Active Processing in Implicit Learning

Vyvyan Huddy

Department of Psychology
University of Glasgow

Submitted for the degree of Ph.D. to the Higher Degree Committee of the Faculty of Social Sciences, University of Glasgow.

August, 2001

Acknowledgements

I express gratitude to my supervisor, Mike Burton, for his guidance and support. I also would like to thank my friends, family and my girlfriend, Jane.

This research was supported by an ESRC scholarship, R00429834476.

Abstract

This thesis is concerned with the phenomenon of implicit learning. Implicit learning occurs "when there is a performance increase on some task, without an associated increase in verbal knowledge about the causes of this performance increase" (Bright, 1993 - p9). In chapter 1, two theoretical interpretations of this type of learning are described. The first proposes that implicit learning reflects the operation of a unconscious learning system (e.g. Reber, 1989). The alternative episodic processing view (Whittlesea and Dorken, 1993) suggests that implicit learning occurs when there is an indeterminate relationship between the explicitly held knowledge acquired during training, and the way this knowledge is used in the test.

Two main experimental findings reported in the chapters 2 - 5 allow for a choice between the two main accounts mentioned above. Firstly, two implicit learning tasks (invariance learning [McGeorge and Burton, 1990] and sequence learning [Nissen and Bullemer, 1987]) demonstrate that learning is dependent on active processing of training stimuli. Secondly, findings from the invariant learning task indicate that episodic knowledge, rather than an abstract rule, is acquired in this type of learning. Both these findings are consistent with the episodic processing account of implicit learning, and not the separate system view. Furthermore, a specific prediction of the episodic account is also confirmed by the data reported in this thesis. This prediction is that the processing demands of the training and test periods must be consistent for successful performance.

This support for the episodic processing account of implicit learning is accompanied by a caveat in chapter 6. It is suggested that the episodic processing view is unsuitable for understanding the type of processes occurring in *all* implicit learning tasks. A possible resolution is offered, in the suggestion that broadening the theoretical scope to a more general consistency model may allow the wider experimental context of implicit learning to be explained.

Table of Contents

Chapter 1 Interpretations of implicit learning

1.1 Introduction

The role of conscious awareness in mediating fundamental learning processes has attracted a considerable amount of interest in recent years. This interest largely stems from the claim that learning can proceed without concurrent awareness of what is being learned, and this has been termed implicit learning. This claim has been interpreted as reflecting the operation of an unconscious learning system that operates separately from learning in more typical situations, where learning does occur with full awareness (e.g. Reber, 1989). An alternative account suggests that it is unnecessary to invoke the idea of an unconscious learning system to explain implicit learning. Instead, implicit learning results from the incidental task demands that are used in experiments that demonstrate learning without awareness. This thesis is concerned with the predictions made by each of these accounts and attempts to test these predictions.

The first chapter of this thesis is concerned with introducing the concept of implicit learning and its theoretical interpretations. The main points of the first chapter are summarised below -

- Firstly, a brief description of early conceptions of the relationship between learning and awareness is presented.

- Then, the concept of implicit learning will be introduced, elucidated by full description of each of the experimental tasks that are claimed to demonstrate implicit learning. This will allow the reader to appreciate the wide experimental context of implicit learning research.

- The experimental tasks will then be drawn together by three main attributes they have in common.

- Next, the methodological problem of determining when learning can be said to occur without awareness is considered.

- Once the methodological context and problems of implicit learning have been described, it is possible to turn to the theoretical interpretations of implicit learning. The three main accounts of implicit learning are considered in turn.

- Finally, a general overview of the main aims of this thesis is presented.

1.2 Learning without insight

Thorndike (1898) studied learning in animals using an apparatus he called the puzzle box, which consisted of a wooden crate with a door in the front panel. In the box there was a latch or a rope that when triggered, allowed the door to open. A bowl of food was placed outside the box that was visible through the slats. When a hungry cat was placed inside the box, it tended to struggle to get out by clawing at the sides of the box in order to reach the food. After some time, the cat would inadvertently pull the lever or rope, triggering the door mechanism, allowing it to escape and eat the food. Some time after, Thorndike placed the cat back in the box once again. It would repeat its struggles to escape and eventually pull the lever as before. There was, however, an important difference to this behaviour. The amount of time, or latency, it took for the correct response to be performed was generally shorter on the second trial than the first, shorter again on the third trial than the second, and so on. The time to escape became progressively shorter on each trial.

This experiment is quoted in virtually every introductory text on Psychology to illustrate the phenomenon of instrumental conditioning. That is, when a response to a stimulus is followed by a reward, that response becomes stronger, such as pulling the lever in the puzzle box. Further to this, Thorndike (1911) suggested that the reward would strengthen any response that preceded it and that no response had some essential property that set it apart from others to the animal, except subsequent reinforcement. In his view, reinforcement occurs automatically. If there had been an insight into this relationship, the cat would have pulled the lever after the initial trial, resulting in an abrupt drop in latency. Contrary to this, such a rapid reduction was not observed, after entering the box the cat would pull the lever with a generally decreasing latency on each successive trial. Indeed, Lieberman (1993) quotes Thorndike (1911) (Pg. 74) as suggesting:

> "The gradual slope of the time-curve . . . shows the absence of reasoning. They represent the wearing smooth of a path in the brain, not the decisions of a rational consciousness."

This conclusion appears to suggest that learning does not require insight, that any association between reward and response can be formed. Indeed, if reinforcement acts automatically, conscious awareness should not be necessary to mediate an association between reinforcement and a response. Greenspoon (1953) tested this prediction in an experiment involving verbal conditioning in college students. In an interview situation the students were given twenty minutes to say all the words that came to mind. Whenever the word produced was a plural noun, the experimenter would provide some reinforcement by saying "mmm-hmm". As the session progressed, the number of plural nouns increased, but the post experimental interviews revealed no awareness that only plural nouns were reinforced. This experiment suggested that reinforcement can proceed with no associated verbalised knowledge, and that the two effects can be dissociated.

Following the Greenspoon (1955) study, concerns were raised over the possibility that the subject could notice the strange attempts the experimenter was making to reinforce the plural nouns. Hefferline, Keenan and Harford (1959) carried out a similar conditioning experiment that was less likely to arouse the suspicion of the subject. In their study, electrodes were attached to various locations on the subjects' body to assess muscular movements. Subjects were told they would be participating in a study designed to assess the effects of stress on body tension, and to this end randomly alternating periods of harsh noise or soothing music were played to subjects. In reality, and unbeknown to the subject, the movement of a small muscle in the subjects' left thumb controlled the alternation of the sound. This movement was so small it could not be detected visually by the subject, but the electrode could detect it. Over the course of the session, there was a marked increase in the contractions made of the thumb by the subject. In the post experimental interview, the subject was unable to verbalise anything about the relationship between the sound and the movement of the thumb. More recently, Lieberman, Sunnucks and Kirk (1998) reported similar learning without awareness using a highly convincing cover task. Subjects were instructed that they would be taking part in a ESP experiment, and the task was to say which of two words the experimenter was thinking about. In fact, the reinforcement contingency was linked to how loudly the subject spoke when responding. The probability of reinforced responses grew as the session progressed, even though the subject reported no knowledge of the relationship between their responses and trial outcomes.

It appears then, that there is some support for Thorndike's suggestion that the association between a reinforcement and response is formed automatically. Furthermore, there is evidence which indicates that certain associations can be hard to verbalise, prompting the suggestion they can be formed without awareness. These experiments involve very simple associations between a response and the reinforcement. This leads to the question of the generality of these findings to more complex learning in humans. The opposite end of the spectrum in terms of complexity, are skilled activities such as driving a car.

1.3 The Acquisition of Skills

When one considers a skilled behaviour it is clear that attempting to describe one's actions is not always possible. This inability to describe the basis of a skilled behaviour occurs because this action can be performed automatically, with no apparent conscious guidance. Hasher and Zacks (1979) have described such automatic processes as effortless, fast, and operating outside of attentional control. Does the existence of these processes represent evidence for learning without awareness of complex behaviours? Simply because skilled behaviour can be performed with so little conscious control does not imply that these behaviours were always carried out in this manner. When learning to drive, for example, a great deal of effort is required to perform actions that will subsequently become automatic. One approach to this problem has been to set a distinction between procedural and declarative knowledge. Whilst declarative knowledge is always reported easily, procedural knowledge is acquired when declarative knowledge is transformed into procedural and the verbalisation of knowledge is lost.

Fitts' (1964) model of skill acquisition incorporates a progression from declarative to procedural knowledge. Within Fitts (1964) account there are three stages of skill development. Knowledge is initially explicit and rule based requiring a large amount of attentional resources, with associated performance tending to be error prone and slow. Refinement of the performance strategy follows, with progression to an associative stage. Here, appropriate strategies are strengthened on the basis of feedback, whereas inappropriate features are weakened. In the final stage, the components of the performance strategy then become so highly practised that they require little attentional resource to guide them. This account has been taken forward by Shiffrin and Schneider's (1977) automatic versus controlled processing view, Anderson's (1982) ACT skill acquisition model and Logan's (1988) instance theory of automisation.

So, it appears that complex behaviour may, in the initial stages at least, be acquired through the application of explicit strategies. This notion of a progression from a set of declarative knowledge that becomes increasingly procedural does not, however, account for a range of findings in the implicit learning literature. This literature suggests that learning, under certain circumstances, can proceed in a manner similar to that originally suggested by Thorndike (1911).

1.4 Implicit Learning

Unlike the progression from declarative to procedural knowledge described above, 'implicit learning' appears to proceed from the very start with no associated explicit knowledge. There is a large and expanding literature claiming to have demonstrated such "implicit learning" (Reber, 1967, Berry and Broadbent, 1984, Nissen and Bullemer, 1987, McGeorge and Burton, 1990).

There are numerous definitions of implicit learning in the literature, each placing emphasis on different aspects of the phenomenon, and these will be considered below. At the most general though,

> "implicit learning occurs when there is an observed increase in performance on some task, without an associated increase in verbal knowledge about the causes of this performance increase" (Bright, 1993 - p9).

Dienes and Berry (1997) contrast implicit learning with explicit learning, suggesting that explicit learning occurs when the stages of the development of a given set of knowledge are readily stated, such as learning how to solve an arithmetic problem by hypothesis testing. Dienes and Berry (1997) claim that implicit learning occurs when the stimulus material is observed or memorised, and the structure is not directly experienced by the subject.

1.5 Experimental Approaches to Implicit Learning

In the aforementioned literature, there are other aspects that have been associated with implicit learning which would be described better following some experimental context. The description below will focus on the initial findings of implicit learning tasks without going into the extensions to these early findings, as these will be dealt with at a later stage. Various tasks have been used to investigate implicit learning, the most prominent of these being Artificial Grammar Learning (from here on AGL).

Artificial Grammar Learning

The earliest evidence Reber used to support the existence of implicit learning was the findings of work on the AGL task (Reber, 1967). In a AGL study, subjects are asked to memorise a series of grammatical strings of letters generated by a finite state grammar (see below for an example).

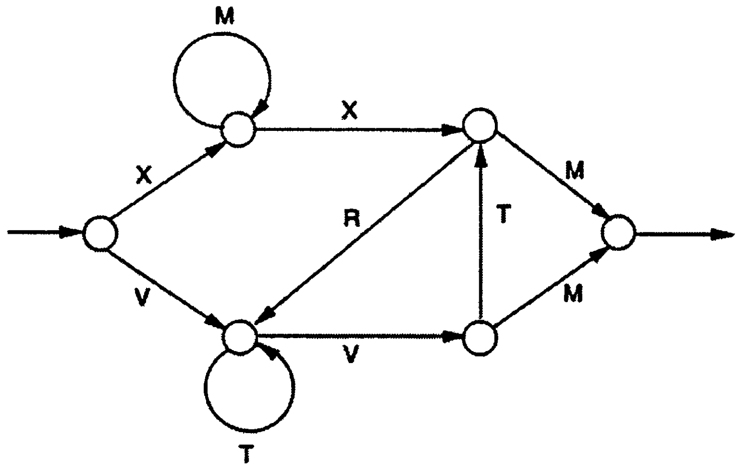

These strings of letters can be generated by this grammar by entering the diagram from the left side and following through pathways until the exit on the right (e.g creating strings such as VTVTM or XMXRVM and so on). This produces strings that conform to the grammar, called grammatical strings. Ungrammatical strings could also be produced by violating the order of working through the system (e.g by producing the string VXM or VRTM). Initially, subjects memorise a series of grammatical strings. Following this, subjects are informed of the existence of a rule structure that constrains the order of the letters and are asked to classify novel grammatical and non grammatical strings. In an initial study, Reber (1967) found that subjects could classify novel grammatical strings significantly above chance levels. In addition, Reber (1967) found that subjects were unable to verbalise the rules of the grammar, and subjects claimed that they had no basis on which to make these decisions.

In a variation of the original task, Reber and Lewis (1977) demonstrated sensitivity to the grammar in a different test task. Following exposure to grammatical strings as before, subjects demonstrated an increased ability to reorder scrambled strings back into their original grammatical order, demonstrating some knowledge of the grammatical structure.

Dynamic Systems Tasks

A second type of task that has provided a substantial amount of data in support of implicit learning is the dynamic systems task. This task presents subjects with a stimulus set that is complex in nature and is also more realistic than the AGL task. The subject typically interacts with a computer program which takes input from the subject and provides a response value that is altered according to an underlying rule.

An example of a dynamic systems task is the "sugar production task", as used by Berry and Broadbent (1984). In this task, subjects take on the role of a manager of a sugar production factory. The task is to reach and maintain the level of sugar output by varying the numbers of workers involved. The subjects enter the number of workers they consider will maximise output, the sugar production total is then updated and this process is repeated over the session. The sugar production level relates to the number of workers by an underlying rule that correlates the current level of production with the previous output and the number of workers. Berry and Broadbent (1984) found that subjects were able to improve their ability in controlling the sugar production task. Like the AGL studies, subjects reported that they were unable to describe the basis of their actions, or they were "unable to put into words" how they went about the task.

The Serial Reaction Time (SRT) Task

A further line of support for the process of implicit learning comes from work by Nissen and Bullemer (1987) and Lewicki et al.. (1987), who introduced a simple but effective technique called the serial reaction time task. There are numerous versions of serial reaction time tasks in the literature, and the most prominent of these will be considered below.

In the Nissen and Bullemer (1987) version of the task, a stimulus is presented in one of four locations and the subject's task is to respond to the location of the stimulus with a corresponding response key as fast as possible. The stimulus moves on to another of the four locations, the subject makes a response, the stimulus again moves to a further location, and so on. The instructions to the subject emphasise only that they should respond as quickly and accurately as possible, but in fact the location changes follow a repeating sequence. Subjects showed a rapid decrease in response time to the location changes as training progressed. Moreover, when the repeating sequence was changed to a random pattern, reaction times increased substantially. This indicates that subjects had become sensitive to the structure of the sequential material. As in the grammar learning studies, subjects were unable to verbalise anything about the sequence and often reported inaccurate information.

Lewicki et al. (1987) used a slightly different procedure involving a visual search task. Subjects had to indicate, by pressing a button, which of four quadrants contained a target digit. The trials were separated into "simple" and "complex" trials. The "simple" trials were arranged so that the target stimulus was the only item on the screen and was easy to detect, while in complex trials the target was presented among a field of 35 distractors hence making a decision more difficult. The task was structured into sections of seven trials, six "simple" trials followed by a seventh "complex" trial. Four of the six "simple" trials predicted the location of the seventh "complex" trial, and each location of the target on the "complex" trial was associated with a unique conjunction of locations on the simple trials. After a period of training on these relations between "simple" and "complex" locations, response time to detect the target on the "complex" trial had decreased significantly. Following training, Lewicki et al. changed the relations between the simple and complex trials so that the target location in the "complex" trial was diagonally opposite to its original location. This change resulted in a large increase in the response time to detect the target object in the "complex" trial. Like the Nissen and Bullemer (1987) task, this indicated learning of the sequential material presented during training. Also like Nissen and Bullemer (1987), subjects did not report any useful information about the rules used to determine the relationship between "simple" and "complex" trials.

The Invariant Learning Task

In an attempt to simplify the relational structure of the stimuli, while maintaining the complexity of individual training instances, McGeorge and Burton (1990) introduced the Invariant Learning task. In the original McGeorge and Burton (1990) study, participants were asked to perform some task (such as arithmetic) on 30 four-digit numbers. Each four digit number contained a "3" digit; this fact was not divulged to the participants. In a subsequent recognition task, 10 pairs of four-digit numbers were presented to subjects one at a time. In each pair, one number contained a "3" (*positive)* and one did not (*negative).* Subjects were then falsely told they had seen one of the numbers in the study phase and they must indicate which of the numbers they had seen. A robust effect was observed for participants choosing the positive over the negative. This effect was seen to persist over different encoding tasks, and when the format of test materials was changed from study to test phase (i.e. the training items were seen in digit format and test items appeared as words). In a post task question session, subjects were unable to report anything pertinent to the invariant digit. They concluded that performance on this task was driven by semantic knowledge of the invariant feature, and that this knowledge was implicit.

Bright and Burton (1993) extended the findings of McGeorge and Burton (1990) by using alternative clock face stimuli in place of the digit strings. The invariance rule was a time range that was consistent over all of the training items so that all of the training items were clock times that varied between 6 and 12 o'clock. As in digit invariance these times were referred to as positives. At test, novel positives were set against negatives (clock times between 12 and 6 o'clock) in a two alternative forced choice. Bright and Burton (1993) demonstrated a preference for positives at test, with no associated verbalisation of the invariance.

1.6 General definitions and attributes of Implicit Learning

It can be seen from these tasks that implicit learning is demonstrated in a number of different experimental contexts. Indeed, there are further demonstrations of effects that have been termed implicit learning (e.g Reber and Millward, 1968;

Shanks, Green and Kolodny, 1994). This proliferation of tasks has resulted in many different perspectives on implicit learning and hence a large number of different definitions. As noted by Frensch (1998), there are "literally dozens of definitions that have been offered and continue to be offered in the literature". Rather than replicating the list of definitions reported by Frensch (1998), definitions will be quoted throughout the current review according to the theoretical context.

The main experimental tasks that claim to demonstrate implicit learning are diverse in methodology. However, there are three main attributes that can be associated with implicit learning, these are that -

- **It occurs without accurate verbalisation of the knowledge used in the test.**

  In all of the experimental tasks described earlier on, subjects performed at above chance levels on the performance test without being able to describe the source of their success. This is the aspect of implicit learning that has provoked a heated debate, and a great deal of effort has been expended on attempting to define criteria for judging whether learning can proceed unconsciously.

- **It is revealed by sensitivity to the structural properties of a stimulus domain.**

  The stimuli used in implicit learning tasks are generated using a rule structure which determines the relationship between the elements that make up the training material. Subjects show sensitivity to this structure by demonstrating facilitation when responding to this structure (e.g sequence learning) or by exhibiting a preference for items that conserve the structure in a test (e.g. AGL). It is important to note that this statement is neutral with respect to the information used by subjects to demonstrate sensitivity.

- **It tends to be associated with incidental training conditions.**

  During training, subjects in implicit learning tasks are not directed towards the underlying structure of the stimuli. The training task usually involves an observation or memorisation task, rather than deliberate hypothesis testing. A caveat should be made at this point that the incidental nature of the training conditions does not imply that subjects are inactive during training. Indeed, Wright and Whittlesea (1998) state that "the absence of hypothesis testing... do not make the learner a passive or unselective recipient of structure". Thus, incidental training should be understood as describing the methodology used in implicit learning tasks, rather than defining the cognitive processes that occur during training.

Of these three characteristics, the first has resulted in the most extensive debate and controversy. This debate centres on the validity of the methodology used to determine the degree of conscious awareness in implicit learning studies. This issue will be addressed in the following section.

1.7 Criteria for verifying when learning occurs without awareness

> "implicit learning experiments have universally adopted the dissociation logic of attempting to demonstrate learning in absence of any detectable degree of awareness"
>
> Shanks and St. John (1994) – p 370

In their review of the implicit learning literature in 1994, Shanks and St John (1994) identified a weakness in implicit learning studies. It had been assumed that if subjects could not verbalise the structure of the stimulus domain to which they have been exposed, and yet they showed sensitivity to that domain in a test, they must have learned the structure of the domain unconsciously (e.g. Reber, 1989). This relies on the assumption that the awareness test accesses any conscious knowledge that the subject may have. However, Shanks and St John (1994) argued that attempting to determine the awareness of subjects *during* training by examining the content of verbal reports *after* testing is not the best way to establish unconscious learning. Nonetheless, they argue that if this methodology is to be used, certain criteria should be met before any findings can be claimed to be valid.

Two criteria for characterising unconscious learning

Shanks and St John (1994) suggested that before making an assumption that the awareness test was accessing all conscious knowledge, a more stringent method for determining the validity of tests of awareness should be used. This method involved assessing the awareness tests used in implicit learning experiments according to two criteria. The first of these concerns the ability of the awareness test to access the critical information that the subject acquired during training. Shanks and St John called this the Information Criterion. The second criterion addressed the problem of the sensitivity of the test of awareness, this was called the Sensitivity Criterion.

According to Shanks and St John (1994), the Information Criterion requires that "it must be possible to establish that the information the experimenter is looking for in the awareness test is indeed the information responsible for the performance changes". That is, if the experiment requires a subject to learn information x, but unbeknown to the experimenter other information y allows the subject to perform at above chance levels, probing the subject about information x may not yield any useful verbal report. This would result in an incorrect conclusion that the subject had unconsciously learned information x. It is necessary, therefore, to determine whether the test of awareness is assessing the same information that the subject uses in the test task.

The Sensitivity Criterion requires that the test of awareness is sensitive to all the conscious knowledge that the subject possesses. This criterion ensures that performance on a test task (e.g. classification test in AGL) is not superior to that of the awareness test simply because the classification test is more sensitive, and therefore results in a higher performance value. Shanks and St John (1994) considered a situation where a single source of knowledge is accessed by a classification test and an awareness test. Above chance performance occurs on classification, with no associated verbalisation on the awareness test. If test sensitivity is not taken into account this would lead to the incorrect conclusion that unconscious knowledge is used on the classification test. A more accurate theory is that conscious knowledge is present and is an influence on the classification test, whereas the less sensitive awareness test fails to detect its presence.

Shanks and St John (1994) propose that this inability of the awareness test to access the relevant information may occur as a result of the difference between the context of the awareness test and the training period. Considering the Nissen and Bullemer (1987) task, the sensitivity to the sequence is normally demonstrated by the change in response time to the location of the stimulus. The sequence is learned by interacting with the sequence in a very similar way, by responding to the stimulus as it moves from location to location. The awareness test used in one version of this task (Lewicki et al., 1987) involved simply asking subjects if they noticed a sequence in the stimulus movements. Since the subject was now interacting with the experimenter and not the stimulus itself, it is clear that there was a substantial change in context. This change in context could contribute to the insensitivity of the awareness test

Shanks and St John (1994) suggest that the Sensitivity Criterion could be met by ensuring that the performance test and awareness test are made as similar as possible, with only the task instructions separating them. This would be more likely to encourage the subject to retrieve as much conscious knowledge as possible. Under these experimental circumstances, a dissociation between performance and verbal report would be more compelling evidence for unconscious learning.

Shanks and St John (1994) extensively reviewed the implicit learning literature with reference to the Information and Sensitivity criteria. They concluded that no study had met these criteria to a level that unequivocally demonstrated the existence of unconscious learning processes. For example, a substantial part of the AGL literature is concerned with determining what knowledge subjects use when making classification judgements. Shanks and St John (1994) report that the evidence for the use of rule knowledge is weak, and other knowledge is more likely to be employed (see later on for a description of these studies [sect 1.8]). This means that asking subjects whether they noticed a rule structure in the stimuli fails to meet the information criterion because this is not the knowledge they used in the test.

Subsequent work carried out by Jimenez, Mendez and Cleeremans (1996) has used awareness measures that are closer to those required to meet the Shanks and St John (1994) criteria. Jimenez et al. (1996) used direct and indirect tests to test the influence of conscious and unconscious influences. Here tasks are matched in all respects, such as context and demands, except instructions differ in the two conditions. In direct tests subjects are instructed to use whatever conscious knowledge they may have. In indirect tests the instructions do not refer to any conscious knowledge. Jimenez et al. (1996) assumed that the direct test would be more sensitive to conscious knowledge and greater sensitivity to some aspect of the stimulus in the indirect test must be a result of unconscious influences.

Jimenez et al. (1996) used these tests in a sequence learning task and claimed to have demonstrated the influence of unconscious learning. However, Shanks and Johnstone (1999) carried out a similar experiment using direct and indirect tests in a sequence learning task which appears to contradict the findings of Jimenez (1996). They claimed that their direct measures of awareness revealed that conscious knowledge of sequences was fully accessible on these objective tests. So it appears that using direct and indirect measures does not result in conclusive evidence of unconscious learning in all cases.

Verifying unconscious learning with a subjective criterion

Dienes and Berry (1997) argued in favour of applying alternative criteria for unconscious learning that are closer to everyday notions of consciousness. They suggest that this leads to more positive conclusions on presence of unconscious learning processes than are allowed by the two criteria proposed by Shanks and St John (1994).

As a starting point Dienes and Berry (1997) suggested that criteria used in the subliminal perception literature may be useful when examining unconscious learning. In a subliminal perception experiment (e.g. Marcel, 1983), subjects are presented with a series of trials where a subliminal stimulus is presented for a few milliseconds preceding a target stimulus on some of the trials but it is absent on other trials. This subliminal stimulus is intended to facilitate the subjects' response to the target, resulting in priming. Following each trial, the subject indicates whether the stimulus was present or not, and how sure they were about the response. Two criteria were developed to assess the subjects' awareness

(Cheesman and Merikle, 1984); the *subjective threshold* and the *objective threshold.* The subjective threshold is met when the subject believes that they are guessing, but still shows above chance performance on discriminating the presence of the stimulus. The objective threshold is met when the subject shows chance performance on discriminating the absence or presence of the stimulus, while they still show priming in response to the target stimulus, indicating that the preceding subliminal stimulus still affected their behaviour.

In the context of implicit learning, Dienes and Berry (1997) suggested that the knowledge used could be said to be below an objective threshold, if a cued recall task reveals chance levels of performance. This assumes that a cued recall task directly measures knowledge that a subject must have used to perform at above chance on classification. To have knowledge that is below a subjective threshold, subjects would have to lack metaknowledge.

Dienes and Berry (1997) proposed that determining a lack of metaknowledge could be divided into two further criteria, the *guessing criterion* and the *zero correlation criterion.* The guessing criterion is met when subjects report that they are guessing and are relying on random selection of the test items. The zero correlation criterion is met when an analysis of the subject's responses reveals that there is no association between their responses and the confidence they have in these responses. Dienes and Berry (1997) suggest that these criteria may correspond to everyday conceptions of unconscious learning.

Like the previous efforts of Shanks and St John (1994), Dienes and Berry (1997) reviewed the implicit learning literature with reference to the criteria they had defined. They claimed that some demonstrations of AGL may reveal the use of knowledge that is below a subjective threshold. They claimed the subjects often classify grammatical strings substantially above chance although they believe that they are guessing. This would appear to meet the guessing criterion described above. Dienes and Berry (1997) presented further evidence from AGL which suggests that the zero correlation criterion may also have been met. An analysis of the confidence judgements of subjects who carried out an AGL task revealed that subjects were just as confident about incorrect decisions as correct decisions. As a result of these findings, Dienes and Berry (1997) argued that construing implicit learning in terms of a subjective threshold criterion may provide evidence for unconscious learning effects.

The limitations of setting criteria for unconscious learning

The assumption of *exhaustivity*, suggested by Riengold and Merikle (1988), addresses the same problem as the Shanks and St John's (1994) sensitivity criterion. A measure of explicit knowledge is said to be exhaustive if it captures all of the relevant knowledge that the subject has. There is a potential problem with this assumption that has been neatly described by Neal and Hesketh (1997) –

> "If the measure of explicit knowledge is not exhaustive, it will always be possible to claim that any observed dissociation is misleading because some unmeasured form of explicit knowledge may be responsible for learning."
>
> Neal and Hesketh (1997) - p 31

This observation leads to the conclusion that the Sensitivity criterion cannot be met if an exhaustive measure of explicit knowledge does not exist. In place of developing such a test of explicit knowledge, Dienes and Berry (1997) propose the subjective threshold criterion, that is described above. This measure of unconscious learning produces more positive conclusions on the presence of unconscious learning in humans than the sensitivity criterion proposed by Shanks and St John (1994). This contradiction highlights a problem with the use of criteria for verifying unconscious learning, to quote Neal and Hesketh (1997) - (p 31) -

> "because different measures produce different results, and there is no objective criterion for choosing a measure in the first place, this type of experimental procedure is empirically indeterminate."

Even attempts to satisfy such criteria are not free of problems. Shanks and St John (1994) suggest that if the context during the training phase and the test phase could be made as similar as possible, the sensitivity criterion has a better chance of being met. However, this causes a further problem because this new measure could be influenced by implicit as well as explicit knowledge. A measure of explicit knowledge that is influenced by implicit knowledge fails to meet the assumption of *exclusivity.* So, experimental approaches that attempt to meet the sensitivity criterion may be flawed by their lack of exclusivity in the same way that verbal report measures are impaired by their deficiency in exhaustivity. Neal and Hesketh (1997) suggest that alternative measures should be investigated to detect separate unconscious and conscious influences on task performance.

Using intention as a route to verifying unconscious learning

Neal and Hesketh (1997) argue that the process dissociation procedure (Jacoby, 1991) could be used to dissociate the effects of implicit and explicit learning. Jacoby's (1991) process dissociation procedure can be shown to separate the influence of intentional and non intentional processes on task performance. This is achieved by examining the effect on task performance when intentional and non intentional processes are allowed to act in concert or where these processes act in opposition. For example, subjects could be presented with a list of words, followed by two tests. In one test, subjects recall as many words as possible from the list, this is the inclusion condition. In the other condition, subjects generate as many words as they can that were *not* on the list, this is the exclusion condition. Any words from the original list that were nevertheless generated in the exclusion condition must be the result of non intentional processes. Although this is a simplification of the process dissociation procedure, it elucidates that non intentional effects of memory can be indexed experimentally. It may be that implicit knowledge may exert a similar non intentional influence on performance in a classification task. Neal and Hesketh (1997) argue that a procedure similar to process dissociation could conceivably detect this. They suggest that using intention in this way to reveal implicit processes does not rely on the sensitivity of the awareness test, and is, therefore, a more appropriate way to verify unconscious learning.

Buchner, Steffens and Rothkegel's (1997) adapted the process dissociation procedure to sequence learning in an attempt to separate conscious and unconscious influences. They revealed that intention to learn the sequence increased performance on an awareness test, but left task performance unaffected. This would indicate that intention can have alternative influences on unconscious and conscious knowledge bases, as suggested by Neal and Hesketh (1997).

In an AGL task, Higham, Vokey and Pritchard (2000) exposed subjects to letter strings produced by two different grammars (GA and GB), and asked subjects to rate strings for grammaticality in two conditions. Firstly, in-concert conditions where the strings were rated as consistent with either grammar, or opposition conditions where the strings were rated as consistent with only one of the grammars (in this case GB). Higham et al. (2000) reported evidence of controlled processing, as subjects could successfully reject items from GA and non grammatical items in the opposition condition. This sensitivity was removed by the introduction of a time deadline at test, as in these circumstances subjects showed a slight tendency to incorrectly classify GA items as GB items. This effect was interpreted by Higham et al. as evidence for automatic influences at test, as subjects under time deadline conditions could not prevent the bias to select GA items.

Redington (2000) questioned these findings on the basis that these results can be interpreted as the influence of a single controlled process. Redington (2000) suggested that the two grammars Higham et al. used were very similar, but the non grammatical items were dissimilar from these. Hence, participants may have mistakenly identified more grammatical items than non grammatical as being from the other grammar as a result of this similarity. Indeed, Dienes et al. (1995) carried out a similar experiment to that of Higham et al., except the non grammatical items were similar to the two grammars presented to subjects. Although this study found controlled influences in grammaticality decisions, no automatic influences were demonstrated.

Perruchet, Gallego and Vinter (1997) offer an alternative interpretation of the process dissociation procedure. They suggest that subjects recall the items in the test, but as a result of impoverished memory, subjects fail to remember the spatiotemporal context of encoding. Thus, even when exclusion retrieval conditions are used, subjects recall the items from the training set and use them in the test.

The process dissociation procedure makes some very strong assumptions, all of which have been criticised in the literature. For example, process dissociation relies on the assumption that task performance is a mixture of implicit or explicit processing components that are assumed to be stochastically independent. Joordens and Merikle (1993) suggested that it would be equally plausible to assume that there is redundancy between explicit and implicit task components, so that situations in which conscious processes operate are a subset of situations where unconscious processes operate. This account makes different estimates of unconscious influences on task performance in comparison to the assumption of independence. There are other models of how conscious and unconscious influences may interact in the process dissociation task (e.g Gardiner and Java, 1993), which make further estimates of unconscious influences.

More generally, Redington (2000) has criticised the use of intention as a route to separating conscious and unconscious influences on task performance. He argues that this position ignores the data from verbal report studies and other measures of conscious awareness. Although these data have been shown to be questionable, Redington (2000) suggests that it is still in need of further explanation and any theory or account of implicit learning that offers no explanation of the lack of verbal reports is unsatisfactory.

