this are found it is uncertain whether certain assumptions made about transfer etc. are in fact correct in relation to a process-pure implicit task. The other possibility is that the digit invariance task is a type of implicit learning task which is different from implicit learning characterised by the grammar task or a sequence learning task. Seger (1994) notes that "implicit learning shows biases and dissociations in learning different stimulus structures" and Frick and Lee (1995) suggest that implicit learning can be divided into three separate paradigms each producing different types of knowledge. One of these paradigms is where subjects process the stimuli which follow a pattern and learn the similarity between the stimuli. This clearly reflects artificial grammar and is the closest of the three of Frick and Lee's paradigms to describing invariance learning. It is known however, that learning in the grammar task is affected by similarity and shows conceptual transfer so one must clearly argue that the digit invariance task represents a fourth type of paradigm. This paradigm could be described as one in which subjects process stimuli which contain an invariance and subjects learn that invariance without demonstrating verbalisable knowledge of it.

Chapter 3 used this paradigm to investigate possible differences in the use of 2-AFC and yes / no recognition measures. Both could be used to test awareness above an "objective threshold" which is to say that knowledge can be said to be implicit if subjects' performance on a forced choice task is at chance. Of course, as Berry and Dienes (1993) note "logically subjects cannot perform at chance on all forced-choice tests otherwise there would be no measure left to indicate knowledge, implicit or otherwise". The digit invariance task of McGeorge and Burton (1990) demonstrates that assuming an objective criteria as a necessary proof of implicit processes does not seem logical in relation to this task as the test phase itself is the more sensitive type of recognition test. Authors such as Shanks and St. John (1994) or Perruchet and Amorin (1992) would argue that above chance performance on this 2-AFC measure demonstrates that subjects have explicit knowledge of the invariant quality present in the display yet verbal reports clearly show that subjects have no such knowledge. The simple nature of the rule underlying successful performance on this task makes the issue even clearer. In tasks such as grammar learning where the rule system is complex one can argue that recognition tasks tap fragmentary knowledge which can still lead to successful performance without necessarily giving 100% accuracy. In the digit invariance task there is one simple rule. Subjects know it or they do not know it and the rule cannot be fragmented. Hence if subjects truly had explicit awareness then they should achieve maximum performance, which they do not. Performance on this task is clearly unconscious based on a subjective threshold but it is unclear how to define the underlying processes based on an objective threshold.

The results of Chapter 3 make this issue somewhat clearer. Several experiments show that subjects can perform at above chance levels on the digit invariance task and the related clocks task (Bright and Burton, 1994) when a 2-AFC task is used. Performance is at chance levels when subjects must perform a yes / no recognition task. This suggests that there is no one ‘objective threshold’ at which all recognition tests exhibit task knowledge. Instead one might have to acknowledge a continuum of tests of differential sensitivity. It may be that the implicit / explicit distinction can be located at one specific part of this continuum or it may be that this is determined according to the type of task used. The digit invariance task seems to require the most sensitive type of recognition task to elicit information whereas above chance performance can be demonstrated in artificial grammar using a yes / no recognition task. Shanks et al (1997) demonstrated that if a 2-AFC paradigm is used in artificial grammar then task performance can reach 80%, which is much larger than that usually observed with yes / no recognition. These results suggest that one cannot adopt a simple definition for an objective threshold. They also point to the digit invariance task as possibly being mediated by implicit processes as one of the criteria for 'implicitness' is for the acquired knowledge to be difficult to elicit. The fact that the invariance knowledge can only be elicited by a 2-AFC and not a yes / no recognition task demonstrates this point.

Chapter 4 attempts to demonstrate real world implicit learning. As the previous chapters seemed to suggest that invariance detection is a robust phenomenon which acquired knowledge unavailable to verbal report or even yes / no recognition tasks, the assumption of this chapter was that invariant properties of the real world should also have been learned by subjects. This follows from the theories put forward by Reber (1993) who suggests that implicit learning is a primitive process and so should be common to everyone with little individual variation in ability. Bright (1994) presented evidence that recalling the direction in which the monarch's head faces on coins can be done only using a 2-AFC task and not by recall. This example of real life invariance detection seems then to bear remarkable similarity in accessibility of knowledge to laboratory experiments in invariance detection. The results reported by Bright (1994) however could not be replicated in Experiment 11. The effect reported by Bright may simply not be robust or it may be that explicit recall may have caused interference in the recall of implicit information. Experiment 12 investigated whether knowledge of this invariance could be elicited under time deadline conditions with an indirect task and using a more sensitive measure of reaction time rather than accuracy. No effect of the invariant property was found. It may be that this type of information is just not capable of being abstracted despite the fact that it is invariant or that the task used was not appropriate to demonstrate such knowledge. Experiment 13 examined whether a different type of invariance (colour) can be learned unconsciously. Using the logo from the British television station 'Channel 4', subjects were given a mock tachistoscopic identification task. They were required to choose which logo out of six options was the one which they had seen subliminally (when in fact they had seen no image at all). Subjects did not perform at above chance levels although when explicitly asked to choose the correct Channel 4 logo, they could do so. This may demonstrate incidental learning of such information but it does not demonstrate implicit learning as subjects could give the correct answer to a question and one cannot rule out episodic memory as the mechanism. This chapter could not offer any convincing evidence for invariance detection in real life. This leads to the possibility that either laboratory tasks demonstrate learning which is not applicable to real life but is some artefact of the laboratory procedure or that learning in real life operates under certain constraints which the experiments in Chapter 4 have violated in some way. In order to determine which may be the correct interpretation, it will be necessary in future research to incorporate into experimental designs constraints of greater ecological validity. For example, we are unlikely to encounter invariance in the environment in a set of items following sequentially on from one another. Experiments which may be useful would examine spaced versus massed presentations of items to simulate this. The other alternative is to discover real world scenarios which do have sequentially presented sets of items containing an invariant quality and attempt to demonstrate such learning in a laboratory setting. As noted however, it is difficult to think of a real world example which fits the constraints of a laboratory implicit learning task.

