Auditory Implicit Learning

Beate Riedel

Department of Psychology
University of Glasgow

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

August, 2003

Aknowledgements

I would like to express my gratitude to my supervisor, Mike Burton, for his guidance and complete support. Further thanks go to my husband, Matthew Hammond, as well as my family and friends, for all their encouragement throughout the years.

Abstract

It has been suggested that much of the information we acquire from our external environment involves processes that do not require conscious awareness (e.g. Reber, 1989; Reber and Winter, 1994). Such knowledge acquisition has been termed implicit learning and this has been put forward as a fundamental process in allowing learning of complex information (e.g. Reber, 1992; Schmidke and Heuer, 1997). It has been proposed that acquisition of the underlying rule structure of stimulus events provides an indication of such a process as being fundamental and general. In contrast, learning bound to more peripheral processes should only be shown when subjects learn, for example, surface features of stimuli or a sequence of motor responses, but not the underlying rules (e.g. Perruchet and Pacteau, 1990; Seger, 1998). The research in this thesis investigates systematically whether implicit learning of sound stimuli behaves any differently to such learning of visual stimuli. This expands the empirical scope of previous studies in the implicit learning field and allows assessment of such processes as fundamental and general.

Chapter 1 provides a background to implicit learning in general and introduces the different concepts involved. Chapters 2 to 4 investigated the generality of findings from visual implicit learning studies in the auditory domain. In particular, they studied the role of rule abstraction in sequence learning (Nissen and Bullemer, 1987) and invariant learning tasks (McGeorge and Burton, 1990). Findings from the sequence learning experiments in Chapters 2 and 3 suggest that subjects were unable to abstract the underlying rule structure of stimuli, as would have been evident from learning of the auditory sequences employed by listening alone. Instead, subjects were only able to learn the relevant associations between their actions (keypress responses) and a set of stimuli. These findings add to evidence from visual implicit learning studies that found peripheral processes involved in such learning. Findings from the invariant learning experiments in Chapter 4 show what types of auditory invariant features

subjects can and cannot learn. This identified for the first time the exact information, or rule, that subjects acquire in such a task in an auditory context. Additionally, it provides some evidence that explicit processes may have been involved. Overall, the findings from the experiments in this thesis put into question that implicit learning is a fundamental process, which involves implicit rule abstraction.

Table of Contents

Chapter 1. Concepts of Implicit Learning

1.1 Introduction

Internalizing the regularities that occur in our external environment plays an important role in our everyday lives. Reber (1989), as well as Reber and Winter (1994), suggested that much of this knowledge acquisition occurs through processes that do not involve conscious awareness. Learning without awareness has been termed implicit and it is this that has been put forward as a fundamental process in allowing acquisition of complex information (e.g. Buchner and Steffens, 2001; Reber, 1992; Schmidtke and Heuer, 1997). Whether learning can proceed without awareness is continuing to generate great interest in cognitive psychology. A particular focus in implicit learning research is on whether such learning can be characterized as an unselective and general process (e.g. Berry and Broadbent, 1988; Buchner and Steffens, 2001; Reber, 1989). Thus, many studies have focused on the nature of the acquired knowledge: if a general process is involved, the knowledge acquired should be of more abstract nature and transfer to different learning contexts, as well as, for example, across domains (e.g. Altmann, Dienes and Goode, 1995). Several studies have explored whether the knowledge acquired in implicit learning tasks consists of the underlying rule structure (i.e. abstract knowledge) or is of more peripheral nature, such as learning of surface features (e.g. Bright and Burton, 1994) or learning of a sequence of motor responses (e.g. Willingham, 1999). Much of the research in implicit learning has been conducted in a visual and visio-motor context. Broadening the scope of implicit learning research using stimuli other than visual would provide empirical evidence to the claim that implicit learning is a fundamental and general process (Buchner, Steffens, Erdfelder and Rothkegel, 1997). The question at the heart of this thesis is whether implicit learning of auditory nonverbal stimuli behaves differently to implicit learning of visual stimuli. By extension, it explores the claim that implicit learning can be described as general by systematically extending it to an auditory context of learning.

In order to provide a background to implicit learning in general and the experimental tasks employed, the first chapter of this thesis introduces the concept of implicit learning and describes some of the main theoretical issues researchers have been concerned with in the implicit learning literature.

1.2 Learning without awareness – early studies

Learning without insight

One of the first experimental studies that claimed to have shown learning without insight was by Thorndike (1911). In his study, cats were rewarded with food when they managed to escape from a so-called puzzle box. A hungry cat, after having spent some time inside the box trying to get out, would eventually pull a lever that opened the door. This happened inadvertently as part of the cats' general attempts of trying to escape to get to the food placed outside the box.

Thorndike observed that when these cats were placed in the box repeatedly, the time taken before pulling the lever would gradually shorten. This instrumental learning set-up is an early example of how a response (i.e. pulling lever) can be strengthened by a reward (i.e. food). Thorndike concluded that the animals' behaviour was mediated by an automatic strengthening of the link between stimulus and response, and that any such association can be formed regardless of the type of reward or response. He concluded that subjects had no conscious awareness of what they were doing. Later experiments conducted with human subjects seemed to support the possibility of rule learning without conscious awareness.

Verbal conditioning procedure

Thorndike (1932) conducted one such study with human subjects. Over a number of days, subjects were presented with several hundred cards. Four lines of equal length were printed on these. Subjects had to indicate which line they thought was the longest and feedback was provided on response. Other distinguishing details could also be seen on most of these cards. These, for example, consisted of a particular number or an ink mark. Thorndike found that subjects seemed to learn an association between a particular feature and a correct response. However, subjects were not always able to report such a correct relationship when asked directly.

Thorndike and Rock (1934) expanded the verbal conditioning procedure to a word association task. Here the experimenter told the subject a particular word and their task was to respond to that word immediately with the first word that came to mind. Subjects were then given a monetary bonus or penalty depending on whether their response had been correct or not. If a response word could follow the experimenter's word in a normal sentence (e.g. 'behind' followed by the subject responding with 'the curtain' or 'the door') this was deemed a correct response. If a response related semantically to the probe word (e.g. 'before' followed with the response 'after') this was deemed an incorrect response. Importantly, subjects were not told of the rule the experimenter employed to decide whether they had made a correct or incorrect response. Thorndike and Rock (1934) found that subjects gradually increased correct responses from the first blocks of trials to the last. They concluded, as with Thorndike's earlier experiments involving cats, that reinforcement (i.e. a correct response followed by a reward) occurred automatically. If subjects had any insight into the relationship between their response and the experimenter's reward, they should have been able to apply the correct response deliberately. This would have resulted in an abrupt rise in correct replies (and in an abrupt drop in latency in the experiments involving cats). However, such a rapid increase in applying the correct response was not observed. Thorndike and Rock (1934) concluded that subjects' gradual increase in correct responses indicated a lack of explicit awareness into the response-reward relationship.

Greenspoon (1955) took up the verbal conditioning procedure and tested subjects in a different context than Thorndike and Rock (1934). Subjects were put into a supposed interview situation and were asked to say as many words they could think of in a limited time span. Whenever they said a plural noun the experimenter provided some reinforcement by saying “mmm-hmm”. Greenspoon (1955) found that with progression of the session the frequency of plural words increased. This occurred despite an apparent lack of awareness of the reinforced production of plural words. However, there have been some criticisms of this study (Dulany, 1961; Hefferline, Keenan and Harford, 1959; Levin, 1961).

The first of these concerns the conclusion that subjects lacked awareness of the reinforcement contingency. Levin (1961) pointed out that the post experimental test of awareness might have lacked sensitivity to pick up any explicit knowledge subjects may have held. Thus, asking subjects whether they had been aware of the purpose of the experiment may have simply been too vague in eliciting any relevant knowledge subjects may have held, despite the apparent lack of verbalizing it.

A second question was raised by Dulany (1961). He investigated whether subjects had learned the actual rule the experimenter had intended, or whether subjects may have learned a different association. In a replication of Greenspoon’s (1955) study, Dulany (1961) found that subjects responded with nouns that belonged to the same category as the actual reinforced word, while maintaining the plural form for all of these responses. An example of this is the word “diamonds” which could result in subjects producing related words such as “rubies, emeralds, pearls, et cetera”. This left a clear question as to Greenspoon’s (1955) hypothesis that subjects lacked awareness while producing seemingly correct plural responses, when the rule they may have learned may have been different from the intended.

The third criticism concerned the possibility that subjects may have become aware of the experimenters' odd behaviour in reinforcing plural words. This could have left the possibility that subjects tried to search for a rule behind the experimenter's behaviour. In an attempt to reduce subjects' suspicion, Hefferline, Keenan and Harford (1959) conducted a conditioning experiment that recorded a particular muscle's movement in a subject's thumb. Importantly, subjects were provided with a very believable cover story: assessing the effects of stress on body tension. Subjects were told that they would be exposed to randomly interchanging intervals of harsh sounds and soothing music. Unbeknownst to subjects, however, these intervals were not played randomly, but were determined by a minimal muscular movement in the subject's thumb. This movement was invisible to the naked eye. Subjects were connected to a number of electrodes during the experiment and one of these recorded the invisible muscle movement in subjects' thumbs. As the session progressed, there was a gradual increase in the contractions of that particular muscle. On post experimental questioning, subjects were unable to verbalize anything relating to the reinforced muscle contractions. One step further in providing a believable cover story, which reduced the possibility of rule searching by subjects, can be found in studies by Lieberman (2000) and Lieberman, Sunnucks and Kirk (1998). Here, subjects were under the impression they were taking part in extrasensory perception (ESP) experiments. Subjects, who were in a different room from the experimenter, were asked to indicate which word out of two the experimenter was thinking about. In both studies subjects gradually increased selection of the correct word, which was tied to a particular rule that had been reinforced (in the former study a 'correct' word was reinforced when it contained a double letter, e.g. apple; in the latter the correct word was tied to the volume of the subject's response). In neither study were subjects able to report any knowledge of the response-reward contingency.

Overall, these early examples of experiments investigating learning without awareness appear to indicate that subjects can form an association between a response and reinforcement without awareness. However, they also provide an early indication of the methodological problems that later studies in implicit learning have been criticised for, such as how to measure whether the acquired knowledge is unconscious and how to operationalize studies investigating learning without awareness. This will be discussed in more detail in later sections in this chapter.

So far, the experiments introduced here involved associations between a response and reinforcement, which can be deemed as fairly simple when one considers some of the complex associations we learn to make in our everyday environments. Human learning does not only consist of these fairly simple associations, but we are used to learning far more complex structures in our everyday environments. The question arises whether learning of simple associations, as exemplified above, can generalize to more complex forms of human learning. One such form of complex learning is the acquisition of natural language.

1.3 Language acquisition as an example of implicit learning

It is generally accepted that the use of language does not require explicit knowledge of the underlying grammatical rules (e.g. Chandler, 1993; Cleeremans, Destrebecqz and Boyer, 1998; Frensch and Rünger, 2003). In fact, most of us learn to recognize and produce grammatical sentences without being able to state the rules of the grammar that underlie the language (e.g. Dienes and Berry, 1997). Although some researchers have argued that acquisition and utilization of information is almost invariably linked to conscious awareness (e.g. Shanks and St. John, 1994), Reber and Winter (1994) disagreed with this argument. They suggested that implicit processes appear to govern large areas of knowledge acquisition required in our everyday lives and one such example is that of language learning and use. Reber and Winter (1994) argued that acquisition of natural language or categories cannot be explained by conscious learning processes. More recently researchers have pointed out connections between implicit learning and psycholinguistics (e.g. Cleeremans et al, 1998; Redington and Chater, 1997; Saffran, Johnson, Aslin and Newport, 1999; Saffran, Newport, Aslin, Tunick and Barrueco, 1997). Research is currently expanding that explores this connection empirically. For example, Saffran et al (1997) investigated how incidental exposure to language-like auditory stimuli (e.g. bupadapatubitutibu...) was enough to enable both children and adult subjects to segment sequences of sounds, which had been played to them continuously, into artificial words (e.g. bupada, patubi, etc.). These artificial words were contained in the original auditory sequence. The ability to segment the continuous stimulus stream was evidenced by above-chance performance on a post experimental recognition test. Saffran et al (1997) proposed that the ability to segment words develops through mechanisms that allow the detection and use of the statistical properties contained in syllable sequences. They suggested this as evidence that this allows language learners to discover words in continuous speech (see also Aslin, Saffran and Newport, 1998). Saffran et al (1997) based their interpretation in the implicit learning literature. The connection can be seen in the fact that language acquisition, like implicit learning, is said to involve incidental learning of complex, structured information (Berry and Dienes, 1993; Cleeremans, 1993; Cleeremans et al, 1998), where incidental has been used as one of the characteristics commonly used to describe implicit learning (this attribute is discussed in detail in section 1.4 below).

It seems that some forms of complex learning in humans, such as language acquisition, involve processes that comprise learning without the intention of the learner to acquire the specific knowledge. Learning of natural languages is such an example, but the question arises as to how implicit learning can be defined and demonstrated in an experimental setting.

1.4 Some definitions and characteristics of implicit learning

The various ways in which implicit learning has been demonstrated in laboratory studies is greatly reliant on the underlying definitions of implicit learning and how these can be operationalized (e.g. Cleeremans et al, 1998; Frensch, 1998; Shanks and St. John, 1994). So far, researchers in implicit learning have failed to provide a satisfactory and unitary definition of implicit learning (e.g. Cleeremans et al, 1998; Frensch, 1998). Frensch (1998) commented that there are "literally dozens of definitions that have been offered and continue to be offered in the literature" (p. 51). This diversity can be seen as symptomatic of the conceptual and methodological problems of implicit learning studies in general. However, there are some common themes that are evident in most definitions and tasks and these will be focussed on here. Although it is beyond the scope of this thesis to provide a unitary definition, it is necessary to be familiar with the main questions that have arisen with regards to defining and demonstrating implicit learning. This will provide a background to the research in this thesis. In line with this, the following section provides a review of the most common attributes implicit learning tasks have been associated with.

Clearly, there are many ways in which implicit learning has been operationalized in experimental studies (examples of the relevant paradigms are provided in section 1.6 below) and the different tasks used reflect different aspects of the phenomenon (Frensch, 1998). Therefore, definitions will mainly be introduced with the main attributes that implicit learning has been associated with:

1) developing a sensitivity to the structural organisation of the stimuli, 2) incidental training conditions and 3) difficulty to express the acquired knowledge verbally.

1.4.1 Sensitivity to the structure of the stimulus domain

Perruchet and Gallego (1997) described implicit learning as a mode in which:

> "subjects' behaviour is sensitive to the structural features of a previously presented situation, without this [...] being due to the intentional exploitation of subjects' explicit knowledge about these features" (p. 124).

Perruchet and Gallego (1997) emphasized two components in their definition: 1) a performance change in subjects' behaviour and 2) failure to mediate any explicit knowledge of this sensitivity. This section focuses on the first part of this description (the second is discussed in detail in section 1.4.3 below).

Measuring implicit learning indirectly

Sensitivity to the structural properties of the stimuli is at the core of demonstrating whether subjects have acquired any knowledge in a learning episode in implicit learning tasks. Such tasks use stimuli that are organized into specific rule structures. This structure governs the relationship between the components making up the learning material. In operational terms it is this that is looked for in implicit learning tasks besides the implicitness of the knowledge acquired. As Perruchet and Gallego (1997) suggested, implicit learning can be demonstrated when subjects show a performance increase to the structure of a complex situation on a behavioural measure. Thus, in a typical implicit learning task, such as artificial grammar learning, subjects are asked to memorize letter strings (e.g. XXRTRXV and QQWMWQP). Unbeknownst to subjects, these are generated according to some rules. Following the memorisation phase, subjects are informed that the strings followed certain rules. Subjects then perform a classification task on a further set of strings, in which they have to class exemplars as following the rules or not. Typically, despite an inability to verbalize the rules, subjects perform above chance in selecting those letter strings that follow the original rules from the memorisation phase (e.g. Reber, 1989). In this case, subjects are said to have become sensitive to the rule structure and this is represented in the preference for items that follow the underlying grammatical rule at test. Another example comes from sequence learning tasks in which subjects are typically asked to respond to some stimuli with keypress responses. Unbeknownst to subjects, stimuli follow a regular, repeating sequence and subjects become sensitive to the underlying regularity as demonstrated in a typical facilitation while responding to this structure (e.g. Nissen and Bullemer, 1987). Invariant learning tasks provide another example of a typical implicit learning task. Here subjects are exposed to a series of stimuli all of which contain an invariant – a stimulus 'quality' that remains constant across trials. Following this exposure subjects are presented with pairs of stimuli, one of which contains the invariant (the positive) and one of which does not (the negative). When subjects are asked to select the item they think they have been exposed to previously they select positive items above chance level, despite the fact that all test items are new to them (e.g. McGeorge and Burton, 1990). This behavioural facilitation in invariant learning tasks is an example of how sensitivity can be measured behaviourally. These tasks can be seen as indirect tests of implicit learning, where the behavioural measure provides information about whether subjects have become sensitive to the underlying structure as suggested by Perruchet and Gallego (1997). It is this sensitivity that is said to indicate that subjects have acquired the underlying rule knowledge in a typical implicit learning task.

1.4.2 Incidental learning conditions

Cleeremans et al (1998) pointed out that one of the most common and, in their view, conceptually impartial description of implicit learning is as follows:

> "...learning is implicit when we acquire new information without intending to do so, and in such a way that the resulting knowledge is difficult to express." (p. 406).

This was also reflected in definitions by, for example, Berry and Dienes (1993) and Seger (1994). Their definitions emphasize the role of intention, as well as the role of accessing what has been learned. A clear division can be seen in this definition: first, the nature of the learning process involved and, secondly, retrieving the acquired knowledge after this process has occurred. According to Frensch (1998) the distinction between such learning and retrieval processes principally affects what has to be demonstrated in a learning episode.

Reber (1967; 1993), as well as Frensch (1998), suggested that the learning process itself is of essence in demonstrating implicit learning, not memory or retrieval. This leads to the second attribute commonly associated with implicit learning: incidental learning conditions. Here subjects are not instructed to learn anything during a learning task and there is no conscious effort to learn the underlying structure of the stimuli, which subjects nevertheless learn. Therefore, subjects cannot have any intention of learning the relevant information and this is what the first part of Cleeremans et al's (1998) definition encapsulates. The attribute is primarily relevant for the operationalization of implicit learning (Frensch, 1998; Cleeremans and Jimenez, 2002). Thus, many researchers go to great length in keeping the true purpose behind an implicit learning experiment hidden from subjects in order to avoid any intentionality in subjects' learning as regards to the stimulus structure (e.g. use of cover stories).

Distinguishing between unconscious and automatic processes

Mathews, Buss, Stanley, Blanchard-Fields, Cho and Druhan (1989) suggested that implicit learning is automatic and occurs without conscious awareness. Frensch (1998) recommended in his review of implicit learning concepts and their operationalization that, in principle, it would be scientifically useful to describe implicit learning as:

> "the nonintentional, automatic acquisition of knowledge about structural relations between objects or events" (p. 48).