Concluding remarks on awareness measures

Redington's (2000) comments about the importance of verbal report data highlights the way in which research on implicit learning, prompted by Shanks and St John (1994), has increasingly focused on finding the methodological solutions to the information and sensitivity criteria. The difficulties in satisfying these criteria have resulted in the introduction of the concept of intention. However, verbal reports can, in some circumstances, produce useful data for understanding implicit learning. For example, Mathews et al. (1989) asked subjects in an artificial grammar learning study to produce verbal reports on how another subject should perform the classification test. These reports were used by naïve subjects in a classification test to some effect showing that verbal reports can produce useful data. Furthermore, the attempts to use intention as a means of isolating conscious and unconscious influences has not yet proved successful (e.g. Redington, 2000 commentary on Higham and Vokey, 2000). Many of the attempts to demonstrate learning in the complete absence of awareness may be unsuccessful because of the nature of the knowledge acquired in implicit learning tasks. In their review, Dienes and Berry (1997) indicate that while the knowledge acquired in implicit learning tasks is not accessible through verbal reports, some knowledge is revealed by more direct tests such as cued report tests. Indeed, Seger (1994) states that knowledge in implicit learning tasks is often on

> "the fringe of consciousness: a field of relatively unarticulated, vague experience, neither fully accessible to consciousness, nor fully separate" (p 421).

Hence, if the knowledge used in implicit learning tasks is not held in separate forms, using the intentional against unintentional manipulation may not yield conclusive results. Moreover, the finding that the knowledge acquired in implicit learning tasks is on the fringe of consciousness, rather than totally inaccessible, is very important for the development of theories of implicit learning. The various theoretical interpretations of implicit learning will be considered in the following section.

1.8 Theoretical Accounts of Implicit Learning

In a review of implicit learning, Reber (1989) concluded that implicit learning tasks such as artificial grammar learning reveal the operation of an unconscious learning system that yields a knowledge base that is "abstract and representative of the structure of the environment" (p 219). The main focus of research on implicit learning has been directed towards either proving or refuting the existence of such an unconscious learning system. Indeed, Mathews (1997) described the current state of implicit learning research as "an ongoing debate between believers and non believers in the existence of a powerful cognitive unconscious" (p 38).

The parallel systems view of implicit learning can be contrasted with two other accounts of implicit learning, the episodic processing account (Whittlesea and Dorken, 1993) and the subjective unit formation account (Perruchet and Gallego, 1998). These frameworks have been taken as representing a view that questions the validity of implicit learning, that of "non believers". However, this is not an accurate characterisation of these accounts. They do not refute the fact that subjects become sensitive to the rule structure of a particular set of stimuli without associated verbalised knowledge. They do question, however, the validity of the claim that these findings suggest that subjects abstract rules from the stimuli they were presented with using a parallel learning system. Instead, they suggest subjects acquire some other, non rule based knowledge that incidentally allows them to become sensitive to the rules in a test. These alternative positions on the basis of implicit learning will be considered below.

In the preceding section, it was reported that much of the experimental evidence suggests that the knowledge acquired in implicit learning tasks is "on the fringe of consciousness" rather than totally inaccessible. If the idea that parallel learning systems are responsible for dissociations between performance measures and verbal reports, one might expect complete inaccessibility of knowledge acquired by the unconscious learning system. As noted earlier, demonstrating learning in the absence of awareness has proved very difficult indeed. For this reason, proponents of parallel learning systems cite other evidence which is taken from two sources, neuropsychology and dual task studies. This evidence will be considered in detail in the next section.

1.8.1 A parallel systems account of implicit learning

The parallel system account can be placed within two main theoretical contexts. Firstly, there is the cognitive evolutionary standpoint, taken by Reber (1989) and Mathews and Rousell (1989). Secondly, some authors consider implicit learning to be similar to the kind of learning that takes place in connectionist models (e.g. Cleeremans, Destrebecqz and Boyer, 1998).

Reber (1989) argued that the evolutionary ancestors of humans showed learning of covariations between stimuli in the environment. This is a primitive form of learning that simply accumulates information about the environment. Reber (1989) suggests that there is "no reason to suppose that these presumably adaptive mental capacities ought to have been lost" (p 230). In other words, this primitive form of learning still exists in humans. According to Reber (1989), these processes predate the evolution of consciousness, as consciousness arrived late on the evolutionary scene. For this reason, these simple learning processes remain impenetrable to conscious inspection. This kind of learning, in Reber's (1989) view, is very similar to the slightly more sophisticated learning that takes place in implicit learning studies (e.g. AGL). As stated earlier, Reber (1989) claims that this system is involved with the acquisition of rules that are represented in a symbolic and abstract form. The evidence for this assumption is presented below.

This position purporting the existence of symbolic rules that are unavailable to consciousness is not tenable according to Cleeremans (1998). Cleeremans (1998) argues that "because these expressions are static and exist independently of the processor that interprets them, they are automatically available to outside inspection" (p 201). So, unlike Reber (1989), Cleeremans (1998) questions the validity of the assumption that implicit learning processes result in the representation of the rules that constrain a given stimulus set (e.g. the grammar in AGL studies).

Instead, Cleeremans (1998) notes that connectionist models may exhibit rule like behaviour, without representing rules themselves, so making these models a better starting point for understanding implicit learning. Connectionist networks are collections of small units that are joined by links. These links strengthen or weaken as each unit becomes active, and this allows the network to process the input and output of information. In this view, "implicit learning is a by product of processing, and involves changes in the very structures that drive processing (the connection weights between units)" (Cleeremans, 1998, p226). Furthermore, these models do not represent knowledge explicitly, it is represented in a distributed manner across the network.

A drawback of this account is that it does not specify how explicit cognition may be represented. As Neal and Hesketh (1997) note, the logical extreme of these arguments is that people never have introspective access to information of any form. Cleeremans (1998) argues that symbolic representations are available to conscious inspection, so it may be inferred that these systems are responsible for explicit learning within this view.

Although the parallel learning system view can be set within a couple of theoretical contexts, the same evidence is used to support the concept of separate learning systems. Apart from dissociations between verbal report and performance measures, which was covered in the previous section, separate systems accounts cite neuropsychology and dual task studies. The neuropsychological evidence will be considered first.

Evidence from neuropsychology

Knowlton and Squire (1992) compared artificial grammar learning in normal and amnesic subjects. Amnesics performed as well as control subjects on grammaticality judgements, but were impaired on classification and recognition when explicit comparison to prior exemplars was emphasised. These deficiencies were attributed to impaired explicit retrieval in the amnesics, with intact grammaticality judgements drawing on implicit knowledge of the grammar. In further papers, Knowlton and Squire (1994, 1996) demonstrated that Amnesic patients were not only impaired on recognition judgements but their knowledge of fragments of grammatical training strings was also impaired.

Subsequently, work by Kinder and Shanks (2000) using a Simple Recurrent Network model of AGL performance, has cast doubt on the supposition that only separate systems models of learning can account for impaired recognition and intact classification in amnesics. The condition was simulated within the SRN by reducing the learning rate parameter, making the assumption that the learning rate of amnesics is slower than controls. The simulation results resemble the experimental data very closely. While classification performance was similar when the learning rate was low or high, recognition performance was considerably better when the learning rate was high. Hence no new mechanisms are required to produce a dissociation between classification and recognition, it can be explained equally well by a slowing down of the learning process in Amnesia.

As explained earlier (sect 1.5), in the sequence learning task subjects become sensitive to sequential movements of a stimulus between locations on a screen, without associated verbalised knowledge. There appears to be some evidence for the activation of non-overlapping brain regions between different levels of awareness in the sequence learning task. In one Positron Emission Tomography (PET) study, Grafton, Hazeltine and Ivry (1995) demonstrated that distinct brain regions are active when subjects report awareness of a sequence compared to a situation where they claim to be unaware. Moreover, as subjects showed greater awareness of the sequence, additional brain areas appeared to become active. This data relies rather heavily on the assumption that the awareness test was suitably detecting the absence of conscious knowledge in the supposed unaware brain state. As explained earlier, assumptions of when subjects are unaware should be made with great caution.

In the first application of the SRT task, Nissen and Bullemer (1987) showed that Korsakoff amnesic patients were able to learn a 10 item repeating sequence. Furthermore, the patients did not reveal any information about the sequence at the awareness test. Shanks and Johnstone (1998) point out that these studies used sequences that did not control for the frequency with which stimulus location was used and other similar factors. That is, the patients need not have learned the sequence transitions to show learning, they could simply learn that location 1 occurred more frequently. Data reported by Keele (1997) shows similar intact learning of sequence information by patients with hippocampal damage. However, there was some evidence that the patients were unable to learn complex information, possibly confirming the criticism of the Nissen and Bullemer (1987) studies. Another study of sequence learning in amnesics was conducted by Reber and Squire (1994). This study appeared to reveal differences between the awareness scores of amnesics and normal subjects. However, they have been questioned by Shanks and Johnstone (1998) on the grounds of low statistical power resulting from small sample sizes.

More recently, Reber and Squire (1998) have demonstrated that patients with medial temporal or diencephalic damage showed normal learning on the task, despite grossly impaired ability to demonstrate learning of the sequence explicitly. Boyd and Winstein (2001) tested patients with unilateral stroke in a series of experiments. The patients responded to the sequence with the hand ipsilateral to the damage, and demonstrated no learning of the sequence even after extended practice. However, when they were given the opportunity to learn the sequence explicitly, a reduction in response time was observed. This data and that of Reber and Squire (1998) shows an interesting dissociation between implicit and explicit learning of sequential information.

In conclusion, although some of the studies reported here are affected by methodological problems, there is evidence that neuropsychological patients are able to learn sequences implicitly but not explicitly (Reber and Squire, 1998), and there is other evidence for the reverse (Boyd and Winstein, 2001). While this data can be taken as support for the existence of conscious and unconscious learning modes, it should be noted that these studies are subject to the same methodological problems of determining when learning is conscious or unconscious as studies with normal subjects. For example, the Grafton et al.. (1995) study does not demonstrate that the unaware subjects are below an objective threshold of awareness, no conclusions can be drawn that implicate separate learning systems. Equally, it is difficult to determine if the subjects in the Boyd and Winstein (2001) study acquire only explicit knowledge of the sequence, it may be that the test is not exclusively sensitive to explicit knowledge and implicit knowledge has an influence. In conclusion, data from neuropsychological studies cannot be taken to be representative of separate learning modes until it is subject to the same methodological constraints that are necessary for studies with normal subjects.

Evidence from dual task studies

The findings of studies with amnesic patients are similar to those where dual task conditions are used with normal subjects. Curran and Keele (1993) found that when subjects had performed a secondary task during training performance was equivalent to situations where single task instructions were used. It was suggested that the secondary task had removed any possible conscious influences from task performance, and task performance revealed only the sole contribution of unconscious learning. However, Neal and Hesketh (1997) note that performance on dual task experiments are subject to the same problems of dissociation as standard SRT task conditions, and cannot be considered as decisive evidence for unconscious processes. Furthermore, Shanks and Johnstone (1998) carried out a dual task sequence learning experiment and demonstrated that appropriate tests of awareness revealed significant amounts of explicit knowledge of the sequence.

Evidence for rule abstraction in implicit learning

Reber's dual system account of implicit learning makes the point that implicit knowledge tends to be in the form of abstract representations of the rule structure that relates stimuli. For example, in the AGL task Redington and Chater (1996) suggest that abstract knowledge is best characterised as information that the subjects possess that is similar to the grammar that generated the letter strings subjects saw during training. Hence, rules go beyond the surface features of the stimuli and are grounded on the conceptual relationships between the letters that form the strings. This is a central aspect of Reber's dual system account, as it allows the conclusion that the proposed unconscious learning system operates using the same symbolic knowledge as does the conscious learning system. As Cleeremans (1998) puts it,

> "there is an unconscious mind that is just the same as the more familiar conscious one, only minus consciousness".

It is important to note that the connectionist account of implicit learning does not specify that rules are represented directly. However, one connectionist model of AGL performance demonstrated by Dienes et al. (1999) models transfer to alternative symbol sets, so findings of transfer are significant for both dual system frameworks.

The strongest evidence for the use of abstract knowledge of this kind in AGL is drawn from studies demonstrating "transfer" to a novel symbol set. The transfer effect in the AGL task occurs in circumstances where elements from which training strings are composed are mapped onto a new vocabulary, so that the underlying syntax is identical but the surface symbols are changed. Subjects perform at above chance levels when judging the grammaticality of the strings in the novel alphabet (e.g. Mathews et al., 1989). This finding, it has been assumed by proponents of the abstraction view, rules out any account which refers to surface features, because in this case, training and test stimuli have different surface forms.

Transfer has been demonstrated many times in the literature (e.g. Altmann et al., 1995; Brooks and Vokey., 1991; Mathews et al., 1989). As stated earlier, this effect has been taken as strong evidence that subjects have acquired an abstract representation of the grammar and are applying this knowledge in the transfer. The original finding has been extended by Altmann, Dienes and Goode (1995), who have demonstrated impressive cross modal transfer. For example, (Experiment 1) one group of subjects were trained on standard letter strings and a second group on a sequence of tones, both sets of stimuli conformed to the same rule structure. Each letter string had an equivalent tone sequence. In the subsequent test, participants performed classification in the same modality (letters to letters or tones to tones) or in opposite modalities (letters to tones or tones to letters). The results indicated that prior exposure to the grammar led to increased performance relative to control groups who had received either random sequences or no training.

While this finding may at face value provide substantial evidence for the use of abstract rule knowledge, it is the extension to this preliminary finding that provides further interest. Altmann et al. (1995) performed post hoc analysis on their data and found that performance was above chance even for sequences where no element was repeated. For example, items such as MTXR do not contain a repeating element, whereas MTTTVT does contain repeating elements. This finding is interesting because it may allow a comparison with another view of transfer performance proposed by Brooks and Vokey (1991). In their view, subjects store whole strings and then use the similarity of those stored strings to new strings presented at test to drive analogical processing. Hence their view is known as the "abstract analogy" account of transfer performance. Repeated elements are important in this view because they allow analogy between strings to be drawn. For example, the string BBGXTR and WWSNPZ are similar on the basis of their repetition structure, this information could not be obtained from strings with no repetitions because the mapping of surface features is entirely arbitrary. Gomez, Gerken and Schvanveldt (2000) make the important point that such repetition structure is a route to testing the validity of the abstract analogies viewpoint. The prediction they made was that transfer can only occur where the items contain repeated elements, with no transfer when strings did not have repeated elements

So, returning to the Altmann et al. (1995) finding that transfer still occurs in the absence of repetitions, this would now seem to be evidence against the abstract analogy viewpoint. However, this finding should not affect the abstract rule based interpretation of transfer. If the grammar is abstracted in the training phase, this process should not occur any differently for repeated or non repeated elements. Hence, the grammar should simply be applied to the new domain.

Gomez, Gerken and Schvanveldt (2000) observed that the ungrammatical strings used for the post hoc comparison in the Altmann et al. study began with illegal elements. They note that these violations have been shown to be particularly salient for learners (see Reber and Allen, 1978; Reber and Lewis, 1977) which may account for the Altmann et al. (1995) findings. To address this issue, Gomez et al. set out to test whether subjects can demonstrate transfer of an identical grammar between two vocabularies where repetition of elements is not allowed by the grammar, and with materials without salient initial bigrams. In these conditions the participants in Gomez et al's experiment did not select grammatical items in the test phase with any greater frequency than non grammatical items. Therefore, it appears that transfer performance is more likely to be dependent on the kind of analogical processing that Brooks and Vokey (1991) originally proposed, rather than a view that emphasises abstraction of rules. Indeed, Gomez et al. make the point that people may have a bias to use information such as repetition patterns. They claim that this occurs because of the perceptual salience of these features.

The acquisition of complex information in sequence learning

There has been some debate about what particular information subjects use to demonstrate performance benefits in the sequence learning task. This debate is similar in some ways to the debate on transfer in the AGL literature, as it centres around the complexity of knowledge acquired by subjects. Reed and Johnson (1994) noted that the sequential movements that subjects respond to may not be balanced in terms of their frequency. For example, the transition from location 1 to location 2 may be less frequent than the transition from location 2 to location 3. In these circumstances, subjects may simply learn that the movement from location 2 to location 3 occurs more frequently and tend to use this information when responding to the targets. No learning of the relationships between the locations is required to show sensitivity. Reed and Johnson (1994) noted a number of other constraints that the sequence has, such as the absolute frequency of each location, for instance the stimulus may simply appear more frequently in location 2. In an attempt to control for these factors, Reed and Johnson (1994) used a procedure that allowed specific conclusions to be drawn on the information that subjects used to perform the task.

Reed and Johnson (1994) devised training and test sequences that differed only in second order condition (SOC) information. A SOC sequence is one in which every location is completely determined by the previous two locations, whereas knowing the previous location alone does not provide enough information to predict the current location. The training and test sequences were equated in terms of the transition frequency, location frequency and other confounding information. The response time at test was significantly slower than the training trials, indicating that subjects were able to learn complex SOC information. Furthermore, results from direct cued generation and recognition tests indicated that the information was implicit. Subsequently, Shanks and Johnstone (1998) have suggested that the design used by Reed and Johnson (1994) was flawed. When these flaws were corrected, subjects were able to reveal sequence knowledge in objective tests, indicating that SOC information cannot be learned without awareness.

Summary of the parallel systems account

Apart from dissociations between performance measures and verbal reports (see sect 1.7), proponents of separate learning systems suggest that the primary evidence in favour of separate conscious and unconscious learning systems is provided by neuropsychological studies. While data from neuropsychological patients provides impressive evidence in favour of separate learning modes, it must be subject to the same tight methodological constraints imposed on studies with normal subjects. In the majority of studies with amnesic patients these controls have not been taken into account. As dual task studies are subject to the same methodological considerations, this problem also applies to these studies. As a result, the evidence from neuropsychology and dual task studies cannot be considered to be conclusive.

The claim that unconscious learning yields abstract knowledge is cast into doubt by the finding that transfer in AGL does not occur when strings do not contain repeated elements (Gomez et al., 2000). This indicates that subjects have to rely on analogical processing of previously seen strings to drive transfer performance. As a result, it seems unlikely that the knowledge acquired in the AGL task is separate from the surface features of the training stimuli and is represented in an abstract form. This conclusion refutes the idea that implicit learning is a process that results in abstract knowledge, as suggested by Reber (1989). Furthermore, the finding that repetition structure is critical to transfer suggests that other knowledge, such as the similarity of training items to those at test may be involved in AGL. This issue will be addressed in the light of alternative accounts of implicit leaning, which are considered in the following sections. As explained earlier, these accounts do not propose the existence of separate learning systems, so the finding that performance is sensitive to the surface form of the stimuli may indirectly provide additional evidence against the idea of separate conscious and unconscious learning systems.

The assumption that the objective structure of the stimulus environment is acquired during training implies that learning proceeds in a unsupervised manner.

The role of the subject is that of a passive learner who has little influence over the acquisition of the stimulus domain. This point has been neatly put by Cleeremans (1993) -

> "a hallmark of implicit learning processes ... is that they proceed in a unintentional way ... any structure emerges as a result of [implicit] processing must be entirely stimulus driven"
>
> Cleeremans (1993) p13
>
> cited by Whittlesea and Wright (1997)

This assumption of passivity contrasts strikingly with the other accounts of implicit learning. In the following discussion, it will become clear that participants in implicit learning experiments are not passive recipients of the stimulus structure, they actively process and manipulate the material they are presented with. It is this interaction with the stimulus that produces sensitivity to the structure. The dual system account fails to take account of the importance of this active processing, and this is why the alternative models that are described below offer a superior account of implicit learning.

1.8.2 Subjective unit formation account of implicit learning

The subjective unit formation account of implicit learning, proposed by Perruchet and Gallego (1998) focuses on how people structure the information that they are presented with. For example, in artificial grammar learning the strings can be thought of as made up from pairs or triplets of letters that frequently occur adjacent to one another. For example, the first two letters in a string may be "TV" in a large proportion of training strings. This division of letters could produce a knowledge base of permissible pairs that can be later used in the test. This account is not limited to AGL, Perruchet and Gallego (1998) argue that this process of division or parsing of stimuli is a fundamental process that shapes the phenomenal experience of the world.

The original motivation for this account of implicit learning originated from work by Dulany, Carlson and Dewey (1984) on the AGL task. In this study, during the grammaticality test subjects were asked to underline the part of the item they believed made the string grammatical or cross out the part of the string that made the string ungrammatical. They performed at above chance on this version of the task.

Perruchet and Pacteau's (1990) experiments were designed to focus specifically on the role of knowledge of permissible pairs of letters and the availability of this knowledge to conscious awareness. Perruchet and Pacteau (1990) demonstrated that grammaticality judgements of subjects initially studying letter strings did not differ from subjects learning from a list of bigrams that made up the strings. Further support for the formation of small chunks of letters is provided by the fact that the verbal reports of subjects who were asked to give instructions to a yoked partner following the training phase most frequently referred to bigrams or trigrams of letters (Mathews et al., 1989).

Data from sequence learning provides further evidence for the idea that subjects acquire constrained fragments or units of knowledge. Perruchet and Gallego

(1998) cite Cleeremans (1993), who reports that even after considerable practice subjects are only influenced by trials four steps away from the current context. Furthermore, there is some further evidence that suggests that the task becomes more difficult as the number of predictive elements increases. Perruchet and Gallego (1990) note that in the study by Cohen, Ivry and Keele (1990), trials which are uniquely predicted by the preceding element (unique relationships) are easier to learn than associations where more elements are required to predict the current context (ambiguous relationships). This indicates that the division of structure into small elements is preferred to more complex longer elements.

It appears that there is evidence from a number of sources that small knowledge units can be used in implicit learning tasks to drive performance. This leaves the question of how these knowledge units are formed. According to Perruchet and Gallego (1998), this process of division or parsing of the training strings occurs by the action of associative processes. The formation of associations occurs as a result of the repeated presence of primitive elements in spatial contiguity (in the case of AGL) or temporal contiguity (in the case of sequence learning). For example, in AGL the letters are the primitive elements which occur in spatial contiguity. Repeated pairings of letters results in the formation of an associative unit, either two or more letters long. Perruchet and Gallego (1998) claim that the limitations of the size of the units may reflect the limited capacity of the perceptual attention system.

The break down of strings into fragments may sound rather like an abstraction of the original stimulus material. In a sense this is correct, what is important however, is that Perruchet and Gallego (1998) argue that these fragments are conscious. Therefore, there is no *unconscious* abstract knowledge base in this model. Perruchet and Pacteau (1990) presented all possible bigrams separately and asked the subjects to rate on a six point scale which they had seen before. The ratings indicated that the subjects had some explicit knowledge of the bigrams that the training strings were composed from. In addition, Dienes, Broadbent and Berry (1991) presented letter string stems and asked subjects what permissible continuations could occur. Again subjects appeared to have some knowledge of the bigrams that made up the strings. In sequence learning, Perruchet and Amorim (1992) demonstrated a correlation between RTs and chunks of trials that subjects had generated in an awareness test.

It may appear that, like Shanks and St John (1994), Perruchet and Gallego (1998) are simply claiming that the knowledge used in implicit learning tasks is explicit. Their view, however, goes further as it provides a role for implicit learning in how these units arrived in explicit memory. It is the process by which information is coded into fragments or units. Perruchet and Vinter (1997) define implicit learning as that which allows -

> "subjects to pass from conscious perceptions and representations to other, generally better structured, conscious perceptions and representations, through the action of intrinsically unconscious mechanisms"
>
> Perruchet and Vinter (1997) – p44

The intrinsically unconscious mechanisms that they refer to here are the associative processes that group the primitives stimuli are made up from. Thus, the subjective experience of a particular stimulus domain changes with greater experience of the domain. Within this view, implicit learning is the structuring of the world from a phenomenological perspective. Since the structure subjects impose on stimuli in implicit learning experiments does not necessarily map onto what the experimenter believes the structure to be, verbal report often yields little useful information from the experimenter's perspective.

There are some researchers who argue that subjects used more than simply fragments of strings in artificial grammar learning. In a recent development, Meulemans and Van der Linden (1997) have built elements of the fragment view into a more traditional abstraction account. By balancing rule knowledge orthogonally to fragment knowledge they were able to demonstrate the use of both kinds of information. Half the test strings contained fragments in common with strings in the training phase (Associated) and half did not (Not Associated). This was coupled with the usual grammatical versus non grammatical items, making four conditions. The degree to which training and test strings had fragments in common was called associative chunk strength (ACS). Meulemans and Van der Linden (1997) also manipulated the amount of training items from a large number (125) to a small number (32). They found that when subjects received a small training set, they were biased to use fragmentary knowledge. However, in the situation where subjects had a much larger set of training stimuli, subjects appeared to use rule like knowledge. This work has been taken to be strong evidence for the use of rule and fragment knowledge under different circumstances in AGL. In response, Johnstone and Shanks (1999) noted that the training strings contained information about the locations of legal fragments and that this information was not included in the ACS measure used by Meulemans and Van der Linden (1997). Using a multiple regression procedure, Johnstone and Shanks (1999) were able to demonstrate that repeated chunks in new positions was a good predictor of performance, whereas grammaticality did not predict performance. These data suggest that the unit formation account is sufficient alone to account for the effect of set size.

It was suggested earlier that transfer to a novel symbol set in AGL can be achieved through the operation of analogical processes. This account of transfer performance is unlike the Perruchet and Gallego (1998) because it refers to the encoding of whole strings. There is some evidence, however, that a fragment account can also explain the transfer effect. Using toy models designed only to provide feasibility proof of performance, Redington and Chater (1996) were able to demonstrate that fragmentary knowledge may be used in transfer. The models provided a match with patterns of observed data. For example, they demonstrate superior performance on the same letter test in comparison with a changed letter test. These models, however, remain silent on the issue of repetition structure which has been shown to be a critical aspect of transfer performance (Gomez et al., 2000).

Concluding remarks on the subjective unit formation account

The account of implicit learning described above involves a substantial shift in understanding what processes are involved in implicit learning tasks in comparison to the dual system account. In place of the unconscious abstraction of rules, the unit formation account proposes that implicit learning reflects the unconscious process of structuring a stimulus domain, and this structuring allows subjects to perform at above chance in implicit learning tasks.

There is a substantial amount of data that supports the claim that subjects build small units of information from the primitives that they are presented with in training. Further evidence suggests that these units are available to consciousness. This finding is significant for the debate on measures of awareness. These knowledge units were only defined as conscious when a direct measure of awareness was used. They are above the objective threshold, as defined by Dienes and Berry (1997), because when subjects see the fragments in a recognition procedure, they select them at above chance. However, when subjects are simply questioned about their metaknowledge, they are unable to report the existence of *rules* indicating that they are below the subjective threshold, as defined by Dienes and Berry (1997).

As noted earlier, some authors cast the status of implicit learning research as reflecting a debate between those who propose the existence of a system that abstracts rules unconsciously, and those who question the existence of such a system. It may be that casting implicit learning in this way isolates the findings of research in implicit learning. Perruchet, Vinter and Gallego (1997) argue that implicit learning should not be an isolated field, instead they argue "implicit learning is at the root of our conscious perception and representation of the world, and its importance for adaptive behaviour is crucial". The important point here is that moving away from the debate on the existence of a unconscious learning system, and instead directing research into attempting to understand what other processes may be occurring has produced a powerful model of implicit learning. This model can generate a number of other testable predictions which will further the understanding of what knowledge is acquired when people encounter complex and structured stimulus domains.

1.8.3 Episodic Processing Account of Implicit Learning

From an early stage, the dual system account of implicit learning has been contrasted with an alternative account of implicit learning, the instance based account proposed by Brooks (1978). This account assumes that all training instances are encoded in full, and the decisions at test are based on the test instances' similarity to those presented during training. No abstract or average form of the training instances is stored in memory.

The episodic processing view is an extension of the instance account, suggesting that subjects encode more than just the training instance itself. This view distinguishes between instances and *experiences of instances*, so that the manner with which the training instances were processed forms a integral part of the representation. This account is now the main competitor to the concept of parallel learning systems, as recent reviews of the field confirm (Dienes and Berry, 1997; Neal and Hesketh, 1997). This view borrows a good deal of its theoretical underpinnings from the instance account, so a description of progression of the instance account is necessary.

An instance based account of implicit learning

The instance account provides a very simple description of what knowledge is acquired in an implicit learning task, that is subjects only encode each exemplar as it is presented. The rule structure of the grammar is not represented in any abstract form, sensitivity to the structure of the grammar arises out the grammatical test items' greater similarity to those presented during the training period. Such use of prior exemplars does not imply that they are explicitly recalled or in fact recallable.

Brooks (1978) gave subjects exemplars that had been generated by two separate grammars. These were used in a paired associate learning task in which these strings were paired with English words. The strings of either grammar could be distinguished by its association to the category of the word to which it was paired. In this case, one grammar was associated with an Old World city and the other a New World city. Subjects subsequently informed of this distinction could classify the exemplars, whereas subjects who were not informed could not discriminate them. Brooks (1978) suggested that subjects were comparing test strings with specific items held in memory, using the specific paired associate as a cue. Reber and Allen (1978) proposed that although this suggests that subjects use exemplar knowledge and not abstract knowledge, the paired associate training conditions stressed close attention to the details of items. However, under conditions where subjects merely observed the strings, the particular details of the stimuli become less important with the emphasis switching to the abstraction of knowledge. Thus, Reber and Allen (1978) propose that the use of abstraction processes versus instance based encoding are balanced in such a way that conditions which promote one discourage the other. Vokey and Brooks (1992) called this view the "dual knowledge" approach. Within this explanation, unconscious processes are the default mode, so when the training task does not encourage instance encoding, rules will be acquired passively (see earlier on).

This "dual knowledge" analysis of performance was questioned by Vokey and Brooks (1992) who suggested that instance based models of memory (e.g Medin and Schaffer, 1978) could account for similarity to specific instances and, more importantly in this case, judgements of grammaticality. In these models, new grammatical items presented at test are likely to resemble a large amount of old grammatical items and so a difference between selection of grammatical versus non grammatical items can be explained by "retrieval time averaging". This process occurs when items' features are, for example, weighted for frequency in such a way that they represent the predictive significance of those features when determining the item's grammatical status. In addition, within these instance based accounts, it is straightforward to account for effects of similarity. Specific items within these models can have a disproportional influence by a "closeness to old" effect. The "closeness to old" effect occurs when a test item is very similar to a training item, so that the particular test item is selected regardless of its grammatical status. Thus, the independent effect of similarity and grammaticality can be accounted for by using "retrieval time averaging" or "closeness to old" processes within instance based accounts, without assumptions of separable knowledge bases. It is important to note that while these models do assume that some abstraction occurs, this abstraction is in no way represented in the knowledge base: it occurs during the retrieval process.

In order to support their claims, Vokey and Brooks (1992) manipulated similarity and grammaticality independently. They found that the effects of the two variables was in the majority of cases, additive, and each variable influenced both judgements of grammaticality and similarity. This lends support to the idea of a single knowledge base because knowledge was applied regardless of the task instructions. This is not the case for a dual knowledge view, as here knowledge of the grammar should be applied in a grammaticality test, with item knowledge being used in a similarity judgement. An even more significant finding was the effect of increasing item individuation by inducing mnemonic training conditions. As noted earlier, Reber and Allen (1978) suggested that increasing item differentiation should reduce effects of grammaticality because it reduces the opportunity for the default abstraction system to operate. Vokey and Brooks (1992) demonstrated that this was not the case, in none of their experiments did item individuation result in a decrease in the effect of grammaticality. In addition, Vokey and Brooks (1992) showed that increasing item individuation decreased the effect of specific similarity, suggesting that items can be too well differentiated to support transfer to new similar items, a finding which is hard to account for under a dual knowledge account.

The validity of the idea that subjects encode whole strings of letters can be questioned on the basis that encoding will not always be maximally efficient. Whittlesea and Dorken (1993) suggested that incomplete encoding of the training stimuli could result in an abstract code of the training stimuli. That is, failure to divert sufficient attention and effort to the training strings may result in only parts of their structure being encoded. Indeed, as Perruchet and Pacteau (1990) have demonstrated, fragments of training strings can be used to select grammatical items at above chance levels, without the need for the whole strings to be experienced. In fact, Vokey and Brooks (1994) have cast the fragment account to be partial encoding or fragmentary retrieval of material. Furthermore, Perruchet and Pacteau (1991) note that the conditions in which AGL performance is maximal, the observation condition of Reber and Allen's (1978) experiment, are also those which are most likely to result in incomplete encoding of particular items. However, the results from the transfer of symbol set experiments conducted by Brooks and Vokey (1991) and Gomez et al. (2000) add support to the position that at some level the representation of whole strings may be important. In addition, there is some evidence that subjects use explicit mini rules generated from their experience with the stimuli (Dulany, 1984), such as "the first letter in the string is T".

The episodic processing view

In response to these apparently contradictory findings of what information is used in the AGL task, Whittlesea and Dorken (1993) propose the acquisition of various types of knowledge; fragments of strings, whole strings or explicit mini rules depending on the circumstances in which the learning occurs. In other words, this view emphasises that variability of the encoding conditions can produce different sorts of knowledge, and that this knowledge reflects the actual experience the subject had with the stimuli. This knowledge is then applied at test according to the extent to which the test task elicits similar processing experiences. Therefore, in this view the properties that the stimulus itself possesses is not necessarily represented or used at test. Different properties of the stimulus will be emphasised by different encoding tasks. This processing view of implicit learning draws on ideas from the transfer appropriate processing framework of memory (Morris, Bransford and Franks, 1977).

The episodic processing account assumes that “memory preserves experiences, not stimulus structures” (Whittlesea and Dorken, 1993 - p230). The nature of these experiences is an interaction between the stimulus structure and the processing performed upon it. Whittlesea and Dorken (1993) suggest that if the encoding task encourages subjects to seek features that are common to all training exemplars, representations are more likely to portray commonalities in the stimuli, as an abstraction account would suggest. If, however, the encoding task induces attention to the details of particular items, this would result in representations of particular items, as in instance based accounts. It is important to note that, in this view, the training task is central to the kind of representations formed. There is no default mode for learning, learning only occurs by active interaction with the stimulus structure.