Chapter 5 moves away from the invariance detection paradigm and examines a very interesting dissociation between implicit and explicit learning. Berry (1991) found that implicit learning could not proceed through observation of stimuli alone. Action during the control of a complex system was necessary. This was not the case when an explicit (using a salient rule) version of the same task was used. Under those conditions, performance was the same for subjects who initially responded to the system or who initially simply observed the responses of another. This is not what was found in another type of implicit learning task, however. Howard et al (1992) demonstrated that subjects observing the stimulus display in a sequence learning task did learn the sequence as well as subjects who responded to it in a reaction time experiment. From Berry's (1991) results, one would not have predicted that learning by observation alone would be demonstrated in other implicit learning tasks. Chapter 5 explores this issue and begins by noting several methodological issues in Howard et al's (1992) procedure which may account for the artefactual finding of learning by observation in an implicit task.

Experiment 14 addressed these possible problems and found that there was no evidence of learning when subjects merely observed the stimulus display. Experiments 15 and 16 demonstrated that this was not due to knowledge being unable to cross domains or observing another caused interruption of the sequence (shown by Stadler (1993, 1995) to affect performance). Experiment 17 tested the hypothesis that subjects in Experiment 14 did not learn because they attempted to discover an underlying pattern which may have had adverse effects on learning implicitly. Experiment 17 shows that subjects given rule discovery instructions do not learn a sequence which 'incidentaly' instructed subjects can learn. Thus it is possible that subjects simply asked to watch another perform the task were suspicious of these instructions and looked for some order within the display. Experiment 18 addresses this problem by presenting subjects with a plausible cover story for simply watching a display. Experiment 18 also attempts to address the other major fault with Experiment 14 in that one cannot be sure that subjects were actually paying attention to the display. Subjects in Experiment 18 were given good reason to attend very carefully to the display and although one cannot state that this equates attention between action and observation subjects, one can be reasonably sure that attention was focused on the display. Despite ensuring that subjects attend to the display and do not search for underlying rule patterns, those subjects who observed the pattern during the learning trials showed no evidence of learning the sequence. The subjects who physically responded to the sequence did show learning. Thus, the implication from the experiments reported in Chapter 5 are that learning cannot proceed by observation alone in an implicit learning task and this finding concurs with that of Berry (1991) in another implicit learning paradigm. This gives additional evidence to advance the suggestion of Dienes and Berry (1997) that qualitative differences demarcate explicit from implicit processes.

Although this dissociation does seem rather strange, the general discussion in Chapter 5 suggest one theory of relations between perception and action which might account for it. Neumann (1990) draws a distinction between 'perception with action' and 'perception used to update an internal representation' and links the latter with conscious awareness. Thus it might be that consciousness is necessary to learn using perception alone. As pointed out this does not mean that action is therefore necessary to learn without awareness. This theory does provide a distinction between the two however and may be the first step in understanding this peculiar finding first reported by Berry (1991). The findings of Chapter 5 lead to one final speculation with respect to the theme of invariance detection from earlier chapters. Chapter 4 could find no evidence for learning a perceptual invariant in the real world yet the preceding chapters presented evidence suggesting that implicit invariance detection was a real and robust phenomenon. The results of Chapter 5 offer an explanation to this apparent conflict of evidence. In laboratory experiments, subjects are required to do something with the stimuli. They respond to them in some manner during the study phase. This is not necessarily the case with the examples used in Chapter 4. Presumably one knows which television channel one is watching or turning over to without waiting for the Channel 4 logo. Similarly, the side of a coin with the monarch's head on is not the side which tells us the denomination of the coin. Thus one could argue that one would simply observe such stimuli without acting in any way towards them. Of course there are many other complex ways in which laboratory tasks and real life tasks differ and I use the terms 'action' and 'responding' in a very loose manner but it is an interesting generalisation of the findings of Chapter 5 and one worth considering in future investigations of real world implicit learning.

References

Abrams, M. (1987). *Implicit learning in the psychiatrically impaired.* Unpublished doctoral dissertation, City University of New York. (cited in Reber, A. S. (1993). *Implicit learning and tacit knowledge: An essay on the cognitive unconscious.* New York, NY, US: Oxford University Press)

Abrams, M., & Reber, A. S. (1988). Implicit learning: Robustness in the face of psychiatric disorders. *Journal of Psycholinguistic Research,* **17** (5), 425-439.

Allen, R., & Reber, A. S. (1980). Very long-term memory for tacit knowledge. *Cognition,* **8,** 175-185.

Altmann, G. T. M., Dienes, Z., & Goode, A. (1995). Modality independence of implicitly learned grammatical knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition,* **21**, 899-912.