This was also reflected by Stadler and Frensch (1994) who argued that learning can be deemed implicit when intention is not involved in the learning process. Frensch (1998) concluded that one of the most common distinctions in conceptualizations of implicit learning is whether the term implicit is taken to be synonymous with unconscious or nonintentional processes. Here, nonintentional is used synonymously with the term automatic and refers to processes that are not intentionally controlled and do not require attention (Cleeremans, 1997; Hasher and Zacks, 1979; Logan, 1990; Seger, 1998). It has also been suggested that automatic processes are fast, involuntary and effortless, and unavailable to conscious awareness (Cleeremans, 1997; Eysenck and Keane, 1990). Frensch (1998) concluded that the term unconscious, in terms of its measurement, does not provide for unambiguous operationalization of the concept of implicit learning, when implicit is taken to mean unconscious. In contrast, conceptualizing implicit as automatic/non-intentional allows an unambiguous operationalization of such learning (Frensch, 1998).

Neal and Hesketh (1997) argued that it is possible to utilize the dissociation between the separate influence of intentional and non-intentional processes on performance based on Jacoby's (1991) process dissociation procedure (PDP). Jacoby's (1991) PDP is a method that allows extraction of separate estimates of conscious and unconscious influences on memory. An example illustrates this procedure (e.g. Destrebecqz and Cleeremans, 2001): Subjects are first trained under incidental conditions on some stimulus structure (training phase). Following this, subjects' explicit knowledge is assessed using two generation tasks, one conforming to an inclusion and one to an exclusion task as proposed by Jacoby (1991). In the inclusion condition, subjects are asked to generate all the elements they had been exposed to in the training phase. They can do this based on conscious recollection or guessing. In the inclusion condition both conscious and unconscious processes can contribute to subjects' performance. In the exclusion condition subjects are explicitly asked to generate elements that are different from those in the training phase. Any items generated from the training phase in the exclusion condition must be due to non-intentional processes. Hence, in contrast with performance in the inclusion condition, conscious and unconscious performances here act against each other. Neal and Hesketh (1997) argued that this adaptation of the PDP could detect non-intentional influence of implicit knowledge on performance. They suggested that detection of implicit processes in this context does not depend on the sensitivity of the awareness test. Therefore, this may be a more appropriate way of confirming unconscious learning.

Buchner, Steffens and Rothkegel (1997), as well as Buchner, Steffens, Erdfelder and Rothkegel (1998), adopted Jacoby's (1991) PDP to a sequence learning context to obtain measures of unconscious and conscious knowledge. They found an association between intention to learn the underlying stimulus structure and increased performance on an awareness test, while task performance was unaffected. This is in line with Neal and Hesketh's (1997) suggestion that consciously and unconsciously stored knowledge can be influenced differently by intention.

Redington (2000), however, argues that the use of intention as a means of separating conscious and unconscious influences on task performance disregards data from verbal reports. Since many definitions and operationalizations of implicit learning include lack of verbalization as an indicator of unconscious processes, Redington (2000) suggested that accounts of implicit learning that exclude an explanation for such lack in verbalization are insufficient.

1.4.3 Verbalization of implicitly acquired knowledge

Focusing on the retrieval of knowledge that has been acquired in an implicit learning episode, this leads to one of the most contentious areas in the implicit learning literature. Seger (1994) described implicit learning as:

> "learning of complex information without complete verbalizable knowledge of what is learned" (p. 163).

The notion of lack of verbalizable knowledge is reflected in many researchers' definitions of implicit learning (e.g. Cleeremans et al, 1998; Berry and Dienes, 1993; Bright, 1993; Reber, 1967). The hypothesis that an inability to verbalize any acquired knowledge in such an implicit task demonstrates unconscious learning (e.g. Reber, 1993) is at the heart of this debate. Thus, as indicated by Seger (1994), as well as in the second part of the definition put forward by Cleeremans et al (1998), the issue here is expression or verbalization of any implicitly acquired information. Much of the evidence for implicit learning representing unconscious processing originates in the apparent dissociation between task performance on the one hand and verbalizable knowledge on the other. This dissociation has been shown and utilized in artificial grammar learning (e.g. Reber, 1989), dynamic systems (e.g. Berry and Broadbent, 1988), sequence learning (e.g. Nissen and Bullemer, 1987) and invariant learning tasks (e.g. Wright and Burton, 1995). All have demonstrated that subjects can perform at above chance level in these tasks without an associated ability to report all the underlying rules (e.g. Berry and Broadbent, 1984; McGeorge and Burton, 1990; Nissen and Bullemer, 1987; Reber, 1967). However, Shanks and St. John (1994), in their major review of implicit learning studies, raised a question over tasks that put forward a lack of concurrent awareness by demonstrating failure to verbalize the knowledge subjects acquired. Since it is operationally difficult to tap into any knowledge during the process of knowledge acquisition, subjects are generally assessed on this point *after* exposure to the relevant stimulus structure, and this commonly consists of some form of verbal report (e.g. McGeorge and

Burton, 1990; Reber, 1967). However, the question arises as to *what* the awareness test in this situation actually taps into: does it access any conscious knowledge subjects may hold? This was one of Shanks and St. John's (1994) major criticisms as they clearly showed a mismatch in attempts to identify the knowledge subjects acquired *during* a particular learning episode with awareness tests *after* stimulus exposure. They did not argue against the use of post experimental awareness tests per se, but suggested that certain criteria should be met if learning in such tasks is to be described as unconscious. Shanks and St. John (1994) suggested the use of two criteria: the Information Criterion and the Sensitivity Criterion.

Two criteria for assessing whether implicitly acquired knowledge is conscious

The first of the two criteria, the Information Criterion, is concerned with the ability of an awareness test to access the relevant information the subject acquired during a learning episode. This criterion focuses on ensuring that the information that is sought in an awareness test is indeed the information responsible for the performance changes that are said to show learning in the first place. This is an important issue, since subjects may have used information other than that sought by the experimenter in the awareness test and it may be this that prevents subjects from expressing the sought knowledge in the first place. If the experimenter failed to tap into the relevant information responsible for incurring a performance change they may falsely conclude that the knowledge subjects failed to verbalize was not conscious, when, in fact, subjects may have simply used different information that may have given rise to the same performance change. Therefore, it is necessary to ensure that a given awareness test accesses the information subjects actually used in the learning task.

The second criterion, the Sensitivity Criterion, is concerned with the sensitivity of the awareness test to tap into *all* relevant conscious knowledge subjects hold. Thus, there is a possibility that subjects have conscious knowledge that is not detected by the awareness test, but this knowledge may contribute to the performance change that is deemed to show learning. Thus, Shanks and St. John argued that the awareness test should be at least as sensitive as the performance test that shows learning. To achieve this sensitivity in the awareness test the performance and the awareness tests should be as similar as possible when it comes to the retrieval context for unconscious knowledge, and may only differ in regards of task instructions (Shanks and St. John, 1994).

Objective tests for measuring implicit learning

Shanks and St. John (1994) considered a verbal awareness test as insufficient for conclusively showing lack of conscious knowledge for reasons encapsulated in the Sensitivity Criterion. Verbal tests will always be dependent on the experimenter asking the correct questions, even in a very detailed and structured interview. Recognition tasks have been put forward as the tasks of choice for assessment of conscious knowledge (Shanks and St. John, 1994). However, Dienes and Perner (1994), as well as Seger (1994), point out that contamination by implicit processes may be responsible for any performance change in these tasks. Therefore, there is a possibility that recognition tasks may tap unconscious as well as conscious processes. This is in line with the notion that no task can be totally 'process-pure', that is tasks can be sensitive to both implicit and explicit influences (Cleeremans, Destrebecqz and Boyer, 1998; Jacoby, Yonelinas and Jennings, 1997).

Jimenez, Mendez and Cleeremans (1996) employed tests that complied with the requirements laid out by Shanks and St. John (1994) for their Information and Sensitivity Criteria in a sequence learning task. In order to do so, they used direct and indirect tests to measure the effects of conscious and unconscious influences. Tasks only differed in terms of the instructions given, but were otherwise matched regarding, for example, task context and demands. Thus, the direct test asked subjects to use any conscious knowledge they were aware of, whereas instructions in the indirect test did not refer to any conscious knowledge. The assumption was that the direct test would show greater sensitivity to conscious knowledge, whereas any sensitivity to aspects of the stimulus in the indirect test must have been due to unconscious influences. Using these tests, Jimenez et al (1996) claimed their findings showed learning without conscious awareness. However, Shanks and Johnston (1999) found conflicting results in a similar sequence learning study using direct and indirect tests. Using these objective tests, their results suggested that conscious knowledge was fully accessible on these test. These contradictory results indicate that the use of direct and indirect tests does not always provide a clear picture of the role of conscious knowledge in these tasks.

A subjective criterion for measuring implicit processes

Reingold and Merikle (1990) suggested that an adequate test of awareness must be *exhaustive*. That is, it must be powerful enough to uncover *all* of subjects' explicit knowledge. This is basically the same as Shanks and St. John's (1994) Sensitivity Criterion. It follows that if a measure of explicit knowledge is not exhaustive, there is always the possibility that some explicit knowledge that remained unmeasured was responsible for learning (Neal and Hesketh, 1997). This leads to the conclusion that the Sensitivity Criterion cannot be achieved if there is no exhaustive test of explicit knowledge. Rather than developing such a measure of explicit knowledge, Merikle (1992) and Dienes and Berry (1997) put forward the subjective threshold criterion in order to distinguish conscious from unconscious processes. This is based in the subliminal perception literature (Cheesman and Merikle, 1984). In a typical subliminal experiment subjects are presented with subliminal stimuli that may or may not have preceded a target stimulus. The subliminal stimulus leads to a primed response to a target stimulus. After each presentation, subjects are required to state whether the stimulus was present, and provide a confidence rating for their response (e.g. Marcel, 1983). In order to verify unconscious learning Cheesman and Merikle (1984) proposed the use of a *subjective threshold* and an *objective threshold*. In the subjective threshold subjects believe they are guessing, although they are performing above chance level. The objective threshold is the point below which subjects score at chance on a recognition or cued recall task of awareness. Dienes and Berry (1997) argued that an implicit task is due to unconscious processes when subjects believe they are guessing (named the *guessing criterion*, Dienes, Altmannn, Kwan and Goode, 1995) or when their confidence ratings are found to be irrelevant to their discrimination accuracy (named the *zero correlation criterion*, Dienes et al, 1995). A second type of metaknowledge category was used by Reber (1993). Reber (1993) distinguished between knowing *that* and knowing *why*, arguing that subjects' ability to respond correctly on a behaviour measure does not necessarily mean the knowledge they used was not implicit as they may be unable to state *why* they chose one letter string as grammatical in, for example, an artificial grammar learning task.

It has become clear that the description of implicit learning as learning without an associated rise in verbalization of what has been learned is problematic in terms of how to access all relevant information at test, as well as finding useful tests that can get to this information. However, as Frensch (1998) argues, this particular definition has predictive value and allows the empirical investigation of attributes thought to be associated with implicit learning.

Investigating implicit learning by extracting any conscious knowledge subjects may hold in combination with operationalizing such learning as non-intentional is related to attempts of dissociating implicit and explicit processes in implicit learning tasks. This can largely be seen in attempts to dissociate implicitly and explicitly acquired knowledge in, for example, the process dissociation procedure described above, or the operationalization of implicit tasks as non-intentional.

1.5 Implicit learning as a general phenomenon

Implicit knowledge acquisition has been characterized as an unselective and automatic accumulation of associated information (e.g. Berry and Broadbent, 1988; Reber, 1989). Reber (1993) suggested that such learning is a fundamental cognitive process, which allows acquisition of complex information that is unavailable to deductive reasoning. Internalizing the regularities underlying the variations in our external environment plays an important role in our everyday activities. Much of implicit learning research utilizes visual stimuli. If implicit learning is an unselective and general process, visual implicit learning studies should transfer to other domains. Widening the scope of implicit learning research using stimuli other than visual would provide empirical evidence to the claim that it is a fundamental and general process (Buchner, Steffens, Erdfelder and Rothkegel, 1997). This could be achieved by employing auditory stimuli in place of visual material.

As already pointed out, acquisition of natural languages is an example of implicit learning in real world environments. Recently, Saffran, Johnson, Aslin and Newport (1999) have extended their previous research to non-linguistic auditory sequences, which were organized into 'tone words' without any phonetic content. Their results indicated that an implicit learning mechanism, which had previously been shown to be involved in word segmentation (Saffran et al, 1997), can also be involved in the segmentation of non-linguistic sequences. Utilizing auditory material, such as simple tones, would widen the scope of empirical research in implicit learning to other domains. Thus, studying implicit learning systematically in the auditory domain will allow evaluation of its generality to other contexts. There are obvious parallels in language and auditory processing, as both have an auditory component in their processing. Like natural languages, learning of, for example, music involves acquisition of highly structured information (e.g. musical grammar, Bigand, Perruchet and Boyer, 1998). Additionally, speech and music consist of a succession of particular sounds occurring in specific orders (e.g. Warren, 1993). A substantial amount of research has been dedicated to the learning processes involved in language, but little to the processes involved in music learning (Tillmann, Bharucha and Bigand, 2000). This extends to a lack of research into auditory implicit learning processes in general. However, there are some implicit learning studies that utilized auditory material. One such example investigated the possibility that acoustic material may impose specific constraints on artificial grammar learning. Bigand, Perruchet and Boyer (1998) used sequences of timbres that were generated according to a typical artificial letter grammar task. Results indicated that subjects were able to learn this auditory artificial timbre grammar (Experiment 1). Importantly, when subjects were presented with auditory material (timbres) they failed to transfer any acquired knowledge to visual letter strings. Bigand et al (1998) presented these results as an indication that exposure to grammatical sequences of timbres primarily leads to knowledge of the surface regularities, rather than to acquisition of the underlying rules or abstract knowledge.

Learning of abstract rules

Tests used to measure abstract learning generally require subjects to make a judgment about a stimulus rather than measuring reaction times. Importantly, rule abstraction or acquisition of abstract knowledge, including that of underlying rules in implicit learning tasks, would not be expected to be bound to surface features of stimuli (e.g. Seger, 1998). Thus, if subjects are able to abstract an underlying rule, rather than rely on matching stimulus features, they should be able to transfer any acquired knowledge across different stimulus formats that follow the same underlying rule structure. Such cross-format transfer has been shown in various artificial grammar learning experiments, in which subjects showed transfer of grammar knowledge acquired on learning strings of letters to different letter sets at test that maintained the underlying grammatical rules (e.g. Gomez and Schvaneveldt, 1994; Knowlton and Squire, 1996; Mathews et al, 1989). Additionally, some studies have shown such transfer to letter-like symbols (e.g. Altmann, Dienes and Goode, 1995, Experiment 4), as well as across sensory modalities (Altmann et al, 1995, Experiments 1 and 2; Manza and Reber, 1997). Further evidence of cross-format transfer comes from invariant learning studies (e.g. Bright and Burton, 1994; Newell and Bright, 2002a/b). Here, subjects were presented with study items that all followed a specific underlying rule (e.g. clock faces in a specific time-range). At test, subjects were asked to make a forced-choice decision between new items, one that followed the rule from the study phase (times falling within the 'invariant' time-range) and one that did not. Importantly, test items were presented in a different surface form (as digital clock faces). Clear transfer effects were found at test when subjects selected those items that followed the specific time-rule they had been exposed to at study. These results provided clear indication that learning was not tied to surface characteristics of the stimuli, but that underlying representation of knowledge was responsible for the cross-format findings (Bright and Burton, 1994).

The findings from visual cross-format transfer studies described above are in contrast to the findings reported by Bigand et al (1998), who failed to show transfer of an artificial grammar rule acquired from exposure to auditory sequences of timbres to visual letter strings. Bigand et al's (1998) results also run counter to findings by Altmann et al (1995, Experiments 1 and 2) who were able to show transfer from the visual (letter sequences) to the auditory modality (pitch sequences) and vice versa in an artificial grammar task (although the latter effects were relatively small).

There are certain implications to findings that show cross-format or cross-domain knowledge transfer. As already pointed out, such transfer would be an indication that the information acquired is not bound to the surface features of the stimuli (e.g. Newell and Bright, 2002b; Seger, 1998), but that the learned knowledge must be of more abstract nature. By extension, such learning could be described as bound to more central cognitive processes, rather than to peripheral ones. It is this that links the described transfer studies to the question of whether implicit learning is a general cognitive process (Reber, 1993). Learning of the underlying rules would be an indication of knowledge acquisition through a more central process. The results from auditory artificial grammar tasks are mixed and do not provide a clear picture on this issue. Results from invariant learning studies indicate a process of rule abstraction, but have not been expanded to include stimuli from other sensory domains that would allow extension of results from the visual domain.

Other auditory implicit learning studies

Perruchet, Bigand and Benoit-Gonin (1997) conducted an auditory implicit learning study using a different paradigm from those described above. This used a sequence learning task constructed exclusively of tones (Experiment 4). Their results failed to show any learning when subjects where asked to respond to each tone with a corresponding keypress response. In contrast, Buchner, Steffens, Erdfelder and Rothkegel (1997) and Buchner, Steffens and Rothkegel (1998) demonstrated facilitation for rule-governed sequences using auditory events.

Here, subjects were required to make corresponding keypress responses to four tones played in a regular, systematic sequence. The findings reflected those originally found by Nissen and Bullemer (1987) and Lewicki, Czyzewska and Hoffman (1987) and confirmed that subjects are able to acquire a sequence composed of auditory material. Buchner et al (1997) suggested that discrimination of tones does not differ from that of visual objects presented in different spatial locations, thus extending sequence learning research to other stimulus domains.

The few results from implicit learning studies that utilized auditory stimuli do not provide a consistent picture and provide, at best, conflicting results. Importantly, they do not allow any insights into whether implicit learning of auditory material behaves any differently to that of visual stimuli. This includes the possibility of transfer effects across domains or learning of underlying rules, which would provide an indication of implicit learning as a general cognitive process.

1.6 Relevant tasks to investigate implicit learning

Acquiring information in the real world is an extremely complex process. Investigating implicit learning using laboratory tasks provides a simpler solution to studying such learning in the real world. Winter and Reber (1994) suggested that artificial learning contexts may serve as a model for understanding the implicit learning processes in natural environments. Artificial material is simpler to control and manipulate for the purposes of experimental investigations than environmental sequences of events. Some of the main paradigms used to investigate implicit learning in the laboratory provide a model for such learning in the real world. Two of these paradigms are introduced here in order to provide a background to the experimental tasks employed throughout the experimental chapters in this thesis.

As already pointed out in Section 1.4, there are a variety of definitions of implicit learning available and this is reflected in the different ways such learning has been operationalized (e.g. Cleeremans et al, 1998; Frensch, 1998; Shanks and St. John, 1994). Seminal studies by Reber (1967) provided the beginnings to what has developed into empirical investigations of the phenomenon through a variety of experimental tasks. The main paradigms are: dynamic systems control (Berry and Broadbent, 1984), artificial grammar learning (Reber, 1967), sequence learning (Nissen and Bullemer, 1987), and invariant learning (McGeorge and Burton, 1990). Typically, implicit learning tasks follow a conceptual design containing the following three parts: 1) subjects are exposed to some complex rule-structured stimulus domain in an incidental learning situation, 2) subjects' performance is measured on the same (online) or on a different task as an indicator of behavioural change and 3) the extent to which subjects' acquired knowledge is conscious is assessed. The experimental tasks utilized in this thesis are sequence learning and invariant learning tasks and these will be focused on here in providing the relevant background to the experimental chapters.

The serial reaction time (SRT) task

The serial reaction time (SRT) task was introduced by Nissen and Bullemer (1987). This task suggested that sequentially organized information can be learned without concurrent awareness of the sequence. Subjects were presented with four, horizontally arranged lights. They were asked to press a corresponding key as fast as possible when its associated light lit up. Unbeknownst to subjects, the order in which the lights were lit followed either a regular, repeating 10-item sequence or followed a random order. Nissen and Bullemer (1987) found that subjects' reaction times decreased with repeated exposure to the regular, but not to the random sequence. Additionally, subjects' reaction times showed a sudden increase when they switched from the regular to the random stimulus display. Associated with this increase in the indirect measure of reaction time was an inability to verbalize anything about the regularity of the sequence.