The underlying principle has been neatly put by Whittlesea and Wright (1997) - p183 "Until the subject encounters the stimuli and processes them in some way for some specific purpose, the effective structure of the set is not real but only potential, it has many potential states, some of which are catastrophically different". This idea guides the understanding of every aspect of what is happening in an implicit learning experiment. The deep structure of the stimulus set is *never directly observed* by the subject. Sensitivity to this deep structure arises, almost by accident, out the *correlation* between the subject's set of experiences and that objective structure, not the objective structure itself. This leads to inaccurate reports on the rule structure of the grammar because people do not directly observe the structure, they instead report whatever elements helped them in satisfying the demands of the task.

Evidence for the use of processing episodes in AGL tasks has been presented by Whittlesea and Dorken (1993). In their experiments, subjects memorised items generated by a grammar, either by pronouncing or spelling the items, then classification was conducted by pronouncing half the test examples and spelling the other half. Test performance was only reliable when study and test processing matched. Thus, the test performance was successful when it cued prior processing episodes. Interestingly, an almost identical pattern of data occurred when the classification and the recognition versions of the test task were used. Whittlesea and Dorken (1993) suggest that this is because subjects use knowledge of processing episodes in both type of test, resulting in the same performance. This finding is consistent with that of Vokey and Brooks (1992) (see earlier on).

Mathews and Rousell (1993) criticised the Whittlesea and Dorken (1993) studies on the basis that subjects' performance was selectively influenced and passive abstraction could not occur under these circumstances. In order to address these concerns, Whittlesea and Wright (1997) presented subjects with two categories that were formally identical but where the items were familiar in one case and unfamiliar in another. Subjects' also performed the same induction task on the stimuli. Therefore, the only manipulation was the familiarity of the category. Subjects' sensitivity to structure was highly influenced by the familiarity of the category showing that processing does not occur in a neutral manner or in isolation from the subjects' experience of the stimuli. Contrary to Mathews and Rousell's (1993) criticism that manipulating task demands artificially produces processing specificity, in this case even when the induction task was held constant, the stimulus was not processed in a neutral manner. This experiment adds further evidence to the view that suggests there is no default mode of processing.

Further support for the episodic account is derived from studies using biconditional grammars. These strings were eight elements long with a full stop separating the first and second halves. The grammar involved three separate rules indicating which letters must occur in the first and second halves of the string. The correspondence rules were X goes with T, P goes with C and S goes with V. According to these rules TPPV.XCCS would be a valid string. Johnstone and Shanks (2000) state that this grammar has three advantages over transitional finite state grammars. First, each of the rules can occur in any of the eight positions. Secondly, because the rule related positions have three intervening letters, it is possible to unconfound rule and fragment knowledge. Finally, the grammatical status of a particular string is easy to determine as all grammatical strings contain three valid rules. This type of grammar is significant because it allows the effect of the training task to have separate effects on different types of structural relationships, fragmentary aspects of the strings and rule aspects of the strings.

Shanks, Johnstone and Staggs (1997) used two induction tasks with a biconditional grammar, one condition which encouraged rule learning (*edit* task) and another which involved more standard orienting conditions (*match* task). Subjects in the *match* condition were at chance in judging grammaticality but they were highly sensitive to the fragment structure of the strings. The data collected by Shanks, Johnstone and Staggs (1997) was unambiguous; when the fragments were familiar, participants showed a greater preference for ungrammatical strings than grammatical. In the *edit* task, subjects were shown flawed examples of strings and were asked to indicate which letters they thought occurred wrongly, after which they were given feedback. This encouraged the use of an hypothesis testing strategy to increase rule learning. Here subjects performed well on judging grammaticality with no effect of whole item similarity, suggesting that the knowledge used was the principles of the grammar. In a further study, Johnstone and Shanks (2000) have shown that subjects in the *edit* condition are much less sensitive to familiar fragmentary elements. Additionally, subjects were shown to have good awareness of the rule structure in the *edit* condition. This not only suggests that subjects were indeed using the rule structure rather than the surface features of the strings, but they may have been using this knowledge with full conscious awareness. Johnstone and Shanks (2000) suggested that their data were consistent with the episodic processing since subjects only acquired rules when the orienting task focused on them, and where subjects attention was left unconstrained, knowledge of fragments became more important.

It is significant that Johnstone and Shanks (2000) demonstrated that rule-like knowledge may be formed with detectable levels of awareness when subjects are directed towards that aspect of the stimulus structure. This finding is in line with four predictions made by Whittlesea and Dorken (1993) on how sensitivity to stimulus structure interacts with conscious awareness -

- One is aware of some knowledge if the task draws attention to that knowledge.
- The incidental learning conditions that are typical in implicit learning experiments tend to mask the fact that knowledge acquired incidentally will be relevant later on.
- One needs some kind of theory about how knowledge that one has is actually relevant to the task in hand in order to verbalise that knowledge.
- Knowledge learning with awareness may be expressed as an unconscious influence rather than an explicit act of remembering. This may be complicated in cases where the relevant knowledge is inherently distributed, as is often the case when participants are exposed to many exemplars.

Whittlesea and Dorken (1993) claim that all four factors are reasons why learning can occur without accurate verbal reports. They maintain, however, that these predictions arise from the general principles of learning, and not from the existence of a separate unconscious learning mode. This implies that the findings of implicit learning studies should not be considered in isolation, as some unique attribute of cognition which is demonstrated under certain circumstances, but as findings that are relevant to the understanding of learning in all circumstances.

The limitations of the episodic account

The encoding of all processing episodes during training is not taken as a plausible knowledge base by all authors. Perruchet, Gallego and Vinter (1997) argue that the life time accumulation of episodes is unrealistic as would lead to a linearly increasing number of episodes from infancy to old age. This position is echoed by Mathews and Rousell (1993), who express concern over the efficiency of a system that stores such a vast amount of information. Neal and Hesketh (1997) respond by suggesting that such criticism is based on intuitive conceptions of the mind, and that there are no empirical data to suggest that storage capacity should act as a constraint on accounts of implicit learning.

Seger (1994) points out that implicit learning is demonstrated through several different dependent measures. According to Seger (1994), these dependent measures reflect different response modalities, conceptual fluency, efficiency and prediction and control. Experimental tasks can be classified according to these modalities, so for example, artificial grammar learning reflects conceptual fluency, while sequence learning reflects efficiency. This classification highlights the necessity for theories of implicit learning to take account of data from different response modalities. A problem with the episodic account is that the experimental evidence which supports this view is drawn solely from classification experiments with artificial grammar strings. However, other evidence may indirectly provide evidence for the episodic processing account if the theoretical foundations of the account are considered in a more broad sense.

The episodic processing account of learning has much in common with the procedural view of cognition, proposed by Kolers and Roediger (1984). Within this framework, information is not represented as an object in memory, but in terms of skill in manipulating symbols. Thus, there is no distinction to be drawn between symbols versus processes or descriptions versus actions, all cognition is described in terms of skills or procedures. For example, the retrieval of the word "bicycle" from memory is no different from the act of riding a bicycle, they both reflect skilled interaction with symbols. This emphasis on interaction is the basis of the episodic model, as subjects actively process and interact with the training stimuli leading to test sensitivity. Thus, the processing view of implicit learning predicts that the stimulus domain must be actively engaged by subjects during training.

Findings from the sequence learning task support this conclusion that the stimulus must be actively responded to during training. Willingham (1999) demonstrated that subjects who merely watched stimuli did not demonstrate learning of the stimulus sequence. This indicates that the subjects have to interact or respond to the stimulus to learn the sequence itself, it is not sufficient to passively observe the stimulus sequence. However, when certain subjects demonstrated a high level of explicit knowledge, indexed by a free recall of the sequence, learning of the sequence by observation was demonstrated by these subjects. It appears that this explicit knowledge is responsible for the learning in this case.

This explicit knowledge may have arisen through mental practice of the sequence movements. Pascual-Leone (1993) demonstrated that mental practice can result in sequence learning without an overt response. However, there are other studies that have failed to demonstrate learning following mental practice only (e.g. Reber and Squire, 1998; Shanks and Cameron, 2000). Shanks and Cameron (2000) claim this discrepancy may be the result of the highly salient sequence elements used by Pascual-Leone (1993). The elusive effect of mental practice is confirmed by the Willingham (1999) study, as only a small number of subjects demonstrated significant explicit knowledge.

Shanks and Cameron (2000) suggested that the processing account of learning (Kolers and Roediger, 1984) may be able to account for the weak effect of mental practice. The processing account proposes that if the operations during training are duplicated by similar operations during the test phase, learning will be revealed. They claim that mental practice may result in mental imagery skill with the sequence, and this mental imagery may not transfer efficiently to the test, resulting in weak effects. It may be that when this mental imagery is highly salient, it may be useful in guiding performance at test. Indeed, when the target stimulus changes colour to make salient certain aspects of the sequence, observational learning can occur (Kelly, Burton and Riedel, 1999).

This understanding of how the episodic processing model can account for sequence learning is incomplete. The episodic processing model rests on the assumption that test performance is driven by reference to previous processing experiences. It is difficult to conceive how episodic information can be useful in, for example, the SRT task. Thus, although the episodic model is a powerful account of performance on classification tasks, it falls short of a complete explanation of all experimental data.

1.9 Summary of theoretical considerations

Early interpretations of implicit learning assumed the existence of a learning system that unconsciously and passively acquires abstract representations of a stimulus environment which are later applied at test (e.g. Reber, 1989). However, the finding that performance in the AGL task does not transfer to a changed letter set provides evidence against the conclusion that abstract knowledge is acquired in this task. Furthermore, the neuropsychological and dual task evidence that has been presented in favour of the operation of parallel conscious and unconscious learning systems is subject to the same methodological problems as single task designs with normal subjects (see sect 1.8.1), rendering this evidence equivocal. As a result, alternative accounts of implicit learning have been put forward as superior explanations.

Perruchet and Gallego (1998) suggest that small units of information are acquired during training in implicit learning tasks. Following training, a large amount of these small chunks of information are stored, which can then be used at test to perform at above chance levels. This account, therefore, assumes that a preferred level of information is acquired. In this sense, the unit formation account is similar to the parallel system view, as both set a default level of knowledge.

In contrast, Whittlesea and Dorken (1993) suggest that there is no preferred level of information acquired during training. Within the episodic processing account, the level information that is acquired is dependent on the processing demands of the induction phase. Therefore, distinct forms of knowledge can be acquired from the stimuli with the same underlying structure. Induction tasks that focus on parts of stimuli may result in the fragmentary kind of knowledge proposed by Perruchet and Gallego (1998). Alternatively, a task that focuses on individual stimuli may result in the representation of whole exemplars in memory.

This is the critical aspect of the processing view. No level of structure is the default for processing; any level can be represented according to the demands of the task. Hence, the use of fragmentary information (as proposed by Perruchet and Gallego [1998]) is perfectly compatible with the episodic processing account. However, the acquisition of abstract information that is held separately from episodic information is inconsistent with the episodic processing view. Therefore, the greatest theoretical contrast is between the episodic processing and parallel systems accounts.

The episodic processing view suggests that the information that subjects use at the test is correlated with the abstract structure of the training items, and since subjects are not aware of this correlation, they are not able to verbalise anything pertinent to the rule structure. If critical aspects of stimuli are stored in this indirect way, then it is reasonable to assume that they may be accessible under certain circumstances. This type of explanation aligns with some observations of what subjects do in fact verbalise following training, given the correct retrieval conditions. For example, Dienes and Berry (1997) have reported that cued tests can elicit some limited awareness following implicit learning experiments. So it appears that the knowledge acquired in implicit learning is on the "fringe" of consciousness (as suggested by Seger, 1994), rather than completely inaccessible.

It has been argued that implicit knowledge and explicit knowledge form the endpoints of a continuum (e.g. Berry and Dienes, 1993; Reber, 1997). If so, knowledge that is on the "fringe" of consciousness may be somewhere in the centre of this domain. However, there is no empirical evidence that such a continuum exists in the human learning system. For this reason, it is more likely that *test relevant* knowledge is on the edges of consciousness as a result of its indirect relationship with information which is directly processed during training. While this position should not be cast as refuting the existence of learning without awareness as a real experimental phenomenon, it does deny the existence of a separate unconscious learning mode to explain it.

A key aspect of the episodic processing account is the constant activity of the learner during training. In contrast, the parallel systems account of implicit learning assumes that unconscious learning is automatic. For example, Reber (1997) states "a system bereft of consciousness is one that operates fully automatically" (p49). It is on this issue that the episodic processing view contrasts most strikingly, as here the operations that are carried out during training directly control performance in the test. The subjective unit formation account remains silent on the issue of whether implicit learning is an automatic process or an active process, and this reinforces the point that the most significant theoretical contrast is between the parallel system and episodic processing views. In the light of this, and the earlier suggestion that the use of fragmentary information is perfectly consistent with the episodic view, from here on, only the parallel system and episodic processing views will be considered.

A concern for the episodic processing account of implicit learning is that the evidence cited in its support is universally taken from classification or recognition tasks. An increasing proportion of recent literature on implicit learning focuses on findings from the sequence learning task, a fact that is problematic for the generality of the episodic processing account. However, the episodic processing account is theoretically close to the Kolers and Roediger (1984) model of skill acquisition. If the episodic account is taken to incorporate predictions made by the Kolers and Roediger (1984) model, then this framework can provide an explanation of sequence learning performance. This indicates that the findings of implicit learning tasks can be integrated into a general account of learning, using a single set of principles to understand the learning of different types of information in varying experimental contexts.

1.10 Preview of the experimental chapters

Research in the field of implicit learning has largely been concerned with attempting to prove or disprove one interpretation of the phenomenon, which casts implicit learning as representing the operation of an unconscious learning system (e.g. Reber, 1989). This has essentially boiled down to a debate on the ability of awareness measures to detect conscious knowledge (Shanks and StJohn, 1994; Dienes and Berry, 1997; Neal and Hesketh, 1997). Alternative definitions of what constitutes unconscious learning are taken by those who support the concept of a separate unconscious learning system (e.g. Dienes and Berry, 1997), and those who question its existence (Shanks and St John, 1994). Thus, the debate on this issue has reached something of an impasse (see sect 1.7). Within this thesis, implicit learning is construed as a experimental phenomenon that may be explained by one of two competing theoretical interpretations, the parallel system account mentioned above, and the episodic processing view (Whittlesea and Dorken, 1993).

The principle aim of this thesis is to discriminate between these two accounts, by focusing on their main points of departure. The first of these is the role of active processing in implicit learning, and the second is the information acquired in implicit learning tasks. An additional theme that will be addressed in this thesis is a specific prediction peculiar to the episodic account. This prediction suggests that the processing demands of the training and test phases must be consistent for successful performance at test. These aspects of the accounts do not specifically address the issue of conscious awareness, which has proved contentious elsewhere.

The main themes recur throughout this body of work, while the general structure of this thesis is guided by more specific experimental questions, which will be addressed in detail at the outset of each chapter.

Chapter 2 The Efficacy of Implicit Invariance Learning

2.1 Introduction

In Chapter 1, recent findings using the Artificial Grammar Learning (AGL) task were discussed. One finding concerned the demonstration of transfer effects in the AGL task. That is, training is carried out on one letter set, with the test items constructed from a different letter set, and appreciable learning can still be revealed on the changed letter set (e.g. Mathews et al., 1989). An understanding of this effect is that the grammar from which the items are constructed has been acquired by the subject, and can be applied regardless of surface form (e.g. Reber, 1989). This understanding has been questioned on the grounds that certain salient items can be used to perform at above chance levels on the transfer test, without the grammar itself being represented in any way (e.g. Brooks and Vokey, 1991). Such a suggestion has implications for findings of transfer in other tasks. One such task is the McGeorge and Burton (1990) invariant learning task. Experiment 1 investigates transfer in this task with reference to the contribution of salient items, which have proved to be important in this task in other work (Wright and Burton, 1995).

A further finding in the AGL literature that may be relevant to the invariant learning task, is the contribution of encoding episodic information (e.g. Whittlesea and Dorken, 1993). The role of this type of information is investigated in further experiments in this chapter. In the course of investigating these processes, it will be seen that the invariant learning task reveals very small learning effects (Experiment 3). These effects may initially raise methodological issues relating to the sensitivity of the test task. However, they represent more interesting points relating to the interaction between test task and training task demands. These issues are addressed in the final experiment in the chapter (Experiment 5), where the role of the training task is investigated.

2.2 Implicit learning of an invariant

In the original McGeorge and Burton (1990) study, participants were asked to perform some task (such as arithmetic) on 30 four-digit numbers. Each four digit number contained a "3" digit; this fact was not pointed out to the participants. In a subsequent recognition task, 10 pairs of four-digit numbers were presented to subjects one at a time. In each pair one number contained a "3" *(positive)* and one did not *(negative)*. Subjects were then falsely told they had seen one of the numbers in the study phase and they must indicate which of the numbers they had seen. A robust effect was observed for participants choosing the positive over the negative. This effect was seen to persist over different encoding tasks, and when the format of test materials was changed from study to test phase (i.e. learning items seen as digits and test items appeared as words). In a post task question session, subjects were unable to report anything pertinent to the invariant digit. They concluded that performance on this task was driven by semantic knowledge of the invariant feature, and that this knowledge was implicit.

Wright and Burton (1995) have questioned this assertion, by suggesting that it is not necessary to learn an invariance rule to perform at above chance performance on this task. They propose that subjects may have been rejecting salient negatives on the basis of information that is more prevalent in these test items. The information to which they refer concerns repetitions within the four digit numbers (e.g. 4667, 2862 etc), these numbers being more distinctive at test. Constraining the positive items to contain "3" reduces the possibility of repetitions within the strings. Therefore, fewer positive items than negative items contained repetitions in the test stimuli. As stated before, items with repetitions would be rejected at test because they are salient. Hence, fewer positive items would be rejected in the test. This would allow participants to score at artificially high levels. To investigate this possibility, they constrained test pairs in two ways, namely where one item contained a repetition and another competing item did not, or where neither contained repetitions *(neutral).*Thus, it is possible to examine the effect of repetitions independently from any effect of the invariant. Where one item contained a repetition, this could occur in the negative (e.g. 2447), so biasing the response to the positive (*towards*), or the positive (e.g. 1138) in turn biasing the response to the negative *(against).* Performance was better predicted by rejection of salient items than selection of positives.

There remained, however, an indication that the invariant still exerted an influence over performance, and there were three pieces of evidence to support this. Firstly, in the *Neutral* condition the effect was marginally significant, with selection of positives at 59% (as chance was at 50%). Secondly, in the *Against* condition, if rejection of salient items was solely controlling performance it would be expected that the salient positive should have biased selection to the negative. However, selection was not significantly biased to the less salient negative. Thirdly, when the three groups were collapsed together, there was a significant trend to select the positive. This would not be expected if rejection of salient items was the only process occurring in this task. Performance should be symmetrical with respect to chance performance across the three groups, yielding chance performance when the three groups are collapsed.

This rejection strategy becomes more important when considering the transfer knowledge of the invariant to different formats (McGeorge and Burton, 1990). That is, subjects presented with training stimuli as digits (3567 etc) would select positives with a greater probability even when test stimuli were words (three five six seven). McGeorge and Burton (1990) suggested that since the effect was not tied to surface features, subjects had acquired semantic information about the invariance. This effect was likened to other transfer effects found in the implicit learning literature, where knowledge does not appear to be tied to the surface form (e.g. Mathews, 1989). It is represented in an abstract manner, and can be applied regardless of surface characteristics.

Importantly, the cross format transfer condition of the McGeorge and Burton study (Experiment 2) reveals a smaller effect than where study and test items were consistent (Experiment 1). It may be that the lowering occurs because subjects are now relying on the rejection strategy alone, with the small effect of invariant sensitivity removed by the change in surface form. That is, the transfer effect demonstrated by McGeorge and Burton (1990) may be an artefact resulting from the use of the rejection strategy, made possible by the presence of salient repetition items.

2.3 Experiment 1

It is the purpose of Experiment 1 to investigate the validity of the cross format transfer effect in implicit invariance learning. The cross format effect occurs when subjects are presented with strings in one format (e.g. words – nine three six one) at study and at test the format is changed to another format (e.g. digits - 2431). Earlier on it was suggested that the rejection strategy demonstrated by Wright and Burton (1990) might be sufficient to account for performance in the cross format conditions. So, it is important to make transparent the contribution of both the rejection strategy and any sensitivity to invariance in this task.

To this end, a set of training strings were constructed that corresponded to the invariance rule, all containing the digit '3'. These strings were written out as words (e.g. three nine seven two). Subjects were then presented with the positive versus negative forced choice test. Here the strings were presented as number strings (e.g. 5671 against 2315), so that the format was different at test in comparison to training conditions. As in Wright and Burton (1995), there were three conditions at test; *Towards, Against and Neutral.* The *towards* test items are constrained so that there is a salient repetition in the negative (e.g. 5661) but not in the positive (e.g. 3298). In the *Against* items, the positive item contains a salient repetition (e.g. 3224), while the negative does not (e.g. 4651). Finally, neither of the *neutral* items contain repetitions (e.g. 2314 against 7659). This design is identical to that of Wright and Burton (1995) except for one aspect, there is a change in surface features between the study and test phase.

In the Wright and Burton (1995) study, when the three conditions were grouped together they displayed a significant but weak bias towards selecting the positive, and a much more powerful bias towards rejecting distinctive repetition items. If the rejection strategy is responsible for the cross format transfer found by McGeorge and Burton (1990), then it should be clearly demonstrated by a bias to select positives in the *Towards* condition, reject negatives in the Against condition, and no bias in towards selecting positives in the *Neutral* condition. More importantly though, there should be no general bias to select positives over the three conditions.

A finding that shows a consistently strong bias to reject salient negative items at test, even when the surface features are changed, would add further evidence to the position that this process occurs in a different way to that of sensitivity to invariance. The data that were reported in the original Wright and Burton (1995) study suggested that this rejection strategy occurred explicitly. A minority of participants (20%) reported using rejection of unseen “doubles” as a strategy when guessing. This was taken as evidence that some subjects were using this strategy to aid performance. However, such assertions must be taken with some caution for two reasons. Firstly, these verbal reports may not be entirely accurate. Secondly, five subjects' responses cannot be taken to represent the whole. Therefore, more evidence is required to dissociate the explicit rejection of negatives from sensitivity to an invariance rule. If such sensitivity to invariance does not transfer across scripts and the explicit rejection of negatives does, these effects may be underlain by different mechanisms. Modifying the Wright and

Burton (1995) study by introducing transfer from study to test could allow this distinction to be demonstrated.

Method

Participants. Thirty undergraduate students from the University of Glasgow took part in the experiment. They were naïve to the experimental hypothesis and procedures.

Materials. Study sets were generated individually for each subject. First, a set of 30 positive items was generated. A positive is defined as a four digit number that contains one "3" digit, a negative being a four digit number that does not contain a "3". There were two other stipulations; none of the numbers included a zero digit and there were no digits that repeated among the four. The study set items were printed as words, e.g. 3756 would appear as "three seven five six". For the test items, four further lists of numbers were generated for each subject: 12 positives with no repetitions, 6 negatives with repetitions, 12 negatives with no repetitions and 6 positives with repetitions. Repetitions within strings did not involve any more than two digits repeating (e.g. 3445), and these are referred to as Doubles. Test pairs were randomly selected so that each subject received 12 pairs, 4 from each of the 3 conditions. In the *towards* condition a Negative Double is paired with a Positive Non Double. The *neutral condition* involves both items as Non Doubles and in the *against* condition the Positive is a Double paired with a Negative Non Double. Test items were constrained so that none were identical to any in the study set.

Design and Procedure. This was a within subjects experiment: All subjects were exposed to the 30 study items, followed by 4 test pairs from each of the 3 conditions; *towards, neutral* and *against.*

The experiment is divided into three phases: a learning phase, a test phase and a test of explicit knowledge. In the learning phase, participants were given 30 positives on slips of paper printed with the 4 digits as words. The slips of paper were removed from the envelope one at a time. The sum of the first two digits was compared with that of the second. If the sum of the first two was larger than the second, then the participants marked the slip with a tick; if the sum was smaller then they marked it with a cross and if the sums were equal they left the slip blank. This orienting task is the same as that utilised by Wright and Burton (1995).

The participants were then presented with a surprise test. Twelve pairs were presented, four from each condition, one at a time, and the participants were falsely told that they had seen one item from each test pair before. They were asked to circle the item they had seen before.

Following the test phase, subjects were given a test of explicit knowledge, which contained three questions. Participants were asked to answer the first of these three questions before turning the paper over to reveal the final question. The questionnaire read as follows:

Could you please answer the following questions in order? Answer all questions even if you feel you are repeating information.

1. Did you notice anything systematic about the numbers in the first part of the experiment?

2. Did any of the numbers appear more frequently than any others?

3. One number appeared more frequently than any other during the first part of the experiment. Could you circle the number that you think it was: If you're not sure then please guess.

One two three four five six seven eight nine

Results

The means for the three conditions are shown below.

Condition	Positives selected		
	Mean No (out of 4)	SD	Percentage
Against	0.83	(0.95)	21%
Neutral	2.17	(1.02)	54%
Towards	3.27	(0.79)	82%

The mean number of positives selected was 6.3 (out of a maximum 12), SD = 1.39. A one sample t-test showed that this was not significantly greater than chance performance of 6, $t(29) = 1.05$. The mean number of positives selected in each condition is shown above. An Analysis of Variance showed a main effect of condition, $F(2, 87) = 52.2$, $p < 0.01$. Tukey HSD t-tests indicated that there were significant differences between *neutral* and *against,* with further differences between *towards* and the other two conditions.

Following this, these scores were tested against chance performance (2.0 for each condition). For the *Towards* condition more positives were selected than would be expected by chance, $t(29) = 8.84$, $p < 0.01$; for the *Neutral* condition the number of positives selected was not different from chance, $t(29) = 0.9$. Finally, in the *Against* condition the number of positives selected was significantly below chance performance, $t(29) = -6.7$, $p < 0.01$.

From the above table it can be seen that in the *Towards* condition items containing repetitions were rejected on 81.7% of occasions, and in the *Against* condition these items were rejected on 79.2% of occasions. Over the two conditions repetitions were classified as unseen on 80.5%. This is far greater than the selection of the Positive containing the invariant (performance at 52%).

The arithmetic task was carried out correctly on 98.5% of the study strings.

Test of Explicit Knowledge. No participant reported anything related to the invariant in response to the question, "Did you notice anything systematic about the numbers in the first part of the experiment?". Four participants mentioned the invariant at the question "did any of the numbers appear more frequently than any others?". Two participants produced the invariant along with other digits and the remaining two produced the invariant on its own. These participants went on to circle the "three" in the response to the question which instructed subjects to indicate which number was most frequent. The mean for these subjects was 6, which was below the mean for the remainder of the group. Two others mentioned the invariant with other digits in response to the second question, but failed to indicate this in response to the last question that was most specific.

Discussion

The results of this experiment indicate that the participants still demonstrate a tendency to reject salient negatives when surface structure is changed. There was no longer any influence of the invariant at test. Firstly, the effect of the positive selection in the *Neutral* condition was no longer marginal. In the Against condition, participants rejected positive at a level that was greater than expected by chance. This is unlike Wright and Burton (1995), where subjects failed to reject the distinctive positive to above chance levels. When all three conditions are collapsed, the bias to selection of positives demonstrated by Wright and Burton (1995) was not found when surface features were changed.

These findings indicate that the effect demonstrated by McGeorge and Burton (1990) in experiment 2 and 3 could be attributed to rejection of salient negatives alone with no contribution from invariance learning. It would appear that invariance learning is sensitive to surface features. This finding indicates that the knowledge used in this task is not held at an abstract level separate from the physical manifestations of the stimuli. The data reported here does not allow any further discussion, although it will be addressed later on in the chapter.

The fact that rejection of salient items is insensitive to surface structure is not surprising if the perspective advocated by Wright and Burton (1995) is followed.

They suggest that negatives are rejected by a metacognitive process, the items' salient repetition structure provides a cue that this particular item is novel and did not occur in the study set. A change in surface structure between training and test should have no effect on this process. The evidence from this experiment would add further weight to this view. The fact that the selection of positives and the rejection of repetitions respond differently to surface changes, suggests that they are separable processes.

Subsequent to this work, Stadler, Warren and Lesch (2000) have carried out a similar experiment on cross form transfer. They demonstrate that the transfer observed in McGeorge and Burton (1990) depended on the use of the rejection strategy. When repetitions were removed from the training and test materials, no transfer was observed. The difference between this study and the current study is that here the effect of rejection of salient items and selection of positives was compared within the same experiment and were shown to be separable processes.

Stadler et al. refer to the lack of cross format effects as "hyperspecificity" of transfer. They suggest that this is similar to the sensitivity to surface feature changes demonstrated in the implicit memory literature. To explain these effects, they cite the processing view of implicit memory (e.g. Roediger, Weldon and Challis, 1999). This view suggests that learning and memory depend upon a match between processing performed during study, and that occurring at test. They suggest that the sensitivity to surface form demonstrated by their data indicates that materials need to be the same between study and test for any effects to occur. This suggests that if the materials are different, processing during study and test do not match, and the result is that performance drops to chance levels. This account would seem to receive further support from the results of Experiment 1.

These data replicate the findings of Wright and Burton (1995), by showing the power of this rejection strategy. Although rejecting salient negatives may be a powerful determinant of performance, it cannot be the only factor at work. Stadler, Warren and Lesch (2000) showed that the learning can be demonstrated with no salient repetition items included in the study materials. In addition, as explained earlier, there was a marginal learning effect evident in the Wright and Burton (1995) data. It may be that the invariant can still exert some influence over and above the rejection of salient repetitions. Alternatively, it may be that the positive versus negative forced choice test could be subject to more artefacts than simply the prevalence of repetitions in the test negatives. For example, subjects may reject numbers made salient for historical reasons (e.g. 1939, 1812 or 1066). Alternatively, numbers can be salient for personal reasons, like birthdays and bank account PIN numbers. It is important to ensure subjects cannot use rejecting negative items as a strategy at any stage. This will allow the effect of selecting positives to be taken seriously as the process underlying performance in this experiment. This issue is investigated in Experiment 2.

2.4 Experiment 2

To date, all research on the McGeorge and Burton (1990) task has used the standard forced choice procedure that places a novel positive against a novel negative. This test configuration relies on the assumption that the positive item is selected against a neutral negative. The data reported by Wright and Burton (1995) would suggest that this reasoning may be flawed and that the negative may be, at the least, equally important in driving performance. Their stimuli were, however, contrived to reveal the strong effects of rejection of salient negative items solely on the basis of their repetition structure.

There is other evidence to suggest that negative instances may not be neutral to the task demands. Ward and Churchill (1998) demonstrated that subjects could be sensitive to negative items over positives in a forced choice test. In their experiment, subjects were presented with a training set consisting entirely of negative items, so that none of the strings contained a '3'. At test, significantly more negative strings than positives were selected. This experiment did not control for repetitions in the study or test strings so it must be considered with some caution. However, what it does illustrate is that the negative can influence judgements towards selection rather than rejection in certain circumstances. It appears then, that the negative can destabilise performance in this task, and this hinders our understanding of what knowledge, if any, is acquired during the training phase.

The next experiment attempts to demonstrate knowledge of invariance indirectly, and without the need for a positive versus negative forced choice. It may be that an invariant feature could exert a dilution effect on a recognition judgement. This could occur by placing a seen positive against a novel positive in one condition, with a seen negative and a novel negative in a second condition. Thus, both items in the test either contain the invariant "3" (in the positive condition) or neither contain the invariant "3" (in the negative condition). The novel positive item would have the invariant in common with the training stimuli and the novel negative would not. The novel positive may, therefore, influence subjects more than the novel negative item, diluting any recognition effect of the old positive in comparison to recognition performance of the old negative.

In the following experiment, the test of invariance is modified so that knowledge is examined by dilution of a response in a recognition test. The aim of this is to rule out the role of rejecting negative items in the McGeorge and Burton (1990) task.

Method

Participants. Twenty undergraduate Psychology students from the University of Glasgow took part in the experiment. They were paid a small fee for their time. They were naïve with respect to the experimental hypothesis and procedures.

Materials. A computer program generated the digit strings individually for each participant. The study items involved 32 four-digit strings of which 26 followed the invariance rule. In this case the invariance rule stipulated the strings must conform to the criterion that the four digit number contains one digit "3". These were known as Positives, if the number did not contain "3" it was classified as a Negative. None of the four digit strings contained zero digits. The strings were constrained so that none of the strings contained repeated digits. A further 6 Positives and 6 Negatives were produced for the test materials. These were constructed to be different from any item in the study set.

Six Positives and Negatives were randomly selected from the study set to make old items. This resulted in 6 test pairs that included one novel positive with one old and another 6 test pairs that contained one novel negative with one old.

Design and Procedure. The experiment had a within subjects design. There were two conditions; in both the subjects were presented with an old item paired with a novel item. In the Positive condition both included the invariant, and in the Negative condition the invariant was absent from the alternatives.

The experiment was divided into three phases: a learning phase, a test phase, and a test of explicit knowledge. In the study phase the subjects were shown the 26 positives and 6 negatives on slips of paper. The orienting task was the same addition and comparison as in Experiment 1.