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge MA: Harvard University Press.

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

Bahrick, H. P. (1954). Incidental learning under two incentive conditions. *Journal of Experimental Psychology* , **47**, 171-172.

Balota, D. (1983). Automatic semantic activation and episodic memory encoding. *Journal of Verbal Learning and Verbal Behavior,* **22**, 88-104.

Berry, D. C. (1991). The role of action in implicit learning. *Quarterly Journal of Experimental Psychology Human Experimental Psychology,* **43A** , 881-906.

Berry, D. C. (1994). Implicit learning: Twenty-five years on. A tutorial. In Carlo Umilta & Morris Moscovitch (Eds.), *Attention and performance 15: Conscious and nonconscious information processing.* (pp. 755-782). Cambridge, MA, US: MIT Press.

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

Berry, D. C., & Broadbent, D. E. (1988). Interactive tasks and the implicit-explicit distinction. *British Journal of Psychology,* **82**, 359-373.

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

Broadbent, D.E. (1977). Levels, hierarchies and the locus of control. *Quarterly Journal of Experimental Psychology*, **29**, 181-201.

Broadbent, D.E., & Aston, B. (1978). Human control of a simulated economic system. *Ergonomics*, **21**, 1035-1043.

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.

Bright, J. E. H. (1994). *Issues in Implicit Learning*. Unpublished Doctoral dissertation. University of Nottingham.

Bright, J. E. H., & Burton, A. M. (1994). Past midnight: Semantic processing in an implicit learning task. *Quarterly Journal of Experimental Psychology Human Experimental Psychology,* **47A** (1), 71-89.

Brewer, W. F. (1974). There is no convincing evidence for operant or classical conditioning in adult humans. In W. B. Weimer & D. S. Palermo (Eds.) *Cognition and the symbolic processes.* Hillsdale, NJ. Erlbaum.

Brody, N. (1989). Unconscious learning of rules: Comment on Reber's analysis of implicit learning. *Journal of Experimental Psychology: General*, **118**, 236-238.

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

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** (3), 316-323.

Buchner, A., Funke, J., & Berry, D. C. (1995). Negative correlations between control performance and verbalizable knowledge: Indicators for implicit learning in process control tasks? *Quarterly Journal of Experimental Psychology Human Experimental Psychology,* **48A** (1), 166-187.

Buchner, A., Steffens, M. C., Erdfelder, E., & Rothkegel, R. (1997). A multinomial model to assess fluency and recollection in a sequence learning task. *Quarterly Journal of Experimental Psychology*, **50A**, 631-663.

Cantor, G. W. (1980). *On symbol string-type and the generalizability of the implicit learning paradigm.* Paper presented at meeting of the Eastern Psychological Association, New York, NY. (cited in Reber, A. S. (1993). *Implicit learning and tacit knowledge: An essay on the cognitive unconscious.* New York, NY, US: Oxford University Press)

Chan, C. (1992). *Implicit cognitice processes: Theoretical issues and applications in computer systems design.* Unpublished DPhil thesis, University of Oxford. (cited in Berry, D. C., & Dienes, Z. (1993). *Implicit learning: Theoretical and empirical issues.* Hove, England: Lawrence Erlbaum Associates, Inc.)

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

Cherry, K. E., & Stadler, M. E. (1995). Implicit learning of a nonverbal sequence in younger and older adults. *Psychology and Aging,* **10** (3), 379-394.

Christovitch, L. A., Fant, G., Serpa-Leitao, A., & Tjernlund, P. (1966). Mimicry of synthetic vowels. (Speech Transmission Laboratory Reports). Stockholm. Royal Institute of Technology. (cited in O. Neumann & W. Prinz (Eds) *Relationships between Perception and Action: Current Approaches.* (1990). Berlin: Springer-Verlag.)

Churchill, E. F., & Gilmore, D. J. (1995) *Knowing what you don't know: Strategies for selection and rejection in invariant learning tasks.* Paper presented at the BPS Cognitive Psychology Section's 12th Annual Conference. Bristol.

Churchill, E. F., Newell, B., & Gilmore, D. J. (1996) *Time after time: Does meaningfulness affect abstraction in implicit learning.* Paper presented at the BPS Cognitive Psychology Section's 13th Annual Conference. Keele.

Claparede, E. (1951). Recognition and 'me-ness'. In D. Rapaport (Ed.), *Organisation and pathology of thought..* (cited in: Schacter, D.L. (1987). Implicit Memory: History and current status. *Journal of Experimental Psychology: Learning, Memory and Cognition* , 13, 501-518.)

Cleeremans, A. (1994a). Awareness and abstraction are graded dimensions. *Behavioral and Brain Sciences,* **17** (3), 402-403.

Cleeremans, A. (1994b). The representation of structure in sequence prediction tasks. In Carlo Umilta & Morris Moscovitch (Eds.), *Attention and performance 15: Conscious and nonconscious information processing.* (pp. 783-809). Cambridge, MA, US: MIT Press.

Cock, J. J., Berry, D.C., & Gaffan, E. A. (1994). New strings for old - the role of similarity processing in an incidental learning task. *Quarterly Journal of Experimental Psychology*, **47**, 1015-1034.

Cohen, A., Ivry, R., & Keele, S. (1990). Attention and structure in sequence learning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **16**, 17-30.