In a similar task, Lewicki et al (1987) presented subjecgts with stimuli that could appear in one of four quadrants on a computer screen. Their task was to indicate which of the four quadrants contained a target stimulus. In the first six trials subjects saw the target only, whereas by the seventh the target was embedded amongst 35 distracters. Unbeknownst to subjects the position of the seventh target was predicted by its position on the first, third, fourth and sixth trial. There were 24 such combinations that functioned as predictors. The measure of learning was subjects' reaction time to the seventh target. If they had acquired the predictive knowledge, their response time on the seventh trial should be significantly faster with practice, when compared to seventh trial in which the target appeared in a different screen quadrant than that predicted. Lewicki et al (1987) found this indeed to be the case. This was an indication that subjects had learned the underlying rule-structure of the stimulus display. Additionally, subjects were unable to report anything about the underlying rules.

The invariant learning task

McGeorge and Burton (1990) introduced the invariant learning task. This was an attempt to simplify the rule-structure of the stimulus display, while conserving relative complexity of learning instances. Subjects were presented with a series of four digit numbers, one at a time. Unbeknownst to subjects, each number contained one digit that remained the same throughout each trial (the invariant). Subjects were asked to perform some task on each number that forced them to process the stimuli, while keeping the true nature of the experiment hidden. This was followed by a surprise recognition test. Here subjects were asked to select one number out of a pair they thought they had seen in the previous phase. What subjects did not know was that one of these contained the invariant (the positive) and the other did not (the negative), while all of the test numbers were new to them. McGeorge and Burton (1990) found that subjects selected the positive over the negative above chance. They suggested this an indication that subjects had acquired the underlying rule of the presence of the invariant. Importantly, when asked what strategy they had used, subjects were unable to report anything that could have accounted for their preference for positives.

A question arose as to the representation of the acquired knowledge. It is possible that subjects abstract some rule across the training stimuli. Alternatively, subjects may be using some simple perceptual similarity between exemplars in order to make their selection. Bright and Burton (1994) tested the possibility of transfer from one stimulus context to a different one using analogue and digital clock faces. The underlying rule was of a more abstract nature than a particular invariant and consisted of a time range which remained consistent across training instances. Subjects selected positives over negatives above chance following transfer from analogue clocks at training to digital clocks at test. This suggested that the information acquired was more abstract than perceptual knowledge of the surface structure. This demonstration of invariant learning also provides an example of the use of different types of stimuli in this context. However, no attempts have been made to investigate whether invariant learning using auditory material behaves differently to visual stimuli in this context.

1.7 Summary and experimental considerations

For researchers such as Shanks and St. John (1994), conscious processes are the default position for information acquisition. However, Reber and Winter (1994) argue that this cannot explain the acquisition of complex information in everyday life, such as those introduced above. It has become clear that the study of implicit learning is complex and so far, no simple unifying solution as to its definition and operationalization has been found (e.g. Cleeremans et al, 1998). Thus, for the purpose of this thesis implicit learning is described at its most general: implicit learning, here, is said to occur when subjects show an increase in performance on some task, without associated increase in verbal knowledge about the basis for this performance change (Underwood and Bright, 1996).

It has been shown that measurement of implicit processes and the resulting knowledge from these may not be a straightforward operation in implicit learning tasks (e.g. Shanks and St. John, 1994; Neal and Hesketh, 1997). A particular problem is the dissociation between implicit and explicit knowledge acquired in an implicit learning task in order to determine implicit processes. It has to be noted that this thesis is not concerned with the assessment of the nature of any acquired knowledge in its implicit learning contexts specifically (i.e. whether it is implicit or explicit), but that implicit learning tasks from the visual domain were utilized in order to operationalize learning without awareness.

A general lack of research into implicit learning using auditory stimuli specifically, as well as a lack in providing consistent results utilizing such stimuli, currently does not provide a clear picture on whether implicit learning of sound stimuli behaves differently to such learning in the visual domain. This question is at the core of this thesis. Hence, stimuli throughout this thesis are exclusively auditory. Utilizing purely auditory material will allow empirical investigation of whether implicit learning of sound stimuli behaves any differently to implicit learning of visual stimuli. By extension, the research in this thesis directly addresses the question of whether subjects abstract the underlying rules or whether learning is tied to more peripheral processes, such as motor response learning. This will allow assessment of whether implicit learning can be described as a fundamental and general process or not. In order to do so, each experimental chapter introduces and addresses a specific issue in visual implicit learning research and investigates the generality of these findings by employing auditory material.

Chapter 2. Learning by listening in an SRT task

2.1 Introduction

It is commonly suggested that implicit knowledge acquisition can be characterized as unselective and automatic accumulation of associated information (e.g. Berry and Broadbent, 1988; Reber, 1989). There is growing evidence that these types of processes may provide the basis for sequence learning (e.g. Frensch and Miner, 1995; Mayr, 1996; Schmidtke and Heuer, 1997). The Serial Reaction Time (SRT) task, or sequence learning task, has become one of the main experimental paradigms employed to investigate implicit learning (e.g. Nissen and Bullemer, 1987; Willingham, Nissen and Bullemer, 1989). Sequence learning tasks typically use visual stimuli, for example an asterisk, which appears in different locations on a computer screen. Subjects' task is to respond to the asterisk according to the spatial location it appeared in by pressing a corresponding key. Each of the response keys corresponds to one of the spatial locations and subjects are asked to make their responses as fast and as accurately as possible. Reaction times are recorded for each response a subject makes. Unbeknownst to subjects, the visual stimuli follow a certain, unchanging spatial sequence, which is repeated over and over. At some point, again unbeknownst to subjects, this sequence is disrupted and a novel or random sequence is introduced to which subjects continue to respond. Learning of the repeating sequence is said to have taken place if subjects' reaction time to the repeating and novel/random sequences is significantly different, with faster responses to the repeating sequence. Thus, subjects are deemed to respond faster to the learning sequence they had been responding to repeatedly, because they are said to be able to anticipate the next item's position. This speeded response should not be observed for random or novel sequences. Overall, it is reaction time that provides an indirect measure of sequence learning in this context. The general view is that such sequence learning can demonstrate learning of complex information without concurrent awareness (Buchner and Frensch, 1997), thus putting this paradigm at the forefront of research into implicit learning.

2.2 Motor response vs conceptual learning

An ongoing debate has arisen as to what subjects learn in a sequence learning task. Subjects are commonly asked to make a motor response (e.g. keypresses) in an SRT task, leaving the possibility that they learn a sequence of responses. This is clearly different from acquiring the underlying rule structure of the sequence, which would make such learning conceptual in nature. This would be shown if subjects acquired the sequence without a motor response tied to the sequence exposure. Sequence learning by observation would provide an instance of such conceptual learning. The evidence to answer the question of what subjects actually learn is somewhat mixed. A study by Howard, Mutter and Howard (1992) is one of the most cited giving evidence of observational learning of a visually presented sequence. Subjects here were required to respond to an asterisk that appeared in one of four boxes, arranged in a row at the bottom of a computer screen. Their task was to press a corresponding key for the box the asterisk appeared in as fast and accurately as possible. Once responded, the asterisk disappeared and appeared in another box. Subjects were told that the asterisk appeared in a random fashion when, in fact, it was shown according to a 10-item or a 16-item systematic sequence. Howard et al (1992) found that those subjects who responded to the first 10% of trials only, and then went on simply to observe the remaining trials, showed as much learning as subjects who responded to all of the trials. In a further experiment Howard et al (1992) had subjects observe the asterisk throughout all trials, without any requirement of responding. Again, they showed a learning effect following observation: a slowdown in reaction time (RT) to a random test sequence when compared to RTs to the systematic repeating sequences. This evidence was supported by Seger (1997) who found observational learning of a 10-item sequence. A slightly different design employed by Mayr (1996) found that subjects could learn a spatial sequence of locations they were not responding to in a dual sequence paradigm. However, other studies by Kelly and Burton (2001) and Willingham (1999) failed to show such learning by observation. Differences in sequence complexity and criticisms regarding the potential level of explicit awareness may provide an explanation to these mixed results. Explicit learning may have allowed sequence learning to occur in those studies that showed learning by observation and this was shown a possibility (Kelly and Burton, 2001; Kelly, Burton, Riedel and Lynch, 2003; Willingham, 1999). In an attempt to replicate Howard et al's (1992) results Willingham (1999) found that only subjects who had been shown to have a high level of concurrent awareness showed any evidence of observational learning. This is supported by evidence from a study by Kelly et al (2003), which attempted to manipulate levels of awareness with a distracter task, as well as salience of the visual stimuli involved. Kelly et al (2003) concluded that effects of sequence learning by observation are mediated by explicit processes and such effects are eliminated under conditions that make it difficult to acquire explicit knowledge of a given sequence.

2.3 Auditory stimuli and the SRT task

The common use of visual stimuli in the SRT task can largely be explained by the ease of experimental manipulations this allows (Buchner and Frensch, 1997). There have been some studies that utilized auditory stimuli in order to expand the empirical scope of sequence learning research, but these have been somewhat unsystematic in their approach to studying auditory sequence learning in its own right (e.g. Buchner, Steffens, Erdfelder and Rothkegel, 1997; Buchner, Steffens and Rothkegel, 1998; Perruchet, Bigand and Benoit-Gonin, 1997; Schmidke and Heuer, 1997). The first study of these was that by Buchner et al (1997). Buchner et al's (1997) primary aim was to evaluate implicit and explicit processes in an SRT task using a variation of the process-dissociation procedure (Jacoby, 1991). As a lesser concern they used auditory stimuli in order to extend the generality of such learning to other experimental conditions (Buchner et al, 1997). Their auditory stimuli were four synthesized tones and subjects' task was to respond to each with a keypress. Prior to exposure to the learning sequences, subjects were trained in the specific key-to-tone mapping. As a cover story they were told that the study would test how well different tones could be discriminated in pitch. Subjects then followed the usual instructions for an SRT task and made their responses as fast and as accurately as possible. Buchner et al (1997) found clear learning effects with subjects showing a clear RT disadvantage to a random test sequence. Experiment 1's design is closely related to Buchner et al's (1997) design providing a starting point for a series of experiments in the current study. The aim was to explore auditory sequence learning in its own right in a systematic manner. Furthermore, this series of experiments investigated some questions that have arisen in the SRT literature in general, utilizing the auditory stimuli in novel ways.

2.4 Auditory sequence learning without a motor response

A particular limitation of previous sequence learning studies that employed auditory stimuli was that all used direct motor response mapping to the stimulus sequences (e.g. Buchner, Steffens, Erdfelder and Rothkegel, 1997; Perruchet, Bigand and Benoit-Gonin, 1997). These studies did not investigate the possibility of learning of auditory sequences without a corresponding motor response. Such learning would be shown if subjects were able to learn an auditory sequence by *listening* alone (equivalent to *observation* in a visual SRT task). The auditory domain may present a class of stimuli that makes it more likely for sequence learning without a motor response tied to the sequence exposure to occur. Auditory material is largely presented sequentially, leaving the listener to process it in a serial fashion. Additionally, auditory stimuli are mostly processed without spatial motor involvement. It is, therefore, reasonable to consider the possibility of auditory sequence learning without a corresponding motor response.

2.5 Sequence structure

All experiments in this study used the same 12-item sequences (one learning sequence and one test sequence). Equal numbers of base items are contained in these. This means that each item can only solely be predicted by its preceding two items. Furthermore, the sequences are constructed so that there are no reversals present (i.e. runs like ABA, where A and B correspond to different stimulus items). Reed and Johnson (1994) suggested the use of such a 12-item sequences, which comply with the Second Order Conditional (SOC), to exclude the occurrence of salient runs (e.g. ABA). They also suggested that these two sequences are matched for a) location frequency (i.e. how often each target location occurs within the sequence), b) transition frequency (i.e. how often each location transition can occur), c) reversal frequency (i.e. the number of times back-and-forth movements occur), d) rate of full coverage (i.e. the mean number of items, ensuring each location has been occupied at least once) and e) rate of complete transition usage (i.e. the mean number of items encountered, ensuring each possible transition occurs at least once). This was taken into account in the current series of experiments and the exact sequences suggested by Reed and Johnson (1994) were used.

The following set of five experiments examines a variation of a common SRT task. Broadly speaking these experiments aim to explore sequence learning effects in the auditory domain. Expanding sequence learning research in the auditory domain in a systematic manner can provide an extended insight into the mechanisms underlying sequence learning in general. Furthermore, the non-spatial nature of auditory stimuli is utilized to investigate whether auditory sequence learning can occur without a corresponding motor response. To do so, stimuli in this study are exclusively auditory and non-verbal in nature.

2.6 Experiment 1

Buchner et al (1997) found clear learning effects in an experiment that used simple auditory tones. In particular, their Experiment 2 showed that auditory stimuli, as the four synthesized tones employed in their study, could provide typical sequence learning effects. Thus, after responding to a repeating auditory 10-item sequence, subjects displayed a significant slowdown in their reaction times when a random sequence was introduced, and a subsequent speedup in their responses when the original repeating sequence was reintroduced. Perruchet, Bigand and Benoit-Gonin (Experiment 4, 1997), on the other hand, failed to show any learning of a 12-item sequence of three tones when subjects were required to respond to these with keypresses. No definite reasons for this lack of learning could be provided. Since Perruchet et al (1997) failed to show auditory sequence learning with corresponding keypress responses, Experiment 1 in the current study established a design that showed whether learning of an auditory sequence could occur with a corresponding motor response. In order to do so, Experiment 1 utilized a design directly based on a previous implicit sequence learning study from the visual domain (Kelly, Burton, Riedel and Lynch, 2003). This had proven reliable and had taken account of previous criticisms of other sequence learning experiments. Thus, it provided a valid basis for the current experiment.

Subjects in Experiment 1 responded to four synthesized tones with corresponding keys. These stimuli were played in a systematic repeating sequence over ten blocks of trials with disruption in block 9 when a novel test sequence was introduced. Learning of the systematic sequence was shown when subjects respond significantly faster to the repeating learning sequence when compared to the novel test sequence.

The aim of Experiment 1 was to create an SRT task that could replicate the results found by Buchner et al (1997) and show whether implicit sequence learning can be found in an auditory SRT task. This will then provide a fresh departure point to investigate sequence learning in the auditory domain systematically in subsequent experiments.

Method

Subjects

16 subjects were recruited from the student population of the University of Glasgow. All were naive to the experimental aims of the study and all received a small payment in return for their participation. All participants reported having normal hearing.

Materials

All stimuli were presented to subjects using an Apple Macintosh G3 computer. The 'SuperLab' experimental package was used to implement the experimental design. All auditory material was presented via headphones to subjects. Responses were made using a common computer keyboard.

All auditory stimuli were generated using Sound Edit 16 (version 2) computer synthesiser software. The sample rate and sample size was set at 44.100 kHz (CD quality) and 16 bits respectively. The four tones chosen for this experiment were 1 = 587.3Hz, 2 = 440.0Hz, 3 = 329.6Hz and 4 = 246.9Hz, where each tone mapped onto sequence positions A, B, C and D respectively (see Design and Procedure section below). These tones had been shown to be reliably distinguishable from one another in a pilot experiment. Each tone was 250ms long.

Design and Procedure

There were three stages in this experiment. In the key practice phase subjects learned the correct key-to-tone mapping. In the sequence learning phase subjects were exposed to a repeating sequence that was at one point disrupted. In the final phase, subjects were given a recognition task to test their implicit learning directly. This last part was unexpected by subjects. The experimenter remained present throughout the experiment.

Subjects were instructed that this experiment would test how well they are able to learn mapping of auditory tones to keys on a computer keyboard. To do so, they were told that they would hear four tones, one at a time. Subjects were informed that their task would be to respond to each tone by pressing a corresponding key on the keyboard as quickly and accurately as they could. In order for subjects to learn the correct key-to-tone mapping, they were given training (see key practice phase below).

The response keys were v, b, n and m on the keyboard mapping onto sequence positions A, B, C and D respectively (where A maps onto the lowest and D onto the highest tone respectively). Subjects were instructed to use the middle and index fingers of each hand for responding. Before starting the key practice phase, subjects were given an example of each key-to-tone mapping.

Subjects were informed that there would be 10 blocks of experimental trials in total with an opportunity to rest between each. They were told that the blocks would consist of the four tones played to them, one after the other, many times in random order.

Stage I: Key Practice

For the key practice phase, subjects were instructed that they would be asked to press each response key to hear its corresponding tone. They had to do this in the order of ABCD DCBA. Following this, subjects were told that they would have to respond to a tone by pressing its corresponding key. They were given onscreen feedback telling them whether they had responded correctly or not. Once subjects had made a response there was a 500ms delay before the next tone was presented. This included the 250ms presentation time of the onscreen feedback. The first set of practice trials consisted of the following sequence: ABCD ABCD DCBA DCBA ABCD DCBA ABCD DCBA. This was followed, without pause, by the remaining 22 out of all the 24 possible sequence orders of ABCD. Piloting had shown that some subjects still made a high number of errors in their key-to-tone mapping after these 22 key practice trials (between 1/2 and 1/3 of errors). For this reason, it was decided to lengthen the key practice phase by presenting the above-mentioned 22 four-tone sequence orders twice, without break and in random order. This reduced the number of errors subjects made overall in the key practice stage, with a progressive improvement throughout. In total subjects were presented with 52 four-tone key practice sequences. The key practice phase took approximately 10 minutes to complete.

Stage II: Sequence Learning

Following the key practice part of the experiment, subjects were told that the experiment would now begin. They were reminded that they would have to respond as fast and accurately as they could. Instead of onscreen feedback, subjects now saw '*****' on the screen once they had made a response. Once a response was made, there was a delay of 500ms (including 250ms display time of '*****'), followed by the next tone.

Two different sequences were used in this experiment, one that represented the learning sequence, which was used in blocks 1 to 8 and block 10, and one that represented the test sequence, which was used in block 9. Each sequence was 12 items long and was taken directly from Kelly, Burton, Riedel and Lynch (2003) to comply with the Second Order Conditional (SOC) detailed by Reed and Johnson (1994). Thus, the learning sequence appeared in the order BDACBADBCDCA and the test sequence in the order of BDBACDABCADC. Since participants were not made aware of the presence of any set sequence order, they were also not made aware of any changes in the repeating sequence from block 8 to 9 and back to the original learning sequence in block 10. The sequences were repeated without pause for eight cycles per block, giving subjects 96 tones to respond to in each block.

Stage III: Direct Measure of Awareness

Following the 10 blocks subjects were presented with a further task. They were informed that the four tones had not been played in random order, but that they had been presented according to a certain sequence that had been repeated over and over. Subjects were told that their knowledge of the sequence would be tested next. To do this they were required to respond to 6-item chunks taken from the 12-item sequences by pressing the corresponding key to the tone they heard as before. They were reminded that they had to respond as fast and accurately as they could. There were 24 chunks in total with half taken from the learning sequence (OLD chunks) and the other half from the test sequence in Block 9 (NEW chunks, see Appendix A.1 for full set of chunks). After responding to each chunk subjects were asked to rate whether the chunk they had just heard was new to them or whether it formed part of the sequence they had been responding to throughout the learning phase (see Appendix A2. for rating scale). There was no time limit for making the decision. Participants were fully debriefed after this final task.