Immediately after subjects had completed this task for all 32 items, they were given a surprise test. Ten slips of paper were presented one at a time, containing one seen item against one novel item. These were divided into two conditions as described above. The left-right positioning of the items was randomly determined for each pair. Subjects were asked to circle the number they had seen before. If they were unsure, subjects were asked to guess. Finally, they were given a test of explicit knowledge, which was identical to the type presented in Experiment 1.

Results

The mean number of seen items in each condition is shown in the table below.

Test Materials	Old strings selected (out of 6)	
	Mean No.	SD
Positive	3.9	1.25
Negative	3.8	0.83

The mean number of seen items selected by the 20 subjects was 3.85 (out of a maximum 6), SD = 1.05. A one-sample t-test showed that this was significantly greater than chance performance of 3, $t\ (19) = 5.11$, $p < 0.01$. A t-test for paired samples showed no effect of condition, $t\ (19) = 0.29$, $p > 0.1$.

The arithmetic task was carried out correctly on 97.3% of the study strings.

Tests of Awareness: No subject reported any information relevant to the task in answer to the first two questions. In response to the third question 2 of the 20 subjects circled the digit "3". These subjects selected seen items at typical levels in comparison with the remainder of the group.

Discussion

The aim of this experiment was to reveal invariant learning in a new form of test that avoids problems of negative rejection. Negative and Positive recognition performance was not significantly different in this experiment. Subjects show that they were no worse at recognising old Positive items than old Negative exemplars. Thus, the positives that were paired with the old positives had no diluting effect on performance. The novel Positive items did not exert the kind of influence on performance that may have been expected if the positive items had acquired some special significance to the subject.

This can be taken as evidence that the selection of positives in the standard forced choice task, where novel negative and positive items are paired, is less important than rejection of salient negatives. Negative items cannot be considered to be neutral in the forced choice test.

However, an aspect of the experimental design may have interfered with sensitivity to invariance. It was assumed that subjects would perform at a similar level with respect to the invariant even though some negative instances were included in the study set. This may not be the case. It is possible that implicit sensitivity to invariance is highly specific and if even a few negative instances could disrupt learning. Thus, the Positive and Negative conditions of this experiment failed to demonstrate a difference on account of the disruption of invariance acquisition. If this is the case, it may have interesting implications for invariance learning. The influence of negative introduction on positive selection may provide interesting insights into the kind of knowledge that is acquired in this task. To establish the effect of negative inclusion it would be necessary to perform an experiment that manipulates the number of negative instances included in the study set, and observe any changes to the extent of positive selection.

This possible effect of negative inclusion indicates that the test constraints used in the current experiment may not be valid for examining learning in this task. If including negatives has an effect on learning which is significant, this implies that any sensitivity to invariance would be altered by the stipulation that old negatives are placed into the study phase. This suggests that it would not be possible to replicate the original conditions of the McGeorge and Burton (1990) study. If so, this test configuration would not appear to be a profitable route for future research.

2.5 Experiment 3

McGeorge and Burton (1990) argued that performance on the invariant learning task could be explained by implicit acquisition of semantic knowledge of invariant characteristics of the study set. Experiment 1 demonstrated that the robust nature of this knowledge in the face of changing surface features is likely to be an artefact of the stimuli used by McGeorge and Burton (1990). The acquisition of an invariant rule may in fact be specific to the same surface features. Implicit learning of invariance has been questioned on other grounds in addition to specificity of transfer. As explained earlier, rejection of salient negatives (Wright and Burton, 1995) has been put forward as another explanation. While this shows that the rejection of salient negative items is indeed an important determinant of performance and can be used in this task, there is still some evidence for an effect of invariant learning (e.g Stadler, Warren and Lesch, 2000). There are two main alternatives which could explain what knowledge underlies this effect of learning. The acquisition of rule-like information about invariance or the encoding of episodic information. It was suggested earlier that these accounts may predict differing patterns of performance as negative items are included in the study materials. The predictions of each account will now be considered.

McGeorge and Burton (1990) have proposed that subjects acquire some semantic knowledge of invariance during this task. In other words, subjects learn a rule that each string contains a "3" digit. The acquisition of this type of rule is dependent on the fact that the study set has this general characteristic that each string has a digit in common. If negative exemplars are presented to subjects along with the positives, then the study set could no longer be described in terms of an invariance rule. If no such rule exists in these circumstances, McGeorge and Burton's position would predict that no learning should take place. This should be the case even with only a small number of negative items.

A highly influential view of implicit learning in the last few years has been the Whittlesea and Dorken (1993) episodic processing account (see sect 1.8.3). Within this view, sensitivity to general properties of the training set (or invariances) occur because rule following test items induce similar processing episodes in comparison with the rule following items from training. This similarity of processing episodes results in the selection of rule following items. Effects of invariance learning can be explained within this account by suggesting that Positives are more prevalent in training, and tend to elicit more similar processing episodes. Positive and Negative items are not theoretically separable within this view. The Positive has not gained any special properties during training; the "3" digit simply makes processing of the positive more similar to strings encountered during training. Hence, introducing negatives should steadily reduce performance as increasing numbers are included in the study materials.

The importance of encoding episodes has been demonstrated previously in this task. Cock, Berry and Gaffan (1994) manipulated the similarity of items at test to those in the study set independently from the invariant (producing *positive similar, negative similar, positive dissimilar and negative dissimilar* items). Cock et al. demonstrated that similarity between test and study item was a stronger determinant of performance than the influence of the invariant. However, manipulating test conditions to gain insight into the information used in this task is problematic. Since the study items in their experiment contained no repetitions, the effect of the invariant would have been similar to the very small marginal effect demonstrated in the Neutral condition of the Wright and Burton study. Hence, the large 75% manipulation of similarity would obscure the small effect of the invariant. That is, the materials encouraged participants to use similarity information, perhaps overriding any residual effect of invariance learning. Introducing negative items into the study phase is a cleaner method of determining the information used in this task as the test task conditions are held constant. The acquisition of a rule against encoding of episodes will reveal alternative patterns of data across the same test, revealing the information used without test bias.

The following experiment attempts to determine how the inclusion of negative items in the study period affects performance. The main prediction is that introducing negatives should reduce the bias to select positives at test, if information from the training period is used at test. More specifically, if McGeorge and Burton (1990) are correct, and semantic knowledge underlies performance in this task, the bias to positives at test should no longer occur where a small number of negatives are included in the training set. However, if the Whittlesea and Dorken (1993) account is followed, the bias to select positives should not significantly decrease with a small proportion of negative items included in the study set.

Method

Participants. Sixty undergraduate students from the University of Glasgow took part in the experiment. They were paid a small fee for their participation.

Materials. Study sets were generated individually for each subject. First, a set of 30 items was produced for the study set. These contained a varying number of Negatives and Positives according to the condition. The All Positive condition included 30 Positive items, the Two Negative condition included 28 Positive strings with 2 Negative strings and the Five Negative condition included 25 Positive with 5 Negative strings. The test items were the same across the three experimental conditions. 10 Positive and 10 Negative strings were produced for the test. As in the study set the strings contained no repetitions in the Negative or the Positives. The test items were also constrained so none had appeared as study set items previously.

Design and Procedure. Subjects were randomly allocated to the All Positives, Two Negatives or Five Negatives conditions. The experiment was again divided into a study phase, a test phase and a test of explicit knowledge.

The study set was presented in the same envelopes and on identical slips of paper as in the previous two experiments. The orienting task utilised the same addition and comparison procedure as earlier. Following the study phase, subjects were given the surprise test. This involved the presentation of ten slips of paper which had a Negative novel item and a Positive novel item printed horizontally. Again, subjects were falsely told that they had seen one of the items in the test pair before. The left - right positioning was randomly determined for each pair. Participants were asked to circle the item they had seen in the study set.

Finally, participants were given a test of explicit knowledge that was identical to the format presented in Experiment 1 and 2.

Results

The means for the three conditions are presented in the table below.

Training Materials	Positives selected (out of 10)	
	Mean No.	SD
All Positives	5.5	(1.40)
Two Negatives	5.8	(1.77)
Five Negatives	4.6	(1.64)

An Analysis of Variance revealed that the main effect of Training set narrowly missed significance, $F(2, 57) = 3.02$, $p < 0.1$. Individual one sample t-tests on each of the conditions revealed no significant effects. The maximum t - value being $t(29) = 2.03$ for Two Negatives (significant on a one tailed test $p < 0.05$), All Positives was $t = 1.60$ and Five Negatives being $t = -1.09$.

The arithmetic task was carried out correctly on 94.5% of the study strings in the All Positives group. This proportion was 96.7% in the Two Negatives group and 95.2% in the Five Negatives group.

Tests of Awareness: No participant reported anything related to the invariant in response to the first question on the explicit knowledge questionnaire. Eight subjects responded, when probed, that the "3" was the most frequent digit. These subjects selected positives at test at the same level as the mean, and could not have been influencing the data at higher levels than subjects who did not report the "3".

Discussion

The data from this experiment were inconclusive. The analysis revealed no significant differences between the different levels of negative inclusion in the study set. What is more important, however, is the fact that although a greater sample size was used in comparison to the Wright and Burton (1995) study, the selection of positives was only significantly different from chance in one experimental condition, Two Positives (the effect was marginal). This suggests that learning in the McGeorge and Burton (1990) invariant learning task is far from reliable. Indeed, such effects may not be theoretically important or interesting if they cannot be easily elicited. However, the fact that these digit string stimuli show such weak effects does not mean that the effect of training negatives cannot be investigated in the invariant learning task. Invariance learning has been demonstrated using stimuli other than the digit strings used by McGeorge and Burton (1990). In experiment 4, clock face stimuli developed by Bright and Burton (1993) are used in a otherwise identical design to that of the current experiment. The reason for this is that learning with clock face stimuli has been shown to reveal a larger learning effect.

2.6 Experiment 4

Bright and Burton (1993) demonstrated implicit invariance learning using the aforementioned clock face stimuli. The invariance rule was a time range that was consistent over all of the training items so that all of the training items clock times fell between 6 and 12 o'clock, as in digit invariance these times were referred to as positives. At test, novel positives were set against negatives (clock times between 12 and 6 o'clock) in a two alternative forced choice. Bright and Burton (1993) demonstrated a preference for positives at test. What is more significant here, though, is that subjects selected 7.8 positives out of 10 test pairs. This is a higher proportion of positives selected than in digit invariance experiments. In the light of this larger effect of learning produced by these stimuli, it may be possible to investigate the effect of negative training items using clock face stimuli.

In experiment 3, the number of negatives that were introduced into the study set was relatively small. This was done because a finding of sensitivity to a small number of negatives would have been evidence against the importance of encoding whole strings in determining performance in this task. This reasoning was undermined, however, by the lack of power in the digit invariance task when repetitions are removed from training and test items. So in the following experiment, larger numbers of negative exemplars are introduced, and the gaps between the numbers of negatives in each of the three groups are enlarged. So in place of the All Positives, Two Negatives and Five Negatives groups; All Positives, Five Negatives and Ten Negatives groups are used.

The predictions made here are similar to those of Experiment 3. If negative items do not influence performance, any effects must be the result of processing conducted during the test period. If negative items during training do affect performance, they may do so in two ways, each of which supports an alternative theoretical interpretation. If rule-like knowledge of invariance determines performance, the selection of positives should be significantly reduced by the inclusion of a small number of negative items during training. On the other hand, if the Whittlesea and Dorken (1993) episodic account is supported, a small number of negative items should not influence performance.

Method

Participants. Sixty undergraduate students from the University of Glasgow took part in the experiment.

Materials. The experimental stimuli were a series of clock faces. Four different shapes of clock face were used circular, square, octagonal, and diamond (see Figure 1 for examples). Four designs of legends 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 straight lines drawn onto the clock faces in black pen.

The times displayed on the clock faces were generated by a computer program. These times conformed to one of two rules, either the times were between 6 and 12 o'clock, these clocks will be referred to as *positives*, or between 12 and 6 o ' clock, these times will be referred to as *negatives*. Examples of the positive or negative clock faces used are shown below.

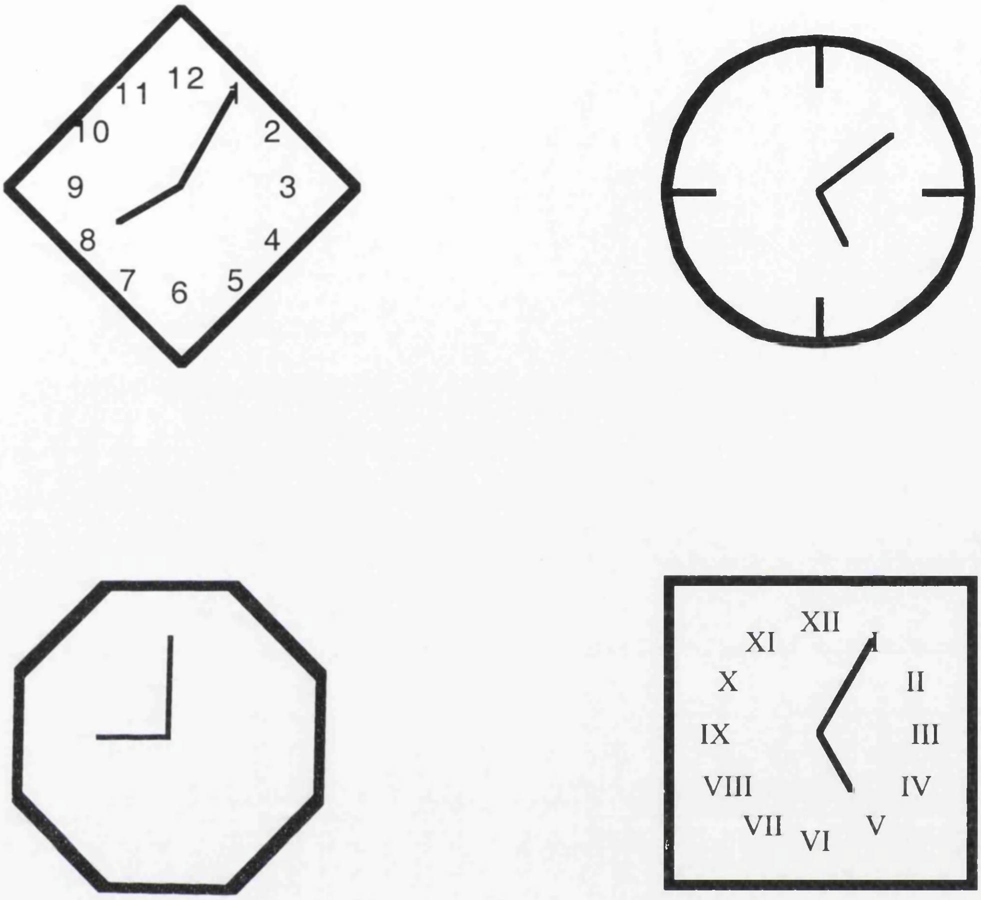

Clock times were generated for the three experimental groups separately. In the All Positives group the materials were composed of thirty positive clock faces for the training items, with ten negatives and ten positives for the two alternative forced choice test. In the Five and Ten Negatives groups, subjects were presented with same test materials as in All Positives. Here though, the training materials were altered by introducing negatives into the set. This was arranged so that the Five Negatives group had five negatives and twenty five positives in the training set, and the Ten Negatives group had ten negatives and twenty positives. There were four sets of stimuli generated for each experimental group.

As seen above, all stimuli in this experiment were presented as analogue times, and the surface features for each face were assigned randomly. The clock times were drawn onto the faces by hand. The training materials and test materials for each set of stimuli were then arranged into two separate booklets.

Design and Procedure. Subjects were randomly allocated to the All Positives, Five Negatives or Ten Negatives conditions. The experiment was again divided into a study phase, a test phase and a test of explicit knowledge.

As described earlier, the materials for each experimental group were divided into four sets. Subjects were given one of the four sets of materials at random. The training clocks booklet were then presented to subjects. They were asked to write down the time displayed on each clock face in a space underneath it. Subjects were free to write out the time in any format they chose. This distractor task is the same as that used by Bright and Burton (1993).

Following the training phase, the first booklet was taken from subjects and they were given the second booklet containing the two alternative forced choice. This consisted of a positive clock time paired with a negative clock time on each page of the booklet. There were ten such pairs made from the ten negatives and ten positives generated earlier. The left/right positioning of the negative and positive clock times was randomly varied. On presentation with the test clocks, subjects were falsely told that they had seen one of the clock times in each of the pairs in the earlier part of the experiment, and that they must indicate with a mark which of the clock times they had seen. If they were unsure they were asked to guess.

Following the test phase, the test materials booklet was taken from subjects and they were presented with a short questionnaire. They were asked to fill out the questions in order, and turn the page over to reveal the last question. The questions subjects were asked were in a graduated format becoming more specific as they progressed:

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

Did you notice anything that the clocks had in common?

The clock times displayed on the clocks all had something in common, what do you think it was? If you don't know just guess.

Results

The mean number of positive clock time selected at test for each group is presented in the table below.

Training Materials	Positives selected (out of 10)	
	Mean No.	SD
All Positives	7.1	(1.7)
Five Negatives	6.5	(1.9)
Ten Negatives	5.2	(2.5)

An Analysis of Variance revealed a significant main effect of Training Set, $F(2, 57) = 4.32$, $p < 0.05$. One sample t - tests showed that All Positives and Five Negatives were both significantly different from chance performance, the lowest t - value was for Five Negatives, $t(19) = 3.4$, $p < 0.05$. Planned comparisons showed that All Positives is significantly different from Ten Negatives, $F(1, 57) = 8.37$, $p < 0.05$; that All Positives and Five Negatives were not significantly different, $F(1, 57) = 0.98$, and that Five Negatives and Ten Negatives are marginally significant, $F(1, 57) = 3.62$, $p < 0.1$.

Tests of Awareness: Three subjects in the Five Negatives group reported that they noticed the majority of clock times appeared to be between 6 and 12 o'clock. These subjects did not deviate substantially from the mean and were not removed from the analysis.

Discussion

The clock face stimuli demonstrated a clear bias to select the Positive item at test. This effect was more convincing than the digit stimuli, which failed to show significant effects. It is now possible to examine the main thrust of this experiment, the effect of negative inclusion on performance.

The effect of Negative inclusion was clear; introducing five negatives clock stimuli did not significantly reduce the bias to select positives. This indicates that introducing a small number of negative exemplars does not break down the bias to select positives. It was suggested earlier that if a invariance rule is acquired, a small number of negative items in the training phase may disrupt learning.

However, this was not the case, as significant sensitivity was demonstrated by the group who received five negatives in the training materials. It appears then, that the alternative possibility, emphasising episodic processing, is a better explanation of the data. These accounts suggest that the fact that the invariant occurs in every string does not result in any abstract representation of invariance. Subjects are biased to select the positive simply because the invariant feature appears more often, or that more positive items are retrieved.

It is somewhat unsatisfactory that this effect could only be revealed by substantially altering the stimuli from digit strings to clock faces. The alteration of stimuli was undertaken because the digit stimuli revealed very small effects of learning, and the effect of negative inclusion could not be investigated. Since this effect has been demonstrated using clock face stimuli, digit stimuli can now be considered once again. It was suggested earlier that if learning of invariant features of digit stimuli cannot be easily demonstrated, these stimuli may not be theoretically helpful. However, the learning effects of these stimuli are not consistent across versions of the task, and these effects may reflect the lower end of this variation. Furthermore, determining the source of the variation may lead to a better understanding of invariance learning, and hence implicit learning in general.

A good example of this is the role of the training task. There is some evidence that learning on this task is not neutral to what processes the subject carries out during training. In their first experiment, McGeorge and Burton (1990) asked subjects to perform arithmetic on the stimuli, and this served as the orienting task.

In a following experiment, they asked subjects to perform a slightly different task. Subjects counted the number of horizontal lines in the stimuli, so for example, "7" would have one horizontal line. This modification significantly reduced performance on this task. It appears then, that it may be possible to vary the sensitivity shown to invariant properties by manipulating the demands of the training period.

Any variation induced by changing the training task is not theoretically neutral. As suggested earlier, the overlap between the training task conditions and test conditions is central to the episodic processing view of implicit learning proposed by Whittlesea and Dorken (1993). Within this view, altering the relationship between the processing demands of the training and test periods varies the overlap between them, resulting in different patterns of performance. Where overlap is high, performance is good, with low overlap resulting in poor performance.

An effect of varying training conditions is not taken into account within a view of implicit learning that assumes that knowledge in these tasks is acquired by an unconscious learning system (e,g. Reber, 1989). This account assumes that information is passively acquired because learning occurs in a stimulus driven manner. The structure (or invariances) of stimuli will be encoded regardless of what task is carried out on the stimuli. Experiment 5 is concerned with clarifying the effect of training task manipulations.

2.7 Experiment 5

Processing accounts of implicit learning emphasise that performance at test is more successful when there is a large degree of overlap between training and test processes. By this understanding, marginally significant learning effects must indicate a small amount of processing overlap between study and test periods. The results of Experiment 3 reveal such marginal effects, so it may be that this is the reason for the small effect of learning revealed. This claim cannot be made without defining the conditions for both high and low degrees of overlap. It is the purpose of the current experiment to investigate conditions where the overlap between study and test is maximal in the invariant learning task. In addition, the conditions where overlap is as little as possible will be demonstrated to provide the purest comparison case.

Wright and Whittlesea (1998) refine the notion of overlapping processing in implicit learning tasks with an account they call "learning without knowing the consequences". They propose that the sensitivity to the general properties of a domain arises from the correlation between these structural properties and the information that subjects acquired explicitly in satisfying the demands of the induction task. Importantly, subjects are not aware of this correlation and are unable to verbalise anything about the structure of the stimulus domain. For example, during the test in the invariant learning task subjects may attempt to remember some of the arithmetic computations they carried out during training. Since an increased number of these calculations contained the digit '3', the positive would be more likely to be selected on the basis of the '3' digit in common with the remembered computations. Since the awareness test probes the subjects as to the occurrence of a regular digit in all strings, the subjects are unlikely to use their knowledge of computations to respond to this question.

One of the strengths of this account is that it can provide some explanation of performance where the induction task is directly related to the structure of the domain, and, conversely, where the induction task is totally unrelated to the stimulus structure. The information that the subject computes in an attempt to satisfy the demands of the induction task can be same as the structure of the stimuli. In this case, the relationship is unambiguous and performance is high. In contrast, the information computed explicitly may have nothing whatsoever to do with the structure or invariances of the stimuli. In which case, this explicitly computed information would not increase performance, and chance performance would be observed.

To test this reasoning, Wright and Whittlesea (1998) carried out an experiment that used an invariant that was designed to relate directly to the arithmetic induction task used by McGeorge and Burton (1990), where the sum of the first two digits was compared with the second two. The digit strings were constructed so that the sum of the first two digits was always the same as the sum of the second two. Since the induction task required subjects to compare the sum of these pairs of digits, the invariant would be made highly salient by the overlap of task and invariance. In the test, subjects performed at ceiling by selecting 98.3% of strings that retained the invariant at test. This level of performance implies that the relationship between task and invariance became transparent.

Wright and Whittlesea (1998) carried out a further experiment which, they claimed, represented the opposite case, where task and invariance were unrelated. The invariant characteristic dictated that the last two digits should always differ by one (e.g. 54, 32, 98 etc). Wright and Whittlesea (1998) claimed that this invariant did not relate to the induction task in any way, and this reasoning was reflected in the data as the subjects performed at chance. This finding is problematic for a number of reasons. Wright and Whittlesea (1998) used only twelve training items, and it is unlikely that learning could occur following exposure to such a small training set. This problem was heightened by the fact that Wright and Whittlesea (1998) did not demonstrate a case where learning of this particular invariant did occur. Furthermore, even if learning could have been successfully demonstrated, it is not clear that this invariant is unrelated to the arithmetic task. Since the last two digits always differ by one, their sum would be odd in each case. This invariance surely correlates with the orienting task, leading to the conclusion that this task does not represent a case where task and invariance are unrelated.

Aside from these methodological concerns, it is important to note that the Wright and Whittlesea (1998) study used novel arithmetic invariant characteristics. It is necessary, therefore, to carry out an experiment using the same invariant characteristics used in previous studies. In addition, the relationship of invariance to induction task will be investigated using cleaner experimental manipulations.

To this end, a set of training stimuli was created that conformed to an invariance rule of the same type that was used in Experiments 1, 2 and 3. The difference is that in Experiment 5 the invariant was low, the digit "2", or high, the digit "8". The induction task in Experiment 5A was to search for the lowest digit in the string and make a key press response. Hence, the invariant was either matched the demands of the orienting task (i.e. 2 – low number) or it did not (i.e. 8 – high number). Experiment 5B was a reversal of Experiment 5A, here subjects were asked to search for the highest digit in the string. This resulted in the opposite relationship of task demands to invariance and served to check that any sensitivity was not the effect of specific stimuli or induction task used in Experiment 5A.

The main prediction made here is that sensitivity to the invariant digit is dependent on the relationship between the induction task demands and the invariant characteristic. That is, when the low numbers are the focus of the orienting task, a low invariant will result in sensitivity, but where the invariant is high, chance performance will be observed. The opposite result is expected where the task is to search for high numbers.

As in previous experiments, a post task questionnaire was presented to subjects. The prediction is that subjects who process the invariant that is related to the induction task demands will consequently report the presence of the invariant in the majority of cases. No clear prediction can be made regarding the report of subjects who process invariant that is unrelated to the induction task demands.

Method - Experiment 5A

Participants. Twenty undergraduate students from the University of Glasgow took part in the experiment.

Materials. The stimuli were constructed in a similar manner to Experiments 1, 2 and 3, although there were some exceptions and for clarity some repetition follows of details reported earlier. Two sets of study and test stimuli were generated, one for each experimental group, Invariant '2' (using an invariant '2') and Invariant '8' (using an invariant '8'). Strings that contain the invariant continue to be referred to as positive items, and negatives are still items without the invariant, the difference here is that the invariant digit can be a '8' or a '2' dependent on the group. The Invariant '2' stimuli for both the study and test phase were created first. These training stimuli consisted of 30 four-digit strings, all of which contained a '2' digit. As in earlier studies, the strings were constrained so that no strings contained a repeating digit. The test stimuli were then generated and consisted of 10 negative items and 10 positive items. No string from the test phase was identical to any string in the study phase. The Invariant '8' stimuli were then generated by replacing all '2' digits in the strings with '8' digits and vice versa, for both the training and test stimuli. The Invariant '8' stimuli were therefore identical to the Invariant '2' group in two important ways. First, they had an invariant digit, in this case '8' in place of '2'. Secondly, the number of '2' digits was the same as the number of '8' digits in the Invariant Low group.

Design and Procedure. Subjects were randomly allocated to the Invariant ‘8’ or Invariant ‘2’ conditions. The experiment was again divided into a study phase, a test phase and a test of explicit knowledge.

The study strings were presented to subjects on a computer screen. Subjects were told that a fixation cross would appear in the centre of the screen and after a short period a four-digit number would replace it. The fixation cross appeared for 300ms. The subjects' task was to indicate using keys 1 – 6 which was the lowest digit of the four. They were instructed to respond as quickly and accurately as possible. After they had made a selection, the fixation cross appeared, followed by the next four digit number and so on. The 30 training strings were randomly presented to each subject.

When they had finished the training phase, some instructions on the test phase were presented. Subjects were told that they would see two four-digit strings on the left and right hand side of then screen. The ten positive and negative test strings were then put together to make ten test pairs. The left right positioning was randomly determined as was the presentation of the pairs. They were told (falsely) that they had seen one of these strings before and to indicate which of the two it was. If it was the right-hand string the instructions were to press the ‘a’ key, and if it was the left-hand string press the ‘l’ key.

Following the test phase, the test of awareness was presented. This involved the same three-question sheet of paper as in Experiments 1, 2 and 3.

Results

The means for the two conditions are presented in the table below.

Training Group	Positive strings selected (out of 10)	
	Mean No.	SD
Invariant '8'	5.2	1.0
Invariant '2'	6.8	1.6

Subjects in the Invariant '2' group selected Positive string in the forced choice test at above chance level (5 out of 10), t (9) = 3.3, $p < 0.05$. This was not the case in the Invariant '8' group, t (9) = 0.58. The means of the two groups were significantly different from each other, t (19) = 2.6, $p < 0.05$.

Tests of Awareness: Subjects provided very little information in response to the first question, and this was not relevant to invariance. Their responses to questions two and three tended to be consistent and are considered together. In the Invariant '2' condition, nine subjects reported that the most frequent digit was '2'. The remaining subject reported that the digit was '7'. In the Invariant '8' group, seven subjects reported that the digit '1' was the most frequent, with two others reporting '6' as the most frequent and the remaining subject noting '2' as most frequent.

Method - Experiment 5B

Participants. Twenty undergraduate students from the University of Glasgow took part in the experiment.

Materials and Procedure. The materials, design and procedure were all identical to Experiment 5A, except in place of searching for the lowest number in the training stimuli during the training phase, subjects were asked to indicate using keys 4 – 9 the highest digit of the four.

Results

The means for the two conditions are presented in the table below.

Training Group	Positive strings selected (out of 10)	
	Mean No.	SD
Invariant '8'	8.3	1.8
Invariant '2'	3.9	2.2

Subjects in the Invariant '8' group selected Positive string in the forced choice test at above chance level (5 out of 10), $t(9) = 5.4$, $p < 0.05$. This was not the case in the Invariant '2' group, $t(9) = -1.5$. The means of the two groups were significantly different from each other, $t(19) = 4.8$, $p < 0.05$.

Tests of Awareness: Subjects provided very little information in response to the first question, and this was not relevant to invariance. Their responses to questions two and three tended to be consistent and are considered together. In Invariant '8' group, nine subjects reported that the most frequent digit was '8', with one subject responding that the most frequent digit was '9'. In the Invariant '2' group, seven subjects claimed that the most frequent digit was '9' and the remaining three subjects all reported that the most frequent digit was '8'.

Discussion

This data from experiment 5 shows very clearly that the relationship of the orienting task to the invariant property can determine performance on the forced choice test. Where the invariant digit was low, and the orienting task induced subjects to search for low numbers, a significantly increased number of positive items were chosen in the forced choice test (Experiment 5A). Similarly, where the orienting task constrained subjects to process high digits, and the invariant was high, performance was significantly above chance (Experiment 5B). However, the opposite was found in the remaining two conditions where the orienting task did not match the invariant characteristic. That is, where the orienting task was either to search for high numbers with a low invariant, or conversely to search for low numbers with a high invariant. Performance was not significantly different from chance in both cases.

These effects were demonstrated by holding the orienting task constant and varying the stimulus in relation to it. Importantly, Experiment 5A and 5B differed in that they demonstrated a reversal of task demands, with a consequent reversal of learning effects. These effects were symmetrical, showing that there are no specific effects of stimuli or task demands. It is the interaction of the orienting task with the invariance that determines the success of learning.

Wright and Whittlesea (1998) have suggested that "what subjects learn is dictated by an interaction between their intentions and the structure of the particular instances they encounter" (p. 415). The data reported from Experiment 5 support this view of what underlies performance in implicit learning tasks. The conception that knowledge is acquired by a unconscious learning mechanism (e.g. Reber, 1989) appears to be refuted by the above data. When the task demands did not relate to the stimulus structure, no learning was observed. This suggests that the task demands have some relationship to the stimuli for learning to take place.

The processing account of implicit learning is able to make some predictions about the degree of learning by referring to the extent of overlapping training and test task demands. The separate systems view of implicit learning can make no prediction about the degree of learning. It would seem that the processing view, such as the Wright and Whittlesea (1998), provides a better account of the data. According to this account, the weak effects of Experiment 3 are the result of a small amount of overlap between the study and test phase.

Wright and Whittlesea (1998) do not specify clearly what is likely to be contained in verbal report data. In Experiment 5, subjects reported that the digits they believed occurred most frequently were those that corresponded to the orienting task they had performed. So, in Experiment 5A low numbers tended to be reported, and in Experiment 5B, high numbers tended to be reported. The verbal reports seem to match the training task demands which goes along with the processing account. However, a more complete explanation of this can be obtained by referring to the particular strategy used by subjects in the awareness test.

In the awareness test, subjects were asked if they noticed anything systematic about the stimuli in the training phase, and whether any of the digits occurred more often than others. Since they were unlikely to have considered anything about these issues during the training phase, they are likely to use available information or an availability heuristic to generate an answer. The availability heuristic (Kahneman and Tversky, 1973) is a strategy for evaluating the frequency of an entity based on the ease with which instances of this entity come to mind. For example, the risk of a train crash would appear greater if one has seen the report of such an incident on the news. In the current context, subjects had processed exclusively, the high digits of the strings (Experiment 5B) or the lower digits (Experiment 5A). Hence, these aspects of the stimuli are the most accessible information, and are therefore more likely to be reported.

The use of an availability heuristic can specify how another aspect of the verbal report data comes about. When the orienting task did not match the invariant, subjects tended to report the "9" or the "1", according to the training task demands. However, this was not the case where the orienting task demands matched the invariant. Here, the majority of subjects accurately reported the invariance from their particular training set, either "8" or "2". This demonstrates that subjects had noticed the invariant when it related directly to their task. The "8" or "2" occurred continuously during training, and therefore served as the most available information.

In summary, the main finding of Experiment 5 was that the interaction of the invariance (or structure) of the stimuli with the training task exerts powerful control over performance, even when using the same invariant characteristics as in previous studies. These effects were shown to be symmetrical when the task demands were reversed, indicating that it was not any special effect of the stimuli used. These data support the processing view of Implicit Learning proposed by Wright and Whittlesea (1998). Furthermore, no learning was demonstrated when the stimuli did not relate the orienting task, suggesting that learning does not proceed in a stimulus driven manner as some perspectives propose (e.g. Reber, 1989). The intentions of the subject appear paramount to the success of learning. This framework provides some explanation for the small amount of learning demonstrated when the orienting task is not directly related to the invariant characteristic, as in Experiment 3. The training task, therefore, not only has theoretical implications; it also has methodological connotations.