Cohen, N.J. (1984) Preserved learning in amnesia: Evidence for multiple memory systems. In L.R. Squire & N. Butters (Eds.) *Neuropsychology of memory* (pp. 83-103). New York: Guildford.

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

Dienes, Z. (1993) Computational models of implicit learning. In D. C. Berry & Z Dienes (Eds.) *Implicit learning: Theoretical and empirical issues*. (pp. 81-112). Hove, England: Lawrence Erlbaum Associates, Inc.

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

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

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

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

Dienes, Z., & Perner, J. (1994). Dissociable definitions of consciousness. *Behavioral and Brain Sciences,* **17** (3), 403-404.

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.

Druhan, B., & Mathews, R. (1989). THIYOS: A classifier system model of implicit knowledge of artificial grammars. *Proceedings of the 11th Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Lawrence Erlbaum Associates. (cited in Berry, D. C., & Dienes, Z. (1993). *Implicit learning: Theoretical and empirical issues.* Hove, England: Lawrence Erlbaum Associates, Inc.)

Dulany, D. E. (1961). Hypotheses and habits in verbal operant conditioning. *Journal of Abnormal and Social Psychology*, **63**, 251-263.

Dulany, D. E., Carlson, R. A., & Dewey, G. I. (1984). A case of syntactical learning and judgment: How conscious and how abstract? *Journal of Experimental Psychology General,* **113** (4), 541-555.

Dunn, R. (1845). Case of suspension of the mental faculties. *Lancet* , **2**, 588-590. (cited in: Schacter, D.L. (1987). Implicit Memory: History and current status. *Journal of Experimental Psychology: Learning, Memory and Cognition* , 13, 501-518.)

Eimer, M., Goschke, T., Schlaghecken, F., & Sturmer, B. (1996). Explicit and implicit learning of event sequences: Evidence from event-related brain potentials. *Journal of Experimental Psychology Learning, Memory, and Cognition,* **22** (4), 970-987.

Erdelyi, M. & Becker, J. (1974). Hypermnesia for pictures: Incremental memory for pictures but not words in multiple recall trials. *Cognitive Psychology*, **6**, 159-171.

Ericsson, K. A., & Simon, H. A. (1980). Verbal reports as data. *Psychological Review*, **87**, 215-251.

Eriksen, C. W., & Yeh, Y. Y. (1985). Allocation of attentionin the visual field. *Journal of Experimental Psychology: Human Perception and Performance*, **5**, 583-597.

Estes, W. K. (1986). Memory storage and retrieval processes in category learning. *Journal of Experimental Psychology: General.*, **115**, 155-174.

Eysenck, M. W. & Keane, M. (1990). *Cognitive Psychology: A Student's Handbook.* Hove UK: Lawrence Erlbaum Associates Ltd.

Fodor, J. (1983). *The Modularity of Mind.* Cambridge, Mass.: MIT press.

Frensch, P. A., Buchner, A., & Lin, J. (1994a). Implicit learning of unique and ambiguous serial transitions in the presence and absence of a distractor task. *Journal of Experimental Psychology Learning, Memory, and Cognition,* **20** (3), 567-584.

Frensch, P. A., & Miner, C. S. (1994b). Effects of presentation rate and individual differences in short-term memory capacity on an indirect measure of serial learning. *Memory and Cognition,* **22** (1), 95-110.

Frick, R. W., & Lee, Y. S. (1995). Implicit learning and concept learning. Quarterly *Journal of Experimental Psychology Human Experimental Psychology,* **48A** (3), 762-782.

Funke, J. & Muller, H. (1988). Eingreifen und prognostizieren als Determinanten von System-identifikation und Systemsteuerung. *Sprache und Kognition,* **7**, 176-186. (cited in Berry, D. C. (1991a). The role of action in implicit learning. *Quarterly Journal of Experimental Psychology Human Experimental Psychology,* **43A** (4), 881-906.)

Gilmore, D. J., & Churchill, E. F. (1996) *Distinctiveness and similarity in implicit invariance detection..* Paper presented at the BPS Cognitive Psychology Section's 13th Annual Conference. Keele.

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** (2), 396-410.

Graf, P. (1990). Life-span changes in implicit and explicit memory. *Bulletin of the Psychonomic* Society, **28**, 353-358.

Graf, P., Squire, L., & Mandler, G. (1984). The information that amnesic patients do not forget. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **10**, 164-178.

Green, R., & Shanks, D. (1993). On the existence of independent explicit and implicit learning systems. *Memory and Cognition*, **21**, 304-317.

Green, D. M. & Swets, J. A. (1966). Signal Detection Theory and Psychophysics. New York: Wiley.

Greenspoon, J. (1955). The reinforcing effect of two spoken sounds on the frequency of two responses. *American Journal of Psychology*, **68**, 409-416.

Haith, M. M., Hazan, C., & Goodman, G.S. (1988). Expectation and anticipation of dynamic visual events by 3.5 month-old babies. *Child Development*.. **59**, 467-479. (cited in Reber, A. S. (1997). Implicit ruminations. *Psychological Bulletin and Review.* **4**, 49-55.)

Haith, M. M., & McCarty, M. E. (1990). Stability of visual expectations at 3.0 months of age. *Developmental Psychology*, **26**, 68-74. (cited in Reber, A. S. (1997). Implicit ruminations. *Psychological Bulletin and Review.* **4**, 49-55.)