Results

Sequence Learning

The measure of interest in this experiment was subjects' reaction times (RTs) to a given repeating auditory sequence. If subjects showed a significant increase in their RTs to a novel test sequence when compared to RTs to the original learning sequence, they were deemed to have acquired some knowledge about the learning sequence.

Figure 1 shows the means across subjects' median RTs for each block. There was a general decrease in RT with practice. Overall error rates across subjects for each block lay below 20 percent. This is relatively high compared to visual SRT tasks, but was not unexpected due to relatively higher difficulty of response mapping in an auditory SRT task. Previous auditory SRT tasks have varied in this respect, allowing rates from 10 up to 25% (e.g. Perruchet, Bigand, Benoit-Gonin, 1997; Buchner, Steffens, Erdfelder and Rothkegel, 1997; Buchner, Steffens and Rothkegel, 1998) reflecting the greater difficulty of auditory response mapping. Importantly, even with an error rate of up to 20% subjects were left with more than 70 responses to make in each block. This was clearly sufficient in providing adequate sequence exposure. This is also reflected in the RT results analysed below.

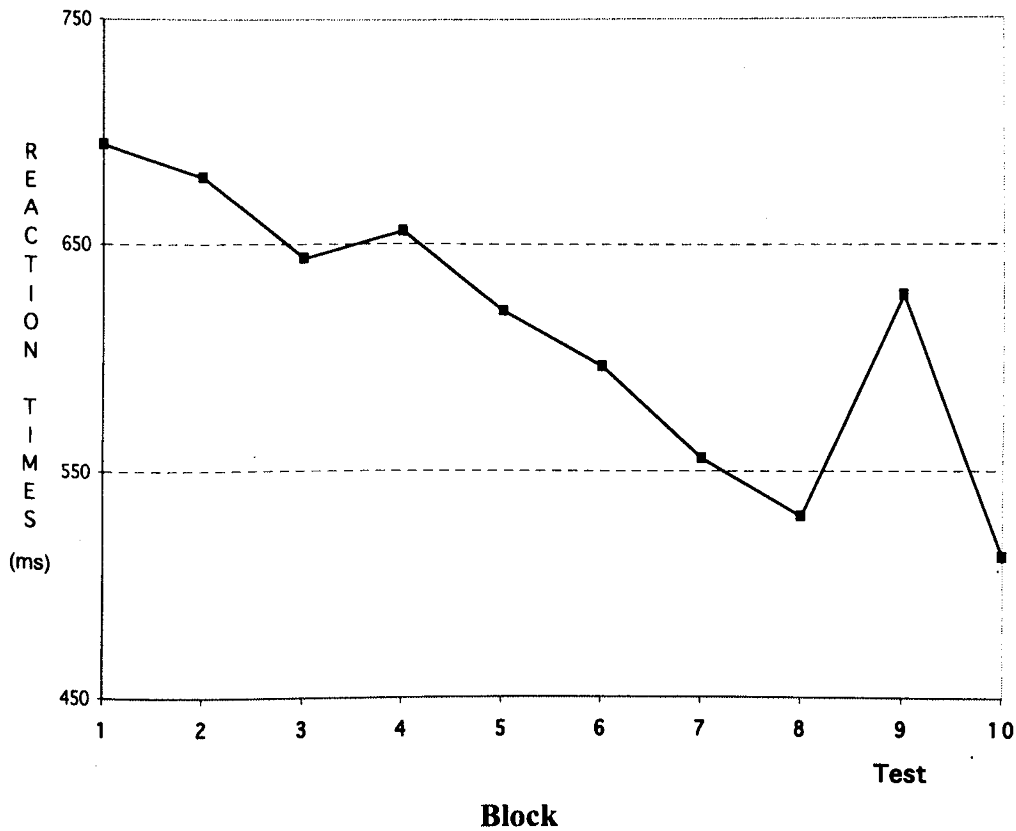

Figure 1. Mean RTs across subjects' median RTs per block. Subjects responded to the repeating tone sequence with keypress responses throughout. A novel test sequence was introduced in block 9. This was reversed to the original learning sequence in block 10.

The three blocks of interest, blocks 8, 9 and 10, show that subjects responded slower when a new sequence was introduced (block 9). This was followed subsequently with a decrease in RT to the original learning sequence (block 10). Table 1 shows the means across subjects' median RTs for these three blocks. There was no significant difference in error rates between the relevant blocks and they will not be referred to hereafter.

Table 1. Mean RTs (ms) and Standard Deviations (in brackets) for Last Three Blocks

Block 8 (learning sequence)	**Block 9 (test sequence)**	**Block 10 (learning sequence)**
530 (193)	629 (133)	512 (166)

For statistical analysis, the average RTs for blocks 8 and 10 were compared to block 9. The statistical comparison of blocks 8 and 9, as well as blocks 9 and 10, were not substantially different and, thus, were not included in the current analysis. In both comparisons there was a significant increase in RTs for the test sequence (block 9) when compared to the original learning sequence (blocks 8 and 10 respectively). In order to see whether there was a significant difference between RTs for the original learning sequence when compared to the novel test sequence overall, a paired sample t-test was performed between the average of blocks 8 + 10 (original) and block 9 (test). There was a significant difference in RTs for this comparison, with $t(15)=3.69$ ($p < 0.05$).

Subjects' speeded response to the learning sequence when compared to a novel test sequence clearly indicates that subjects are able to anticipate which item of the sequence will follow, thus decreasing their reaction times to the learning sequence. These results indicate that subjects have learned the repeating auditory sequence.

Direct Measure of Awareness

In table 2 the results for the recognition awareness test can be seen. There were a total of 24 chunks, made up equally of 12 OLD and 12 NEW items. The overall mean number of chunks identified correctly is presented. Also, correctly identified OLD and NEW items are presented individually.

Table 2. Means for chunks correctly identified by subjects in direct awareness test (recognition paradigm) with standard deviations in parentheses.

Total Chunks Correct	**OLD chunks (learning sequence)**	**NEW chunks (test sequence)**
10 (3.3) (out of 24)	6.3 (2.7) (out of 12)	3.7 (1.3) (out of 12)

A one-sample t-test compared the number of overall chunks correctly identified (both OLD and NEW) to chance performance of 12 (out of a total of 24). The overall performance in this recognition task was below chance, with $t_{(15)} = -2.38$; $p < 0.05$.

The data clearly shows that subjects were unable to identify chunks of items they had been exposed to previously in the sequence learning task. This indicates that subjects where not aware of the sequential nature of the repeating sequence throughout the learning task.

Discussion

The results from Experiment 1 show a clear sequence learning effect in a design that used auditory stimuli to which subjects responded with keypress responses. Additionally, the direct test of awareness indicated that subjects were unaware of the systematic sequence structure of the auditory sequence, supporting the conclusion that the learning effect found here was based on implicit knowledge. This is clearly in line with the original findings by Buchner et al (1997) and establishes that auditory sequence learning can occur in the current context. However, the question of *what* subjects have learned arises. By using corresponding keypress responses to the sequence elements, subjects may have learned a sequence of motor responses. This is clearly distinct from learning of the underlying rule structure of that sequence, which would make such learning conceptual in nature. It is this question of whether subjects learn a systematic repeating sequence at the motor response level or conceptually that is of ongoing concern in sequence learning research (e.g. Kelly, Burton, Riedel and Lynch, 2003; Willingham, 1999). Those studies that have investigated this question directly used visual stimuli. It is useful to expand sequence learning research to other experimental contexts and explore whether learning of auditory sequences can occur without a corresponding motor response tied to the sequence exposure. Experiment 2 explored this issue further.

2.7 Experiment 2

Experiment 2 was an adapted version of Experiment 1 and used the same basic design of 10 blocks of trials. The aim of Experiment 2 was to explore whether subjects could learn a given systematic auditory sequence by listening alone. As pointed out in the discussion to Experiment 1, such learning would be indicative of conceptual learning of the underlying rule structure, rather than learning of a sequence of responses. In addition, subjects in the current experiment were required to engage in a distracter task. Dual task conditions have been shown to lessen any explicit learning effects found in single task learning, making it unlikely that explicit awareness develops throughout this task (e.g. Goschke, 1998; Keele, Jennings, Jones, Caulton and Cohen, 1995; Seger, 1997). At the same time, evidence shows that sequences can be learned under conditions of attentional load (Hsiao and Reber, 1998 for a review). The secondary task in dual task studies commonly takes the form of tone counting. This is clearly not a possibility in the current task, since the auditory nature of the stimuli would prevent subjects from performing such a secondary task. An alternative task was found, which required subjects to do simple Maths calculations on a sheet of paper. This also provided a cover story, keeping the true nature of this experiment hidden: subjects in this instance were told that this experiment would test how well they could do simple Maths equations while being exposed to some random background noise.

Experiment 2 tested two groups of subjects in a design that saw them listening to a repeating systematic auditory sequence for the first seven blocks of trials.

Following this, subjects responded with a keypress response in the last three blocks as had subjects in Experiment 1. This allowed recording of RTs providing a measure to test for any learning effects.

Acquiring the correct key-to-tone mapping has been viewed as relatively more difficult than mapping to visual stimuli in equivalent visual sequence learning tasks (e.g. Buchner et al, 1997; Buchner et al, 1998; Perruchet et al, 1997). This poses a potential difficulty in the current design: it was inappropriate to provide the key-to-tone mapping training prior to the listening phase in order to avoid covert responding. Therefore, this training phase was placed just prior to the last three response blocks, following listening exposure in the first seven. Additionally, it was important to keep the interruption from listening exposure to responding as brief as possible. For this reason a further adjustment was made to the previous experiment and the key practice phase shortened to an absolute minimum. A pilot experiment was run and the shortened key practice phase was not found to affect RTs or error rates. This established the briefer key practice phase as sufficient in the current context.

Subjects started using the keypress response in the last three blocks of trials. This left the possibility that any reaction time differences between these blocks could be due to a continued practice effects. In order to take account of this possibility, there were two conditions in this experiment. In the first condition the same systematic sequence was maintained throughout the last three blocks of trials (no-change). In the second condition, a novel test sequence was introduced in block 9 in the same way as in Experiment 1 (with-change). Thus, if subjects in the with-change condition learned the systematic auditory sequence they had been listening to, they should show a significant RT decrease to the original learning sequence in block 10 when compared to RTs to the novel test sequence in block 9. However, if subjects in the no-change condition also showed a significant decrease in their RTs from block 9 to block 10, this decrease would represent a continued practice effect of the keypress responses and not a learning effect. This would clearly put into doubt whether any reaction time differences in the with-change condition could be due to learning effects. In this way, the no-change condition provided a baseline to the with-change condition. Overall, testing subjects in these two conditions allowed assessment of any sequence learning effects through listening exposure represented in reaction time decreases to the systematic repeating sequence.

Method

Subjects

32 subjects were recruited from the student population of the University of Glasgow. All were naive to the experimental aims of the study and all received a small payment in return for their participation. All participants reported having normal hearing.

Materials

Materials were generated and presented in the same way as in Experiment 1 and all were the same as in the previous experiment except where stated.

Design and Procedure

The 12-item sequences were repeated without pause for eight cycles per block, making a total of 96 trials per block. This was the same as in Experiment 1. There were three parts to this experiment: In Stage Ia subjects listened to the auditory sequence while engaged in the secondary task. In Stage II subjects were given the appropriate key-to-tone practice. In Stage Ib subjects responded to the auditory stimuli with keypress responses.

There were two conditions in this experiment: a *no-change* condition and a *with-change* condition. The sequence in the *no-change* condition consisted of one tone sequence only and was the same as the learning sequence in Experiment 1 (mapping spatially onto BDACBADBCDCA). In this condition, no novel test sequence was introduced in block 9, but the sequence remained constant throughout all ten blocks of trials. The same sequence was employed in the *with-change* condition. However, here the novel test sequence from Experiment 1 was introduced in block 9 (mapping spatially onto BDBACDABCADC). Subjects were not made aware of the sequential nature of the auditory stimuli. Subjects were randomly assigned to each condition in equal numbers.

This was a mixed design, *with-change* presenting a between subjects factor and *block* a within subjects factor.

The experimenter remained present throughout the experiment. Instructions and the procedure were the same for subjects in both conditions.

Stage I.a: Sequence Learning

Subjects were informed that the purpose of the experiment was to explore whether their ability to do simple Maths calculations was affected by listening to auditory material via headphones. Thus, they were told they had to calculate a list of Maths equations as fast and accurately as they could (see Appendix B for full list) while listening to some random sequences of tones. They were advised that there would be 10 blocks of auditory material in total and that they could rest between each.

Subjects were presented with the first seven blocks of trials. These will be referred to as 'listening' blocks, since subjects simply listened to the tone sequences via headphones while doing the Maths task. Since subjects did not respond to the tones as in Experiment 1, ISIs for the tones had to be set. It became clear that with tones of 250ms length, an ISI of 250ms between each tone was sufficient for presentation.

Following completion of the first seven blocks, subjects were asked to respond to the auditory stimuli by pressing corresponding keys for each of the four tones. To do so, they went through a key practice phase.

Stage II: Key Practice

The instructions and set-up for this phase remained the same as those in Experiment 1's key practice phase, except for the following changes: In order to keep the interruption between the listening phase and the response phase to a minimum subjects were exposed to the following practice trials only: ABCD ABCD DCBA DCBA ABCD DCBA ABCD DCBA. As a further assurance of maximum key-to-tone mapping learning, the four tones used in Experiment 1 were adjusted so that they were acoustically further apart, making them more distinguishable. The tones used in this experiment were 1 = 123.5Hz, 2 = 196.0Hz, 3 = 440.7Hz and 4 = 523.3Hz (mapping onto the sequence positions ABCD respectively, where A was the lowest and D the highest tone). In this instance, the key practice phase took approximately one minute.

Stage I.b: Sequence Learning - continued

Following the key practice phase subjects resumed with the remaining three blocks (blocks 8, 9 and 10). Subjects were informed that they now had to respond to each tone pressing its corresponding key as fast and as accurately as possible.

The current and subsequent experiments did not focus specifically on the question of whether any sequence learning effects were implicit or explicit in nature. For reasons of economy the direct test of awareness from Experiment 1 was, therefore, no longer employed.

Results

The measure of interest in this experiment was subjects' RT. If subjects were able to acquire the auditory sequence by listening, a significant interaction between sequence change and block would be expected. In this case, subjects in the with-change condition would show a significant decrease in RT from block 9 (novel test sequence) to block 10 (original sequence), whereas subjects in the no-change condition would not show any RT difference between blocks 9 (original sequence) and 10 (original sequence), or such a decrease would be significantly less than that found in the with-change condition.

The three blocks of interest were those in which subjects made keypress responses (blocks 8, 9 and 10). Figure 2 shows the means across subjects' median RTs for each response block. There was a general decrease in RT throughout these blocks in both conditions. Error rates across subjects for the response blocks lay, again, just below 20 percent. There was no statistical difference in these between the response blocks and they will not be referred to hereafter.

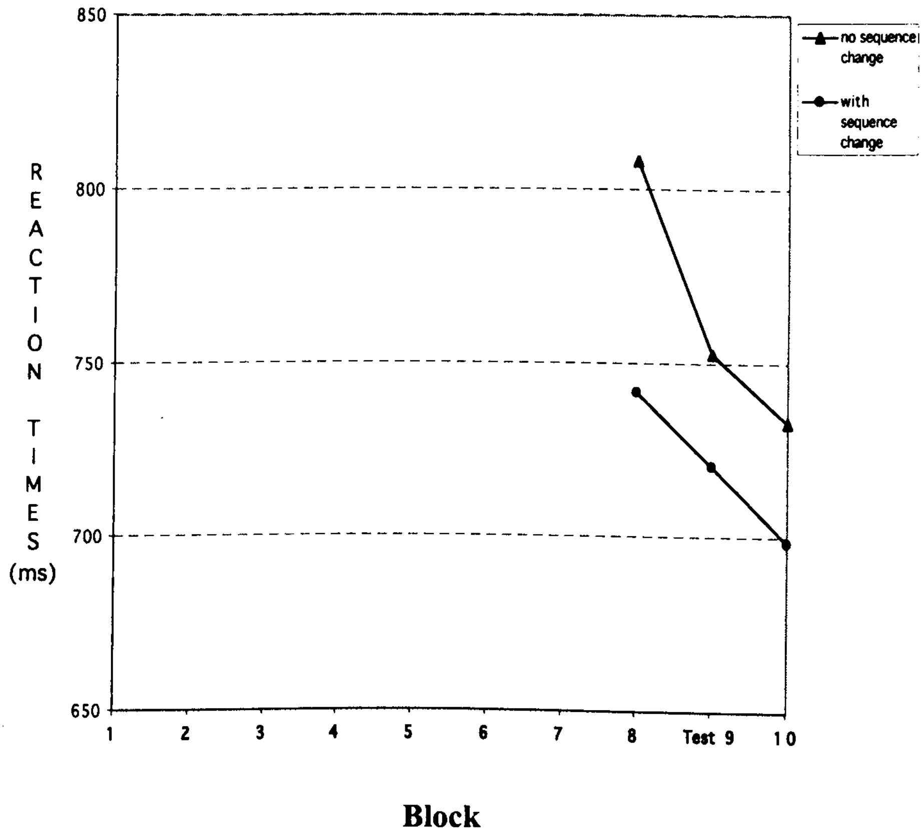

Figure 2. Mean RTs across subjects' median RTs per response block (8, 9 and 10). Subjects listened to the repeating auditory sequence throughout blocks 1 to 7. Subjects switched from listening to a keypress response in the final three blocks. In the no-change condition the same learning sequence was maintained throughout all ten blocks. In the with-change condition a novel test sequence was introduced in block 9 and block 10 reversed back to the original learning sequence.

Table 3 shows the means across subjects' median RTs for the relevant blocks. The blocks of interest were block 9 and 10. Block 8 was deemed unreliable in providing RTs reflecting anything other than practice effects, since this was the first block in which subjects made keypress responses.

Table 3. Mean RTs (ms) and Standard Deviations (in brackets) for Last Two Blocks

Condition (sequence change)	Block 9	Block 10
no-change	753 (140)	733 (148)
with-change	721 (178)	699 (171)

A mixed design two-way ANOVA on sequence change (no-change or with-change) X block (block 9 or block 10) revealed no main effect of sequence change ($F(1, 30) = 0.38$, $p > 0.05$) and no main effect of block ($F(1,30) = 1.91$, $p > 0.05$). Importantly, there was no interaction ($F(1, 30) = 0.005$, $p > 0.05$).

Subjects did not show a significant decrease in RTs in block 10 when compared to block 9. These results reveal neither practice nor learning effects, as would have been encapsulated in an RT reduction from block 9 to block 10. Importantly, this is also reflected in the failure to show any interaction between sequence change and block. Thus, subjects in the with change condition did not show a smaller decrease in their RTs between blocks 9 (test) and 10 (original) when compared to subjects in the no change condition in these blocks. These results indicate that subjects in the two conditions did not behave differently. Overall, these findings suggest that subjects failed to learn the repeating auditory sequence.

Discussion

The aim of this experiment was to assess any sequence learning effects represented in reaction time decreases to the systematic repeating sequence when subjects had merely listened to it. The results failed to show any evidence for such learning effects. They also did not show any practice effect for the keypress responses in the last two blocks of trials. This indicates that subjects failed to learn the underlying rule structure of these stimuli when they merely listened to the systematic sequence, ruling out conceptual learning of the auditory event sequence. It has to be noted that individual variability proved high in this experiment and this may have been the reason for the non-significant effect. It is a possibility that subjects varied considerably in their keypress responses due to difficulties in applying the correct key-to-tone mapping. However, there may be other reasons why the current design failed to capture any sequence learning effects.