2.8 General Discussion

A consistent theme of this chapter has been the efficacy of the learning effects in the McGeorge and Burton (1990) task. The learning effects have been shown to be sensitive to surface features (Experiment 1), insensitive to a small number of negative items in the training period (Experiment 4) and modulated by the training task demands (Experiment 5). Throughout, the effects of each of these factors have been understood in the light of processing accounts of Implicit Learning and Memory (Whittlesea and Dorken, 1993. Wright and Whittlesea, 1998. Roediger, Weldon and Challis, 1989). The findings of each experiment will now be related to the others, with reference to processing accounts.

Sensitivity to invariant characteristics was shown to be eliminated when the surface form of the stimuli changed from the training phase to the test phase (Experiment 1). This finding is in line with subsequent data reported by Stadler, Warren and Lesch (2000), and therefore it is unlikely that the information acquired in this task is held at an abstract or conceptual level as McGeorge and Burton (1990) suggested. It is more likely that performance is driven by a match between processing events that occurred during the test phase, and processing events that occurred during training.

In an attempt to place these findings in a theoretical context, Stadler, Warren and Lesch (2000) compare specificity of surface form in the McGeorge and Burton (1990) task to similar effects in the implicit memory literature. At this point it should be noted that implicit learning and implicit memory have been defined as distinct phenomena (Buchner and Wippich, 1998). While implicit learning refers to the acquisition of relationships between stimuli, implicit memory reflects the facilitation of a response following single exposures of familiar stimuli (e.g. Jacoby, 1983). Implicit learning can be revealed in test conditions where the surface features are altered (e.g. Mathews et al., 1989) and implicit memory effects do appear to be sensitive to surface form (Jacoby and Hayman, 1987). As performance in the invariant learning task is specific to surface features, Stadler et al. suggest that it reflects processing similar to that occurring in implicit memory tasks rather than implicit learning tasks. Classifying invariance learning in this manner implies that its theoretical interpretation may be distinct from other forms of implicit learning.

This classification may not be required if there is a single theoretical interpretation of both implicit learning and implicit memory. Wright and Whittlesea (1998) have compared the findings of implicit learning tasks with those in the implicit memory literature, in this case using the processing account as a point of convergence. They cite the seminal paper by Jacoby (1983) to illustrate that the processing account of implicit memory is akin to the processing account of implicit learning, and describe how this account can explain both sets of findings. Jacoby (1983) demonstrated that subjects asked to generate the antonym of *hot* encode the meaning of the word *cold*, but not its visual properties, whereas reading *cold* selectively triggers visual properties but not conceptual aspects of the word. Aspects that are not required to satisfy the demands of the task remain unprocessed. This encoding variability reveals different patterns of performance on different tests. There is no need to assume that performance on one test is underlain by a different type of knowledge in comparison to another test. In an implicit learning task, the training phase requires certain mental operations to be carried out on stimuli. These may be related to the test task demands, or unrelated to the test task. Performance is dependent on the test task bearing some relationship to the training task demands.

This account is supported by the findings of Experiment 5. Where the invariant digit was low, and the orienting task induced subjects to search for low numbers, a significantly increased number of positive items were chosen in the forced choice test (Experiment 5A). This was replicated where the orienting task involved searching for high numbers. Here the training task induced processing of information that would later become important in the test task. According to Wright and Whittlesea (1998), during the training stage the invariant property had no special status, it is only when this property later correlates with the test task demands that it can assist in satisfying those demands. Conversely, where the orienting task was to search for high numbers with a low invariant, and vice versa, performance was at chance. This occurs because the information that is encoded during the test phase cannot be used at test to satisfy the demands of the test task. The same kind of processing occurs when, as in Jacoby (1983), a subject has experienced the visual properties of *hot*, and the semantic properties are required at test, performance is likely to fail because the encoded information is of no use to the test demands. If the test requires the visual properties of *hot*, the subject would be able to draw on the previous visual processing of the word *hot*, resulting in improved performance. Hence, the data reported in implicit memory and implicit learning tasks can be understood within a single set of principles.

Returning to the findings of Experiment 1, the test task demands require subjects to indicate which of two digit strings are most familiar to them. According to the processing account (Roediger, Weldron and Challis, 1989), once an item has been encountered, this results in later presentations of the item being processed with greater ease (i.e. more fluently), biasing it towards selection at test. When the surface features change, this fluency mechanism can no longer operate, giving the positive item no advantage over the negative. Hence, specificity of surface form is a consequence of this view of performance. This account assumes that the application of similarity based mechanisms in tandem with fluency. That is, at test subjects compare the items in the forced choice test to previously seen exemplars in training, so selecting the item which is most similar to an exemplar from training.

Experiment 4 suggests that the use of similarity or a episodic mechanism is a reasonable assumption to make concerning performance in this task. Following on from work by Cock, Berry and Gaffan (1994), who revealed the use of similarity at test, Experiment 4 demonstrated that when the training set no longer followed an invariance rule, it merely had a majority feature, performance did not deteriorate significantly. This indicates that the above chance selection of positive items does not have to occur when there is an invariant, it can also occur in conditions when the positive item is more prevalent in the study set. This occurs because, at test, the positive item will be similar to a larger number of episodes of stimuli than the negative, as a result of the presence of the invariant feature. It would not make any difference if the feature occurred in every of the strings, this invariance has no special status.

Along with the acknowledgement of specificity of surface features, these findings are a problem for an account of invariance learning that purports the abstraction of an invariance rule from the training materials, such as McGeorge and Burton's

(1990) original understanding of performance. Furthermore, in Experiment 5, chance performance is observed when the orienting task does not relate to the invariant property, demonstrating that the sensitivity to invariance is not an automatic consequence of interacting with the stimuli. This causes further problems for any account suggesting that subjects acquire knowledge of invariance that is a special quality acquired from the stimuli, regardless of surface form or processing conditions, held separately from the stimuli themselves.

The experiments reported in this chapter have been primarily concerned with manipulations of the training task conditions. In experiment 3, it was observed that certain training tasks used on stimuli can produce relatively small learning effects at test. This was interpreted as demonstrating a small degree of low overlap of study and test processes. Nevertheless, it may be that these training conditions can reveal more significant learning effects on a test task that is more appropriate to the processes carried out during the test. It is the purpose of the following chapter to explore the role of test task demands in the invariant learning task.

Chapter 3 The Effect of Alternative Test Task Demands

3.1 Introduction

In the previous chapter, it was demonstrated that the information acquired in the invariant learning task is tied to the surface features of the training stimuli. It was argued that this finding shows that subjects select items at test according to their fluency or ease of processing. Such fluency, however, is only one mechanism for making decisions about stimuli. Whittlesea and Leboe (2000) suggest that people can use another heuristic in remembering and classification tasks, the generation heuristic. According to Whittlesea and Leboe (2000) the generation heuristic is "based on the production of information about a prior experience with a stimulus, information that is not available in the current stimulus display" (p85). False memory effects provide some evidence for the use of this heuristic. For example, when given a word "sleep" following a set of items that are in a particular context (e.g. dream, bed, night, snore), subjects falsely claim to have seen a target word in that context (Roediger and McDermott, 1995). Hence, subjects are generating information about the training period that the test stimulus does not itself possess, and are using this information to aid responding.

Studies on invariance learning, and indeed implicit learning in general, have tended to rely on classification of novel stimuli revealing knowledge about a given stimulus domain. It may be that these tests of knowledge rely mainly on the fluency with which the test items are processed to guide decisions about stimuli. This leads to the question of whether the fluency of processing of test stimuli is critical for any test task to reveal learning effects. It may be that sensitivity to invariance, for example, can be revealed in other ways. As noted above, subjects have been shown to generate aspects of the training stimuli which are not present in the test stimuli to make judgements. It may be possible to devise a test task that focuses on this generation aspect of test processing rather than the fluency of processing test items alone. If learning effects cannot be revealed by the generation of information, it is likely that knowledge of invariance can only be revealed by an increased fluency of processing novel test strings.

3.2 Generation of an Invariant Feature

In the following experiments, a test task is introduced that encourages subjects to generate features of stimuli at test that they have encountered in the study phase. To this end, we apply a modification of the fragment completion task (e.g. Tulving, Schacter and Stark, 1983) to induce generation of an invariant feature. In the Tulving et al. study participants are presented with incomplete words (e.g. _O_O_GA_) and asked to fill in the blanks to make up a word. In this task, half the solutions were words from a target list, but the subjects were not informed of this. Fragment completion therefore served as a test of implicit memory. Tulving et al. found that subjects were more likely to complete fragments correctly when the solution corresponded to a target word. They found that correctly recognised words were not completed any better than words that subjects failed to recognise. Therefore, fragment completion could serve as a valid implicit memory test.

This test of memory can easily be applied to the invariance learning task. The first phase of the experiment would be carried out in an identical manner as in previous invariant learning studies. The test materials presented to participants would consist of three digits and space (e.g. 7_26, 928_ etc). The space can be placed in any of the four positions. These numbers could be constrained to include the invariant or it could be absent from the string (e.g. invariant present = 43_1 or invariant absent = 291_). The subject's task is to complete the fragment with a digit to make up a four digit number they saw in the study phase.

Altering the demands of the test task may result in a shift in the processing resources used at test in comparison to forced choice materials. Therefore, in addition to investigating generation of knowledge (experiments 6 and 7), the following experiments examine the role of processing in this task. The fragment completion task induces subjects to generate potential candidates to complete the string. In a forced choice task subjects are asked to choose which of two numbers they have seen before. These tasks differ in that the former involves articulation of digits whereas the latter does not. Experiment 8 examines the effect of suppressing articulation during the study phase. A reduction in performance would indicate that the overlap in processing resources between study and test must be an important determinant of performance. This finding would be consistent with processing theories of implicit learning (see Neal & Hesketh, 1997; Whittlesea & Dorken, 1993). Stadler, Warren and Lesch (2000) demonstrated that invariance learning is sensitive to surface form using forced choice materials. If the test task demands are changed so that subjects generate information about the stimuli, it is likely that surface feature changes will no longer influence performance. Experiment 9 investigates this possibility.

To date, studies of invariance learning have increasingly supported a processing basis of successful performance, rather than the application of abstract knowledge. Cock, Berry, & Gaffan (1994) demonstrated that similarity to instances in the study phase was a more important factor than apparent knowledge of invariance. Thus, reference to specific episodes during study can determine performance. More recently, Stadler, Warren and Lesch (2000) showed that a change in surface features can reduce performance to chance levels. This finding is again consistent with the processing view of implicit learning.

In summary, the following experiments are primarily intended to demonstrate that knowledge of invariant features can be demonstrated in a generation task rather than a forced choice test. The secondary aim of these experiments is to investigate the effect of changing processing demands on fragment completion performance.

3.3 Experiment 6

The purpose of this experiment is to determine whether it is possible to demonstrate learning of invariant features in the McGeorge and Burton (1990) task using a fragment completion task. To this end, the test task materials were contrived in two ways. Absent strings were created by removing the invariant from a invariant string (e.g. 3451 would become _451). Present strings were created by removing one of the three other digits from a invariant string (e.g. 3451 could become 3_51). Present strings were included to prevent response bias. If participants produced an invariant in literally every test string then this may induce them to inhibit production of invariant strings in later trials.

Certain predictions can be made about performance on this test. If the study set encountered contains an invariant digit in every string then it would be expected that subjects may produce more invariant digits than other digits in the test. That is, participants will produce more invariant digits in the Absent condition when they have received an invariant digit in the study set materials (Invariant Group), than where they have received four naturally occurring numbers (Random Group) study set materials.

A weaker prediction can be made about performance in the Present condition, if subjects have implicitly acquired information about a single invariant digit then they may inhibit production of invariant digits if one is already present in the string. Hence, this test may provide a more comprehensive assessment of knowledge acquired during the study phase by potentially examining production and inhibition of an invariant digit.

However, this prediction is made with some caution because the strings are constrained to contain no repetitions which have been shown to be highly salient (Wright and Burton, 1995). Consequently, subjects may not produce repeated digits in the strings because these doubles did not occur in the study phase. This hypothesis can be tested by examining the production of repetitions of all digits and comparing this score to the production of invariant digits in the test phase.

Method

Participants. Twenty undergraduate students from the University of Glasgow took part in this experiment. They were paid a small fee for their participation.

Materials. Study sets were generated individually for each subject. The sets were generated separately for the Invariant and Random study set materials groups. In the Invariant group a set of 30 items was produced for the study set. These were constrained so the strings contained an invariant "3" digit. In the Random group a set of 30 four digit numbers were generated so all numbers 1-9 would appear naturally. The strings in both groups were constrained in order to contain no repetitions. A further 10 Positive and 10 Negative items were generated for each subject in both groups as the test items. As in the study phase, none of these items contained repetitions. The test items were constrained so none had appeared as study items previously. These test items were then altered to form fragments in two forms for the Present and Absent condition. In the Present condition, the 10 Positives for each subject were selected and one digit removed at random which was not the invariant and replaced with a "_". For example, 4369 would become one of "_369, 43_9 or 436_". In the Absent condition, each of the 10 Negatives for each subject was selected and again one digit was removed and replaced with a "_". Since no invariant is present in the Negatives one of the four digits was selected. For example, 2815 could become one of " _815, 2_15, 28_5 or 281_". The study and test items were presented with each string appearing on a page of a small booklets.

Design and Procedure. As in experiment 1 – 5, the experiment was divided into three phases: a study phase, a test phase and a test of explicit knowledge.

Subjects were randomly allocated to receive Invariant or Random study materials which are explained above. Apart from the constraints imposed on the four digit strings the experimental procedure was identical for both groups. The four digit study strings were presented on individual pages of the small booklet. Subjects were asked to perform the same arithmetic task as in experiment 1 - 3.

They were then presented with the surprise test phase consisting of 10 string fragments from each of the 2 conditions; *Present* or *Absent.* The items were presented on individual pages of the small booklets. These items were randomised, so that no systematic order of presentation occurred. Participants were falsely told that the fragment represents three quarters of a whole number they had seen in the study phase. They were asked to fill in the gap with a digit so that the fragment would make a whole number that they had seen in the study phase. Like the study phase, they were told to complete each page of the booklet without referring back or forward to other test strings. Following the test phase subjects were given the awareness questionnaire.

Results

The means for the two factors are presented in the table below.

Study set	Invariant produced in test string (out of 10)			
	Present		Absent	
	Mean No.	SE	Mean No.	SE
Invariant	0.2	0.13	2.30	0.30
Random	0.3	0.21	1.1	0.38

An Analysis of Variance revealed a significant main effect of study materials, $F(1, 18) = 4.76$, $p < 0.05$. It also demonstrates a main effect of test materials, $F(1, 18) = 24.82$, $p < 0.01$. The Study set x Test materials interaction was also significant, $F(1, 18) = 4.99$, $p < 0.05$.

In the Absent condition participants in the Invariant group produced significantly more invariant digits than in the Random group, $F(1, 36) = 6.17$, $p < 0.05$. This difference was not replicated in the Present condition, $F(1, 36) = 0.16$.

In the Invariant group, the mean number of invariant digits produced in the Absent condition was significantly higher than in the Present condition, $F(1, 18) = 26.03$, $p < 0.01$. The difference was not found in the Random group, although the effect was marginal, $F(1, 18) = 3.78$.

Analysis of repetition production - The means for digit repetitions were divided into two types, repetition of invariants (43_2 completed as 4332) or repetition of other digits (43_2 completed as 4322). Using these means it is possible to determine if " 3 " is produced in present strings with a lower frequency than other digits that are repeated.

Subjects produced repetitions of digits other than "3" in 2.5% of Present test strings. In comparison, the mean number of "3" digits produced in the Present condition was 2%. A paired samples t-test revealed that the production of invariant digits and repetitions of other digits are not significantly different, $t\ (19) = 1.0$.

The arithmetic task was carried out correctly on 95% of the study strings in the random group and 95.3% of strings in the invariant group.

Test of Awareness: No participant reported anything related to the invariant in response to the first question on the explicit knowledge questionnaire. One participant circled "three" in response to the final question but did not mention it in response to the first two. No subject reported anything pertinent to repetitions in response to the first question.

Discussion

Constraining the materials received by subjects in the Invariant group to contain an invariant digit resulted in the predicted difference in relation to the Random group. Participants produced a significantly larger number of invariant digits when they had received an Invariant study set than when the study set they received contained random naturally occurring numbers. This main finding confirms the prediction made that subjects would produce the invariant with higher frequency under the conditions described above.

It was suggested in the introduction that when the invariant was already present in the test string, implicitly acquired information may prevent subjects completing the string with a repeated invariant digit. Every string subjects encountered contained an invariant digit but never more than one, therefore it is possible subjects may have acquired this information. The analysis of invariant production seemed to imply that subjects were indeed selectively inhibiting invariant digits in the Present condition. Analysis of the group that received invariant material revealed that production of invariant digits was significantly lower in the Present than the Absent condition, whereas this difference was not replicated in the group that received random materials. In addition, the difference between the Invariant group and the Random group also did not significantly differ in the Present condition. Therefore, it is possible that the invariant is inhibited when it is present in the string. But, as aforementioned, the study set materials were constrained so that they did not contain repetitions. Analysis revealed the production of repetitions in general was not significantly different from the production of invariant digits in the Present condition. This suggests that the inhibition of "3" in the Present condition may reflect a floor effect which is a result of subjects inhibiting production of repeated digits as a matter of course.

Inhibition of repetitions may occur because subjects did not see any repetitions in the study set and therefore did not complete strings in the test phase with repetitions. Strings with repeated numbers tend to be salient (as in experiment 1 and Wright and Burton, 1995) and would therefore be remembered by subjects.

Some subjects could have noticed the repetition constraints of the study set and used it as a strategy to aid their memory.

3.4 Experiment 7

In experiment 6 it was shown that the fragment completion test can reveal varying effects when subjects have been exposed to training strings that all contained an invariant digit compared to where all the digits appeared naturally in the training strings. This sensitivity was primarily revealed by the increased production of '3' digits in the absent test materials (e.g. those that did not contain '3' – 72_1). A potentially significant secondary finding was that subjects appeared to inhibit production of the invariant '3' in the present test materials (e.g. those that did contain the '3' – 13_5). It may be that this effect is the result of implicitly acquired information that each training contained a single '3' digit. This possibility was undermined by the fact that subjects appear to produce repetitions of all other digits apart from '3' at the same level that they produced the invariant '3' in the present strings. This indicates that the test materials may bias subjects against generating any digits that are already present in the particular test string (e.g. subjects biased against completing 42_1 as 4221 etc). This effect would generalise to the invariant, resulting in an apparent bias to inhibit the '3' in the present strings.

It was suggested above that this bias to inhibit the production of repetitions in the test strings may be the result of the absence of items containing repetitions occurring during the training phase. Since subjects did not see any training items with repetitions, they did not complete the test strings with repeated digits.

Alternatively, this effect may be the result of strategic processing during the test phase, with no reference to the training items. Test items, like training items, appear salient when completed with a repeated digit. Hence, subjects may be biased against generating digits that produce salient items regardless of what items they saw during training.

The most accessible way to differentiate between these possibilities is to include items with repetitions in the training phase. If the floor effect of producing repetitions results from the fact that subjects did not see training items containing repetitions, subjects will now have seen a number of items with repeating digits and may complete test strings with repetitions. On the other hand, if subjects fail to generate repetitions as a result of a test bias against generating digits that result in salient strings, the repetition items presented during training should have no effect.

In addition, the strings will be constrained by preventing the '3' digit from repeating in the training strings, so that only one '3' could occur in the training items. Hence, if subjects show a tendency to repeat digits other than '3' in the test materials, while inhibiting production of the '3' digit in the present strings, it may be that they have become sensitive to the singular occurrence of the '3' digit in the training strings. However, if the '3' is repeated in the present test strings (e.g. 23_1 is completed 2331) at the same level as other repeats (e.g. 2_59 is completed as 2559), subjects do not become sensitive to the fact that the '3' digit occurs only once in each string.

Method

Participants. Twenty undergraduate students from the University of Glasgow took part in the experiment. They were paid a small fee for their participation.

Materials. In the Invariant group a set of 30 items was generated for the study set. These were all Positive items with certain constraints regarding repetitions. All digits apart from "3" appeared naturally, that is they were not prevented from repeating. In the Random group a set of 30 four digit numbers were generated so all digits 1-9 would appear naturally. There were no constraints on repetitions within the strings. A subset of ten of the thirty training strings could be -

3254, 3621, 7437, 3622, 1138, 1326, 9349, 9312, 3682, 1653

The test items were generated in exactly the way as Experiment 6.

Design and Procedure. The experiment followed the same three phase procedure using an identical orienting task on study items as in Experiment 6. To clarify; the only alteration was the constraints on repetitions in the study set.

Results

The means for the two factors are presented in the table below.

Study set	Invariant produced in test string (out of 10)			
	Present		Absent	
	Mean No.	SD	Mean No.	SD
Invariant	0.8	1.14	2.0	1.25
Random	0.7	0.67	1.0	0.94

An Analysis of Variance revealed a significant main effect of Invariant, $F(1, 18) = 6.39$, $p < 0.05$. There was no significant interaction or main effect of Group, the largest value for $F = 2.49$.

In the Absent condition, participants did not produce significantly more invariants in the Invariant group than in the Random, although the effect was marginal, $F(1, 18) = 4.09$, $p < 0.1$. Again in the Present condition there was no significant effect, $F = 0.06$.

In the Invariant group, participants produced more '3' digits in the Absent condition than in the present, $F = 8.18$, $p < 0.01$. This difference was not found in the Random group, $F = 0.5$.

The arithmetic task was carried out correctly on 99.6% of the study strings in the random group and 98.3% of strings in the invariant group.

<u>Analysis of Repetitions.</u> This analysis was carried out in the same way as in Experiment 4. Repetitions of digits present in the test strings were produced in

both the absent and present conditions in 8.6% of test strings. This was significantly larger than zero, $t(19) = 6.4$, $p < 0.05$. A paired samples t-test reveals that the difference between production of repetitions in general and production of invariants in the present condition are not significantly different, $t(19) = 0.23$.

Test of Awareness. No participant reported the invariant in response to the first question on the explicit knowledge questionnaire. Two participants circled "three" in response to the final question. They had mentioned the invariant in answer to the second question, but among other digits and were therefore not removed. No subject mentioned the repetition of numbers in any sense.

Discussion

The indication of this experiment is that repeated digits are produced in fragment test strings when the training strings contained repeated digits. Subjects placed repetitions of digits already present in the test strings significantly above zero in this experiment. When comparing this finding to that of experiment 6, the conclusion must be that the presence of repetitions during training is critical for the production of repetitions during the test. Consequently, this rules out the possibility that subjects are biased against producing repeated digits regardless of what items they are exposed to during the training phase. The floor effect in the present condition of experiment 6 was conceivably due to the lack of repetition items in the training set. Subjects did not see any repetition items in the training phase, and therefore did not complete any strings with a repeated digit. This inhibition of repeating digits could have occurred by explicit or implicit recourse to the training strings, and this issue will be investigated in experiment 8.

In the Invariant group, the '3' digit was constrained so that it would occur in every string and it would occur only once in every string. This constraint did not result in a selective inhibition of the '3' digit in the Present condition. The '3' digit was produced as often in the present condition as repeated digits were produced in general. Furthermore, there was no difference in the production of '3' digits in the Present condition by subjects in the Invariant or Random groups. Therefore, the presence of an invariant '3' had no effect on the production of '3' digits in the Present condition when repetitions were introduced into the training phase. The apparent inhibition effect that was found in the Present condition in experiment 6 clearly due to the lack of items with repetitions in the training phase, leading to a general bias against the production of repetitions at test.

It would seem then, that the effect of the invariant '3' during training is only detected in the Absent test condition. The Present condition does not show any sensitivity to the invariant '3', even when the training materials are constrained so that it does not repeat. Hence, the sensitivity is demonstrated only where the digit is produced at greater levels than usual. Although this finding has refined the locus of sensitivity of the fragment completion test, it does not provide an explanation of the processes that are involved in this increased production of digits. The processes that give rise to this effect are investigated in experiment 8.

3.5 Experiment 8

In experiment 6 it was shown that the fragment completion task can show differential effects when subjects have performed an orienting task on strings that all contain an invariant digit and where the strings are random. The next experiment aims to define the mechanism which underlies performance on the fragment completion task.

The increased production of the invariant '3' may occur because subjects have acquired some conceptual knowledge of the invariant as suggested by McGeorge and Burton (1990). Fragment test strings may be completed by the application of this abstract conceptual knowledge. Accounts of implicit learning that emphasise abstraction assume that acquisition and application of knowledge occurs passively. Therefore, processing conditions during the study phase should have little influence on the passive acquisition of stimulus structure. This provides a means of testing the applicability of this account to performance on this task.

In the previous chapter, the episodic processing account of implicit learning has been put forward as a more comprehensive account of performance in this task. The importance of overlapping processing constraints may provide an alternative interpretation of performance in this task. In the completion task, subjects must articulate potential candidates for completion of the strings. According to the episodic view, this pool of candidates must be triggered by a previous learning episode which involves similar processing to that occurring in the test phase. The phonological process during test may then rely on the phonological articulation of digits during the orienting task. If this articulation is prevented, then the useful overlapping information used to generate digits may be removed resulting in a reduction in performance to levels comparable with the control group.

In addition to the influence of verbalisation during orienting, an aspect of Experiment 6 that requires further scrutiny is the floor effect that was demonstrated in the Present condition of both groups. Analysis revealed that subjects were inhibiting production of all repeated digits at test and not just invariant digits. It was suggested that, since the study set did not contain repetitions, subjects would avoid producing digits that were already in the fragment. Experiment 7 demonstrated that if the training strings do contain repetitions, then subjects complete some test strings with repeated digits. It is surmised that this process is under subjects' conscious control or it may even be an implicit effect.

A questionnaire was presented to participants with the intention to examine knowledge of repeated digits. Firstly, subjects were asked if they used strategies during the test phase; they then commented on strategies that involved avoiding the numbers present in the test string and, finally, they were given examples of numbers containing repetitions and were instructed to indicate how they were different from those in the study set.

An additional aspect of this experiment investigated the generation of numbers in neutral conditions. That is, what do people produce without the cue of the digits in the test string? Subjects were asked to generate random numbers between 1 - 9 and place them in the gap between three other addition signs " e.g. + + _ + or +

_ + + ". By removing the digits from the test string no context is available for subjects to use. Symbols were used to make this section as close as possible to the standard test. This was self paced and followed the fragment completion test.

Method

Participants. Forty three undergraduate students from the University of Glasgow took part in this experiment. They were paid a small fee for their participation.

Materials. Study set and test materials were generated in an identical fashion to that of Experiment 1. In place of the invariant "3" an invariant "5" was used in this experiment.

Design and Procedure. The experiment was divided into four phases: a study phase, a test phase, a neutral generation phase and a test of explicit knowledge.

The experiment was a mixed design with two between subjects factors: Suppression vs No Suppression during the study phase, and Invariant vs Random study materials. Suppression conditions were performed by the repetition of a phase "alpha beta gamma" out loud by participants. Before commencing, subjects were trained in the correct manner of articulation. They were encouraged to keep articulation as continuous as possible, and endeavour not to leave gaps between the words. Participants were asked to practise this phrase to the above specifications, when they had become proficient the experiment began.

Subjects performed the same test as in Experiment 7 except that the strings were contrived to have a "5" digit present or absent. Again, the subjects were asked to complete the string with a digit that they believed made up a four digit number observed in the first part of the experiment.

Following this, the subjects were presented with a booklet containing 10 random generation materials. These consisted of three plus signs with a underscore randomly among them (e.g. + + _ +). Subjects were asked to place a digit 1 - 9 randomly in the gap. The generation was self paced.

In addition to the standard post task questionnaire, a second sheet of additional questions were presented, which followed a graduated format. The final question which presented subjects with numbers with repetitions was not visible until the paper was turned over. The question are presented below.

4. Were most of your responses guesses?

5. Did you use any strategies when filling in the spaces?

6. Did you have a strategy for avoiding numbers depending on the other digits in the test string? If so, how?

7. How are the numbers listed below different from those you did the addition and comparison with earlier?

4423

7355 and so on.

Results

The means and standard error for the three factors are presented in the table below. Three subjects in the invariant group showed good awareness of the invariant digit, and these subjects were removed from the analysis.

		Invariant produced in test string (out of 10)			
Suppression	Training set	Present		Absent	
		Mean	SE	Mean	SE
No Supp	Invariant	0.1	0.10	3.0	0.37
	Random	0.2	0.13	1.0	0.30
Supp	Invariant	0.2	0.13	1.7	0.37
	Random	0.2	0.13	1.6	0.31

An Analysis of Variance revealed a significant main effect of study materials group $F(1, 36) = 5$, $p < 0.05$ and test materials $F(1, 36) = 99.5$, $p < 0.05$. The three-way interaction of articulation, study materials and test materials was significant, $F(1, 36) = 5$, $p < 0.05$.

In the Absent condition, subjects in the No Suppression group who were presented with invariant materials produced significantly more invariant digits than subjects in the No Suppression group who received random materials, F (1,

36) = 27.2, $p < 0.05$. This difference was not replicated in the Suppression group, $F(1, 36) < 1$. Examination of the effect of Suppression in the Absent condition reveals that when subjects received invariant study materials they produced more invariant digits under No Suppression conditions than under Suppression conditions, $F(1, 36) = 11.4$, $p < 0.05$. This difference did not occur when subjects received random materials, $F(1, 36) = 2.4$.

The arithmetic task was carried out correctly on 94% of the study strings in the Suppression group and 97.6% of strings in the No Suppression group.

Production of digits in neutral strings. The means for invariant production in each condition are shown below.

Study set	Invariant digits produced in the neutral test strings			
	Suppression		No Suppression	
	Mean No.	SE	Mean No.	SE
Invariant	1.2	0.20	1.1	0.23
Random	0.7	0.26	1.0	0.21

An Analysis of Variance revealed no significant effects of any of the factors. The difference in invariant production between subjects who received random materials and those who received invariant materials failed to reach significance $F(1, 36) = 1.7$.

Test of Explicit Knowledge. Three participants were removed from the analysis because they verbally reported knowledge of the invariant digit in question 2 of the explicitness test and reinforced this by indicating the invariant digit as the most frequent in question 3.

In response to question 4, all subjects admitted to guessing most of the time during the completion test. In question 5, 12 subjects reported using memory for the results of arithmetic done in the orienting task as a cue to generate digits for the test. In answer to question 6, 43% of participants said they had avoided the digits that were already present in the string as a strategy. The remaining participants did not mention anything pertinent to repetitions in response to this question. However, 83% of participants noted that the numbers presented to them in question 6 were different from those in the study set because they contained repetitions.

Discussion

Participants who did not perform a phonological suppression task produced similar performance in the test phase in comparison to participants in Experiment 6. That is, constraining the study strings to contain an invariant digit in the Invariant group produced the predicted difference in relation to the Random group. This difference was not replicated, however, when participants had performed a phonological suppression task during encoding. This main finding suggests that phonological encoding is required for subjects to demonstrate increased invariant production following exposure to invariant strings. Therefore, it is clear that invariant properties are not encoded passively but phonological articulation during study is necessary for completion of test strings.

The floor effect that was demonstrated in the Present condition of both Experiment 6 and 8 appears to be the result of an explicit bias to avoid production of a digit that is already present in the test string. In the questionnaire, 43% of subjects reported that they used the strategy of avoiding digits which already occurred in the test string. When presented with strings containing repetitions, 83% of the sample were able to verbalise the point that these test strings contained repetitions while the items in the study set did not. On the more specific question, subjects were more easily able to compare these items with instances they saw in the study set and the difference became apparent to a greater extent.

Regarding the neutral test conditions, participants showed no bias to production of invariant digits when the test items did not provide any cue to the strings in the study set. Therefore, this process is not simply a passive production of digits based on some very crude frequency priming of digits that were processed in the study phase.

3.6 Experiment 9

From the data presented in Experiment 8, the mechanism underlying performance on the completion task appears sensitive to phonological interference during the study phase. Thus, the overlap between the processing that occurs at study and test is an important determinant of performance in the invariance learning task. This finding is consistent with episodic processing views of implicit learning.

One limitation of this explanation is that it assumes that a different process underlies performance on the completion test in comparison to the standard forced choice procedure. This assumption cannot be made solely on the basis of the data from Experiment 8. A dissociation between performance on this completion task and existing forced choice data needs to be demonstrated in order to clarify that distinct processing occurs in these test tasks. A potential candidate for this is the transfer of the effect across different surface features. If the phonological information determines string completion performance then a change in surface features should not interfere with the effect. In the case of the forced choice test, performance was reduced to chance levels when surface features were changed (Stadler, Warren and Lesch, 2000). Thus, in Experiment 9, subjects were presented with study materials in which the four digit numbers were written out as words (e.g. four three nine one). The test materials were the same as those presented in Experiment 6 and 7.

Method

Participants. Twenty seven undergraduate students from the University of Glasgow took part in this experiment and they were paid a small fee for their participation.

Materials. Study and test materials were generated in an identical fashion to that of Experiment 6 with the exception that the study set strings were written out as words rather than as digits (e.g. three nine two one).

Design and Procedure. The experiment was carried out in the same way as Experiment 6 with the exception that the same extended questionnaire was presented to participants as in Experiment 8. This constituted the third phase of the experiment.

Results

The means and standard error for the two factors are presented in the table below. Three subjects in the invariant group were shown to have good awareness of the invariant digit and were removed from the analysis.