Hartman, M., Knopman, D. S., & Nissen, M. J. (1989). Implicit learning of new verbal associations. *Journal of Experimental Psychology Learning, Memory, and Cognition,* **15** (6), 1070-1082.

Hay, J. F., & Jacoby, L.L. (1996). Separating habit and recollection: memory slips, process dissociations and probability matching. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **22**, 1323-1335.

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

Hendrickx, H., De Houwer, J., Baeyens, F., Eelen, P., & Van Avermaet, E. (1997). Hidden covariation detection might be very hidden indeed. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **23**, 201-220.

Hill, T., Lewicki, P., Czyzewska, M., & Schuller, (1990). The role of learned inferential encoding rules in the perception of faces: Effects of nonconscious self-perpetuation of a bias. *Journal of Personality and Social Psychology*, **26**, 350-371.

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

Holender, D. (1986). Semantic activation without conscious identification in dichotic listening, parafoveal vision and visual masking: A survey and re-appraisal. *The Behavioral and Brain Sciences,* **9**, 1-23.

Holland, J., Holyoak, K., Nisbett, R., & Thagard, P. (1986). *Induction: Processes of inference, learning and discovery.* London: MIT press.

Homa, D., Sterling, S., & Trepel, L. (1981). Limitations of exemplar-based generalisation and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning and Memory*, **7**, 418-439.

Howell, D.C. (1992). *Statistical Methods for Psychology (Third edition).* Belmont, California: Duxbury Press.

Howard, D. V. (1988). Implicit and explicit assessment of cognitive aging. In M. L. Howe & E. J. Brainerd (Eds.) *Cognitive development in adulthood: Progress in cognitive development research.* (pp. 3-37). New York: Springer-Verlag.

Howard, J. H., & Ballas, J. A. (1980). Syntactic and semantic factors in the classification of nonspeech transient patterns. *Perception and Psychophysics.* **28**, 431-439.

Howard, J., Mutter, S., & Howard, D. (1992). Learning by event observation. *Journal of Experimental Psychology: Learning, Memory and Cognition,* **18** (5), 1029-1039.

Irwin, F. W., Kauffman, K., Prior, G. and Weaver, H. B. (1934). On learning without awareness of what is being learned. *Journal of Experimental Psychology:,* **17**, 823-827.

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

Jacoby, L. L. & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General,* **110**, 306-340.

Jacoby, L. L. & Hay, J. F. (1997). Age-related deficits in memory: Theory and application. To appear in M. A. Conway (Ed.) *Theories of Memory II.*

Jacoby, L. L., Toth, J. P. & Yonelinas, A. P. (1993). Separating conscious and unconscious influences of memory: Attention, awareness and control. *Journal of Experimental Psychology: General,* **122**, 139-154.

Jacoby, L. L., Yonelinas, A. P., & Jennings, J. M. (1997). The relation between conscious and unconscious (automatic) influences: a declaration of independence. In J.D. Cohen & J. W. Schooler (Eds.) *Scientific Approaches to Consciousness.* Mahwah, NJ: Lawrence Erlbaum Associates Ltd.

Jones, G. V. (1990). Misremembering a familiar object: When left is not right. *Memory and Cognition*, **18**, 174-182.

Jones, G.V., & Martin, M. (1992). Misremembering a familiar object: Mnemonic illusion, not drawing bias. *Memory and Cognition*, **20**, 211-213.

Jordan, M.I. (1995). The organisation of action sequences: Evidence from a relearning task. *Journal of Motor Behavior,* **27**, 179-192.

Keele, S.W., Jennings, P., Jones, S., Caulton, D. & Cohen, A. (1995). On the modularity of sequence representation. *Journal of Motor Behavior,* **27**, 17-30.

Keppel, G, Saufley, W. H., & Tokunaga, H. (1992). *Introduction to Design and Analysis: A Student's Handbook.* New York: W.H. Freeman & Co.

Knopman, D. S., & Nissen, M. J. (1987). Implicit learning in patients with probable Alzheimer's disease. *Neurology,* **37** (5), 784-788.

Knowlton, B., Ramus, S., & Squire, L. (1992). Intact artificial grammar learning in amnesia: Dissociations of classification learning and explicit memory for specific instances. *Psychological Science*, **3**, 172-179.

Korsakoff, S.S. (1889). Etude medico-psychologique sur une forme des maladies de la memoire. *Revue Philosophique,* **28**, 501-530. (cited in: Schacter, D.L. (1987). Implicit Memory: History and current status. *Journal of Experimental Psychology: Learning, Memory and Cognition* , 13, 501-518.)

Kushner, M., Cleeremans. A., & Reber, A. (1991). Implicit detection of event interdependencies and a PDP model of the process. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, 215-220. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Levin, S. M. (1961). The effects of awareness on verbal conditioning. *Journal of Experimental Psychology:,* **61**, 67-75.

Lewicki, P. (1985). Nonconscious biasing effects of single instances on subsequent judgments. *Journal of Personality and Social Psychology*, **48**, 563-574.

Lewicki, P. (1986a) *Nonconscious Social Information Processing*. New York: Academic Press.

Lewicki, P. (1986b). Processing information about covariations that cannot be articulated. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **12**, 135-146.

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

Lewicki, P., & Hill, T. (1989). On the status of nonconscious processes in human cognition: Comment on Reber. *Journal of Experimental Psychology General,* **118** (3), 239-241.