Kelly and Burton (2001) pointed to the possibility that subjects observing a visual event sequence may fail to transfer the acquired sequence knowledge to a keypress response, because of the possibility of modality specific learning. They suggested that previous findings in this area are inconclusive. Cohen, Ivry and Keele (1990), as well as Keele et al (1995) found RT savings when subjects transferred their responses between arm muscles and finger muscle effectors. Furthermore, Keele et al (1995) also demonstrated some RT savings when subjects transferred their responses from manual keypresses to verbal responses.

These data suggest that the learned contingencies in an SRT task can transfer between modalities. However, results by Ziessler (1994) run opposite to these findings. His results indicated that subjects making differential responses to targets showed learning. Those who responded to the majority of targets, however, did not. These findings point to a need for specific motor responses for sequence learning to occur. These contradictory results are given further fuel by Nattkemper and Prinz (1997) who identified some problems with those studies that showed transfer of learning between different modalities. In the context of the current experiment this leaves the possibility that subjects listening to the systematic sequence may not have transferred the learned to a manual keypress response. Employing auditory stimuli in the current sequence learning context is novel and there is a possibility that the auditory stimuli used may require a different motor response, which would be more appropriate for indirectly capturing any learning effects through a reaction time measure. A voice response is more closely related to the auditory domain than keypresses, since they are directly linked to the auditory domain. This makes the use of a voice response potentially more suitable for securing any auditory sequence learning effects. A voice response will allow capture of reaction times, thus providing an appropriate measure of any auditory sequence learning effects. Experiment 3 explored further whether subjects could learn an auditory sequence by listening utilizing a different kind of motor response than the current experiment.

2.8 Experiment 3

Experiment 3 was a replication of the design in Experiment 2 except for substitution of the keypress response used then with voice responses in the current experiment. Voice responses have been successfully employed to capture RTs in previous sequence learning studies (e.g. Keele, Jennings, Jones, Caulton and Cohen, 1995; Kelly and Burton, 2001). Those studies generally used verbal voice responses (e.g. "left", "right" or "centre"). Such verbal voice responses are clearly inappropriate in an auditory SRT task like the current experiments, since subjects would have to learn mapping of a verbal response onto each individual stimulus. However, as long as the voice response is meaningful to subjects, the content should not affect RTs obtained. Asking subjects to replicate the pitch of a tone they just heard is such a meaningful voice response. Additionally, it was necessary to keep the voice responses constant between subjects. For this reason subjects were asked to replicate the pitch of a given tone by saying/singing the syllable "Da". Moreover, it decreased any disruption between the listening and response stages further than the shortened key practice phase in the previous experiment.

Experiment 3 tested two groups of subjects in a design that saw them listening to a repeating systematic auditory sequence for the first seven blocks of trials. Following this, subjects responded with voice responses in the last three blocks of trials. Again, there were two conditions: a no-change and a with-change condition. If subjects in the with-change condition learned the systematic auditory sequence they should show a significant RT decrease to the original learning sequence in block 10 when compared to RTs to the novel test sequence in block 9. However, if subjects in the no-change condition also showed a significant decrease in their RTs from block 9 to block 10, this decrease could only represent a continued practice effect of the voice response and not a learning effect. This would clearly put into doubt whether any RT differences in the with-change condition would have been due to learning effects. Thus, the no-change condition provided a baseline to the with-change condition. Overall, testing subjects in these two conditions allowed assessment of any sequence learning effects through listening exposure, which would be represented in RT decreases to the repeating systematic sequence.

Method

Subjects

36 subjects were recruited from the student population of the University of Glasgow. All were naive to the experimental aims of the study and all received a small payment in return for their participation. All participants reported having normal hearing.

Materials

Materials were generated and presented in the same way as in Experiment 2 and all were the same as in the previous experiment except where stated.

Since subjects were now required to make a voice response, the actual tones employed were altered. There was no requirement for the tones to have maximum frequency separation in this context, as long as they remained clearly distinct. Learning the key-to-tone mapping response in the previous experiments had been considered easier with tones that were further apart in frequency. The main requirement of the tones in the current experiment, however, was that both male and female subjects could easily replicate them with an accurate voice response while perceiving them as clearly distinguishable. Therefore, the four tones used here were closer in frequency than in the previous experiments, but also remained clearly distinct from on another. The tones used in this experiment, then, were A = 130.8Hz, B = 146.8Hz, C = 196.0Hz and D = 246.9Hz (where A was the lowest and D the highest tone). All voice responses were recorded using a standard Apple Macintosh microphone positioned on top of the computer screen. Tones were 250ms in length and were separated by a 250ms ISI.

Design and Procedure

The 12-item sequences were repeated without pause for eight cycles per block, making a total of 96 trials per block. This was the same as in Experiment 2. There were three parts to this experiment: In Stage Ia subjects listened to the auditory sequence while engaged in the secondary task. In Stage II subjects were given the appropriate key-to-tone practice. In Stage Ib subjects responded to the auditory stimuli with voice responses.

There were two conditions in this experiment: a *no change* condition and a *with change* condition. The sequence in the *no change* condition consisted of one tone sequence only and was the same as the learning sequence in Experiment 1 (mapping spatially onto BDACBADBCDCA). In this condition, no novel test sequence was introduced in block 9, but the sequence remained constant throughout all ten blocks of trials. The same sequence was employed in the *with change* condition. However, here a novel test sequence was introduced in block 9 (mapping spatially onto BDBACDABCADC, as in the previous experiments). Subjects were not made aware of the sequential nature of the auditory stimuli. Subjects were randomly assigned to each condition, totalling 20 in the no-change and 16 in the with-change condition.

This was a mixed design, with-change presenting a between-subjects factor and block a within-subjects factor.

The experimenter remained present throughout Stage Ia and Stage II of the experiment only.

Stage I.a: Sequence Learning

This was an exact replication of the sequence learning stage in Experiments 2. Subjects were given instructions as before. This stage comprised blocks 1 to 7.

Stage II: Voice Response Instructions

Following block 7, subjects were instructed that there would be a change in the experiment. They were informed that they would now have to make a response to each tone and that there would be no further Maths calculations. They were told that they would hear, as before, four tones played to them, one at a time, in random order. Their task was to respond to each tone by singing it back into the microphone positioned above the computer screen. They were instructed that this 'singing' had to be done in a specific way: subjects had to sing the tone back using the syllable 'da' in the corresponding pitch of a given tone. Using the syllable 'da' kept responses uniform between all subjects and allowed concise recording of subjects' RTs. Subjects were positioned approximately 30cm away from the microphone and told that they would have to direct their response directly at the microphone. Subjects were asked to make their voice responses as fast and accurately as they could. The RTs for each voice response were recorded. Accuracy as to the pitch was not recorded. Subjects were given an example of each tone and had to make an accurate response as instructed to each in the presence of the experimenter. The experimenter was present until stage I.b below.

Stage I.b: Sequence Learning - continued

Following the voice practice phase subjects resumed the remaining three blocks (blocks 8, 9 and 10). Subjects were informed that they now had to respond to each tone making the appropriate voice response as fast and as accurately as possible.

Results

The measure of interest in this experiment was subjects' RT. If subjects were able to acquire the auditory sequence by listening to, a significant interaction between sequence change and block would be expected. In this case, subjects in the with-change condition should show a significant decrease in RT from block 9 (test sequence) to block 10 (original sequence), while subjects in the no-change condition would not show any significant decrease or a significantly lesser decrease between these blocks than subjects in the with-change condition.

The three blocks of interest were those subjects made voice responses to (blocks 8, 9 and 10). Figure 3 shows the means across subjects' median RTs for each response block. There was a general decrease in RT throughout these blocks in both conditions.

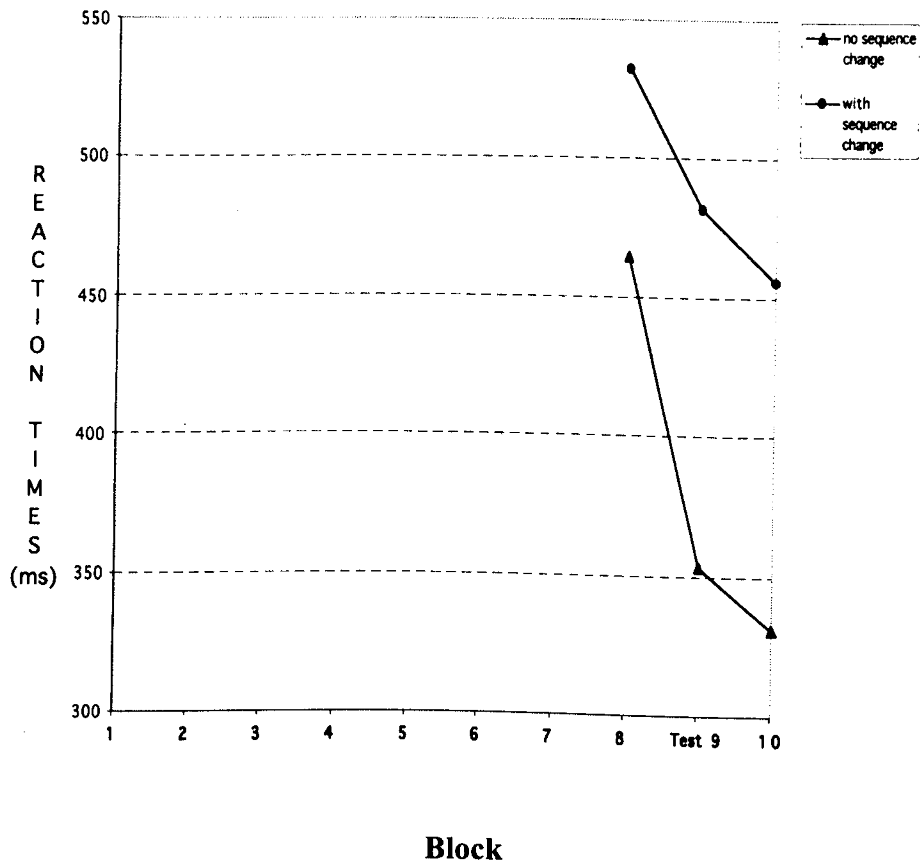

Figure 3. Mean RTs across subjects' median RTs per response block (8, 9 and 10). Subjects listened to the repeating auditory sequence throughout blocks 1 to 7. Subjects switched from listening to a voice response in the final three blocks. In the no-change condition the same learning sequence was maintained throughout all ten blocks. In the with-change condition a novel test sequence was introduced in block 9 and block 10 reversed back to the original learning sequence.

Table 4 shows the means across subjects' median RTs for the relevant blocks. The blocks of interest were block 8, 9 and 10. In this case, data for block 8 was analysed alongside the other blocks, since there were noticeably larger RTs in the with-change condition when compared to the no-change condition.

Table 4. Mean RTs (ms) and Standard Deviations (in brackets) for Last Two Blocks

Condition (sequence change)	Block 8	Block 9	Block 10
no-change	465 (221)	354 (198)	332 (173)
with-change	532 (202)	482 (191)	456 (193)

A mixed design 2x3 ANOVA on sequence change (no-change or with-change) X block (block 8 or block 9 or block 10) revealed no main effect of sequence change ($F(1, 34) = 2.93$, $p > 0.05$) but a main effect of block ($F(2,34) = 16.7$, $p < 0.05$). Importantly, there was no interaction ($F(2, 34) = 1.61$, $p > 0.05$).

Lack of any interaction between sequence change and block suggests that subjects performed equally when these two conditions were compared. Thus, subjects in the with-change condition did not show a decrease in RTs from block 9 (test) to block 10 (original) that was significantly different from subjects' RT decrease between these blocks in the no-change condition. Therefore, the main effect of block cannot be due to learning effects but represents practice effects. Additionally, the seemingly large overall RT differences between subjects in the two conditions did not prove to be significantly different from one another suggesting that overall performance did not differ between these two groups. Again, this was supported by lack of finding any interaction effect.

Overall, these results indicate that subjects failed to learn the repeating auditory sequence.

Discussion

The aim of this experiment was to assess any sequence learning effects captured in voice responses to a systematic auditory sequence subjects listened to. Learning of an auditory systematic sequence, here, would have been evidence for learning independently of motor responses. This would have given support to previous results from visual SRT tasks that showed learning by observation (e.g. Howard, Mutter and Howard, 1992; Seger, 1997). The results failed to show such learning effects and add further to evidence that failed to show learning by observation (e.g. Kelly and Burton, 2001; Willingham, 1999) and extend these results to the auditory domain. It has to be noted that individual variability proved high in this experiment and this may have been the reason for the non-significant result. It is a possibility that subjects varied considerably in making the voice responses and may have been a difference in skills that was responsible for the high individual variability.

However, there is a possibility that subjects failed to learn the auditory sequence because of the distracter task subjects engaged in, both in Experiment 2 and 3. The nature of the auditory stimuli made it necessary to find a different secondary task than the common tone counting. Attending to and calculating simple Maths equations was deemed a suitable task in this context. However, this did not take account of the focus of subjects' attention in these experiments. Kelly and Burton (2001) pointed out that observation tasks (equivalent to listening tasks in the current study) might be different from previous studies that have shown learning of sequences under conditions of attentional load. They postulated that subjects in observational studies might not have been attending the stimuli at all. This would be in contrast to conditions of attentional load in which subjects commonly have some attention focused on the event sequence. The distracter task in Experiments 2 and 3 clearly did not allow subjects to aim any attention at the auditory event sequence and this might have prevented any learning by listening to occur. By extension, it is possible that subjects did not listen to the auditory sequence at all in the current context. This possibility should clearly be taken into account in any future experiments investigating sequence learning by listening.

There also remains a question as to the suitability of the voice response for capturing relevant RT differences. This is particularly emphasized by the large standard deviations and the noticeable difference in RTs between the two experimental conditions in the current experiment (although these were not found to be significantly different from one another). Thus, it is still a question whether subjects failed to learn the auditory sequence by listening or whether the motor responses employed failed to capture any real differences in RTs. As already discussed in the introduction to the current experiment Cohen et al (1990) and Keele et al (1995) found transfer effects in an SRT task. In Keele et al's (1995) study, subjects had to respond to a visual sequence by either making spatial keypress responses or by making verbal responses to visual stimuli during the learning phase. During the test phase subjects switched their responses from keypress to verbal and vice versa. Learning was shown for both transfer conditions, although the effect was much smaller in the case of switching from keypresses to verbal responses. Experiment 1 in the current study has already established that subjects are able to learn a systematic auditory sequence when measured through a set of spatial motor responses (keypresses). An experiment that utilizes spatial keypress responses for the auditory sequence exposure and then switches to a non-spatial voice response should show learning at transfer according to the results found by Keele et al (1995). This would also verify whether the use of voice responses is sufficient in picking up any sequence learning effects in the current context. Experiment 4 investigated this further.

2.9 Experiment 4

Experiment 4 aimed to explore transfer effects in an auditory SRT task. In line with Keele et al's (1995) results, changing from a keypress to a voice response in an auditory sequence learning task should provide evidence of transfer of such learning. Should such transfer be shown, it would add evidence for sequence learning as independent from a spatial component of sequence exposure and would extend Keele et al's (1995) findings to the auditory domain. This would provide evidence that auditory sequence learning does not require motor response mapping, but can occur by learning of event contingencies, which would make such learning conceptual in nature.

Method

Subjects

20 subjects were recruited from the student population of the University of Glasgow. All were naive to the experimental aims of the study and all received a small payment in return for their participation. All participants reported having normal hearing.

Materials

The first part of this experiment used the same materials as Experiment 1 except where stated. This included a key practice phase (Stage I) and sequence learning phase (Stage II). Stage III used the same materials as the voice response instruction stage in Experiment 3. Although subjects used a keypress response throughout the first part of this experiment, a voice response in the final three blocks made it appropriate to use the tones from the previous voice response experiment.

Design and Procedure

The spatial sequences used for the learning sequence and the test sequence were the same as Experiment 1 (mapping spatially onto BDACBADBCDCA and BDBACDABCADC respectively). Subjects were not made aware of the sequential nature of the auditory stimuli. The sequences were repeated without pause for eight cycles per block, making a total of 96 trials per block. This was the same as in previous experiments.

Stage I: Key Practice

This was the same as Experiment 1.

Stage II.a: Sequence Learning

This was an exact replication of the sequence learning stage in Experiment 1. Subjects were given instructions as before. This stage comprised blocks 1 to 7. The experimenter remained present throughout this part of the experiment.

Stage III: Voice Response Instructions

This was the same as Experiment 3's voice response instructions.

Stage II.b: Sequence Learning - continued

Following the voice response instructions subjects resumed with the remaining three blocks (blocks 8, 9 and 10) using the voice responses as instructed.

Results

As in the previous experiment, the measure of interest in this experiment was subjects' RTs to a given repeating auditory sequence. If subjects showed a significant increase in their RTs to a novel test sequence when compared to RTs to the repeating learning sequence, they were deemed to have acquired some knowledge about the repeating sequence.

Figure 4 shows the means across subjects' median RTs for each response block. There was a general decrease in RT throughout the first seven blocks in which subjects used a keypress response. This general decrease continues in the last three blocks in which subjects used a voice response. Overall error rates across subjects for each keypress response block lay below 20 percent. The relevant blocks for measuring any sequence learning were the final three blocks to which subjects made voice responses.

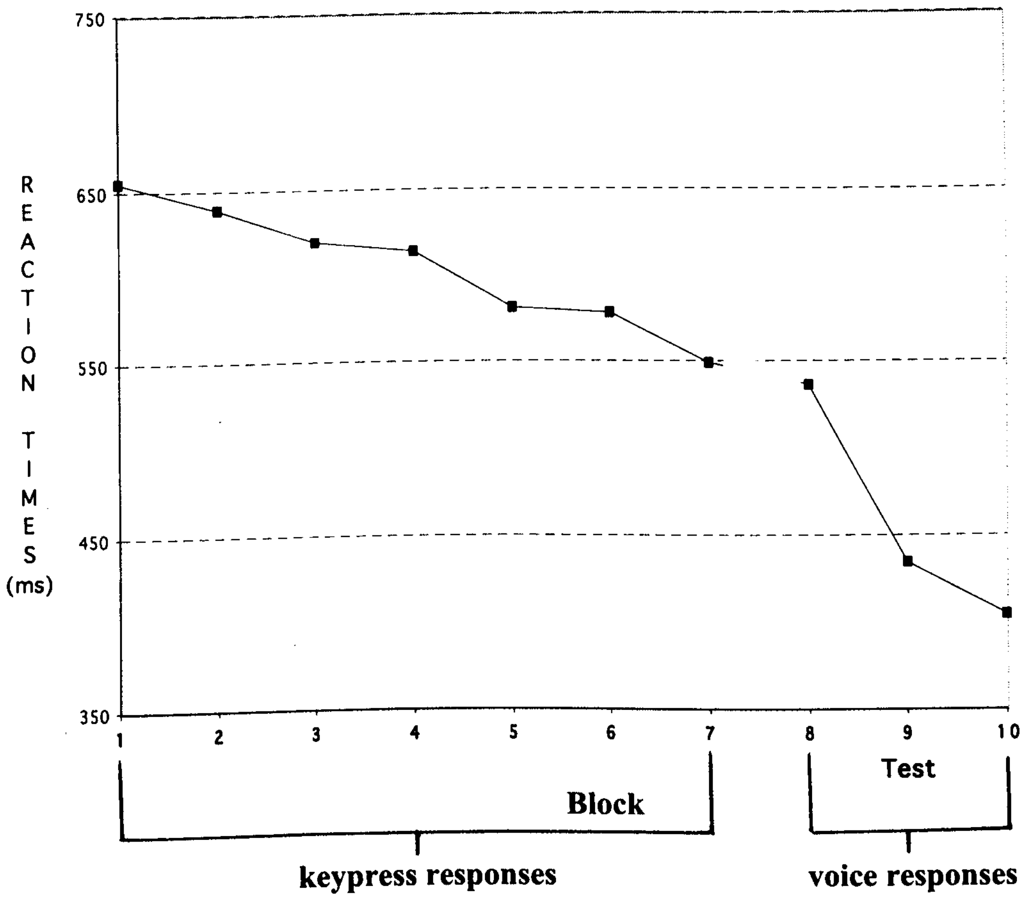

Figure 4. Mean RTs across subjects' median RTs per response block. Subjects responded to the repeating auditory sequence throughout blocks 1 to 7 with keypress responses. Subjects switched to a voice response for the final three blocks. A novel test sequence was introduced in block 9. Block 10 reversed back to the original repeating sequence.