Study set	Invariant produced in test string (out of 10)			
	Present		Absent	
	Mean No.	SE	Mean No.	SE
Invariant	0.4	0.2	2.2	0.3
Random	0.3	0.2	1.4	0.3

An Analysis of Variance revealed a significant main effect of test materials $F(1, 22) = 37.5$, $p < 0.05$. The main effect of study materials and interaction of test and study materials failed to reach significance. However, in the absent condition the production of "3" was significantly higher in participants who had received an invariant study set as opposed to those participants who had received random study materials, $F(1, 44) = 4.15, p < 0.05$.

The arithmetic task was carried out correctly on 97% of the study strings in the random group and 96.7% of strings in the invariant group.

Test of Awareness. Three participants were removed from the analysis because they reported knowledge of the invariant in response to question 2 of the awareness test and reinforced this by indicating "3" as the most frequent in response to question 3.

In response to question 4, all subjects admitted to guessing most of the time during the completion test. In question 5, arithmetic was cited by 50% of participants as the main strategy they used. In answer to question 6, 33% of participants reported that they had avoided placing repetitions of numbers already present in the test string. In response to question 7, 75% of participants noted that the numbers in the example were different to those in the study set because they contained repetitions.

Discussion

Participants who had received study materials in which every string contained a "three" produced significantly more "3" digits when completing the Absent strings than participants who received four naturally occurring numbers. This replicates the findings of Experiments 6, 7 and 8. It also demonstrates that changing surface features does not substantially affect performance on the completion task. It is suggested above that the completion task involves a different type of processing in comparison to the forced choice procedure that has been used in other studies on invariance learning. This position is supported by the data as changes in surface features do not substantially affect performance, unlike the findings of Stadler et al. (2000) where sensitivity was apparent.

Experiment 8 demonstrated that repetition structure is an aspect of stimulus structure that subjects may use as an explicit strategy; this finding is replicated in experiment 9. In the questionnaire, 33% of subjects reported the use of a strategy for avoiding digits that already occur in the test string. This is slightly less than the level in Experiment 8. Despite this, when presented with strings of digits containing repetitions, 75% of participants noticed and reported the difference between the numbers presented and those in the study set.

3.7 General Discussion

The data reported here demonstrate that it is possible to elicit knowledge of invariant properties of digit stimuli using a fragment completion task. In experiment 6, subjects who processed strings containing an invariant digit completed test strings with an invariant digit significantly more often with this invariant than people who processed strings with naturally occurring numbers. This effect was revealed only by an increased production of invariant digits in the absent test fragments. Invariant production in the present test fragments was somewhat obscured in experiment 6 by a bias against producing repeating digits. When this bias was accounted for in experiment 7, there was no tendency for subjects to show sensitivity to the *singular* occurrence of the training invariant in each string. This suggests the representation of the invariance is not very specific. Instead, the information may be more general; referring to the frequency of occurrence of the digits in the training strings or about specific strings themselves. However, claims on the information used in this task cannot be made on the basis of experiment 6 and 7.

The effect was replicated consistently across the four experiments reported in this chapter. Studies of invariant learning that use forced choice test materials have exhibited inconsistent patterns of performance (see experiment 3). It could be that these forced choice tasks require different processing resources than are used in completing strings with an invariant digit. Indeed, considering the nature of what the two test tasks require people to do may aid the understanding of the difference in their characteristics of performance. This focus on processing demands is consistent with episodic processing accounts of implicit learning (Whittlesea and Dorken, 1993). The overlap in processing resources engaged during study and test is the mechanism by which episodic processing theories of implicit learning operate. This account is supported by the data presented here. The completion test task requires people to articulate potential candidates to fill the fragments that are presented at test. When articulation is prevented in the study phase, articulation at test has no similar processing episode from the study phase to map onto. From the episodic framework, it follows that performance in the test is inhibited, this was demonstrated in the data from Experiment 8.

The difference in processing used by fragment completion in comparison to forced choice test materials was evident in Experiment 9, where performance was consistent over changing surface features. This finding is unlike Stadler et al. (2000), who demonstrated no sensitivity across surface form with forced choice test materials.

The apparent transfer revealed in Experiment 9 can also be explained within the processing account. Data from Experiment 8 indicates that articulation of digits is the resource required for performance in the completion task. The change of surface features should not interfere with articulation, hence performance in the completion task is not affected. Paradoxically, it is the consistency of processing resources that results in apparent transfer in the completion task. The same argument can be applied with regard to the lack of transfer shown in studies using forced choice test materials. The forced choice task requires participants to compare whole strings and make a decision on which string was present in the study set. Thus, the processing that this task draws on is more likely to be visual representations of strings seen in the study phase. If surface features change then consistency of processing resources will not occur and performance drops to chance.

This explanation for the pattern of data reported here is consistent with the recent processing account of Implicit Learning put forward by Whittlesea and Wright (1997). They call this account of Implicit Learning, "Learning without knowing the consequences". This view of implicit learning emphasises that subjects are unaware that " processing a particular item this way rather than that way... they are in fact exercising an option to prepare for the future in a specific way" (Pg 196). Considering the fact that subjects will process many aspects of the stimuli in many ways, it seems that they are prepared to perform in a number of different ways to a number of test tasks. In fact Wright and Whittlesea (1998) have made a similar point: "At test knowledge is better understood as a set of resources to be drawn on for many different purposes in interacting with the world" (p415). After encountering the study materials in the McGeorge and Burton (1990) task, subjects appear to use knowledge of the study materials differently according to the demands of the test materials. They adapt to test circumstances, rather than apply knowledge passively.

The experiments reported in this chapter demonstrate that knowledge of invariance can be generated given the correct test circumstances. This reinforces the point that subjects are actively involved in satisfying the demands of the test. When presented with the test stimuli, this triggers the generation of potentially useful experiences of the training phase which subjects apply during the test. In the experiments reported in this chapter, the fragment completion test captures this process of generation. Furthermore, this use of generation indicates that the fluency of processing test stimuli is not the only route to revealing knowledge of invariance. Indeed, subjects can apply different heuristics or strategies during the test phase according to the demands of the test task. These heuristics focus on different experiences of the training stimuli, resulting in differing sensitivity to training task manipulations.

On the issue of awareness, Wright and Whittlesea (1998) note that information can become explicit in implicit learning tasks when the task carried out during encoding maps onto the same information as the task carried out at test. They comment that in carrying out the demands of the training task, information may be encoded that is only indirectly related to the training task and this knowledge may be used to aid performance at test. Since this knowledge is not the focus of the training task, it is not processed in conscious awareness. Taking the task presented here as an example, one of the questions in the awareness test probed knowledge of strategies used during the test phase. In responding to this the majority of subjects reported using strategies involving arithmetic as an aide to memory. This is not surprising because subjects were induced to perform arithmetic as the orienting task. The key point is that none of these strategies would help them perform in the completion test or come up with useful information in the awareness test. Similar to a suggestion by Whittlesea & Dorken (1997) concerning tests of awareness, subjects were operating under the wrong theory of what knowledge is relevant and hence they fail on the awareness test. However, they can use some of the indirectly encoded information to perform at above chance levels in the completion task. This knowledge is not available to consciousness because it was not the focus of the training task. Such indirectly encoded information may be compatible with certain operations that are carried out in the test phase, in which case, it is revealed in the test phase. For example, in the completion task, digits were encoded phonologically in the study phase and this aided performance in a test phase that required such phonological information.

This analysis of the information used in the generation test does not precisely specify which aspects of the stimulus are used to perform at above chance levels.

The next chapter will attempt to address the question of what information is used during the test.

Chapter 4 The Information Acquired in the Invariant Learning Task

4.1 Introduction

Reber (1989) suggested that the information acquired in implicit learning tasks is in the form of abstract rules. For example, in artificial grammar learning (see sect 1.5), Reber (1989) proposes that subjects acquire knowledge of the rule structure that determines the structure of the training items. However, this understanding of the knowledge acquired in artificial grammar learning has proved contentious. One view is that it is more likely that subjects classify test strings by making analogies between individual training and test strings (e.g. Brooks, 1978). This analogical processing may be successful because of greater similarity of grammatical training and test strings (e.g. Vokey and Brooks, 1992). There are further perspectives considered in detail elsewhere (see sect 1.8.2 and 1.8.3). This disagreement on the basis of artificial grammar learning has led to the use of the simpler forms of implicit learning, such as the learning of invariant features, in order to study this type of learning.

The main focus of study in the preceding chapters has been the effect of manipulating processing conditions in the invariant learning task. In other research, there is an ongoing effort to determine the information used in the invariant learning task (e.g. Cock, Berry and Gaffan, 1994. Churchill and Gilmore, 1998. Ward and Churchill, 1998); therefore, the information acquired in the invariant learning task will be the main focus of this current chapter.

The training stimuli in the digit version of the invariant learning task are constrained, so that each string contains a digit in common. This constraint allows the possibility that subjects may become sensitive to the higher frequency of this invariant feature, rather than acquiring abstract knowledge of the invariance, as suggested by McGeorge and Burton (1990). There is some evidence that the acquisition of frequency information can be acquired automatically and without awareness (e.g. Hasher and Zacks, 1979), so it is possible that similar frequency acquisition may explain implicit invariant learning effects.

In response to this, Bright and Burton (1993) presented subjects with less rigidly defined invariant feature that did not give rise to increased occurrences of a specific feature. As described earlier (sect 2.6), Bright and Burton (1993) demonstrated implicit invariance learning using clock face stimuli. The invariance rule was a time range that was consistent over all of the training items so that all of the training items clock times fell between 6 and 12 o'clock. At test, subjects showed a preference for clock faces that displayed times within this range. This preference indicates that subjects can become sensitive to an invariant characteristic that is not simply a highly frequent invariant feature. Thus, the operation of a crude frequency counting mechanism is not an essential mechanism for invariance learning. However, these data do not rule out the operation of frequency counting in acquisition of an invariant feature of digit strings.

In Chapter 3, it was explained that if subjects produced more invariant digits in the present condition of the test strings (i.e. those fragments that contain the invariant), it may indicate that they are simply primed to produce more digits from the study set, regardless of test context. This kind of increase in production of digits that occurred very frequently in training during the test is akin to a frequency counting process. The production of invariant digits in the present strings could not be observed as a result of the bias to inhibit the production of repeated digits. In experiment 8, the production of digits in neutral strings was investigated, which allows the production of digits without test context to be examined. Using these conditions, there was no increased production of invariant digits from training. It would appear that there needs to be some context from training for any increase in invariant production to be revealed. This finding indicates that it is unlikely a crude frequency priming mechanism is operating.

The neutral test strings offer a different form of test in comparison to the standard conditions as subjects do not recollect any aspects of the training stimuli. Therefore, these data from digit production in neutral strings are not sufficient to conclude that digit frequency plays no role in digit invariant learning. More convincing evidence against the use of frequency information could be provided if the frequency of all digits (i.e. not just the invariant) in the training phase can be shown to have no relationship with the production of digits during the test phase.

4.2 Further analysis of the completion data from experiment 6

In experiment 6, the analysis focused on the production of the invariant '3' digit only. The production of the other digits was not included in the analysis. In this section, the level of production of the digits other than the invariant (i.e. 1, 2, 4 –

9) in the test fragments will be considered in relation to the frequency that these digits occur in the training set.

As each subject was presented with a different set of study materials, the frequency of each digit within the study sets varied considerably. It is meaningless, therefore, to examine the production of particular digits averaged across participants. Instead, a ranking system was developed so that the digit that occurred most frequently in particular training materials was ranked '8', the digit that occurred least frequently was ranked '1' and the other digits in between occupied the other rankings. For each subject, therefore, different digits corresponded to each of the ranking positions. Hence, the level of digit production for each rank was averaged across subjects, rather than the level of production of specific digits. This resulted in a production score for each rank. This score pooled the values of present and absent test fragments to produce a single production score.

The training stimuli were characterised according to a further measure, the frequency score. Each subject had different frequencies of digits that resulted in each rank. For example, subject 1 had fifteen '8' digits that were ranked as the most frequent. However, subject 2 had sixteen '4' digits ranked as the most frequent. The average of these frequencies for all subjects produced the frequency score for each rank.

The association between training set frequency and test production can be examined by plotting the production score against frequency score of each rank (see fig 1) There are two sets of rankings plotted below which were derived from the invariant and random groups separately. It was considered that these groups should yield separate ranks because the presence of an invariant alters the frequency of the other digits in the training strings.

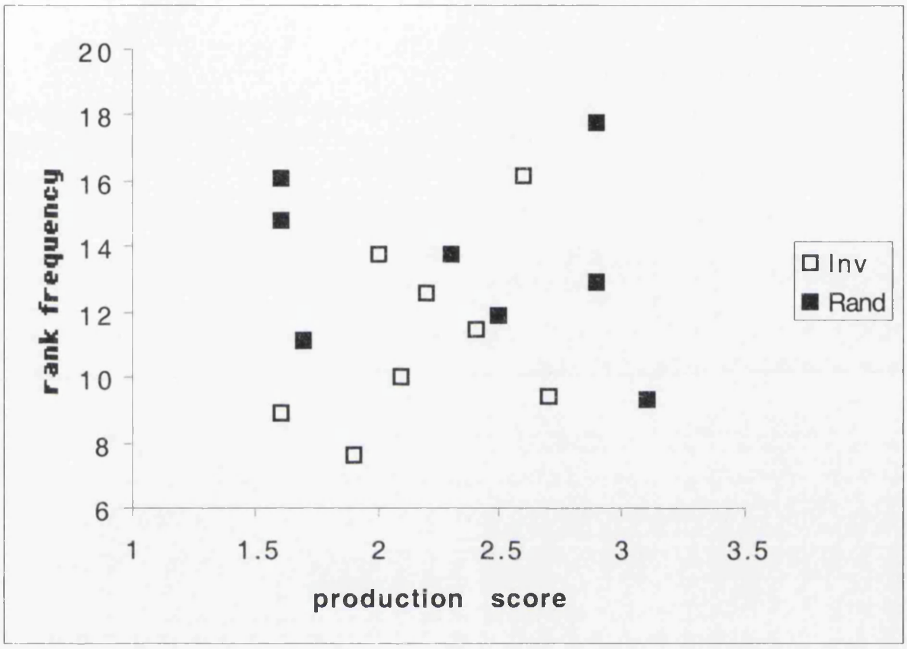

Figure 1

A Pearson correlation analysis of these scores reveals a value of r (16) = 0.077, which is not significant. This indicates that there is no association between the frequency of occurrence of digits in the training phase and their production in the test fragments. Therefore, the increased production of the invariant '3' digit in experiment 6 is not simply a result of its higher frequency within the training materials. One point of concern is the degree of influence that small variations in the training frequency could have on digit production at test. The frequency of even the highest rank digits in the training set is an average of eighteen occurrences, far less frequent than the thirty occurrences of the invariant. It is possible that the small variation in frequency during the training would not have a discernible effect on the production of digits during training. In comparison, the invariant is clearly far more frequent than the other digits, and this may result in a sufficient increase to be detectable. However, the analysis reported above is unambiguous; there is no association between training frequency and subsequent digit production. A weak effect of frequency may be expected to reveal a hint of association, but this is not the case.

These data suggest that a frequency counting mechanism does not control the production of digits in the completion test. An alternative possibility is that subjects are responding on the basis of abstract knowledge of the invariance. When considering the suppression data reported in chapter 3, it is unlikely that such abstract knowledge does control performance. It was argued in chapter 3 that the dependence of performance on articulation indicates that episodic knowledge of the training phase is used in this task. That is, during the test subjects' decisions are based on memory for specific items encountered during the training phase.

The importance of memory for instances in invariance learning has been investigated by Cock, Berry and Gaffan (1994). They developed four indices to measure the similarity between the training and test strings. For example, one measure was the sum of digits that each test string shared with each of the thirty training strings. Interestingly, selected negative strings scored higher than rejected negatives on the matrices, although they did not discriminate between selected and rejected positive strings. Hence, there are some difficulties with the way similarity was scored. As a result, Cock et al. manipulated the similarity of items in the forced choice test to those in the study set independently of the invariant. The data revealed that if the test item was similar to a training item, this similarity would bias the response to a greater extent than if the test item was positive or negative. Thus, there is good evidence that subjects use the familiarity of the test items to guide decisions.

In Chapter 3, subjects completed fragments of novel strings at test with a single digit to make up a string they believed to occur in the test phase. It is quite conceivable that subjects bring to mind instances or experiences of strings from the training phase when completing the test strings. It is the aim of experiment 10 to investigate the role of memory for instances in determining performance on the invariant learning task when the fragment completion test is used.

4.3 Experiment 10

As explained earlier, Cock et al. demonstrated that memory for whole training strings can exert an influence over performance in the forced choice test. It is the aim of the current experiment to determine the role of memory for particular items when the completion test is used.

If performance on the completion test is underlain by memory for whole instances, the invariant digit has no special status. That is, completing fragments with any digit, including the invariant, accesses the same knowledge base – memory for particular strings. The aim of the test of knowledge should be to access the knowledge of the invariant aspect of the training, and the knowledge of other aspects of the strings. If performance is no different for these two aspects of the strings, it would be reasonable to assume that the invariant digit has no special status, and is generated at test in the same way as other aspects of the strings. To investigate this possibility, a set of test stimuli was constructed which separately accessed knowledge of both the invariant digit and other aspects of the strings.

In the training phase subjects encountered thirty positive test items as before. However, during the completion test subjects were presented with twenty fragments of strings that were derived from items that occurred in the training set. Furthermore, these fragments were made by removing the invariant digit, in this case invariant '5', from the invariant training items (e.g. so that 2541 could become 2_41) or by removing another digit (e.g. 2541 could become 254_). Since all test items were derived from positive items, the conditions were comparable. These test conditions represent the same structure as the present and absent conditions used in the experiments in chapter 3. However, in place of measuring the degree of production of the invariant digit, the accuracy of completing the fragment correctly was measured.

The main prediction is that if subjects use memory for whole exemplars during the test, fragment completion performance will not differ in the absent and present strings. However, if some knowledge of invariance is acquired, fragment completion performance will be higher in the absent than the present fragments, as a result of an increased number of invariant digits produced. In order to be sure that any increase in the absent condition is not an artefact of the test materials, a second experimental group was included, the negatives group. In the negatives group, subjects were presented with twenty positive test items and ten negatives during training. Therefore, no knowledge of invariance could be acquired. If subjects in the absent condition were to perform higher than those in the present, it would prove that any effect in the invariant group is an artefact of the test materials. In addition, the same awareness test was administered to participants in order to replicate the findings of experiment 8 and 9 on the awareness of repetitions.

Method

Participants. Twenty one undergraduate students from the University of Glasgow took part in this experiment. They were paid a small fee for their participation.

Materials. Study sets were generated individually for each subject. The sets were generated separately for the Invariant and Negatives groups. In the Invariant group a set of 30 items was produced for the study set. These were all Positive items and contained no repetitions within the strings (as before). In the Negatives group a set of 30 four digit numbers was generated so that 20 of the strings were positives (contained an invariant 5) and 10 were negatives (did not contain an invariant 5).

The test materials fragments were generated by taking the 20 positive items from the study set and removing a digit. In the Present condition a digit other than "5" was removed (e.g. a string 4531 would become 45_1). In the absent conditions a 5 digit was always removed (e.g. a string 7591 would become 7_91).

Design and Procedure. The experimental procedure was carried out in an identical manner to that of experiments 6 and 7.

Results

Means and standard error for the two factors are presented in the table below. One subject was removed from the analysis for having awareness of the invariant.

Study set	Number of strings completed correctly (out of 10)			
	Present		Absent	
	Mean No.	SE	Mean No.	SE
Invariant	2.1	0.31	2.4	0.43
Negatives	2.2	0.41	2.7	0.40

An Analysis of Variance revealed that neither of the main effects nor interaction reached significance. It would be unwise to compare the means here with a purported level of chance. This is because it is unclear what that level of chance would be. Experiments 6, 7 and 8 have demonstrated that the invariant influences subjects' responses to a large extent. Therefore, it would not be possible to determine a fair level across the two groups.

The arithmetic task was carried out correctly on 96.3% of the study strings in the Invariant group and 97.2% of strings in the Negatives group.

Tests of Explicit Knowledge. One subject reported knowledge of the invariant digit in the awareness test. Again, all subjects said they felt they were guessing most of the time. In response to question 5 subjects reported using similar arithmetic as in Experiment 6. In answer to question 6, 40 % of subjects reported the use of avoiding repetitions as a strategy. When they were given strings containing repetitions in question 7, 90% noted that these items were different from those in the study set because they contained repetitions.

Discussion

In the invariant group, completion performance was equivalent in the absent and present test strings. There was no difference if the subject was completing the string with an invariant digit or another digit. In addition, the verbal reports of subjects in the awareness test mirrored those of the subjects in experiments 8 and 9. That is, the majority of subjects, when prompted, noticed that the strings in the training phase did not contain repetitions and a large minority of subjects reported using this as a strategy to aid performance.

This finding implies that the invariant '5' digit had not acquired any special status as a result of occurring in every training string; it was treated as any other digit during the test phase. Moreover, this suggests that subjects use memory for individual items when completing the absent test strings and not semantic knowledge of the invariance itself. This claim may be premature, however, if the required responses made in the test are considered. In the absent condition, there were ten test strings in which the required response was the invariant '5'; yet the responses in the present condition could be any of the eight other digits. Subjects in the absent condition may have been discouraged from entering the same response in the absent strings more than a couple of times, especially if they failed to notice that every training string contained a ‘5’ digit. It is necessary, therefore, to conduct a further experiment that takes this potential bias into account.

4.4 Experiment 11

In the previous experiment, it was demonstrated that completion performance is equivalent in conditions either where the fragment is completed with a digit that was invariant in the training materials, or if the fragment is completed with any other digit. Hence, it appears that the invariant does not acquire any special quality during training. However, the test materials may have been subject to a response bias that stems from the fact that the correct response to the absent strings was always the same, a ‘5’ digit, while responses to the present strings were varied across the other eight digits.

It may be possible to counteract this bias using the subjects' tendency to inhibit the production of repeated digits. In the current experiment, a number of novel present fragments were placed into the test set, in order that the proportion of present fragments at test was much higher than absent fragments. The natural bias against repeating the ‘5’ digit in the present test fragments would mean that the number of ‘5’ digits produced in the present fragments would be close to zero. Since the present fragments are in the majority, this would reduce the number of instances when the invariant ‘5’ would be a potential response. This limits the possibility that subjects may consider their responses to be unrepresentative of the training digit frequencies, as the number of occasions '5' is a potential response is reduced. Apart from the introduction of novel present strings, the experimental procedure was identical to experiment 10.

Method

Participants. Twenty-four undergraduate students from the University of Glasgow took part in this experiment. They were paid a small fee for their participation.

Materials. Study sets were generated in an identical fashion to Experiment 7 with the alteration in the Negatives group that 16 positive items occurred and 14 negatives,

The test materials fragments were generated by taking the 16 positive items from the study set and removing a digit. In the Present condition a digit other than "5" was removed, and in the absent condition, a "5" digit was removed. A further 14 novel positive items were generated and the invariant removed, as in the present condition. None of these items were the same as those in the study phase.

Design and Procedure. The experimental procedure was carried out in an identical manner to that of experiments 10.

Results

Means and standard error for the two factors are presented in the table below.

Study set	Number of strings completed correctly (out of 8)			
	Present		Absent	
	Mean No.	SE	Mean No.	SE
Invariant	1.1	0.28	1.7	0.48
Negatives	1.8	0.41	1.8	0.35

An Analysis of Variance revealed that neither of the factors main effects nor interaction reached significance. Since the level of chance is undefined, the experiment does not allow a comparison with chance (see above).

The arithmetic task was carried out correctly on 93.9% of the study strings in the Invariant group and 94.3% of strings in the Negatives group.

Tests of Explicit Knowledge. No subjects reported any knowledge of the invariant digit in the awareness test. Four subjects mentioned the invariant among other digits in response to the second question. None of these subjects went on to circle the invariant in the third question.

Discussion

The results of experiment 11 mirrored those of experiment 10; completion performance was equivalent in the present and absent test fragments. In both the experimental groups, there was no difference in performance if the string was completed with the invariant '5' or any of the other digits. It was suggested that the findings of experiment 10 may be influenced by the fact that the correct response in the absent condition is always the same. If subjects were to complete fragments correctly, then a number of '5' digits would have been produced; this may appear odd to subjects, especially if they failed to notice that all training strings contained a '5' digit, so resulting in inhibition of further '5' digits. In the current experiment, the test materials were controlled, in order that fewer '5' digits would be produced, and this did not appear to affect performance. It seems then, that the number of correct responses in the absent condition is the same as that in the present condition because the same process underlies completion in both conditions. It is likely that subjects refer to memory for previously encountered strings when completing the test fragments. This memory based account makes no distinction between digits that occur in every string and those which occur less frequently, as the knowledge base is only the stored exemplars seen in training.

4.5 General Discussion

Following exposure to a set of training items that conformed to an invariance rule, Cock, Berry and Gaffan (1994) demonstrated that subjects at test tend to select an item that is similar to one seen in training in a forced choice test over one that is dissimilar, even if it violates the invariance rule. Cock et al. suggested that this finding shows that subjects rely on a sense of familiarity with the test items to guide judgements during the test. The use of memory based processing has been revealed as a powerful determinant of performance in the artificial grammar learning task (e.g. Vokey and Brooks, 1994). It appears that memory for previous training instances has good explanatory power across different experimental tasks that are claimed to demonstrate implicit learning.

In chapter 3, an alternative route to accessing sensitivity to an invariant feature was introduced; the completion test. In the completion test, subjects are asked to complete a fragment of a four digit string with a digit to make a number that they think occurred earlier in the training set. In the experiments reported in the current chapter, subjects were presented with fragments of old training strings and were given the same test instructions. Under these circumstances, subjects showed no difference in their ability to complete these strings with a digit that was invariant in the training materials, or, for that matter, any other digit. This implies that the invariant did not acquire any special quality during training. Instead, responding in this task is dependent on memory based processing of particular training strings rather than abstract information about invariant characteristics.

Memory for particular instances, in this explanation, refers to memory for specific individual training strings that occurred during training. It assumes that subjects have remembered at least some of the training strings in full. In the experiments reported in chapter 3, the test fragments were derived from strings which did not occur in the training set. In this case, it must be assumed that subjects rely on the same memory based processing, except here, fragments are completed in order to be as close as possible to remembered strings. Since subjects feel they are guessing the majority of the time, it is quite conceivable for such a process to operate. It is important to note that this mechanism is equivalent to the notion of specific similarity to training strings that was proposed by Cock, Berry and Gaffan (1994).

Cock et al. expand on this notion of specific similarity of test strings to particular training strings with the idea of *general* similarity. They claim general similarity is the "summed similarity between a test string and all members of the learning set" (p1031). Cock et al. comment that it is difficult to apply such a conception of general similarity to the invariant learning task. Indeed, if the training items are described only by a simple invariance rule, such summed similarity of the test string to all training strings is akin to a frequency based account. The analysis that was performed on experiment 6 did not support this notion of a frequency counting account. For this reason, the memory for specific strings is put forward as a superior account of performance.

It is straightforward to explain a mechanism behind the use of similarity processing in the forced choice test of invariance learning. Subjects select strings that share a number of features with a particular item they believe they saw before. In the completion test, subjects must be cued by the test fragment to recall some item from training. They would then complete the test string with a digit which is consistent with the remembered item. The suppression data from experiment 8, in tandem with the transfer data from experiment 9, would seem to suggest that subjects are using memory for articulation of strings during training, rather than a visual representation of the strings.

The problem with this account is that the verbal reports of subjects in experiments 11 and 12 failed to provide any evidence that subjects were using a strategy of recalling particular items. As in chapter 3, the main verbal reports of strategies that subjects used concerned memory for arithmetic computations. This may not be a problem, however, if it is considered that the correct responses to the later fragment completion task are correlated with these computations. By remembering the processing operations carried out on the training strings, subjects are able to perform at above chance levels on the test, and yet remain unable to verbalise any of the individual training strings.

Throughout the work on this thesis, there has been continuous reference to the role of active processing in task performance; at test subjects use information that was generated during the active processing of the training stimuli. This is consistent with the episodic processing account of implicit learning (Whittlesea and Dorken, 1993). However, there are some findings on the sequence learning task that are inconsistent with the idea that subjects have to actively attend to the training materials, and these will be considered in chapter 5.

Chapter 5 Response Relevance in the SRT task

5.1 Introduction

Subjects in implicit learning experiments have been described as passive learners who have little influence over the acquisition of the structure of a stimulus domain (Lewicki and Hill, 1989). This means that implicit learning is directed entirely by the structure of the stimulus domain, and any variability in learning is the result of variations in the stimulus itself and not the task demands. In contrast, there exists a large amount of data that suggest subjects do not process the stimulus structure in a passive manner. For example, Whittlesea and Wright (1997) demonstrated that familiarity of the stimulus domain can have a powerful impact on the acquisition of sensitivity to its structure. In experiment 5 (sect 2.7), it was shown that sensitivity to the invariant property was not demonstrated if the task demands did not direct the subject to process the invariant characteristic of the stimuli. This indicates that the structure of stimuli is not acquired automatically.

However, this suggestion is not supported by other studies on the sequence learning task. Sequence learning is typically demonstrated when subjects are asked to react to items presented in different locations on a screen in a choice reaction time task. Unbeknown to the subject, the successive locations follow a sequence which repeats several times. Subjects are shown to demonstrate a speedup in their response times, without being able to describe the cause of this reduction (e.g. Nissen and Bullemer, 1987). In a development of this single sequence task, Mayr (1996) demonstrated that subjects were able to learn two sequences simultaneously. Importantly, subjects did not make an overt response to both sequences, and yet showed appreciable learning on the sequence to which they did not respond. This finding is significant because it implies that some automatic learning of sequence information can occur.

5.2 Simultaneous Learning of Two Independent Sequences

Mayr (1996) demonstrated learning of two independent sequences, one of which represented a sequence of location changes, and the other a sequence of object changes. The stimulus object varied across four locations positioned in the corners of the screen. Unlike Nissen and Bullemer (1987), four different object types were presented, a black square, a black circle, a white square or a white circle. Subjects were instructed to respond to these object changes. Importantly, this meant that the location changes did not require a response and thus spatial locations were not significant for the selection of a goal related motor response. Mayr (1996) demonstrated learning of both the spatial and object changes, and, in a second experiment, displayed how joint learning of object and spatial sequences was as efficient as the learning of single sequences.

Mayr (1996) interpreted these data as reflecting the operation of separate sequence learning systems, and suggested that the spatial sequence is learned by a system that acquires the sequence of spatial orientations to the successive locations of the stimulus object. This system operates independently from the system that acquires the sequence of object changes or the non spatial sequence. According to Mayr, the non spatial sequence is acquired by a system that is involved with the selection of a motor response. This system will acquire any information relevant to responses, so that location, colour, size and so on may be learned as long as they relate to the response demands. From this point of view, non-spatial regularities cannot be acquired when they are not relevant for the goal related response. This possibility was not tested, however, by the original Mayr (1996) study because subjects were instructed to respond to the object changes or the non-spatial dimension. A clearer picture of dual sequence learning using the Mayr (1996) design could be obtained by using the same spatial and non-spatial sequence dimensions, with responses made to the spatial dimension rather than the non-spatial. Using this design, it would be possible to ascertain if non-spatial sequences can be acquired without a direct response being made to them, and this possibility will be tested in experiment 12.

5.3 Experiment 12

In experiment 12, a replication of the Mayr (1996) study will be undertaken except that instead of responding to the object changes, subjects will be instructed to respond to the location changes of the object. The aim of this is to demonstrate that sensitivity to non spatial sequence material is not acquired when subjects do not respond directly to it.

As in Mayr (1996), participants were presented with four different objects that appeared at four different locations, and the succession of the location and object changes followed regular, but uncorrelated sequences. The same eight and nine item sequences were used as in Mayr (1996), and these were repeated nine and eight times respectively to make a 72 trial block, and to ensure that each element was paired with each other once. A finding of no learning of the object changes would indicate that non spatial sequences need to be related to task demands in order to be learned.

Method

Participants. Thirty two undergraduate students from the University of Glasgow took part in the experiment. They were paid a small fee for their participation.

Stimuli. Stimuli were presented on a 15in monitor and occurred at four different locations. These stimuli occurred in four boxes of side length 3.5cm which were horizontally positioned across the screen, each separated by 1.2cm. Four objects could occur in each of the locations; a black square, a white square, a black circle and a white circle. Object width and height was 2.2cm. During the practice trials an asterix, also of side length 2.2cm, was used. The response to stimulus interval (RSI) was zero.

Procedure. Before the main training trials began, subjects performed a 72 trial practice block where the stimulus object was an asterix. These 72 trials were organised so that they followed a random pattern through the whole block so there was no systematic sequence. No element occurred in the same location on successive trials. Throughout the experiment, subjects were instructed to respond to the location of the object. If the object occurred in the leftmost box, subjects pressed the "d" key on the keyboard, the "f" key for the left middle box, the "j" key for the right middle box and the "k" key for the rightmost box. The instructions emphasised both speed and accuracy.

Following this, subjects were instructed that the object to which they were to respond would alternate between black circle or square, and a white circle or square. They were told to ignore these changes and to continue responding to location, as in the practice block. Two sequences were used that corresponded to the sequences used by Mayr (1996). Subjects were not informed that the appearances of the objects would follow a regular pattern. Sequence A had eight elements: DBDABCAC and sequence B had nine: CDADBCABA. Here the letters corresponded either to the objects (A = black square, B = white square, C = black circle and D = white circle) or the locations (A for leftmost moving across the screen to D for rightmost). For each subject, one sequence was applied to locations and the other to objects and this assignment was counterbalanced. The training blocks consisted of 16 blocks of 72 trials each. This resulted in eight repetitions of the nine element sequence and nine repetitions of the eight element sequence. Thus, each element in either sequence was paired only once with each element of the other. In blocks 9, 12 and 15, the location or object sequence could be random, or both sequences could be random. For half the subjects, the location sequence was random in block 9 and the object sequence random in block 12, and for the other subjects the reverse assignment was used. All subjects saw random object and location sequences in block 15.