Lewicki, P., Hill, T., & Czyzewska, M. (1994). Nonconscious indirect inferences in encoding. *Journal of Experimental Psychology: General*, **123**, 257-263.

Lewicki, P., Hill, T., & Czyzewska, M. (1997). Hidden covariation detection: a fundamental and ubiquitous phenomenon. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **23**, 221-228.

Lewicki, P., Hill, T., & Sasaki, I. (1989). Self-perpetuating development of encoding biases. *Journal of Experimental Psychology: General*, **118**, 323-337.

Lieberman, D. A. (unpublished manuscript). Reinforcement without awareness: The role of relaxation in implicit learning.

Manza, L., & Reber, A. S. (1992*). Inter and intra modal transfer of an implicitly acquired rule system.* Unpublished manuscript. (cited in Reber, A. S. (1993). *Implicit learning and tacit knowledge: An essay on the cognitive unconscious.* New York, NY, US: Oxford University Press.)

Marcel, A. (1980). Conscious and pre-conscious recognition of polysemous words: Locating the selective effects of prior verbal context. In R. Nickeron (Ed.) *Attention and Performance VIII.* Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Marcel, A. (1983). Conscious and unconscious perception: An approach to the relation between phenomenal experience and perceptual processes. *Cognitive Psychology*, **15**, 283-300.

Marescaux, P-J., Luc, F., & Karnas, G. (1989). Modes d'apprentissage selectif et nonselectif et connaissances acquises au control d'un processus: Evaluation d'un modele simule. *Cahiers de Psychologie Cognitive*, **9**, 239-264. (cited in Berry, D. C., & Dienes, Z. (1993). *Implicit learning: Theoretical and empirical issues.* Hove, England: Lawrence Erlbaum Associates, Inc.)

Massaro, D.W. (1990). An Information Processing Analysis of Perception and Action. In O. Neumann & W. Prinz (Eds) *Relationships between Perception and Action: Current Approaches.* (1990). Berlin: Springer-Verlag.

Mathews, R. C. (1997). Is research painting a biased picture of implicit learning? The dangers of methodological purity in scientific debate. *Psychonomic Bulletin and Review,* **4,** 38-42.

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 Human Experimental Psychology*, **40** (1-A), 135-165.

Mathews, R. C., Buss, R. R., Stanley, W. B., & Blanchard Fields, F. (1989). Role of implicit and explicit processes in learning from examples: A synergistic effect.

Journal of Experimental Psychology Learning, Memory, and Cognition, **15** (6), 1083-1100.

Maybery, M., Taylor, M., & O'Brien Malone, A. (1995). Implicit learning: Sensitive to age but not IQ. Special Issue: Cognitive development. *Australian Journal of Psychology,* **47** (1), 8-17.

Maylor, E. A. (1995). *Preserved repetition priming in old age. A world record attempt on the longest interval.* Paper presented to the Experimental Psychology Society. Cambridge.

Mayr, U. (1996) Spatial attention and implicit sequence learning: Evidence for independent learning of spatial and nonspatial sequences. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **22**, 350-364.

McAndrews, M. P., & Moscovitch, M. (1985). Rule-based and exemplar-based classification in artificial grammar learning. *Memory and Cognition,* **13**, 469-475.

McDowall, J., Lustig, A. & Parkin, G. (1995). Indirect learning of event sequences: The effects of divided attention and stimulus continuity. *Canadian Journal of Experimental Psychology*, **49**, 415-435.

McGeorge, P. (1990). *Implicit Learning and Human Memory: Theoretical and Practical Issues.* Unpublished Doctoral dissertation. University of Nottingham.

McGeorge, P., & Burton, A. M. (1990). Semantic processing in an incidental learning task. *Quarterly Journal of Experimental Psychology Human Experimental Psychology,* **42** (3-A), 597-609.

McGeorge, P., Crawford, J. R., & Kelly, S. W. (unpublished manuscript). The relationship between explicit memory, age, psychometric intelligence and an incidental learning task.

McGeorge, P., Crawford, J. R., & Kelly, S. W. (1997). The relationships between psychometric intelligence and learning in an implicit and explicit task. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **23**, 239-245.

McClelland, J.L. (1988). Connectionist models and psychological. *Journal of Memory and Language,* **27**, 107-123. (cited in Dienes, Z. (1993) Computational models of implicit learning. In D. C. Berry & Z Dienes (Eds.) *Implicit learning: Theoretical and empirical issues.* (pp. 81-112). Hove, England: Lawrence Erlbaum Associates, Inc.)

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.

Millward, R.B. (1981) Models of concept formation. In R. E. Snow, P. A. Frederico, & W. E. Montague (Eds.) *Aptitude, learning and instruction.* Hillsdale, NJ: Erlbaum. (cited in Reber, A. S. (1993). *Implicit learning and tacit knowledge: An essay on the cognitive unconscious.* New York, NY, US: Oxford University Press).

Millward, R.B., & Reber, A. S. (1972). Probability learning: Contingent-event sequences with lags. *American Journal of Psychology.* **85**, 81-98.

Millward, R.B., & Reber, A. S. (1968). Event-recall in probability learning. *Journal of Verbal Learning and Verbal Behaviour.* **7**, 980-989.

Morton, J. (1967). A singular lack of incidental learning. *Nature,* **215**, 203-204.