Table 5 shows the means across subjects' median RTs for the relevant blocks. The blocks of interest were those to which subjects made a voice response (8, 9 and 10). Block 8 was deemed unreliable in providing RTs reflecting anything other than continued response practice effects, since this was the first block in which subjects used a voice response.

Table 5. Mean RTs (ms) and Standard Deviations (in brackets) for Last Two Blocks

Block 9 (test sequence)	**Block 10 (learning sequence)**
434 (149)	404 (166)

In order to see whether there was a significant difference between RTs for the original sequence when compared to the test sequence, a paired sample t-test was performed between block 9 (test sequence) and block 10 (original sequence). No significant difference in RTs for this comparison was found (t-value of $t(19)=1.586$, $p > 0.05$).

These results do not show transfer of sequence learning from spatial keypress responses to non-spatial voice responses.

Discussion

The results of this experiment do not provide any indication of sequence learning transfer from a spatial motor response to a non-spatial voice response. This is in disagreement with findings by Cohen et al (1990) and Keele et al (1995), which found such transfer in visual tasks. Experiment 1 had shown that subjects could acquire an auditory sequence when they responded to it with keypress responses. This implies that subjects in the current experiment acquired sequence knowledge but failed to transfer it to the non-spatial voice response used here.

It is, therefore, a possibility that the voice response utilized here does not carry the sensitivity to procure any reaction time decreases to the original learning sequence. As with the previous experiment, it has to be noted that individual variability proved high in this experiment and this may have been the reason for the non-significant result.

A further possibility is that lack of explicit processes in the sequence acquisition may be responsible for failure to show any sequence learning. Willingham, Wells and Farrell (2000), as well as Kelly and Burton (2001), suggested that transfer of learning found in Keele et al's (1995) study could have been due to explicit processes. Indeed, when explicit effects were removed, no sequence learning by observation could be established. Subjects in Experiment 1 underwent exactly the same procedure as subjects in the current experiment until block 8. Experiment 1 established that subjects acquired the auditory sequence without any concurrent awareness. Therefore, and in line with Kelly and Burton (2001) and Willingham et al's (2000) findings, the results found here point to the possibility that subjects failed to show any transfer of learning, because they lacked the necessary explicit knowledge. However, this was not tested for in the current experiment.

In order to establish whether the voice response may have lacked sensitivity in the current and previous experiments, it was useful to carry out a final experiment to establish whether it was unsuitable for capturing any RT differences.

2.10 Experiment 5

Experiment 5 aimed to establish whether it was failure of the voice response to procure any sequence learning effects in Experiments 3 and 4. This employed a voice response throughout all 10 blocks of trials. Ziessler (1994) found some learning when subjects used a single motor response to a visual sequence, although this was weaker than when each stimulus called for a distinct motor response. Using a voice response that requires subjects to sing back each tone at its correct pitch should provide distinct responses to each stimulus. However, even if these voice responses fail to provide a distinct response to each stimulus, making them similar to the single motor response employed by Ziessler (1994), sequence learning, if weaker, should still be shown. When subjects respond to a systematic repeating sequence, which is at some point disrupted, a clear reaction time advantage should be found to the repeating sequence. Experiment 5 investigated whether subjects show such a reaction time advantage to a systematic auditory sequence they responded to with the voice response throughout. If such reaction time advantage was shown, it would indicate that the voice response is adequate in obtaining sequence learning effects in this context.

Method

Subjects

20 subjects were recruited from the student population of the University of Glasgow. All were naive to the experimental aims of the study and all received a small payment in return for their participation. All participants reported having normal hearing.

Materials

All materials were generated and presented in the same way as in Experiments 3and 4.

Design and Procedure

The sequences used for the learning and novel test sequences were the same as in Experiment 1. Subjects were not made aware of the sequential nature of the auditory stimuli. The blocks were shortened from the previous eight sequence cycles to three for each block, making a total of 36 trials. A pilot experiment had shown that this allowed accurate completion of all blocks using a voice response, without fatiguing subjects.

Subjects were instructed that this experiment would test how well they are able to replicate a set of four tones presented to them via headphones. To do so, they were told that they would hear four tones, one at a time. Subjects were informed that their task would be to respond to each tone by making a voice response as quickly and accurately as they could. They were also told that they would be instructed on the correct nature of the voice response required of them.

They were informed that there would be 10 blocks of experimental trials in total with an opportunity to rest between each. Subjects were told that these would consist of the four tones being played to them, one after the other, many times in random order.

Subjects' task was to respond to each tone by singing it back into the microphone positioned above the computer screen. The actual voice response set-up and basic instructions were the same as in the previous experiments containing voice responses. Subjects were informed that they would have to respond as fast as they could, as well as being as accurate in replicating the given pitch of a tone as they could. They were given an example of each tone and had to make an accurate response as instructed to each tone in the presence of the experimenter. The experimenter was not present throughout the experimental trials.

Results

The measure of interest in this experiment was subjects' RTs to a given repeating auditory sequence. If subjects showed a significant increase in their RTs to a novel test sequence when compared to RTs to the repeating learning sequence, they were deemed to have acquired some knowledge about the learning sequence.

Figure 5 shows the means across subjects' median RTs for each block. There was a general decrease in RT with practice. This flattens out towards the latter half of blocks.

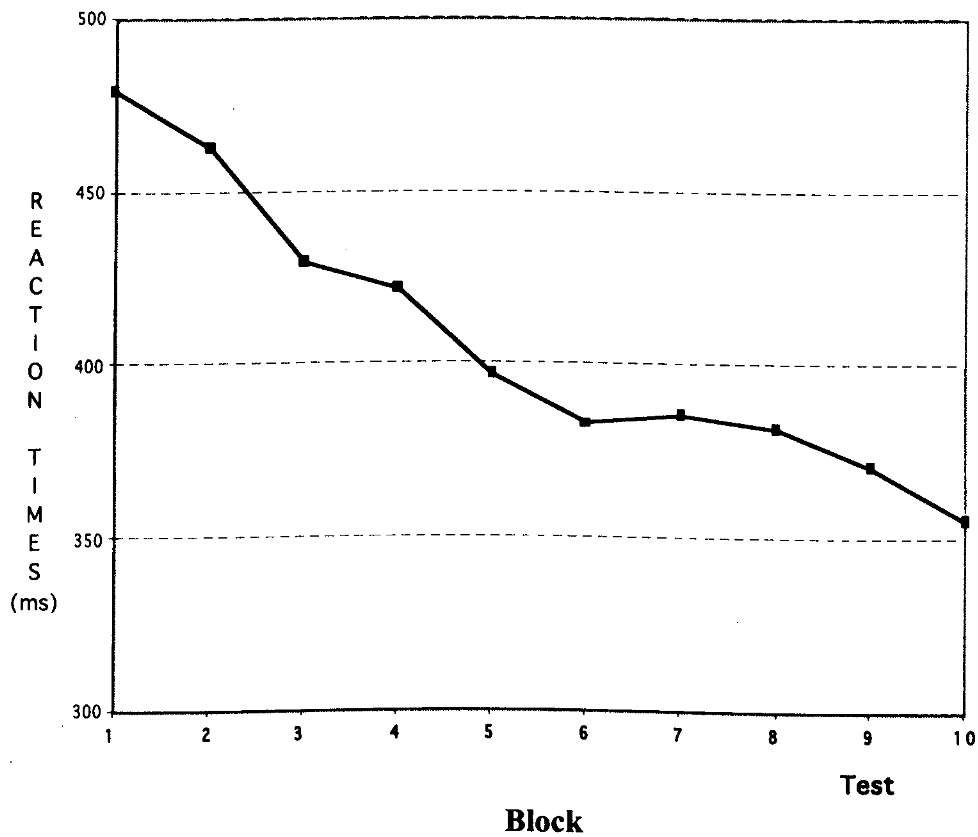

Figure 5. Mean RTs across subjects' median RTs per block. Subjects responded to the repeating tone sequence with a voice response throughout. A novel test sequence was introduced in block 9. This was reversed to the original repeating sequence in block 10.

The three blocks of interest (blocks 8, 9 and 10) show that subjects responded slower when a new sequence was introduced (block 9). This was followed by a further decrease in RT to the original learning sequence (block 10). Table 6 shows the means across subjects' median RTS from these last three blocks.

Table 6. Mean RTs (ms) and Standard Deviations (in brackets) for Last Three Blocks

Block 8 (learning sequence)	**Block 9 (test sequence)**	**Block 10 (learning sequence)**
381 (152)	371 (144)	356 (152)

In order to see whether there was a significant difference between RTs for the original sequence when compared to the test sequence, paired sample t-tests were performed between blocks 8 (original) and 9 (test) and blocks 10 (original) and 9 (test).

There were no significant differences in RTs for these comparisons. The comparison of block 8 and 9 provided a t-value of $t(19)= 1.802$ ($p > 0.05$) and the comparison between blocks 10 and 9 showed a t-value of $t(19)=1.194$ ($p > 0.05$).

These results fail to show any learning of the auditory sequence.

Discussion

These results do not indicate any sequence learning effect when subjects responded to the auditory stimuli with a non-verbal voice response. Keele et al (1995) found learning when subjects used a verbal voice response (Experiment 3). Although it has been suggested that this learning might have been due to explicit processes (e.g. Kelly and Burton, 2001; Willingham, 1999), there is also the possibility that the verbal nature of the responses may have added a conceptual level the tone sequences themselves could not provide here. In Keele et al's (1995) study subjects responded to three visual stimuli, presented in different spatial locations. In one condition, responses consisted of the words 'left', 'middle' and 'right' throughout all trial blocks. A clear sequence learning effect was found, represented in RT differences recorded through the voice responses. In contrast, Experiment 5 used a non-verbal response and it may have been this difference in responses that lead to failure of finding any effects in the current experiment.

Another possibility is that no auditory sequence learning took place in this context. Furthermore, the failure to show sequence learning in this context may also have been due to ceiling effects in the voice response itself. It seems that, overall, the voice response employed to capture auditory sequence learning effects here was insufficient to procure any effects represented in reaction time differences.

2.11 General Discussion

The main question under investigation in this series of experiments was whether subjects could learn an auditory sequence by listening. Such learning is akin to learning by observation in a visual SRT task, which is viewed as learning at the conceptual level. Learning at the conceptual level is in opposition to learning at the motor response level where subjects learn a sequence of motor responses, rather than the underlying rule structure. Evidence for learning by observation has been mixed (e.g. Howard, Mutter and Howard, 1992; Seger, 1997; Willingham, 1999; Kelly and Burton, 2001) and it was hoped that using auditory sequences may help to shed new light on this ongoing debate.

Results from Experiment 1 are clear evidence that the SRT task chosen was robust in providing a design in which subjects lack any explicit awareness. This makes this task clearly implicit in nature. However, this was not the focus of interest in this study and the remainder of experiments looked at the possibility of auditory sequence learning by listening more directly. Experiment 1 clearly showed that subjects are able to learn an auditory sequence of tones they responded to with a spatial motor response. Experiments 2 and 3 failed to provide any evidence for learning by listening when subjects responded with spatial keypress (Experiment 2) or non-spatial voice responses (Experiment 3). There were some questions as to the sensitivity of the voice response, however, and subsequent experiments tested the suitability of the actual voice response employed. Experiment 4 failed to show any transfer of a learning effect from a spatial keypress response to a non-spatial voice response. Evidence from

Experiment 1 clearly showed that subjects learned the sequence when coupled with a keypress response, so that failure to show any learning effect in Experiment 4 must have been due to failure of transfer or lack of sensitivity in procuring any reaction time differences through the voice response. No clear answer could be provided in regards of the sensitivity of the voice response. In order to settle the question of sensitivity of the voice response used here, a final experiment was conducted. Experiment 5 saw subjects responding with a non-spatial voice response throughout all sequence trials. Again, this failed to provide any evidence of auditory sequence learning.

Willingham et al (2000) suggested that sequences are encoded in egocentric space, which is privileged to the motor system and implicit in processing. Evidence from their study suggested that a sequence of response locations must be retained to allow for sequence learning to transfer from one response to another. Mayr (1995) found that subjects could learn a sequence of stimulus locations without a corresponding motor response. He also suggested that non-spatial stimuli could only be learned when a spatial motor response was tied to their sequence exposure. Evidence from a study by Willingham (1999) supports the critical role of responses. Failure to show learning of an auditory sequence in the current series of experiments is also in line with this view. Overall, this series of experiments shows that subjects are able to learn a repeating auditory sequence when they respond to it using spatial motor responses (Experiment 1). However, no evidence of such learning was found when subjects used a non-spatial voice response (Experiment 5). The combined evidence from this series of experiments indicates that subjects have to be engaged with the repeating auditory sequence via a spatial motor response for learning to occur. Overall, no sequence learning by listening was shown. It has to be noted, however, that individual variability proved high in those experiments in which subjects listened to the auditory sequences first (Experiment 2) and those in which they made voice responses (Experiments 3, 4 and 5) and this may have been the reason for the non-significant results. It is a possibility that subjects varied considerably in making their responses and it may have been a difference in skills that was responsible for the high individual variability.

In conclusion, the five experiments reported here did not support the existence of auditory sequence learning through listening. This provides an extended insight into the mechanisms underlying sequence learning in general and adds further evidence that purely conceptual sequence learning, without response exposure, does not arise. In future it would be interesting to test whether spatially presented auditory stimuli give rise to sequence learning, as the spatial locations did in the Mayr (1996) study. Additionally, it is possible that a different aspect of non-verbal auditory stimuli, such as rhythm, may prove to show similarly special properties as spatial attributes have shown in the visual domain.

Chapter 3. Learning with or without a motor response

3.1 Introduction

The mechanisms underlying sequence learning have been described as general and non-selective. Thus, it has been suggested that it is possible to characterise implicit sequence learning as the acquisition of all relevant associations between a person's actions and the stimuli (Schmidtke and Heuer, 1997). Sequence learning is typically shown when subjects are asked to respond to visual items presented in different locations on a computer screen in a choice reaction time task. Unbeknownst to subjects, the successive stimulus displays follow a systematic, repeating sequence. Subjects show a speedup in their response times, while lacking the ability to describe the cause of this (e.g. Nissen and Bullemer, 1987). However, the question has arisen whether it is possible for subjects to learn a repeating sequence without a corresponding motor response tied to the sequence exposure. Thus, if subjects were able to learn a visual repeating sequence without a motor response tied to the sequence exposure, they should be able to learn that sequence by mere observation. By introducing a motor response, it is possible that subjects learn the actual response sequence itself (i.e. a sequence of key presses). This has to be considered as separate from learning of the underlying rule structure of a given sequence, which would make such learning conceptual in nature. Whether sequence learning takes place at the motor response level or the conceptual level is of ongoing debate in the serial reaction time (SRT) task literature and the results of these studies are contradictory. Studies by Howard, Mutter and Howard (1992) and Seger (1997), for example, showed learning of visual sequences that subjects observed, but did not respond to, which is akin to learning at the conceptual level. However, Willingham (1999) and Kelly and Burton (2001) failed to show such learning by observation alone. One of the arguments to answer this discrepancy in findings focuses on subjects' different levels of awareness of the observed sequences, and it has been suggested that subjects are only able to learn sequences with a concurrent high level of awareness (Willingham, 1999; Kelly and Burton, 2001).

The argument that visual sequences can potentially be learned at a conceptual level, rather than at the motor response level, has its origins in the fact that much of sequence learning research uses the visual domain for ease of motor response mapping onto stimulus events. This allows easy recording of reaction times. There have been some studies that used auditory stimuli in an SRT task (e.g. Schmidtke and Heuer, 1997; Buchner, Steffens, Erdfelder and Rothkegel, 1997; Buchner, Steffens and Rothkegel, 1998) in an attempt to extend sequence-learning studies to the auditory domain. Showing sequence learning in the auditory domain would extend the generality of such learning to other experimental conditions. With one exception (Perruchet, Bigand and Benoit-Gonin, 1997), auditory sequence learning was shown in these studies. However, only the Buchner, Steffens, Erdfelder and Rothkegel (1997, Exp. 3) and Schmidtke and Heuer (1997) studies can somewhat be related to learning of an auditory sequence without a corresponding motor response, since all the other studies employed motor response mapping. Here subjects learned uniquely related response stimulus intervals (RSIs) between sequences of tones they responded to. The aim of the previous chapter was to establish whether subjects could learn a systematic auditory sequence *without* a corresponding motor response tied to the sequence exposure, i.e. by *listening* alone (which is akin to observation in the visual domain). The series of experiments conducted did not establish such an effect. Although this shows lack of proof of auditory sequence learning as conceptual, that series of experiments may not provide the full picture. As already pointed out, some studies have shown learning of a visual sequence that was simply observed by subjects (Howard et al, 1992; Seger, 1997), and another finding such sequence learning when the sequence not responded-to consisted of spatial locations (Mayr, 1996). Overall, this indicates that learning of such sequences is possible when there is a spatial component to the systematic sequences to be learned. However, despite the failure to find any sequence learning in the previous set of experiments, it is possible that the measure employed to do so was simply not sensitive enough to pick up any appreciable sequence learning effect. The current chapter aims to investigate further whether auditory sequences can be learned conceptually, without a corresponding motor response tied to subjects' sequence exposure, i.e. by listening alone.

3.2 Simultaneous learning of two uncorrelated sequences

Mayr (1996) used a dual sequence paradigm in which subjects were exposed to two independent visual sequences concurrently. One sequence contained four different objects, while the other sequence consisted of four spatial locations where objects were displayed (the four corners of the computer screen). Subjects made responses to the objects by pressing corresponding keys for each of the four objects, whereas responses were not tied to the spatial locations. Mayr (1996) found learning of both the object and the spatial location sequences. The importance of this finding is that learning of the spatial location sequence did not require a corresponding motor response. This made the learning of the not responded-to sequence special, since acquisition occurred without a motor response to the sequence items (as in the learning of the object sequence). Mayr (1996) argued that the learning of the location sequence was possible due to a spatial orienting system, which is distinct from a system required to learn the response-related sequence. This view suggests that in order for learning of a non-spatial sequence to occur (e.g. objects), that sequence must have a motor response tied to the sequence exposure.