Following block 16, subjects were questioned about their metaknowledge of the sequential structure of the objects and location changes. Like the Mayr (1996) study, subjects were falsely led to believe they had been randomly assigned to one of four different conditions. Subjects were told there were four conditions, Condition 1, where the object and location changes followed a specific pattern, Condition 2 where the location changes were regular and object changes random, Condition 3 where the object changes were regular, and the location random, or Condition 4, where both object and location changes were random. Subjects were asked which of these four conditions they felt they had been assigned to.

When they had indicated which condition they thought that they had been assigned to, they were told that both the sequences were regular and a short test would follow which would probe their knowledge of the sequences. This test involved generation of items of the sequence based on three cue items. In this respect, the test differed from the Mayr (1996) study because three, rather than two, elements were used as a cue. In the Mayr (1996) study, the generation test was made as short as possible. The reason for this is that testing of both sequences could easily induce contamination from one to the other. It is possible that increasing the number of cue items from two to three may improve the sensitivity of the test, without substantially enlarging the generation test. Subjects saw three cue items and at presentation of the third item, they were asked to predict the next item in the sequence. This was repeated until the prediction of each sequence item had been attempted. Like the Mayr (1996) study, the order of presentation of the object and location generation test was counterbalanced.

Results

Overall Learning Effects: Median RTs were computed per subject and block, and are shown in Figure 1.

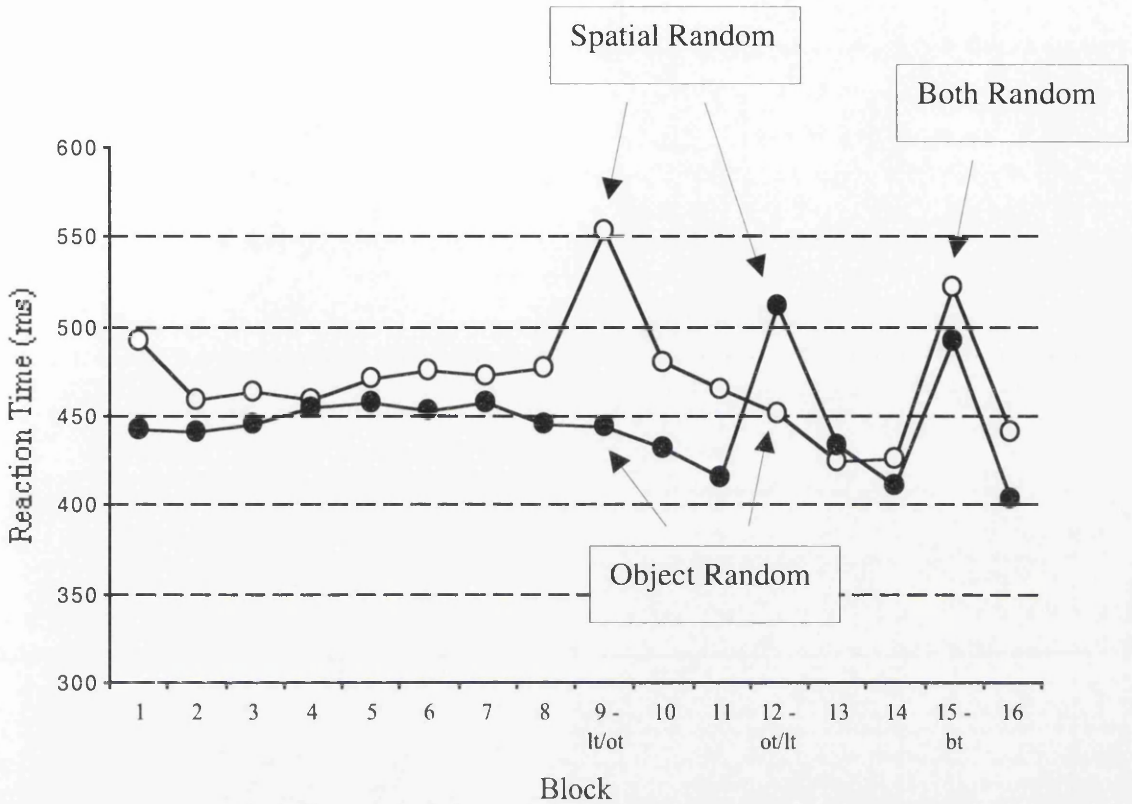

Figure 1 Reaction times as a function of training, separately for participants who were exposed to random spatial sequence in block 9 and a random object (non-spatial) sequence in block12 (white circles), and for the participants for whom the reverse order was used (black circles). All subjects responded to the location changes of the stimulus object.

There is no practice effect, that is, the RTs do not appear to decrease as the Blocks progress. This does not, however, indicate that no learning took place because in the blocks where location sequence was replaced by random movement or where both location and object changes became random, an increase in RTs occurred. This increase in RTs is determined by comparing the RTs on random blocks with those on adjacent sequenced blocks. This yielded learning scores for the location sequence alone, object sequence alone or where both sequences were tested together. The means and standard deviations of these scores are presented in the table below.

Sequence Test	Learning Score	
	Mean	SD
Location	80	39
Object	6	26
Both	86	40

For the location sequence, the learning score was significantly greater than zero, $t\,(31) = 11.4$, $p < 0.05$, as was the learning score for both sequences together, $t\,(31) = 12$, $p < 0.05$. However, the object sequence learning score failed to reach significance, $t\,(31) = 1.2$. From these data, it is clear that subjects show learning of the location sequence but because there is no corresponding increase in RTs when the object sequence ceases, this indicates no learning of the object sequence. Comparing the difference scores between the object and location reveals that they are significantly different, $t\,(31) = 8.6$, $p < 0.05$.

Comparing learning of the Eight and Nine element sequences: The data reported above represents learning effects of the eight and nine element sequences collapsed together. When the learning of these sequences is considered separately, the pattern of performance is no different. A two factor Sequence Type X Sequence Length Analysis of Variance reveals, importantly, a main effect of Sequence Type, $F(1, 30) = 84.9$, $p < 0.05$, but no main effect of Sequence Length, $F(1, 30) = 2.71$, $p > 0.1$. Furthermore, these factors did not interact, $F(1, 30) << 1$. From this, we can assume that when a location response is used, the 8 and 9 item sequence are comparable.

Subjects' Awareness of the Sequences: The post task question prompted subjects to declare whether they noticed any structure in either 1) object and location changes, 2) location changes alone, 3) object changes alone or 4) whether no sequence was present in either. The number of subjects who endorsed each of the responses is shown in the table below, collapsed across subjects who responded to the eight or the nine item sequence.

Allocation Response	Number of Responses
1. Location and Object	7
2. Location	16
3. Object	5
4. Neither	4

A Chi-Square test reveals that there is a difference in the numbers of responses in each category, $X^2(3) = 11.3$, $p < 0.05$. A location (2) versus others (1, 3 and 4)

Chi-Square test indicates that significantly more subjects responded "location" than any other response, X^2 (1) = 10.5, $p < 0.05$.

These data suggest that subjects noticed the location sequences more than they noticed object sequences. The next question is, do they have greater awareness of the location sequence than the object sequence? The generate test of object and location sequence was intended to probe subjects' conscious knowledge of the sequences. The means for the object and location generation scores are given in the table below.

Sequence	Generation Score (out of 8.5)		Percentage
	Mean No.	SD	
Object	2.5	1.0	30
Location	3.5	1.7	41

The generation score represents the number of correct predictions subjects made as the changes cycled through the sequence. The Generation Score for Location changes was significantly larger than that for Object changes, t (31) = 2.57, $p < 0.05$.

Discussion

The results of this experiment indicate, as predicted, that a non spatial sequence does not exert any influence on performance when it is not relevant for the response. That is, when an aspect of the stimulus is not critical for the response, it does not influence performance. In this task, the only factor that influenced performance was the location sequence, to which the subject responded. The object sequence was not responded to, and the test revealed no increase in RTs, suggesting that this aspect of the stimulus remained unprocessed. The reaction time data was mirrored by the verbal reports and generation scores. More subjects reported that the location changes followed a sequence than any of the other categories. In addition, the generation score of the location sequence was also significantly higher than the object sequence. Taken together, these findings indicate that subjects tended to have greater awareness of the location sequence than the object sequence.

These data suggest that an alteration of the task demands from responding to object changes to responding to location changes can have a substantial impact on performance and awareness. Mayr (1996) demonstrated learning of both sequences when the response was to the object changes, whereas in the current experiment, only the sequence that required a response showed any learning. The dual system account proposed by Mayr (1996) may explain this difference. Mayr (1996) suggested that spatial sequences are learned by a spatial attention system which directs orienting to the succession of location changes. Non spatial or object changes are not processed by this system, instead they are processed by the system that selects a response. This view suggests that non spatial sequences require a response in order to influence performance at test, and the data appear to support this position.

There seems to be a contrast between this response based non-spatial sequence learning, and Mayr's (1996) demonstration of non-response based spatial sequence learning. Mayr (1996) suggests that this indicates that learning of a spatial sequence can occur without an overt response, while a non spatial sequence does require a response. This question of whether sequence learning can occur without an associated response is explored in further detail in Experiment 13.

5.4 Experiment 13

In the Mayr (1996) study, the spatial sequence was not relevant for the response, but appreciable learning was still demonstrated. In experiment 12, the non spatial sequence was not relevant for the response demand, and here, no sensitivity to this sequence was acquired. This implies that spatial and non spatial sequences are processed by alternative mechanisms. Perhaps more important, however, is the finding of spatial sequence learning which is not directly related to the response demands of the task. This finding suggests that an overt response is not required for learning of a spatial sequence to occur. Evidence for pure perceptual learning of a sequence has been provided using a single sequence task by Howard, Mutter and Howard (1992) prior to Mayr's work. Howard et al. presented a group of subjects with the standard sequence learning task, except no response was required during the first three blocks of trials. On the fourth block, subjects responded as normal to the stimuli, and on the fifth block the random pattern was introduced. Subjects showed the standard increase in response time when the random block was introduced, even though they had only observed the sequence for the majority of the training phase. Howard et al. concluded that subjects had learned the sequence in a purely perceptual manner. More recently, Willingham (1999) questioned these findings on the basis of the use of explicit knowledge in the Howard et al. study. When the explicit effects were removed, no effect of learning by observation was present.

There is a further methodological problem with studies that attempt to demonstrate perceptual learning of sequential learning by pure observation. Although subjects are watching the screen, it is impossible to be sure they are, in fact, attending to the stimulus changes. This is the advantage of the dual sequence task; subjects are actively following the object changes of the stimulus, and to see the next object change they must pay attention to the location of the object. This methodology may have an advantage over the standard Nissen and Bullemer (1987) task, in that it is possible to investigate learning of sequential material that is attended to but does not require an overt response, such as the spatial sequence used in Mayr (1996). This type of sequence learning would be better described as *non response relevant* rather than *observed,* since both stimulus dimensions are clearly observed, but only one is response relevant.

A problem for the Mayr (1996) dual sequence task is that the spatial and non-spatial sequence dimensions require different responses. This brings the further question of spatial versus non-spatial sequence learning into the experimental interpretation, thus clouding the issue of non response relevant sequence learning. Furthermore, since this finding of non response relevant spatial sequence learning is demonstrated using different response demands (i.e. object response) in comparison to the standard demonstration of spatial sequence learning (which use location response), it may be that these findings are not comparable. A superior test of non response relevant spatial sequence learning is possible by using the dual dimension task with two *spatial* dimensions. Hence, the response demands would be the same as the single spatial sequence learning task, with the advantage of greater certainty that the non response relevant dimension is processed.

A dual sequence task could be constructed in the spatial domain by varying stimulus movements along two spatial dimensions. The stimulus object could move between the squares of a grid, rather than four locations in a row. Movements between the columns of the grid could follow one sequence, while movements between the rows of the grid could follow a second sequence. Hence, the stimulus would be moving independently along two separate dimensions, according to two independent sequences.

Importantly, this design allows a simple mapping to be formed between either the vertical or the horizontal movements, and the key press response. As in the Mayr (1996) study, it is possible to have one response relevant dimension and one dimension that is not relevant to the response. The alteration here is that both dimensions occur in the same spatial domain. Although Mayr (1996) has demonstrated learning of a spatial sequence that is not directly response relevant, learning of a spatial sequence that is not response relevant has never been demonstrated where the response is made to another spatial dimension.

Rather than using a dual sequence design, only a single sequence was presented to subjects, with movements along the second dimension following a random pattern. Since this was the first test of this experimental set up, a single sequence was used to simplify the experimental design as much as possible. As a result, subjects would be exposed to a single sequence, and hence knowledge of only one sequence is tested.

To summarise the experimental design, subjects were presented with a four by four grid, and the stimulus object, a black square, could move between any of the sixteen grid positions. One group of subjects responded to the movements of the square between the columns of the grid, and another group responded to the movements of the square between the rows. The movements between the grid positions were simultaneously determined by two independent patterns; a repeating sequence and a random sequence. If the repeating sequence determined the square's movements between the rows of the grid, the random sequence determined movements between the columns, and vice versa. The subjects responded to the dimension of the grid that followed the random pattern so that the other dimension followed a repeating sequence. For the Vertical sequence group, the repeating sequence determined movements between the rows of the grid, while the subjects responded to the column in which the stimulus appeared, and this followed a random pattern. The opposite relationship was used in the Horizontal sequence group.

From here on, the vertical and horizontal groups will be considered together in terms of two dimensions. In both cases, the dimension of the grid to which the subject responds will be referred to as the response relevant dimension (or RR) and the dimension of the grid that the subject does not respond to will be referred to as the non response relevant or (non-RR). Hence, the repeating sequence occurs on the non-RR dimension in both groups, while the subjects respond to the RR dimension of the grid.

Following five blocks of training, the non-RR dimension changed from a repeating sequence to a random sequence, so that both dimensions of the grid now followed a random pattern. Any sensitivity to the sequence that had been acquired on the non-RR dimension should now be revealed as a increase in response time. Hence, this served as the first test of knowledge.

The next phase of the experiment is referred to as the Switch phase, as here, subjects switch their responses from one dimension to the other. Hence, if the movements of the square between the rows of the grid had been the RR dimension, the movements of the square between the columns now becomes the RR dimension and vice versa. After two blocks of completely random trials, the third block returns to the same sequence constraints as in the training blocks. Since the RR dimension has switched from horizontal to vertical or vice versa, this allows the subject to directly respond to the sequence for the first time, thus allowing a further test of sequence knowledge. Any facilitation in comparison to the adjacent random blocks would reveal sensitivity to the repeating sequence.

Method

Participants. Twenty two undergraduate students from the University of Glasgow took part in the experiment. They were paid a small fee for participation.

Stimuli. Stimuli were presented on 15in monitor and occurred at sixteen different locations within a square grid formation, so that each side was four boxes long. The grid was arranged so that the whole area of the screen was covered by the grid. The stimuli that appeared in the boxes were black squares of side length 2cm. The response to stimulus interval (RSI) was 300 ms.

Design and Procedure. The position of the stimuli within the grid will be described by its horizontal and vertical location in the grid. These horizontal and vertical locations are each labelled A, B, C and D, so that A on the horizontal dimension was the right column of squares, moving across to D on the left. On the vertical dimension, A was the top row of squares, moving down to D at the bottom. So, for example, BD would be the inside left location on the bottom row.

The experiment was divided into two phases, first the Training phase, followed by the Switch phase. In the Training phase, the movement of the square along each of the dimensions could either be determined by the eight element sequence from Experiment 12 - DBDABCAC, or it could be a random pattern. As in Experiment 12, this sequence was repeated nine times within one block of trials to make a total of 72 trials. This sequence determined movements on one dimension and location changes would be random on the other. Following from this, subjects were randomly allocated into two groups, the Horizontal Sequence group and the Vertical Sequence group. In the Horizontal Sequence group, the sequence determined movements on the horizontal dimension, and the vertical movements were random. In the Vertical Sequence group, the sequence determined movements on the vertical dimension, with location changes on the horizontal dimension following a random pattern. No element in the grid was repeated on successive trials. This arrangement was used for the first five blocks of trials. In the sixth block, the movements of the square along both dimensions were determined by a random pattern, again for 72 trials, and this served as the first test of any knowledge acquired in the training phase.

Throughout the experiment, subjects responded to the stimuli using their index finger. Subjects assigned to the Horizontal Sequence group responded to the row the stimulus appeared in during the training phase (by pressing "9", "6", "3" and "." on the numeric keypad). Those assigned the Vertical Sequence group responded to the column in which the square appeared during training (by pressing "4", "5", "6" and "+" on the numeric keypad). Hence, the stimulus object moved randomly between locations on the RR dimension, with the non-RR dimension varied according to the eight item sequence.

The Switch phase consisted of four blocks of trials. During the first two blocks of trials, the movements of the square along both dimensions varied according to a random pattern. In the third block of the Switch phase, location changes reverted back to the same sequences used during the Training phase. In the final block of the Switch phase, the location changes of both dimensions would again follow a random pattern.

In the Switch phase, the responses were switched to the opposite dimension to that used in the training phase, so that subjects who responded to the column in which the square appeared during the training phase, now responded to the row in which square appeared, and vice versa. This meant that subjects now responded to the dimension of the stimulus movements that followed the sequence (i.e. the *responses* followed the sequence presented during the training phase) in the third block of the transfer phase. This served as an opportunity for the subjects to respond to the sequenced material for the first time, and constituted the second test of the sequence.

Throughout the experiment, when a incorrect response was made, a tone of duration 200ms was presented immediately following the response. The square moved to the next location in the sequence to which the subject made a response, and so on.

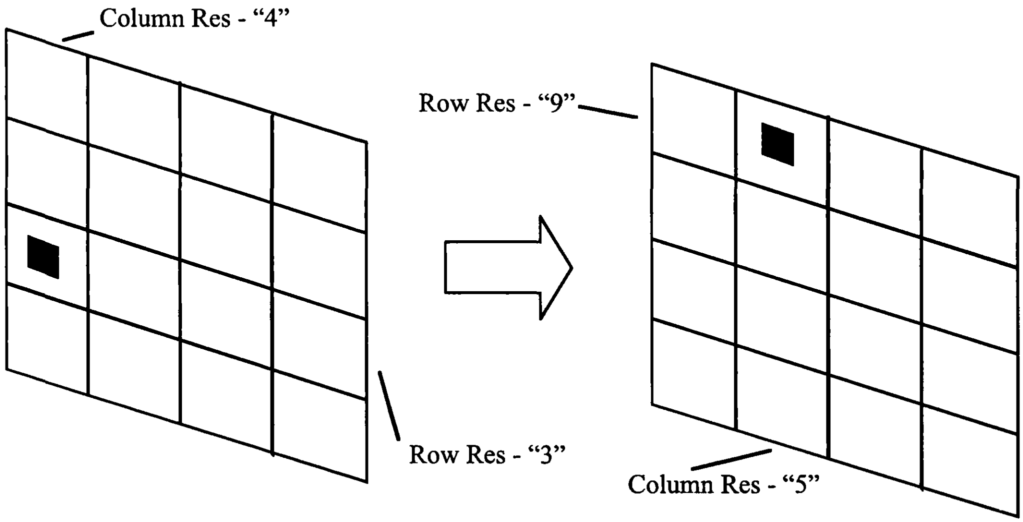

This diagram represents the responses on the keypad made on two successive trials by subjects responding to the row in which the square appears, or the subjects responding to the column in which it appears.

Results

Learning Effects: Median RTs were computed per subject and block, and are shown in Figure 2.

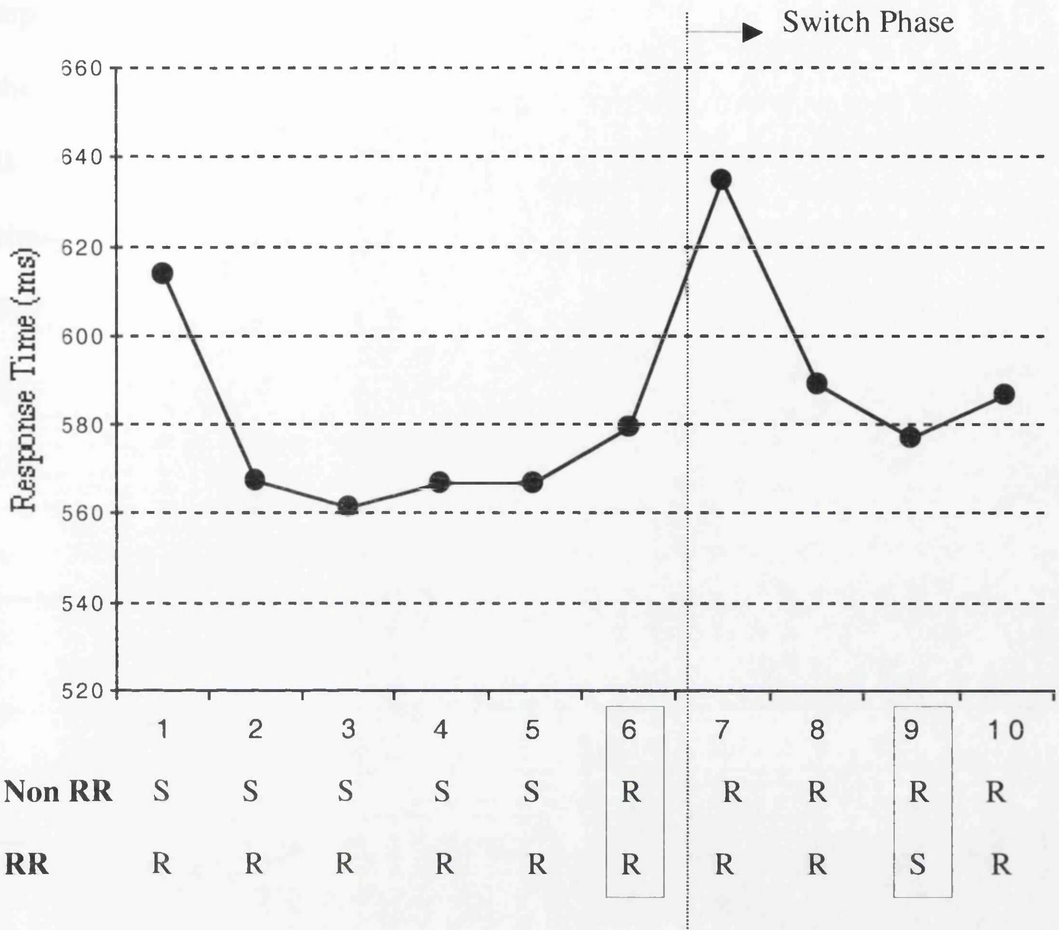

Figure 2. Reaction times to the response relevant dimension as a function of training, collapsed across the Horizontal and Vertical sequence groups. The Non RR and RR are marked for each block according to the sequence type that occurred on each block, so that if a random sequence was used the block is marked "R" and the repeating sequence is marked "S". The test blocks, blocks 6 and 9, are in the boxes.

The crucial RT data are those obtained in blocks 5, 6, 9 and 10, where the learning scores are determined. There are two learning scores; the non RR score which represents the change in RTs when the non RR dimension changes from a repeating sequence (block 5) to a random sequence (block 6). Following block 6, the response task switched from responding to the column in which the square appeared, to responding to the row in which the square appeared, and vice versa. In block 9, the subjects responded directly to the repeating sequence, as the repeating sequence and random sequence constrained the square's movements in the same way as blocks 1 - 5. Hence, the learning score derived from these trials is called the Direct learning score. The Direct learning score represents the change in RTs when the RR dimension changes from a repeating sequence (block 9) to a random sequence (block 10). The non RR, Direct learning scores and standard deviations are presented in the table below.

	non RR		Direct	
	Mean	SD	Mean	SD
Learning Score (ms)	12.8	52.5	9.5	46.4

Neither the non RR learning score [t (21) = 1.1] nor the Direct learning score [t (21) = 0.9] were significantly larger than zero.

Discussion

As no increase in response time occurred at the test blocks in comparison to the previous transfer blocks, this indicates that sensitivity to the sequence was not revealed in either the non RR or Direct learning scores.

The aim of the current study was to demonstrate, in the simplest possible conditions, that a sequence of spatial locations can be learned when the response is made to a second spatial dimension. The data suggest that subjects did not acquire any sensitivity to the spatial sequence that was present on the non-RR dimension. In the Mayr (1996) task, there was clear learning of a spatial sequence that was unrelated to response selection. It may be that the discrepancy between these findings is related to the response task that each study asked subjects to perform. That is, spatial sequence learning of stimulus movements on a non-RR dimension does not occur when it is in the same domain as the RR dimension. However, there are some methodological concerns that need to be addressed before considering this point further.

Willingham, Nissen and Bullemer (1989) carried out an experiment where the stimuli were a series of colour patches that occurred at four locations on the screen. The location changes followed a specific sequence, while the colour patches varied according to a random pattern. The responses were made to the colour patches, in order that the location changes were not response relevant. No learning of the location changes was demonstrated by Willingham et al. (1989). Mayr (1996) supposes that this finding can be explained in that the separation between the locations was not very large, and it may be that the small separation of the locations would show only a small performance advantage. In the current study, the locations were spread out to cover the whole screen area in an attempt to maximise separation. However, as in the Willingham (1989) study, this separation may not have been sufficient to show any advantage.

The response in the current study required subjects to press the keys using their index finger. Unlike the standard situation where each location is given a specific finger press, inducing a mapping between the index finger response and the spatial location is not a trivial task. Furthermore, even when the mapping has been induced, there are four possible locations for the response to be made, which makes responses difficult with the single index finger. Following each response, the subject is faced with locating the position of the next response from four possible locations. This is not necessary when each finger is allocated with a single response location.

In summary, it appears that a non response relevant spatial sequence is not acquired when the response dimension is a further spatial sequence. This contrasts with findings reported by Mayr (1996) who demonstrated that a spatial sequence *can* be acquired without a direct response. There are two possible reasons for the failure to demonstrate learning in the current study. Firstly, the separation of the locations between which the object moves may not have been sufficiently large to reveal a significant performance benefit. Secondly, the response mapping between the index finger and the four response locations may not have allowed a simple response mapping to be obtained. In Experiment 14, these issues are addressed by simplifying the design of the experiment in an attempt to reveal an effect of learning a spatial sequence without a specific response.

5.5 Experiment 14

Experiment 13 demonstrated that sensitivity to a spatial sequence does not occur when the non-RR and RR dimensions are both spatial. It was suggested that this failure to demonstrate learning may be the result of methodological problems. The first of these was the small separation of the stimulus locations in the grid, which may have made any performance benefit difficult to detect. The second problem arose from the singular use of an index finger to respond to four locations. This response set up is not as intuitive as the standard SRT task where a separate finger is assigned to each stimulus location.

It may be possible to address both these points with a single change in the experimental design. Reducing the number of stimulus locations on each dimension from four to three would increase the separation of each stimulus location, if the screen area is maintained. It would simplify the response if three response locations, rather than four, were used. This would make the movement between responses faster and more efficient.

There is some precedent for using the index finger with three response locations. Cohen, Ivry and Keele (1990) used such a configuration when investigating the degree of effector independence in the SRT task. In addition, they used very simple five item sequences in their studies. These sequences were shown to reveal strong learning effects across training.

In the following experiment, the basic design of Experiment 13 is retained with the major modification that three stimulus locations are used on each dimension in place of the four used previously. In addition, the same simple five item and three element sequences that were used by Cohen et al. were employed here, as they have been shown to reveal clear learning effects. Moreover, as they are shorter than the eight element sequences used in Experiment 13 they should be acquired more easily.

Method

Participants. Twenty four undergraduate students from the University of Glasgow took part in the experiment. They were paid a small fee for participation.

Stimuli. The stimuli were presented within a grid formation as in Experiment 13. Here though, the grid contained nine locations. The side length of the grid was 22 cm, making the side length of each quadrant 7.5cm. As in Experiment 13, the stimuli occurring within the boxes were small squares, with the exception that the side length of the square was 1.5cm. The response to stimulus interval (RSI) was 200ms.

Design and Procedure. The design of this experiment was identical to that used in experiment 13. There were three major differences in the procedure; the first of these was the use of shorter five item sequences (see below for a description). Secondly, the number of training blocks was increased from five to seven, so that there were twelve experimental blocks. Finally, here there were less grid locations to which the square could move (from sixteen locations in Experiment 13 to nine locations in the current experiment). This lowered number of total locations reduces the number of elements that any sequence could use, so, in place of each dimension having four elements: A to D, here, three elements are used: A to C. The organisation of these elements remains the same, with A referring to the left column through to C on the right. The rows are labelled in order that A represents the top row of locations moving down to C as the bottom row of locations. For example, AB is the leftmost middle location of the nine.

The response keys used were "4", "5" and "6" on the numeric keypad for the horizontal dimension, and "8", "5" and "2" on the numeric keypad for the vertical dimension.

The five element sequences were presented in six different orderings. The six structured sequences were ABCBC, ACBCB, BACAC, BCACA, CABAB and CBABA. Cohen et al. (1990) has suggested that these sequences represent all possible structures of five item sequences with three elements. For each block of 100 trials, the sequence was repeated 20 times. No feedback was given on the reaction time task because the three element sequence was a simple mapping to three response locations.

Results

Learning Effects: Median RTs were computed per subject and block, and the average for each block across subjects is displayed in Figure 3.

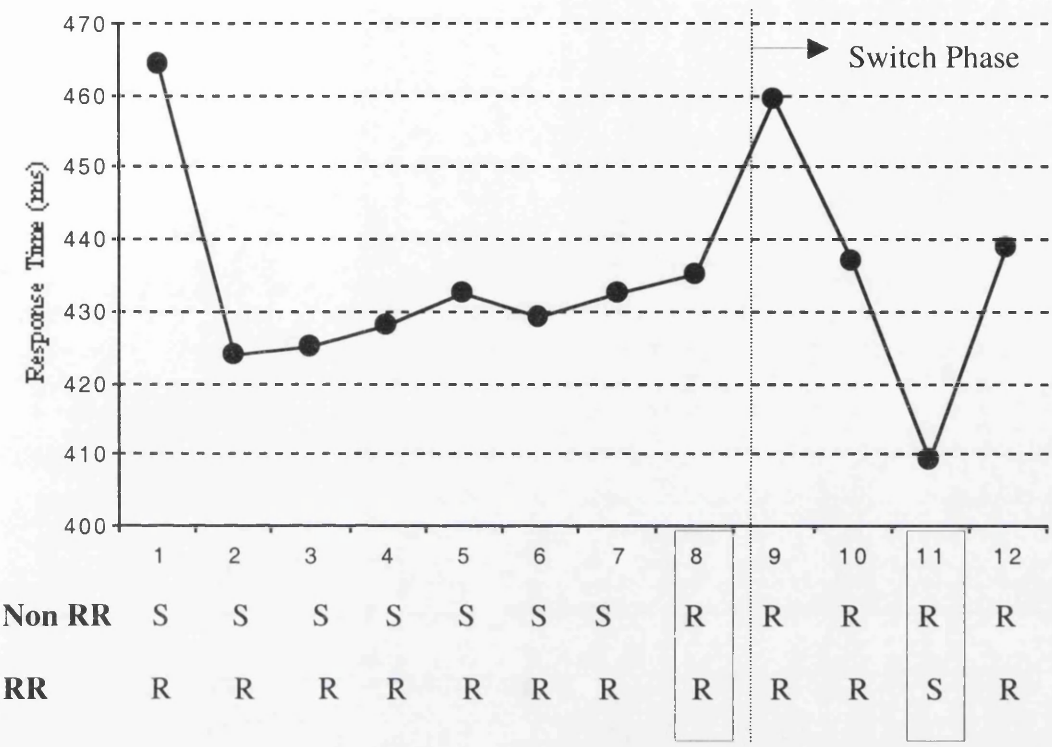

Figure 3. Reaction times to the response relevant dimension as a function of training, collapsed across the Horizontal and Vertical Sequence groups. As in figure 2, the non RR and RR dimensions are marked according to the sequence type that was used in a particular block. The test blocks, blocks 8 and 11, are in the boxes.

The learning scores were determined from the RTs in blocks 7, 8, 11 and 12, in the same way as the previous experiment. Again, this yielded non RR and Direct learning scores, displayed in the table below.

	non RR		Direct	
	Mean	SD	Mean	SD
Learning Score (ms)	2.6	33.8	29.4	55.8

The non RR learning score is not significantly different from zero [t (23) = 0.4]. In contrast, the Direct learning score was significantly different from zero [t (23) = 2.53, $p < 0.05$].

Discussion

The results of Experiment 14 are similar to those of Experiment 13. However, unlike Experiment 13, there is some evidence of sensitivity in the Direct test, as the Direct learning score is significantly different from chance. In contrast, the non-RR learning score did not reveal any learning effects.

This difference between these two tests may indicate that the apparent significant learning score is a artefact of learning within the Direct test block. In the second test block (block 9), subjects are no longer simply observing the sequence, they are responding to the sequence itself, and sequence learning by response should now function. The block is one hundred trials long and the sequence only five items, which results in twenty cycles of the sequence within the test block. Considering that the many sequence learning studies show an appreciable effect after twenty four cycles of a longer sequence, twenty cycles should be sufficient for a significant degree of learning to occur. Therefore, an effect of learning within the block is likely to be responsible for any effect of learning.