Muller, H.J., & Findlay, J. (1987). Sensitivity and criterion effects in the spatial cuing of visual attention. *Perception and Psychophysics*, **42**, 383-399.

Musen, G., & Squire, L. R. (1993a). Implicit learning of color-word associations using a Stroop paradigm. *Journal of Experimental Psychology Learning, Memory, and Cognition,* **19** (4), 789-798.

Musen, G., & Squire, L. R. (1993b). On the implicit learning of novel associations by amnesic patients and normal subjects. *Neuropsychology,* **7** (2), 119-135.

Myers, C., & Conner, M. (1992). Age differences in skill acquisition and transfer in an implicit learning paradigm. *Applied Cognitive Psychology,* **6** (5), 429-442.

Nelson, T.O. (1978). Detecting small amounts of information in memory: Savings for nonrecognized items. *Journal of Experimental Psychology: Human Learning, Memory and Cognition* , **4**, 453-468.

Neves, D. M. & Anderson, J. R. (1981). Knowledge Compilation: Mechanisms for the Automatisation of Cognitive Skills. In J. R. Anderson (Ed.) *Cognitive Skills and their Acquisition.* Hillsdale NJ: Erlbaum Associates.

Neumann, O. (1990). Visual Attention and Action. In O. Neumann & W. Prinz (Eds) *Relationships between Perception and Action: Current Approaches.* (1990). Berlin: Springer-Verlag.

Nickerson, R. S. & Adams, M. J. (1979). Long-term memory for a common object. *Cognitive Psychology*, **11**, 287-307.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, **84**, 231-259.

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

Nissen, M.J., Knopman, D. S., & Schacter, D. L. (1987). Neurochemical dissociation of memory systems. *Neurology,* **37**, 789-794.

Nissen, M. J., Willingham, D., & Hartman, M. (1989). Explicit and implicit remembering: When is learning preserved in amnesia? *Neuropsychologia,* **27** (3), 341-352.

Perruchet, P. (1994). Learning from complex rule-governed environments: On the proper functions of nonconscious and conscious processes. In Carlo Umilta & Morris Moscovitch (Eds.), *Attention and performance 15: Conscious and nonconscious information processing.* (pp. 811-835). Cambridge, MA, US: MIT Press.

Perruchet, P., & Amorin, M. (1992). Conscious knowledgeand changes in performance in sequence learning: Evidence against dissociation. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **18**, 785-800.

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

Perruchet, P., & Pacteau, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology General,* **119** (3), 264-275.

Perruchet, P., Vinter, A., & Gallego, J. (1997). Implicit learning shapes new conscious percepts and representations. *Psychonomic Bulletin and Review*, **4**, 43-48.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology.* **77**, 353-363.

Posner, M. I., & Keele, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology.* **83**, 304-308.

Postman, L. (1962). Rewards and punishments in human learning. In L. Postman (Ed.) *Psychology in the making.* New York: Knopf.

Rathus, J., Reber, A. S., & Kushner, M. (1990). *Implicit and explicit Learning: Differential effects of affective states.* (cited in Reber, A. S. (1993). *Implicit learning* *and tacit knowledge: An essay on the cognitive unconscious.* New York, NY, US: Oxford University Press)

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

Reber, A.S. (1976). Implicit learning of synthetic languages: The role of instructional set. *Journal of Experimental Psychology: Human Learning and Memory,* **2**, 88-94.

Reber, A. S. (1976). Implicit learning of synthetic languages: The role of instructional set. *Journal of Experimental Psychology Human Learning and Memory,* **2** (1), 88-94.

Reber, A. S. (1989a). More thoughts on the unconscious: Reply to Brody and to Lewicki and Hill. *Journal of Experimental Psychology General,* **118** (3), 242-244.

Reber, A. S. (1989b). Implicit learning and tacit knowledge. *Journal of Experimental Psychology General,* **118** (3), 219-235.

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

Reber, A. S. (1997). Implicit ruminations. *Psychological Bulletin and Review.* **4**, 49-55.

Reber, A.S. & Allen, R. (1978). Analogy and abstraction strategies in synthetic grammar learning: A functionalist interpretation. Cognition, **6**, 189-221.

Reber, A. S., Allen, R., & Regan, S. (1985). Syntactical learning and judgment, still unconscious and still abstract: Comment on Dulany, Carlson, and Dewey. *Journal of Experimental Psychology General,* **114** (1), 17-24.

Reber, A. S., Kassin, S. M., Lewis, S., & Cantor, G. W. (1980). On the relationship between implicit and explicit modes in the learning of a complex rule structure. *Journal of Experimental Psychology: Human Learning, Memory and Cognition*, **6**, 492-502.

Reber, A. S., & Lewis, S. (1977). Implicit learning: An analysis of the form and structure of a body of tacit knowledge. *Cognition,* **5** (4), 333-361.

Reber, A. S. & Millward, R. B. (1965). Probability learning and memory for event sequences. *Psychonomic Science*, **3**, 431-432.

Reber, A. S. & Millward, R. B. (1968). Event observation in probability learning. *Journal of Experimental Psychology:*, **84**, 85-99.

Reber, A. S. & Millward, R. B. (1971). Event tracking in probability learning. *American Journal of Psychology:*, **77**, 317-327.