In order to investigate further whether learning of a given auditory sequence is possible without a corresponding motor response, Mayr's (1996) first experiment was adapted to the auditory domain. This addresses the question of whether there are any other stimulus contexts that may allow learning of a sequence without corresponding motor responses. It is possible that there are stimulus classes other than those requiring a spatial orienting system that may elicit sequence learning without a motor response tied to the sequence. The auditory domain may present such a class of stimuli. Auditory material is largely presented sequentially and the listener is used to processing such material in a serial fashion. Furthermore, auditory material is usually processed without spatial motor involvement. Therefore, it seems reasonable to consider the possibility that auditory material may give rise to sequence learning without a corresponding motor response. The previous series of experiments, using a single sequence learning paradigm, addressed this issue to some extent. Here, learning of an auditory tonal sequence was only shown when subjects had to respond to the items of that sequence with a spatial motor response (key presses). In two particular manipulations subjects *listened* to sequences of tones without responding to the items, or subjects used a non-spatial voice response to each sequence item. No discernible learning was found in either condition. These results failed to support the possibility that auditory sequence learning may be possible by listening alone. Auditory sequence learning was only shown when a spatial response was relevant for the stimulus selection in a given sequence, which is in line with Mayr's (1996) suggestion. However, questions have arisen as to the sensitivity of the previous experiments: it is possible that the design was not sensitive enough to pick up any sequence learning by listening alone. Failure to elicit sequence learning using a voice response may have been due to a lack of sensitivity in the voice response to capture any appreciable reaction time gain to the systematic sequence. It is possible, that subjects performed at ceiling using the voice response. The current study used the dual sequence paradigm employed by Mayr (1996), replicating his design by adapting it to the auditory domain. The underlying premise of the dual sequence paradigm is that subjects' reaction times will be affected in the responded-to sequence when the not responded-to sequence is disrupted. This could provide a more accurate way of obtaining any sequence learning effects, rather than the single sequence paradigm employed in the previous set of experiments. By adopting Mayr's (1996) design to the auditory domain, reaction time becomes a suitably sensitive measure to study sequence learning effects.

The current experiment used two different auditory sequences: one of four speakers' voices (2 male and 2 female speakers) and one of four colour names (blue, green, red and white). Choosing these two auditory stimulus classes (i.e. one mostly consisting of different auditory surface qualities and one semantic) allowed simultaneous presentation of two uncorrelated auditory sequences (i.e. speakers saying the different colour names). Additionally, there is a clear qualitative difference in the two stimulus classes used here: On the one hand, there are the colour names that can be easily distinguished by their semantic meaning. On the other hand, there are the different voices that are distinguishable by their different surface qualities alone. A clear difference in difficulty with which the stimuli in the semantic category and the surface feature category can be distinguished should be expected. Thus, it should be easy to distinguish the colour names from one another, whereas this should be harder to do for the four voices. Although subjects underwent a learning phase to acquire the relevant key-to-stimulus mapping and were only included if they had reached a preset learning criterion, there is an argument to say that the colour names should be more easily mapped to their corresponding keypress responses than the voices. Thus, the colour sequence should be more readily learned than the voice sequence.

The current experiment used a fully counterbalanced design. This addresses an issue that arose with Mayr's (1996) design in which only one stimulus dimension was used for motor responses (namely subjects responded to the objects, but never the spatial locations). Overall, Mayr (1996) interpreted his results as evidence for the involvement of separate sequence learning systems. In line with this, Mayr argued that the non-spatial sequence (here objects) is acquired by a system that entails the selection of motor responses. In contrast, he argued that a system independent from motor selection could acquire the sequence of spatial orientations. The system that selects the relevant motor responses can acquire any regularities relevant to responses. Thus, as long as the material is related to the response demands, stimulus entities such as colour, size and form may be acquired. It follows from this interpretation that acquisition of non-spatial regularities will not occur when they are not relevant to the response selection. In terms of Mayr's (1996) design learning of the object sequence would not have been expected if the items in that sequence had not been coupled with a motor response. However, this was not tested. By using a fully counterbalanced design in the current experiment, this will be addressed for the two stimulus classes used here. This is of particular importance, since the domain the stimuli were chosen from differs greatly from the visual stimuli used by Mayr (1996).

Providing a fully counterbalanced design will allow assessment of any independent effects of the different auditory stimulus sequences used. This seems necessary, since it is possible that the two different auditory stimulus sequences pose different processing demands. There is a possibility that one may be learned without a corresponding motor response, whereas the other may not. In this way, the whole picture will be provided and any stimulus domain differences between colour names and speakers' voices can be explored. The question of whether there are other types of stimulus domains that give rise to learning without an overt response tied to the sequence exposure is at the heart of this study.

3.3 Experiment 6

Experiment 6 is a replication of Mayr's (1996) study. However, in place of the object and location sequences, a sequence of four different voices and a sequence of four colour names was used. The aim of this study is twofold: to show that both types of auditory sequences can be learned in a sequence learning set-up when they are tied to a corresponding motor response, and to investigate whether auditory sequences of the types used here can be learned without a corresponding motor response, essentially by listening alone.

Method

Subjects

96 participants were tested in this study. All participants were recruited from the student population of the University of Glasgow. All were naïve as to the experimental aims of the study and received a small payment in return for their participation. All participants reported having normal hearing.

Materials

All stimuli were presented to subjects using an Apple Macintosh iMac computer. The "SuperLab 1.74" experimental package was used to implement the experimental design. All auditory material was presented via headphones to subjects. Responses were made using a common computer keyboard.

Four voices were recorded saying the colours "blue", "green", "red" and "white", which were then spliced into individual stimuli constituting each of the colour names. Two of the speakers were male and two female and the voices were tested for their distinctiveness in a pilot experiment. A total of 16 individual stimuli were created. A response-stimulus interval of 500ms was set. All auditory stimuli were recorded and manipulated using "Sound Edit 16" (version 2) computer synthesiser software.

Design and Procedure

There were two conditions in this experiment, namely a voice condition and a colour condition. In the voice condition subjects responded to the voices by pressing corresponding keys for each of the four voices. In the colour condition subjects responded to the colour names by pressing corresponding keys for each of the four colours. The colours "blue", "green", "red" and "white" mapped onto A, B, C and D respectively. The voices were equally mapped onto A, B, C and D (A = first male voice, B = first female voice, C = second male voice, D = second female voice). The response keys were v, b, n and m on the keyboard mapping onto the voice stimuli or the colour stimuli A, B, C and D respectively. Subjects were instructed to use the middle and index fingers of each hand for responding. Before starting the key practice phase, subjects were given an example of each key-to-audio stimulus mapping.

Depending on the condition, subjects were instructed that this experiment would test how well they are able to learn mapping of voices or colour names onto keys on a computer keyboard. They were told that they would hear four auditory stimuli (voices or colour names), one at a time. Subjects were informed that their task would be to respond to each stimulus by pressing a corresponding key on the keyboard as quickly and accurately as possible.

There were three stages in this experiment. In the key practice phase subjects learned the correct key-to-voice or key-to-colour name mapping. In the sequence learning phase subjects were exposed to the repeating sequences, which were disrupted at specific points. In the final phase, subjects were given an awareness test to test any learning directly. This last part was unexpected to subjects. The experimenter remained present throughout the key practice phase and the awareness test. Subjects were at no point made aware of the existence of any repeating sequences or any disruption of these.

Phase I: Key Practice

At the beginning of the key practice phase, subjects were instructed to press each response key to hear its corresponding auditory stimulus (voice or colour name, depending on the condition). They had to do this in the order of ABCD DCBA ABCD DCBA ABCD DCBA ABCD DCBA. The stimuli used here were the same as those used in the following experimental phase. The training phase differed from Mayr's (1996) training phase in that subjects were presented with both stimulus constituents (voices and colours) in the training phase without a separation of the two at any stage. Learning of the auditory key-to-stimulus mapping was expected to be harder than the mapping of visual stimuli in the Mayr (1996) study. This is especially true for the different voices in the auditory set-up. Exposure to both stimulus conditions at this stage is expected to halt any confusion in the sequence learning phase when subjects will already be familiar with the presentation mode, lessening the possibility of any confusion at that stage. Following familiarisation with the key-to-stimulus mappings, subjects performed a 96-trial practice block of the stimulus responses. These 96 trials followed a random order. However, the run of these trials was constrained so that neither of the voice or colour sequences contained any repetitions. Subjects were told that they would hear a stimulus and they had to respond to it by pressing its corresponding key to the stimulus condition they were allocated to. They were given onscreen feedback telling them whether they had pressed the correct key or not. The feedback was there to enable subjects to adjust their responses should they make any mistakes. Once subjects had made a response there was a 500ms delay before the next auditory stimulus was played. This included the 250ms presentation time of the onscreen feedback. Subjects who had reached a preset learning criterion went on to complete the experiment. Subjects who made more than 10% errors in the 96 trials finished here and did not complete the experiment. The key practice phase took approximately 10 minutes to complete.

Eight subjects' data was replaced in the voice condition and nine in the colour condition. The replaced subjects had made more than 10% errors in their keypress responses across all blocks and were eliminated from analysis.

Phase II: Sequence Learning

Following successful completion of the key practice phase, subjects were told that the experiment would now begin. Subjects were informed that there would be 16 blocks of experimental trials in total with an opportunity to rest between each. They were told that the blocks would consist of the four stimuli they had to respond to, played one after the other in random order. They were reminded that they would have to respond as fast and accurately as they could. There was no further feedback to subjects' responses with the screen remaining blank throughout the sequence learning trials.

Two sequences were used that corresponded to the sequence orders used by Mayr (1996). Sequence A followed the order of DBDABCAC consisting of eight items. Sequence B followed the order of CDADBCABA consisting of nine items. These sequences corresponded either to the voices (where A = 1st male voice, B = 1st female voice, C = 2nd male voice and D = 2nd female voice) or the colour names (A = blue, B = green, C = red and D = white). The two sequences were organized so that they made up a total of 72 trials in each of the 16 sequence learning blocks. These followed nine repetitions of the eight-item sequence and eight repetitions of the nine-item sequence. In this way, each element in either sequence was paired only once with each item in the other sequence. In blocks 9 and 12 either the voice or colour sequence was random, and in Block 15 both sequences were random. In the voice condition, half of the subjects were presented with the voice sequence random in Block 9 and the colour sequence random in Block 12. For the other half of subjects the reverse assignment was in place. In the colour condition, half of the subjects were presented with the colour sequence random in Block 9 and the voice sequence random in Block 12. For the other half of subjects the reverse assignment was used.

Stage III: Direct Measure of Awareness

Following Block16 subjects in both conditions were tested on any knowledge they held about both sequences. Subjects were falsely led to believe they had been randomly allocated to one of four experimental conditions:

Condition 1 –both voice and colour sequence followed a regular pattern

Condition 2 – the voice sequence was regular and the colour sequence random

Condition 3 – the voice sequence was random and the colour sequence regular

Condition 4 – both voice and colour sequence followed a random pattern.

Subjects had to indicate which of the four conditions they thought they had been assigned to. Once they had indicated their choice, they were informed that, in fact, both sequences had followed independent regular patterns. A short test that assessed any knowledge subjects held of both sequences followed. This test was a generation task that provided subjects with three cue items of a given sequence to which they had to generate the next item in the sequence. This generation task differed from Mayr's (1996) task in that it used three items, rather than two items as in the Mayr (1996) study. The aim of using just two items was to keep the test as short as possible without contaminating subjects' knowledge of a given sequence with knowledge of the other. This was a possibility, since subjects were tested on both sequences. However, it is reasonable to assume that increasing the cue items from two to three will provide an increase in the test's sensitivity without considerably lengthening the generation test.

In the voice condition, subjects were played three cue voices, all remaining constant in the colour they said (i.e. "blue"). They were then asked to generate the next voice item using the same key-to-voice mapping they had used before. Following this, subjects were asked to indicate how many colour names they had heard and if they recalled all four of them correctly, they were played three cue colour items to which they had to predict the next colour name saying it out loud to the experimenter. In the colour condition subjects were played three cue colour names using a neutral voice. They were then asked to generate the next colour name using the same key-to-colour mapping as before. Following this, subjects were asked to indicate how many voices they had heard in the previous phase. If they recalled the use of four voices they were asked to indicate a distinguishing feature of each of these (i.e. male, female, accent, etc). Subjects were presented with three cue voices to which they had to predict the next voice to the experimenter. This process was repeated for both sequences until the prediction of each sequence item had been attempted. The order of presentation for the voice and colour name generation test was counterbalanced.

Results

Overall Learning Effects: Median reaction times (RTs) were calculated for each participant and block. Reaction times > 2000ms were excluded from analysis. Figure 1 displays subjects' reaction times in the voice condition, in which subjects responded to the voices with keypress responses. This shows reaction times across blocks separately for those participants who had the voice sequence random in Block 9 and the colour sequence random in Block 12, as well as for the participants for whom the reverse order was applied.

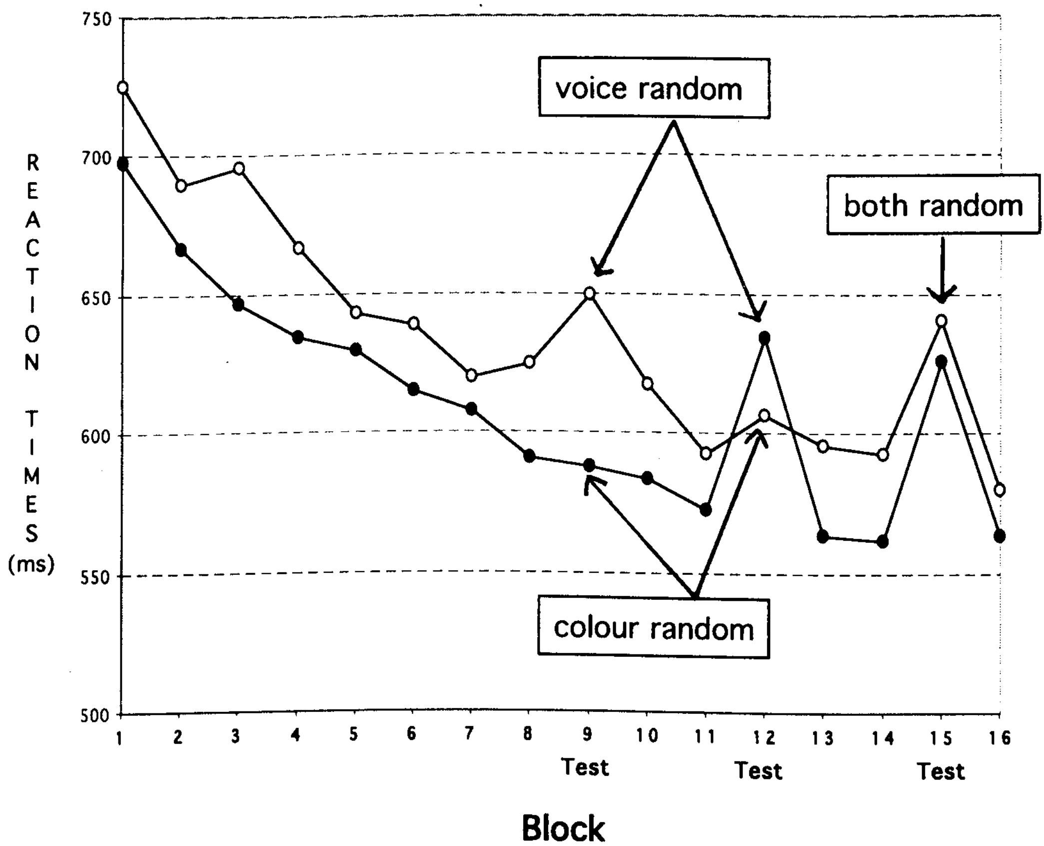

Figure 1. Voice Condition. Reaction times as a function of training blocks, separately for participants who were exposed to a random voice sequence in block 9 and a random colour sequence in Block 12 (white circles), and for the participants for whom the reverse order was used (black circles). Both sequences were random in Block 15. All participants responded to the voice sequence.

Figure 2 displays subjects' reaction times in the colour condition, in which subjects responded to the colour names with keypress responses. It shows reaction times across blocks separately for those participants who had the colour sequence random in block 9 and the voice sequence random in block12, as well as for the participants for whom the reverse order was employed.

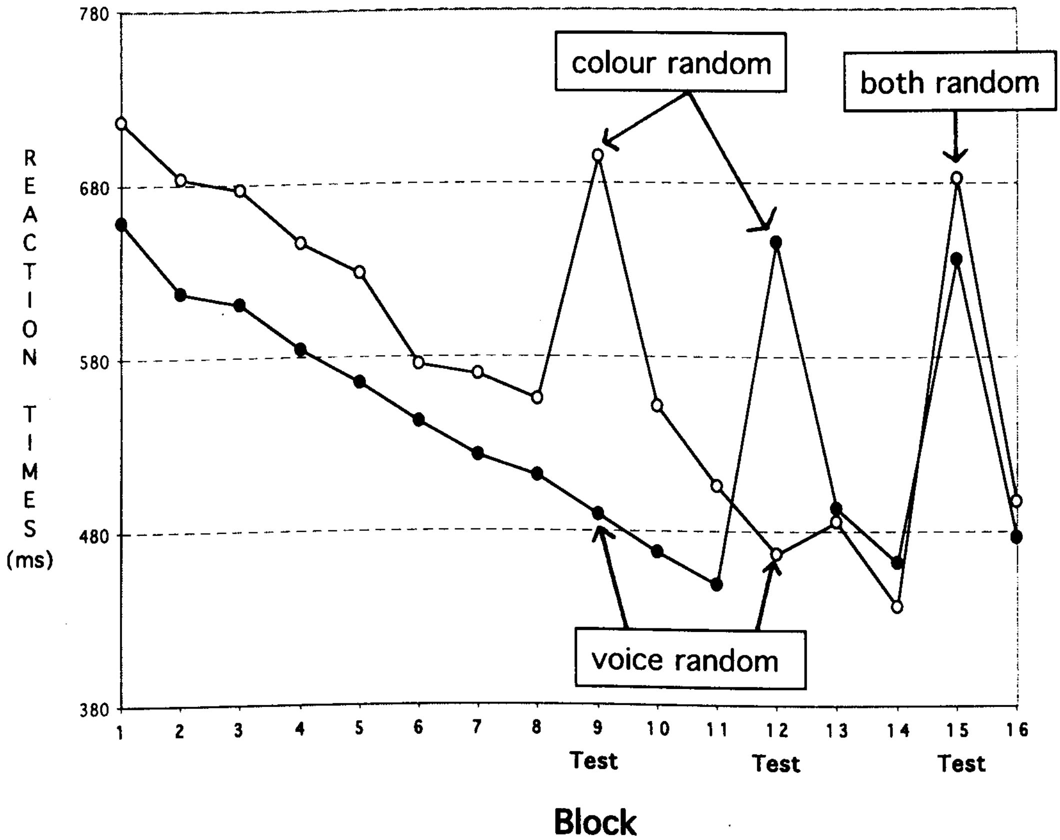

Figure 2. Colour Condition. Reaction times as a function of training blocks, separately for participants who were exposed to a random colour sequence in block 9 and a random voice sequence in Block 12 (white circles), and for the participants for whom the reverse order was used (black circles). Both sequences were random in Block 15. All participants responded to the colour sequence.

A clear practice effect can be seen in both the voice and colour condition. Furthermore, an increase in RTs in each of the random blocks is shown for those sequences that were responded to in blocks 9 and 12. Additionally, an increase in RTs was found for the block in which both sequences became random

(Block 15). The analysis involved the comparison of the mean RTs (across subjects' median RTs) between a random sequence block and its adjacent regular sequence blocks. Related-sample t-tests revealed that learning for both the colour and voice sequences occurred when they were responded-to. No learning of either the voice sequence or the colour sequence was shown when these were not responded-to.