5.6 General Discussion

The main aim of this chapter is to investigate the role of active responding in determining sequence learning performance. The question was addressed using non spatial (experiment 12) and spatial (experiment 13 and 14) sequences. The main finding is that sequence learning appears to be dependent on subjects having an opportunity to respond to the sequence itself. This effect was consistent in non spatial and, in contrast to other experimental findings, spatial sequences. These effects will now be considered in further detail.

It is not greatly surprising that sensitivity to a non spatial sequence can only be acquired if the sequence was directly responded to during training. In experiment 12, subjects performed a location detection task, and the object changes could not have aided performance in any way. For this reason, it is likely that they remained unprocessed during the location detection task. Mayr (1996) proposes that non spatial or object sequences may be acquired through a response based sequence learning system that picks up any stimulus dimension which is correlated with responses, regardless of whether it is a shape or colour variable. The object sequence in experiment 12 did not correlate with the response in any way and therefore could not be acquired by such a mechanism.

Mayr's (1996) finding of spatial sequence learning in the absence of an overt response would seem to indicate that a response based interpretation is incomplete. However, the response demands of this study are different from those in previous spatial sequence learning studies. Subjects responded to the changes in the shape of the stimulus object, with the spatial dimension non response relevant. In experiment 13 and 14, the response was made to the spatial location of the object, with the sequence occurring on a non response relevant spatial dimension. In this case, no effect of the non response relevant dimension was detected. This position is in line with other research which supports the critical role of responses in sequence learning (e.g. Willingham, 1999), with the added certainty that the non response relevant or observed sequence has been processed. So, it appears that the Mayr (1996) study may be a special case of non response relevant sequence learning, resulting from the spatial/non spatial dual dimension task.

The degree of awareness of the subjects in the Mayr (1996) study may explain why these data present a special case. If the explicit measures reported by Mayr (1996) are examined, fifty five of the sixty participants in the study performed at above chance levels on the prediction test of awareness of the spatial sequence. This indicates that subjects had good explicit awareness of the spatial sequence, and this may account for the effect of learning of the non response relevant dimension. It could be argued that prediction tests of the type used in experiment 12 and Mayr (1996) can be influenced by implicit knowledge of the sequence, that is, they are not process pure (Willingham, Greeley and Bardone, 1993). While this may be the case, a large number of subjects reported detecting some sequence in the stimulus movements or the object changes (forty five out of sixty), indicating that these significant prediction scores are accompanied by a feeling of metaknowledge, implicating the use of explicit knowledge. Taken together, these explicit effects are not merely minor consideration when interpreting this data.

In section 1.8.3, it was noted that the procedural view of learning (Kolers and Roediger, 1984) can be a useful framework for understanding the sequence learning task. Within this framework, learning is construed as the accumulation of procedures for interacting with symbols. These procedures are associated with specific tasks, so, in the case of sequence learning the stored procedure would involve responding to the location of a stimulus. Hence, within this account, information about non response or task irrelevant sequences should not be acquired. This prediction is supported by the data presented in this chapter, as sensitivity was only shown to response relevant sequences. This suggests that subjects are not passive receivers of structure in the training stimuli; instead, it appears the critical aspects of the structure must be processed during the training phase for learning to occur. The importance of the active processing of the stimulus material is akin to findings of other implicit learning tasks reported within this thesis and elsewhere. In the next chapter, a summary of the main findings of all the experiments reported in this thesis will be presented.

Chapter 6 General Discussion

The first part of this chapter provides a brief summary of the main findings reported in each of the experimental chapters of this thesis. Following this, these findings will be drawn together to assess their significance with respect to the theoretical issues set out in chapter 1.

Chapter 2 - The efficacy of invariant learning

The experiments reported in chapter 2 investigated several manipulations of the invariant learning task (McGeorge and Burton, 1990). In this task, subjects are exposed to training stimuli which consist of 30 four digit numbers containing an invariant "3". In subsequent forced choice test subjects tend to select novel numbers containing this invariant over numbers without it. In experiment 1, sensitivity was eliminated when the surface form of the stimuli changed from the training phase to the test phase. It appears that decisions at the forced choice test in the invariant learning task are made on the basis of processing fluency. When surface features change, this fluency cannot be used, giving no advantage to the items that preserve the invariant.

The use of fluency of processing at test implies that subjects are relying on the familiarity of the items presented at test to make decisions (as suggested by Cock, Berry and Gaffan, 1994), and experiment 4 provided further evidence supporting this position. When the training set was defined by a majority feature, rather than an invariance rule, performance did not significantly deteriorate. This indicates that the invariant does not acquire some special quality that sets it apart from the other digits in the training strings.

Sensitivity to an invariant characteristic is not only influenced by the surface features, the demands of the training task are also critical for successful test performance. In experiment 5, sensitivity to the invariant property was shown only when the training task demands matched the invariant property (e.g. search for low numbers with a invariant "2"), while performance was at chance when the training task was unrelated to the invariant (e.g. search for low numbers with an invariant "8"). This demonstrates that if the training task does not relate in any way to the structure of the stimuli, performance at test will be at chance levels. Conversely, when the training task maps directly onto the structure of the training strings, performance is maximal.

Chapter 3 - The effect of alternative test task demands

Investigating the generation of information in the invariant learning task was the major aim of chapter 3. In experiment 6, a digit string fragment completion test was used to demonstrate increased generation of an invariant digit following exposure to training exemplars that all contained this invariant. This finding indicates that in addition to increased fluency of processing test items, sensitivity to invariance can be revealed by generation of invariant information.

It was predicted in chapter 3 that since the completion task requires the articulation of digits, this process may lead to the use of phonological information preserved from the training period. This possibility was tested in experiment 8 where subjects performed phonological suppression during training. Performance on the completion task dropped to chance levels in the suppression group, suggesting that articulation of previously experienced strings is a key aspect of generation performance. Generation performance was seen to be insensitive to the surface features of the training period, as successful performance was demonstrated regardless of the surface form of the training stimuli.

Chapter 4 - The information acquired in the invariant learning task

In the completion test reported in chapter 3, subjects are given a fragment of a four digit number and are asked to complete the fragment to make a number they saw during training. In experiments 10 and 11 of this chapter, subjects were given fragments of items which had appeared in the training set. These fragments were constructed so that there were two alternative correct responses, either where the invariant completed the string correctly or another digit was required. Under these circumstances, subjects showed no difference in their ability to complete these strings with a digit that was invariant in the training strings, or any of the other digits. This finding suggests that the invariant did not acquire any special quality during training. Instead, subjects refer to particular items that they can remember from the training period.

Chapter 5 - Response relevance in the SRT task

This chapter investigated the role of active responding in sequence learning performance. In the sequence learning task, subjects are asked to react to a stimulus object presented in different locations on a screen in a choice reaction time task. The succession of locations follow a sequence which repeats ten times or so within a block of trials. A speedup of reaction time is demonstrated across training trials, although subjects cannot describe the cause of this speedup. In a development of this task, Mayr (1996) presented two sequences on two separate stimulus dimensions, a spatial dimension and a non spatial dimension (the shape of the stimulus object). Simultaneous learning of both sequences was demonstrated, although subjects did not respond directly to the spatial sequence. This suggests that the spatial sequence was learned automatically, contrasting with the processing view which emphasises constant activity of the learner.

The experiments reported in chapter 5 contradict the findings of Mayr (1996). In place of the dual sequence task, a dual dimension task was used where the stimulus object moved along two spatial dimensions (experiments 13 and 14). A repeating sequence occurred on only one of the dimensions for simplicity. In this case, learning of the spatial sequence was not demonstrated when the repeating sequence was presented on the dimension that was not response relevant. Furthermore, in experiment 12, no sensitivity to a non spatial sequence was evident when the response was made to a spatial dimension. This finding suggests that learning can only occur when a response is made to the training stimuli, it does not occur automatically. In addition, subjects were shown to be aware of the stimulus dimension they responded to, the location sequence, and not the object sequence. This indicates that task demands during training can determine the extent of conscious awareness of a sequence.

Discussion

Two main accounts of implicit learning were described in chapter 1, the parallel system and the episodic processing views. The parallel system account (e.g. Reber, 1989) holds that the sensitivity to structure of stimuli that occurs in a typical implicit learning experiment is the result of the operation of an implicit learning system. This system abstracts the structure of a given stimulus domain without awareness, and this occurs in a passive manner. In contrast, within the episodic account, particular experiences of stimuli are preserved from the training phase, and the similarity of these experiences to the test experiences determines test performance.

It was suggested that the episodic processing account provides a superior explanation of implicit learning. The episodic processing account may provide a superior explanation of experimental data as it places more emphasis on active processing, and this position is in tandem with a large amount of experimental data that was described in chapter 1. These data suggest that subjects actively process the training stimulus they encounter according to the demands of the training task, and this processing is reflected in later test performance.

Several experiments reported in this thesis are pertinent to the issue of active processing in implicit learning. This question was addressed indirectly in experiment 8. When subjects performed phonological suppression during training, no effective learning of the invariance was demonstrated, indicating that learning is not an automatic consequence of interacting with the stimulus. Similarly, in experiment 5, sensitivity to the invariance did not occur as an automatic consequence of perceiving the stimuli. When the induction task focused on aspects of the stimulus structure that were not related to the invariant information, no sensitivity was shown to the invariance. In contrast, when the training task directed subjects to the invariant property, maximal sensitivity was revealed. This experiment demonstrates the importance of processing the appropriate aspects of the stimulus structure that will later become important for the test.

In experiment 12, similar findings were demonstrated in a very different experimental context. Here, subjects only became sensitive to a sequence when they were given an opportunity to respond to the sequence. When the sequence was not response relevant, no sensitivity was demonstrated. This is similar to the situation in experiment 5 where successful test performance was only demonstrated when the training task correlated with invariant properties. The post task verbal reports from experiments 5 and 12 both indicate that when the training task directs subjects to the critical aspects of the training material, greater awareness of the structure is demonstrated. These findings are significant because they demonstrate that training task demands interact with conscious awareness.

Further evidence for the absence of learning in circumstances where the sequence was non response relevant was presented in experiments 13 and 14. In this task, the stimulus object moved between the squares of a grid according to a random sequence and a repeating sequence. Subjects responded to the movements of the square between the rows or columns of the grid, so that if one was relevant to the response the other was not. During training, the sequence always occurred on the stimulus dimension that was non response relevant. The only circumstances where sensitivity was revealed was where the subject responded directly to the sequence. Within twenty cycles of the sequence, a significant decrease in reaction time was observed in comparison to the adjacent random blocks. No such reduction was observed when the sequence occurred on the non response relevant dimension. A potential problem for this experiment was that the subject had already seen the sequence on the initial training trials where it was non response relevant. The obvious route by which to test whether this exposure made any difference to performance on the response relevant block could be to use a comparison group where the training trials were random on both dimensions.

Up to this point, the focus has been on how processing operations during training determine subsequent test performance. In chapter 3, the emphasis is reversed from how processing during training determines test performance, to determining how changed test processing demands influence performance. The point is made that the test should not be conceived as accessing some specific set of knowledge; instead the test accesses a pool of knowledge of the training stimuli. This view is taken from Wright and Whittlesea's (1998) interpretation of the knowledge acquired in the training phase as a set of resources that can be drawn upon for many purposes when interacting with the world. Thus, different tests should access distinct aspects of this pool of knowledge.

This view is supported by the findings of experiments 8 and 9. In these studies, subjects completed digit fragments in the test rather than choosing between exemplars in a forced choice test. Unlike the forced choice test, the completion test was not sensitive to a change in surface form between the training and test periods. This indicates that distinct knowledge bases are used in each test. The fact that performance on the completion test was sensitive to phonological suppression during training (experiment 8) would seem to point to the use of phonological information. The important point is that the training period yields different types of knowledge that can be applied in different ways according to the demands of the test task. The consistency between the test task demands and the processing that occurred during training determines success on the test of knowledge.

A further line of investigation in this thesis was to specify the type of information that is acquired in the invariant learning task. In experiment 4, it was shown that introducing a few negative items into the training period did not disrupt test performance. This indicates that no special or abstract knowledge of invariance is acquired. This finding was replicated in experiments 10 and 11, where completion of digit fragments was no more successful when the fragment was completed with the invariant from training or any other digit. Again, this implies that the invariant from training was not weighted above the other digits encountered during training. Taken together, these findings are good evidence against the McGeorge and Burton (1990) position that abstract knowledge of the invariant feature is acquired. Instead, it is more likely that knowledge of individual experiences is applied at test, as suggested by the episodic processing account. In chapter 4, it was put forward that subjects use memory for specific strings at test in the invariant learning task. These strings are not available to conscious inspection because they are not directly represented. The strings seen in training correlate with the computations carried out on them, and subjects report computations most often as their strategy in the test. So it is likely that using explicit memory for these computations allows above chance performance on this test.

It is important to note that the episodic account predicts that while this conclusion is possible with this particular induction task, another induction task may result in alternative information used in the test. For example, in experiment 5 subjects were asked to search for the highest or lowest digit in the string. When the invariant was high and the training task induced subjects to search for high numbers, performance was maximal. An important secondary finding was that subjects' verbal reports consistently and accurately reported the invariant as the most frequent digit. It is possible to speculate that, at test, subjects were using some strategy of searching for high numbers, or indeed the invariant. Using this induction task in place of computation, a different pattern of performance may be expected in experiments 10 and 11, with the invariant being produced most often in the digit fragments.

In this thesis, several predictions of the episodic account of implicit learning were tested. First, the episodic processing account predicts that the subject must perform some active task on the training stimuli for learning to occur, and this was reflected in the experimental data reported in this thesis. Secondly, the processing carried out on stimuli during the training and test periods must be consistent for successful performance at test. This prediction was also confirmed by the experimental data reported within this thesis.

In contrast, none of the experiments appear to provide evidence for automatic acquisition of knowledge. Learning does not proceed in an automatic manner where the structure of the stimulus environment is "absorbed" without supervision, as some authors suggest (e.g. Lewicki and Hill, 1989). Furthermore, the findings of experiments 4, 10 and 11 do not support the acquisition of abstract knowledge in the invariant learning task. Instead, it appears more likely that episodic knowledge underlies performance on this task. This conclusion is in line with the majority of findings with the artificial grammar learning task (e.g Vokey and Brooks, 1992; Whittlesea and Dorken, 1993; Johnstone and Shanks, 2000).

In sum, converging evidence from several sources indicates that the episodic processing account (Whittlesea and Dorken, 1993) provides the best available explanation of implicit learning. This view of implicit learning is not, however, free of problems. It was suggested in chapter 1 that the evidence cited in its support is taken only from classification or recognition studies. Indeed, the phrase "episodic processing" does not appear appropriate for the sequence learning task as the emphasis it places on the events does not map onto the experimental demonstration of sequence learning. In sequence learning, the increasing efficacy of learning is noticeable as a gradual decrease in response time. This type of learning is more akin to the predictions of the Kolers and Roediger (1984) procedural view of learning. Within this framework, learning is perceived as an accumulation of procedures over time. This build up of procedural knowledge is tied strongly to the processing that conducted on the stimuli during training; a prediction that is very close to those made by the episodic processing account of implicit learning. Indeed, it may be better to consider these types of models under the umbrella term of processing accounts of implicit learning, rather than "episodic" processing. This is quite reasonable since they share the fundamental principle of processing consistency. Using this broader interpretation of processing consistency, the wider experimental context of implicit learning research (i.e. classification, generation or reaction time tasks) can be encompassed by this account.

Broadening the theoretical context of implicit learning research in this way does not come without its costs. The chief concern is that the underlying mechanisms of the processing account are rather hard to pin down. Processing based terms that have been used frequently in this thesis, such as the "consistency of processing" or "training task demands", are not easily specified.

This difficulty with specifying the mechanism of processing type frameworks has been highlighted previously. For example, Baddeley, (1978) suggested that the levels of processing framework for memory could only be described as "a useful rule of thumb for predicting the outcome of certain types of experiments" (p148). Furthermore, in a critique of Kolers and Smythe's (1984) procedural approach to cognition, Allport (1984) asked the question "how should we *specify* the dependence of skills on circumstances" (p323), to which he claimed Kolers and Smythe (1984) "offer no hint of an answer" (p323). These points are no doubt valid criticisms of the processing account in general, but in the case of implicit learning, the alternative structural account lacks supporting evidence (see chapter 1). In comparison, the processing view of implicit learning provides a simple interpretation of implicit learning, which can integrate new findings such as the response of the completion test to various training task manipulations (experiments 8 and 9). The parallel system view is not flexible enough to provide an explanation of these types of effects.

From this we must conclude that, while the processing account provides the best account of implicit learning, one aim of future research should be to specify further the mechanism of aspects of the processing account such as "consistency" or "fluency". Fortunately, some progress on this problem has been made by Whittlesea and Leboe (2000). As reported in chapter 3, they suggest that classification decisions can be based on the use of heuristics. Although there is no space to consider their work in detail, suffice it to say that this work specifies the mechanism of the processing account in more detail than previous analysis.

This processing view of memory and learning is a general view of these functions of the mind, which can provide an explanation of learning in standard circumstances, as well as those in implicit learning experiments. If so, why should there be a demarcation point between learning that occurs in so called implicit learning experiments from other more standard explicit learning contexts? In terms of the cognitive structures that underlie these tasks there is no distinction. However, the experimental finding that people have difficulty verbalising the knowledge they use in implicit learning experiments is unequivocal, so there must be some distinction. This distinction lies in the task demands of implicit learning and explicit learning contexts rather than knowledge structures that are proposed to underlie them. Implicit learning represents circumstances when there is an indeterminate relationship between the explicitly held knowledge acquired during training and the way this knowledge is used at test (as argued by Whittlesea and Dorken, 1993). This indeterminacy does not arise, however, from some specific function of the mind, it results from the incidental task demands of implicit learning experiments. In more typical laboratory learning situations, it is the aim of the experimenter to make the structure of the stimuli unambiguous in order to maximise learning effects, hence subjects can better verbalise the underlying rules.

This interpretation is subtly distinct from processing views proposed elsewhere (e.g. Whittlesea and Dorken, 1997), as it pays attention to the content of verbal reports. In the past, proponents of the processing view have suggested that verbal reports should be abandoned, and research should solely focus on processing manipulations. More generally, Kolers and Roediger (1984) make the point that cognitive processes should be described by what they can do rather than what they know. This position is made untenable by the fact that verbal reports can reveal useful information under certain circumstances. For example, in experiments 5 and 12 the training task focused on the critical aspects of the stimulus structure, resulting in high performance at test and significantly, reasonably accurate verbal reports. These findings are in line with those of Johnstone and Shanks (2000) (see sect 1.8) where subjects directed to the rule structure of biconditional artificial grammars showed both sensitivity at test and good awareness of the rules.

These demonstrations of conditions where verbal report is accurate may point towards a new research strategy in implicit learning research. Experiments where learning is seen to occur in the absence of verbal report should be contrasted with demonstrations of learning where verbal report was accurate. Examining the main differences in the training task demands in each case will lead to a better understanding of how training task demands interact with verbal reports. Hence, conditions where learning can occur without verbal report can be specified more accurately.

References

Allport, A. (1984). Alternatives to the computational view of mind: The baby or the bathwater? *Journal of Verbal Learning and Verbal Behavior*, **23**, 315 - 324.

Altmann, G., Dienes, Z., & Goode, A. (1995). Modality independence of implicitly learned grammatical knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **21**(4), 899 - 912.

Anderson, J. (1982). Acquisition of cognitive skill. *Psychological Review*, **89**, 369 - 406.

Baddeley, A. (1978). The trouble with levels: A reexamination of Craik and Lockhart's framework for memory research. *Psychological Review*, **85**(3), 139 - 152.

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

Berry, D., & Cock, J. (1998). Implicit learning of invariant features. In M. F. Stadler, PA (Ed.), *The handbook of Implicit Learning* . Thousand Oaks London New Delhi: SAGE.

Berry, D., & Dienes, Z. (1993). *Implicit learning: Theoretical and empirical issues*. Hove, UK: Erlbaum.

Boyd, L., & Winstein, C. (2001). Implicit motor sequence learning in humans following unilateral stroke: the impact of practice and explicit knowledge. *Neurosci. Lett.*, **298**, 65 - 69.

Bright, J. (1993). *Issues in Implicit Learning.* Unpublished PhD Thesis, University of Nottingham.

Bright, J., & Burton, A. (1994). Past Midnight: Semantic Processing in an Implicit Learning Task. *Quaterly Journal of Experimental Psychology*, **47A**(1), 71-89.

Bright, J. E. H., & Burton, A. M. (1998). Ringing the changes: Where abstraction occurs in implicit learning. *European Journal of Cognitive Psychology*, **10**(2), 113-130.

Brooks, L. (1978). Non-analytic concept formation and memory for instances. In E. Rosch & B. Lloyd (Eds.), *Cognition and Concepts* . Hillsdale, NJ.: Erlbaum.

Brooks, L., & Vokey, J. (1991). Abstract Analogies and Abstract Grammars: Comments on Reber (1989) and Mathews (1989). *Journal of Experimental Psychology: General*, **120**(3), 316 - 323.

Buchner, A., Steffens, M., Erdfelder, E., & Rothkegel, R. (1997). A multinomial model to assess fluency and recollection in a sequence learning task. *The Quarterly Journal of Experimental Psychology*, **50A**(3), 631 - 663.

Bucher, A. & Wippich, W. Differences and commonalities between implicit learning and implicit memory (pp 3 - 46). In D. Berry (Ed.), *How implicit is implicit learning* (pp. 196 - 234). Oxford: Oxford University Press.

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

Churchill, E., & Gilmore, D. (1998). Selection through rejection: Reconsidering the invariant learning paradigm. *Quaterly Journal of Experimental Psychology*, **51A**(1), 1-17.

Cleeremans, A. (1998). Principles for Implicit Learning. In D. Berry (Ed.), *How implicit is implicit learning* (pp. 196 - 234). Oxford: Oxford University Press.

Cleeremans, A., Destrebecqz, A., & Boyer, M. (1998). Implicit learning: news from the front. *Trends in Cognitive Sciences*, **2**(10), 406-416.

Cock, J., Berry, D., & Gaffan, E. (1994). The role of similarity processing in an incidental learning task. *Quaterly Journal of Experimental Psychology*, **47A**, 1015-1034.

Cohen, A., Ivry, R., & Keele, S. (1990). Attention and structure in sequence learning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **16**(1), 17 - 30.

Curran, T. (1997). Higher order associative learning in amnesia: evidence from the serial reaction time task. *Journal of Cognitive Neuroscience*, **9**(522 - 533).

Curran, T., & Keele, S. (1993). Attentional and non attentional forms of sequence learning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **19**, 189 - 202.

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

Dienes, Z., Altmann, G., Kwan, L., & Goode, A. (1995). Unconscious knowledge of artificial grammers is applied stategically. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **21**, 1322-38.

Dienes, Z., & Berry, D. (1997). Implicit Learning: Below the subjective threshold. *Psychonomic Bulletin and Review*, **4**(1), 3-23.

Dienes, Z., Broadbent, D., & Berry, D. (1991). Implicit and explicit knowledge bases in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **17**(5), 875-887.

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

Fitts, P. (1964). Perceptual-motor skill learning. In A. Melton (Ed.), *Categories of human learning* (pp. 243 - 285). New York: Academic Press.

Frensch, P. (1998). One concept, multiple meanings. In M. Stadler & P. Frensch (Eds.), *Handbook of implicit learning* . London: SAGE Publications.

Gardiner, J., & Java, R. (1993). Recognising and remembering. In A. Collins, S. Gathercole, M. Conway, & P. Morris (Eds.), *Theories of memory*. Hove, England.: Lawrence Erlbaum Associates.

Gomez, R., Gerken, L., & Schvaneveldt, R. (2000). The basis of transfer in artificial grammar learning. *Memory and Cognition*, **28**(2), 253 - 263.

Grafton, S. T., Hazeltine, E., & Ivry, R. B. (1998). Abstract and effector-specific representations of motor sequences identified with PET. *Journal of Neuroscience*, **18**(22), 9420-9428.

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

Hasher, L., & Zacks, R. (1979). Automatic and effortful processes in memory. *Journal of Experimental Psychology: General*, **108**, 356 - 388.

Hefferline, R., Keenan, D., & Harford, R. (1959). Escape and avoidance conditioning in human subjects without their observation of the response. *Science*, **130**, 1338 - 1339.

Highham, P., Vokey, J., & Pritchard, J. (2000). Beyond dissociation logic: Evidence for controlled and automatic influences in artifical grammar learning. *Journal of Experimental Psychology: General*, **129**(4), 457 - 470.

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

Jacoby, L. (1983). Remembering the data: Analyzing interactive processes in reading. *Journal of Verbal Learning and Verbal Behavior*, **22**, 485 - 508.

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

Jacoby, L., & Hayman, C. (1987). Specific visual transfer in word identification. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **13**, 456 - 463.

Jimenez, L., Mendez, C., & Cleeremans, A. (1996). Comparing direct and indirect measures of sequence learning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **22**(4), 948 - 969.

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

Johnstone, T., & Shanks, D. (in press). Abstractionist and processing accounts of implicit learning. *Cognitive Psychology*.

Joordens, S., & Merikle, P. (1993). Independence or redundancy? Two models of conscious and unconscious influences. *Journal of Experimental Psychology: General*, **122**(4), 462 - 467.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, **80**, 237 - 251.

Kelly, S., Burton, A., Riedel, B., & Lynch, E. (2000). Sequence learning by action and observation: evidence for separate mechanisms. *unpublished manuscript.*

Kinder, A., & Shanks, D. R. (in press). Amnesia and the declarative/procedural distinction: A recurrent network model of classification, recognition, and repetition priming. *Journal of Cognitive Neuroscience.*

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.

Kolers, P., & Roediger, H. (1984). Procedures of mind. *Journal of Verbal Learning and Verbal Behavior*, **23**, 425 - 449.

Kolers, P., & Smythe, W. (1984). Symbol manipulation: Alternatives to the computational view of mind. *Journal of Verbal Learning and Verbal Behavior*, **23**(289 - 314).

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

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., & Bizot, E. (1988). Acquisition of procedural knowledge about a pattern of stimuli that cannot be articulated. *Cognitive Psychology*, **20**, 24 - 37.

Lieberman, D. (1993). *Learning, Behaviour and Cognition.* Pacific Grove, California: Brooks/Cole.

Lieberman, D., Sunnucks, W., & Kirk, J. (1998). Reinforcement without awareness: I. Voice level. *Quarterly Journal of Experimental Psychology*, **51B** (4), 301 - 316.

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

Marcel, A. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology*, **15**, 197 - 237.

Mathews, R. (1997). Is research painting a biased picture of implicit learning? The dangers of methodological purity in scientific debate. *Psychonomic Bulletin and Review*, **4**(1), 38 - 42.

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

Mathews, R., & Roussel, L. (1993). Automatic abstraction of stimulus structure from episodes: Comment on Whittlese and Dorken (1993). *Journal of Experimental Psychology: General*, **122**(3), 397 - 400.

Mathews, R., & Roussel, L. (1998). Abstractness of implicit knowledge: A cognitive evolutionary perspective. In D. Berry (Ed.), *How implicit is implicit learning?* . Oxford.: Oxford University Press.

Mayr, U. (1996). Spatial attention and implicit sequence learning: Evidence for independent learning of spatial and non spatial sequences. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **22**(2), 350 - 364.

McGeorge, P., & Burton, A. (1990). Semantic processing in an incidental learning task. *Quaterly Journal of Experimental Psychology*, **42A**, 597-610.

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

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

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

Neal, A., & Hesketh, B. (1997). Episodic knowledge and implicit learning. *Psychonomic Bulletin and Review*, **4**(1), 24-37.

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

Pasual-Leone, A., Grafman, J., Clark, K., Stewart, M., Massaquoi, S., Lou, J., & Hallet, M. (1993). Procedural learning in Parkinson's disease and cerebellar degeneration. *Annals of Neurology*, **34**(594 - 602).

Perner, J., & Dienes, Z. (1999). Higher order thinking. *Behavioral and Brain Sciences*, **22**(1), 164 (7 pages).

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

Perruchet, P., & Gallego, J. (1998). A subjective unit formation account of implicit learning. In D. Berry (Ed.), *How implicit is implicit learning?* (pp. 124 - 161). Oxford: Oxford University Press.

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., & Pacteau, C. (1991). Implicit acquisition of abstract knowledge about artificial grammars: some methodological and conceptual issues. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **120**, 112 - 116.

Perruchet, P., Vinter, A., & Gallego, J. (1997). Implicit learning shapes new conscious percepts and representations. *Psychonomic Bulletin and Review*, **4**(1), 43_48.

Reber, A. (1967). Implicit learning of Artificial Grammers. *Journal of Verbal Learning and Verbal Behavoir*, **5**, 855-63.

Reber, A. (1989). Implicit Learning and Tacit Knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **118**(3), 219 - 235.

Reber, A. (1997). Implicit Ruminations. *Psychonomic Bulletin and Review*, **4**(1), 49 - 55.

Reber, A., & Allen, R. (1978). Analogic and abstraction strategies in synthetic grammar learning: A functionalist interpretation. *Cognition*, **6**, 193 - 221.

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

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

Reber, A., & Squire, L. (1994). Parallel brain systems for learning with and without awareness. *Learning and Memory*, **1**, 217 - 229.

Reber, A., & Squire, L. (1998). Encapsulation of implicit and explicit memory in sequence learning. *Journal of Cognitive Neuroscience*, **10**, 248 - 263.

Redington, M. (2000). Not evidence for separable controlled and automatic influences in artifial grammar learning: Comment on Higham, Vokey and Pritchard (2000). *Journal of Experimental Psychology: General*, **129**(4), 471 - 475.

Redington, M., & Chater, N. (1996). Transfer in Artificial Grammar Learning: A Reevaluation. *Journal of Experimental Psychology*, **125**(2), 123 - 138.

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

Rheingold, E., & Merikle, P. (1988). Using direct and indirect measures to study perception without awareness. *Perception and Psychophysics*, **44**, 563 - 575.

Roediger, H., & McDermott, K. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **21**, 803 - 814.

Roediger, H., Weldon, M., & Challis, B. (1989). Explaining dissociations between implicit and explicit measures of retention : A processing account. In H. C. Roediger, FIM. (Ed.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 3-41). Hillsdale: Erlbaum.

Seger, C. (1994a). Criteria for implicit learning: Deemphasize conscious access, emphasize amnesia. *Behavioral and Brain Scienes*, **17**(3), 421 - 422.

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

Shanks, D., & Cameron, A. (in press). The effect of mental practice on performance in a sequential reaction time task. *Journal of Motor Behavior.*

Shanks, D., Green, R., & Kolodny, J. (1994). A critical examination of the evidence for nonconscious (implicit) learning. In C. Umilta & M. Moscovitch (Eds.), *Attention and Performance XV.* . Cambridge, Mass.: MIT Press.

Shanks, D., & Johnstone, T. (1998). Implicit knowledge in sequential learning tasks. In M. Stadler & P. Frensch (Eds.), *Handbook of Implicit Learning* . London: SAGE Publications.

Shanks, D., & Johnstone, T. (1999). Evaluating the relationship between explicit and implicit Knowledge in a sequential reaction time task. *Journal of Experimental Psychology: Learning, memory and cognition.*, **25**(6), 1435 - 1451.

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

Shanks, D., & St John, M. (1994). Characteristic of dissociable human learning systems. *Behavioural and Brain Sciences*, **17**(3), 367 - 447.

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

Stadler, M., Warren, J., & Lesch, S. (2000). Is there cross format transfer in implicit invariance learning? *Quaterly Journal of Experimental Psychology*, **53A**(1), 235-245.

Thorndike, E. (1911). *Animal Intelligence*. New York: McMillan. (As cited in Lieberman, D. (1993). *Learning, Behaviour and Cognition*).

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

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

Vokey, J., & Brooks, L. (1994). Fragmentary Knowledge and the Processing-Specific Control of Structural Sensitivity. *Journal Of Experimental Psychology: Learning, Memory and Cognition*, **20**(6), 1504 - 1510.

Ward, G., & Churchill, E. (1998). Two tests of instance-based and abstract-rule based accounts of invariant learning. *Acta Psychologica*, **99**, 235-253.

Whittlesea, B., & Dorken, M. (1993). Incidentally, things in general are particularly determined: an episodic-processing account of Implicit Learning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **122**(2), 227-248.

Whittlesea, B., & Dorken, M. (1997). Implicit learning: Indirect, not unconscious. *Psychnomic Bulletin and Review*, **4**(1), 63-67.

Whittlesea, B., & Leboe, J. (2000). The heuristic basis of remembering and classification: fluency, generation and resemblance. *Journal of Experimental Psychology: General*, **129**(1), 84-106.

Whittlesea, B., & Wright, R. (1997). Implicit (and Explicit) Learning: Acting adaptively without knowing the consequences. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **23**(1), 181-200.

Willingham, D., Greeley, T., & Bardone, A. (1993). Dissociation in a serial response task using a recognition measure: Comment on Perruchet and Amorim (1992). *Journal of Experimental Psychology: Learning, Memory and Cognition*, **19**, 1424 - 1430.

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

Willingham, D. B. (1999). Implicit motor sequence learning is not purely perceptual. *Memory & Cognition*, **27**(3), 561-572.

Wright, R., & Burton, A. (1995). Implicit learning of an Invariant: Just say no. *Quaterly Journal of Experimental Psychology*, **48A**(3), 783-796.

Wright, R. L., & Whittlesea, B. W. A. (1998). Implicit learning of complex structures: Active adaptation and selective processing in acquisition and application. *Memory & Cognition*, **26**(2), 402-420.