Reber, A. S., Walkenfeld, F. F., & Hernstadt, R. (1991). Implicit and explicit learning: Individual differences and IQ. *Journal of Experimental Psychology Learning, Memory, and Cognition,* **17** (5), 888-896.

Reber, A. S. & Winter, B. (1994). What manner of mind is this? *Behavioral and Brain Sciences,* **17** (3), 418-419.

Reed, J., & Johnson, P. (1994). Assessing implicit learning with indirect tests: Determining what is learned about sequence structure. *Journal of Experimental Psychology Learning, Memory, and Cognition,* **20** (3), 585-594.

Reingold, E. M. & Merikle, P. M. (1990). On the inter-relatedness of theory and measurement in the study of unconscious processes. *Mind & Language,* **5**, 9-28.

Repp, B.H. & Williams, D. R. (1985). Categorical trends in vowel imitation. *Speech Communication.*, **4**, 105-120.

Richardson, J. T. E. (1993). The curious case of coins. *The Psychologist*, **6**, August, pp 360-366.

Rousel, L. & Mathews, R. (unpublished manuscript). THIYOS: A synthesis of rule based and exemplar based models of implicit learning. (cited in Berry, D. C., & Dienes, Z. (1993). *Implicit learning: Theoretical and empirical issues*. Hove, England: Lawrence Erlbaum Associates, Inc).

Rozin, P. (1976). The evolution of intelligence and access to the cognitive unconscious. *Progress in Psychobiology and Physiological Psychology*, **6**, 245-280.

Rubin, D.C., & Kontis, T.C. (1983). A schema for common cents. *Memory and* Cognition, 11, 335-341.

Sanderson, P. (1989). Verbalisable knowledge and skilled task performance: associations, dissociations and mental models. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **15**, 729-747.

Schacter, D.L. (1984). Toward the multidisciplinary study of memory: Ontogeny, phylogeny and pathology of memory systems. In L. R. Squire & N. Butters (Eds.), *Neuropsychology of Memory.* (pp. 13-24). New York: Guildford Press.

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 Science*, **608**, 543-571.

Seger, C. A. (1994). Implicit learning. *Psychological Bulletin,* **115** (2), 163-196.

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., Green, R. E. A., & Kolodny, J. A. (1994a). A critical examination of the evidence for unconscious (implicit) learning. In Carlo Umilta & Morris Moscovitch (Eds.), *Attention and performance 15: Conscious and nonconscious information processing.* (pp. 837-860). Cambridge, MA, US: MIT Press.

Shanks, D. R., Johnstone, T., & Staggs, L. (1997). Abstraction processes in artificial grammar learning. *The Quarterly Journal of Experimental Psychology,* **50A** (1), 216-252.

Shanks, D. R., & St John, M. F. (1994b). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences,* **17** (3), 367-447.

Sherry, D. F., & Schacter, D. L. (1987). The evolution of multiple memory systems. *Psychological Review*, **94**, 439-454.

Shimamura, A. P. (1986). Priming effects in amnesia: Evidence for a dissociable memory function. *Quarterly Journal of Experimental Psychology*, **38A**, 619-644.

Sloman, S. A., Hayman, C. A., Ohta, N., & Law, J. (1988). Forgetting in primed fragment completion. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **14**, 223-239.

Smith, J.D., & Kemler-Nelson, D.G. (1984). Overall similarity in adults' classification: The child in all of us. *Journal of Experimental Psychology: General*, **113**, 137-159.

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

Stadler, M.A. (1989). On learning complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **15**, 1061-1069.

Stadler, M.A. (1993). Implicit serial learning: Questions inspired by Hebb (1961). *Memory and Cognition,* **21**, 819-827.

Stadler, M.A. (1995). Role of attention in implicit learning. *Journal of Experimental Psychology: Learning, Memory and Cognition,* **21**, 674-685.

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

Squire, L. R. (1986). Mechanisms of memory. *Science*, **232**, 1612-1619.

Squire, L. R. (1987). *Memory and Brain..* New York: Oxford University Press.

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

Thorndike, E. L. (1932). *The Fundamentals of Learning.* Teachers College, Columbia University: New York.

Thorndike, E. L., & Rock, R.T. (1934). Learning without awareness of what is being learned or intent to learn it. *Journal of Experimental Psychology,* **11**, 1-19.

Tulving, E. (1983). *Elements of Episodic Memory.* Oxford: The Clarendon Press.

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

Turner, C., & Fischler, I. (1993). Speeded test of implicit knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **19**, 1165-1177.

Vokey, J. R., & Brooks, L. R. (unpublished manuscript). Taming the clever unconscious: Analogic and abstraction strategies in artificial grammar learning. (cited in McAndrews, M. P., & Moscovitch, M. (1985). Rule-based and exemplar-based classification in artificial grammar learning. *Memory and Cognition,* **13**, 469-475.)

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

Vokey, J. R., & Higham, P. A. (manuscript submitted for publication). Abstract analogies and positive transfer in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory and Cognition.*

Whittlesea, B. W., & Dorken, M. D. (1993). Incidentally, things in general are particularly determined: An episodic-processing account of implicit learning. *Journal of Experimental Psychology General,* **122** (2), 227-248.

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

Wright, R. L., & Burton, A. M. (1995). Implicit learning of an invariant: Just say no. *Quarterly Journal of Experimental Psychology Human Experimental Psychology,* **48A**, 783-796.