In the voice condition, comparing RTs between the repeating and random sequence blocks, significantly faster RTs in the repeating sequence blocks were shown when these were the blocks in which the responded-to voice sequence was disrupted (collapsed across the 8-item and 9-item sequences; $t = 4.022$, $p < 0.05$). There was no significant difference in RTs between the blocks in which the not-responded-to colour sequence became random and the adjacent repeating sequence blocks ($t = 0.812$, $p > 0.05$). A significant difference was found in RTs for Block 15, in which both the responded-to voice sequence and not responded-to colour sequence became random, and the adjacent repeating sequence blocks ($t = 6.311$, $p < 0.05$).

For the colour condition, the same pattern of results emerged: a significant difference in RTs between the random and repeating sequence blocks could be shown when these were the blocks in which the responded-to colour sequence was disrupted (collapsed across the 8-item and 9-item sequences; $t = 8.734$, $p < 0.05$), showing significantly faster RTs for the repeating sequence. There was no significant difference in RTs between the blocks in which the not-responded-to voice sequence became random and the adjacent repeating sequence blocks ($t = 1.615$, $p > 0.05$). A significant difference was found in RTs for Block 15, in which both sequences were random, and the adjacent repeating sequence blocks ($t = 8.895$, $p < 0.05$).

These results indicate that subjects were able to learn the sequence they responded-to, whether this was made up of voices or colour names. No evidence of learning the not responded-to sequence, whether it be voice or colour names, was found.

Learning of the eight-item and nine-item sequences: The above results used data that was collapsed across the eight-item and nine-item sequences. For ease of computing whether there was any difference in performance between subjects responding to the eight-item and nine-item sequences, difference scores between random blocks 9 and 12 and the adjacent blocks were calculated for each subject and block. When considering the learning of the two sequence lengths separately, performance did not differ. A 2x2 (Sequence Length x Sequence Type) ANOVA revealed a main effect of Sequence Type, $F(1,92) = 29.4$ ($p < 0.05$), but no main effect of Sequence Length, $F(1, 92) = 0.934$ ($p > 0.1$). Also, there was no interaction between these factors, $F(1, 92) = 2.0$ ($p > 0.1$). From this we can infer that the use of the eight- and nine-item sequences in both conditions is comparable.

Learning of the voice and colour sequences: The above ANOVA also gives an indication of any differences in performance between the voice and colour conditions. From the main effect of Sequence Type we can assume that subjects in the colour condition showed greater learning of the colour sequence, than subjects showed of the voice sequence in the voice condition.

Assessment of Awareness of the Sequences: the post-experimental questions asked subjects to declare whether they noticed any regularity in either 1) voice and colour sequences, 2) voice sequence alone, 3) colour sequence alone or 4) neither sequence. The number of subjects who subscribed to each category is shown separately for the voice and colour conditions in the tables below. The data is collapsed across subjects' responses for the eight- and nine-item sequences.

Voice Condition (those subjects who responded to the voice sequence)

Category Chosen	**Number of Responses**
1. Voice and Colour	10
2. Voice	24
3. Colour	3
4. Neither	11

A Chi-Square test reveals that there is a difference in the number of responses in each category for the voice responders, $X^2(3) = 18.8$ ($p < 0.05$). A voice (category 2) versus others (categories 1, 3 and 4) Chi-Square test shows that significantly more subjects responded "Voice" than any other response, $X^2(1) = 20.0$ ($p < 0.05$). The data suggests that subjects noticed the regularity in the voice changes more than those in the colour when they responded to the voices in the experiment.

Colour Condition (those subjects who responded to the colour sequence)

Category Chosen	**Number of Responses**
1. Voice and Colour	19
2. Voice	1
3. Colour	26
4. Neither	2

A Chi-Square test reveals that there is a difference in the number of responses in each category for the colour responders, $X^2(3) = 38.8$ ($p < 0.05$). A colour (category 3) versus others (categories 1, 2 and 4) Chi-Square test shows that significantly more subjects responded "Colour" than any other response,

$X^2(1) = 21.8$ ($p < 0.05$). This data suggests that subjects noticed the regularity in the colour changes more than those in the voices when they responded to the colours in the experiment.

The generate task probed any knowledge subjects held of the sequences. This helps to determine whether subjects held more knowledge of one or the other sequence. The means for the voice and colour generation scores are given in the tables below, separately for the voice and colour conditions, collapsed across the eight- and nine-item sequences.

Voice Condition (those subjects who responded to the voice sequence)

Sequence	Generation Score (out of 8.5)		
	Mean Number	**SD**	**Percentage**
Voice	4.4	1.9	52
Colour	3.3	1.4	39

In the voice condition, subjects did not perform significantly above chance of 4.25 in generating items for the voice sequence ($t_{(47)} = 0.54$; $p > 0.05$). Subjects performed significantly below chance in generating the items of the colour sequence ($t_{(47)} = -4.65$; $p < 0.05$). These data suggest that, in the voice condition, subjects did not acquire any appreciable knowledge of either the voice sequence they responded to, nor the colour sequence they did not respond to.

Colour Condition (those subjects who responded to the colour sequence)

Sequence	Generation Score (out of 8.5)		
	Mean Number	**SD**	**Percentage**
Colour	6.3	2.3	74
Voice	2.2	1.9	26

In the colour condition, subjects generated items for the colour sequence significantly above chance performance of 4.25 ($t_{(47)} = 6.11$; $p<0.05$). Subjects performed significantly below chance in generating the items for the voice sequence ($t_{(47)} = -7.40$; $p<0.05$). These data suggest that, in the colour condition, subjects acquired appreciable knowledge of the colour sequence they responded to, but did not gain such knowledge about the voice sequence they did not respond to.

3.4 General Discussion

The results of this experiment suggest that auditory sequences of the types used here do not exert any influence on performance when they are not relevant for making a 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 whether a sequence was responded-to or not. Both voice and colour sequences showed a significant increase in reaction times (compared to a random sequence) when they were responded-to, but no such increase in reaction times was shown for either sequence when it was not responded-to. This suggests that the underlying sequential aspect of the stimulus sequences remained unprocessed. Previous studies have shown that a sequence can be learned when a motor response is tied to the exposure and this included auditory sequences of tones (e.g. Buchner, Steffens, Erdfelder and Rothkegel, 1997; Buchner, Steffens and Rothkegel, 1998). These results were clearly replicated here. However, the main question here was whether an auditory sequence could be learned, without a corresponding spatial motor response tied to the sequence exposure, by listening alone. The results found here extend Mayr's (1996) original findings to the auditory domain. It seems there is nothing 'special' about either auditory colour name sequences or sequences of voices that would give rise to learning these without an overt response tied to the sequence exposure. Using a fully counterbalanced design showed this to be the case for both auditory sequences. This is of importance, since Mayr (1996) showed learning of the not responded-to spatial location sequence and learning of the responded-to object sequence, but he did not investigate whether object sequences could be learned when they were not responded-to.

A clear difference in learning of the colour sequence when compared to the voice sequence was found. The colour sequence was acquired more readily than the voice sequence. This is not surprising, since the colour-to-key mapping was acquired more easily, as the voices were less distinguishable from one another when compared to the colour names. Although subjects accomplished the voice-to-key mapping in this task, this was done less easily. This is reflected in the longer response latencies in the voice condition. The faster overall response times in the colour condition show that the key-to-colour mapping was accomplished more readily than the voice-to-key mapping. This can also be taken to indicate that the colour sequence was learned more easily than the voice sequence. Additionally, a pattern of results emerged in the colour condition that supports this assertion, but was unexpected. When looking at the reaction time data more closely it emerged that some subjects in the colour condition had extremely fast response times to the repeating sequence overall. When looking at response latencies that were faster than 200ms, it was found that nine subjects responded below this threshold more than 10% of the time across all blocks. Additionally, a further six subjects responded more then 5% below this 200ms threshold across all blocks. Some of these reaction times were in the 1-50ms region. However, this was not accompanied by keypress errors for those trials. Overall, this indicates that subjects in the colour condition were able to anticipate which item in the sequence would come next. No such pattern was found in the voice condition. Taking this into account, it was particularly important to show that the colour sequence, although learned more easily, was not learned in the voice condition when it was not responded-to. Thus, the results show that even the easier-to-learn auditory sequence of colour names did not give rise to learning by listening alone.

The awareness data shows that subjects were aware of some regularity in the sequence they responded to, both in the voice and the colour condition. However, the generate task shows that subjects in the voice condition did not acquire any appreciable knowledge of the voice sequence they responded to. The pattern of results is different in the colour condition in which subjects performed significantly above chance in generating the colour sequence they responded to. A similar pattern of results was found in both conditions for the secondary sequence subjects did not respond to: subjects failed to gain appreciable knowledge of the secondary colour sequence when they responded to the voice sequence and subjects failed to acquire appreciable knowledge of the secondary voice sequence when they responded to the colour sequence. Overall, this data suggests that subjects acquired definite knowledge of the colour sequence when they responded to that sequence, but no such pattern was found in the voice condition. It seems that the colour sequence was more salient, giving rise to higher levels of awareness in this context. Willingham (1999) and Kelly and Burton (2001) suggested that learning of visual sequences by observation alone may have been due to high levels of concurrent awareness.

Awareness was lacking for the secondary sequence subjects did not respond to for both conditions. This is of interest, since there was a possibility that subjects would become aware of the more salient secondary colour sequence in the voice condition. This was obviously not the case and it is possible that failure to learn the secondary sequence may have been due to a lack of concurrent high levels of awareness of that sequence.

The results clearly show that learning of the type of auditory sequences employed here occurs when a motor response is tied to the sequence exposure. However, no learning of the secondary sequence, which participants simply listened to, was shown for either type of sequence. Although failure to show such learning does not preclude the possibility that it exists, taking into account evidence from the previous series of experiments and studies in the visual domain, the continued failure to show sequence learning at the conceptual level does indicate that for such learning to occur, a motor response is required. The only exception so far has been the learning of spatial location sequences as found by Mayr (1996). In future, it would be interesting to see whether similar sequence learning could be shown with spatial auditory stimuli.

In conclusion, in this context there appears to be nothing special about the auditory domain when it comes to sequence learning by listening alone. Mayr's (1996) finding that a location sequence gave rise to learning without a corresponding motor response indicates that a spatial sequence holds some special property that gives rise to such learning. This was not found for the auditory stimuli employed here.

Chapter 4. Invariant Learning of Auditory Features

4.1 Introduction

Berry and Dienes (1993) suggested that implicit learning could be defined in terms of unintended learning of a fairly complex stimulus coupled with difficulty in expressing the acquired knowledge. However, implicit learning may not be limited to learning of complex structures, but may be replaced by learning of simple underlying rules, as long as these remain unavailable to conscious awareness. McGeorge and Burton (1990) employed such a simple type of rule in an incidental learning task, in which subjects were exposed to different exemplars that all contained an invariant. In their invariant learning task subjects were presented with 30 four-digit numbers, one at a time. Subjects were required to do some arithmetic task on these and did not notice the presence of the invariant "3" in each exemplar. Following this phase, subjects performed a forced-choice recognition task. They were presented with pairs of four-digit numbers and were falsely led to believe they had seen one number in a pair in the previous phase. Subjects' task was to choose which one they thought they had seen previously. What they did not know was that they had not seen either number before. However, one contained the invariant "3" (the positive) and the other did not (the negative). McGeorge and Burton (1990) found that subjects chose the positive significantly above chance performance when compared to selection of negatives. This was coupled with apparent lack of awareness for applying knowledge of the invariant. McGeorge and Burton (1990) interpreted this as evidence for implicit learning. In a further experiment, McGeorge and Burton (1990) found transfer of this implicit knowledge across stimuli surface forms (i.e. number strings expressed in digits at study; number strings expressed in words at test). This suggests that the invariant rule knowledge subjects acquire is at the conceptual level. Other studies have also found implicit knowledge transfer when the underlying structure remains unchanged, but the surface structure varies, between study and test (e.g. Altman, Dienes & Goode, 1995; Bright and Burton, 1994; Huddy and Burton, 2002; Mathews, Buss, Stanley, Blanchard-Fields, Cho and Druhan, 1989; Newell and Bright, 2002; Reber, 1969).

4.2 Auditory Stimuli and the Invariant Learning Task

In the original study by McGeorge and Burton (1990) digits in the form of number strings and as written words were used. These type of visual stimuli were also employed in various forms in other invariant learning studies (Churchill and Gilmore, 1998; Cock, Berry and Gaffan, 1994; Huddy and Burton, 2002; Newell and Bright, 2002a, 2002b; Stadler, Warren and Lesch, 2000; Ward and Churchill, 1998; Wright and Burton, 1995; Wright and Whittlesea, 1998). Other visual stimuli employed consisted of clock faces (both analogue and digital clock faces; Bright and Burton, 1994 and Newell and Bright, 2002). Lastly, Kelly, Burton, Kato and Akamatsu (2001) investigated learning of real-world regularities using visual stimuli such as coins and logos. Although one particular invariant learning study investigated phonological issues in conjunction with digit stimuli (Newell and Bright, 2002a), all of the studies in the invariant learning literature used visual stimuli. None employed purely auditory stimuli. Ward and Churchill (1998) posited the question whether other stimuli than digits can prove successful in distinguishing between different processing accounts of invariant learning. The current study introduces a new class of stimuli in an invariant learning context and aims to extend invariant learning research to new experimental contexts. In turn, this will provide new insights into issues in invariant learning in general.

Auditory stimuli, such as tones, may present different processing demands than visually presented stimuli, such as digits. At present, we do not know whether the basic phenomenon found by McGeorge and Burton (1990) can be replicated with auditory material. It is possible that there are stimulus classes other than visual material that would allow such learning to occur. The auditory domain may present such a class of stimuli. Investigating the possibility of invariant learning in the auditory domain will provide a fresh departure point and possibly allow us to gain new insights into the mechanisms underlying invariant learning in particular and implicit learning in the auditory domain in general. Additionally, the potentially different processing demands of auditory stimuli (such as the tones used in the current experiments) may affect awareness of the underlying invariant rule. Thus, extending the invariant learning literature to the auditory domain would also be useful for exploring awareness issues in implicit learning in general.

4.3 Experiment 7

Investigating invariant learning in the auditory domain covers new ground and it is necessary to employ a suitable design. In order to do this it is useful to remain close to the original design that provided a departure point for other invariant learning studies. For this reason Experiment 1 from the original McGeorge and Burton (1990) study was adapted to accommodate auditory material that could be mapped directly onto the nine digits used initially. Simple synthesized tones provide auditory material that is suitable for this purpose and these easily replaced the digits used in the original design. Thus, Experiment 7 is a replication of Experiment 1 in the McGeorge and Burton (1990) study, except for substitution of the original nine digits (1 – 9) with auditory tones. This provides a starting point that will establish whether the general invariant learning phenomenon found by McGeorge and Burton (1990) can be replicated in the auditory domain. McGeorge and Burton (1990) found that subjects selected test items containing an invariant digit "3" above chance when they had been exposed to a study set of items all containing the invariant, even though all test strings were new to them. On post task questioning, subjects did not seem to have been aware of the presence of the invariant rule. McGeorge and Burton (1990) interpreted this as evidence that subjects' bias towards selecting strings that contained the invariant at test was based on an implicitly acquired rule. At present, we do not know what auditory regularities, if any, can be learned in an invariant learning paradigm. The current experiment is the first attempt to find an auditory regularity that may be acquired in this context.

In the current auditory adaptation of the original McGeorge and Burton (1990) experiment, subjects were exposed to a learning set of four-tone sequences, all of which contained the invariant tone. They performed a distracter task and were not made aware of the common invariant feature between each sequence. This distracter task saw the greatest change from the original design: subjects in the McGeorge and Burton (1990) study performed some arithmetic task on the learning set of digits. This ensured that subjects attended each digit in a learning sequence, as well as distracting from the true nature of the experiment. In order to provide an equivalent distracter task in the current experiment, subjects were asked to decide which of the tones in a given sequence had been the highest in pitch. This ensured that they attended to each individual tone in a learning sequence, as well as providing a suitable cover story to the purpose of the experiment. Following this distracter task, subjects were given a surprise forced-choice recognition test in which they were presented with pairs of four-tone sequences. They were falsely led to believe that they had heard one of the pair previously when, in fact, these were all completely new. One sequence in the pair, however, contained the invariant tone G (the positive), whereas the other did not (the negative). Subjects had to choose the sequence they thought they had heard in the previous set of 30 items. If subjects acquired the underlying invariant knowledge (i.e. the given tone), they should show selection of positives at test significantly above chance, when compared to negatives.

Method

Subjects

20 subjects were recruited from the student population of the University of Glasgow. All received a small payment in return for their participation. All participants reported having normal hearing.

Materials

All stimuli were presented to subjects using an Apple Macintosh G3 computer. The "SuperLab" experimental package was used to implement the experimental design. All auditory material was presented via headphones to subjects. Responses were made using a common computer keyboard.

All auditory stimuli were generated using Sound Edit 16 (version 2) computer synthesizer software. The sample rate and sample size was set at 44.100 kHz (CD quality) and 16 bits respectively. The nine tones in this experiment were taken from the 12-tone scale and were as follows: 1) C1 (130.8Hz), 2) D1 (146.8Hz), 3) E1 (164.8Hz), 4) F1 (174.6), 5) G (196.0Hz = invariant), 6) A (220.0Hz), 7) B (246.9Hz), 8) C2 (261.6Hz) and 9) D2 (293.7Hz). Each tone was 500ms in length.

The experimental material consisted of four-tone sequences, drawn randomly from the nine tones above. The tones in each sequence were separated by a 250ms ISI. Study sets of these were generated individually for each subject.

Firstly, a set of 40 positive items was generated. A positive is defined as a four-tone sequence that contains at least one tone G. Secondly, a set of ten negative items was generated, with a negative constrained to contain no G. The study sets were further restricted so that there were no repetitions of a particular four-tone sequence within a subject's study set.

Design and Procedure

This was a within-subjects design, in which subjects were presented with 30 study items drawn from a set of 40 positive sequences. This was followed by ten pairs of sequences made-up of the remaining ten positive and ten negative sequences.

The experiment consisted of two stages: a learning phase and a test phase. In the learning phase subjects were presented with 30 positive sequences. Subjects were told that this experiment would test their ability to distinguish tones they heard via headphones. They were informed that they would hear 30 four-tone sequences, one at a time, and that they had to decide which of the four tones had been the highest in pitch in a given sequence. They were instructed that they had to record their response by pressing 1, 2, 3 or 4 on the computer keyboard (each representing tone 1, tone 2, tone 3 and tone 4 in a sequence respectively). This served as the distracter task, leaving the true nature of this task unrecognizable to subjects. The learning phase was followed by an unexpected test phase. Subjects were given ten pairs of sequences, presented with a 1500ms pause between the two sequences in a pair. Each pair consisted of a positive and a negative sequence with the order of these randomized throughout the ten test trials. Neither of these two sequences had been heard in the learning phase previously. Subjects were falsely told that they had heard one sequence in a pair in the previous phase of the experiment. They were asked to decide which one they had heard before by pressing "1" on the keyboard if it had been the first sequence, or "2" if it had been the second one in the pair. Subjects were permitted to replay a pair of sequences one time, or make their response immediately after first play. If they were unsure as to which sequence they had heard previously, they were asked to guess.

This was a self-paced task and each subject took approximately 10 minutes to complete the entire task.

Results

The mean number of positives selected in the test phase was 5.65 (out of a possible of 10), with a standard deviation (SD) of 1.6. A one-sample t-test compared this with chance performance of 5. This showed that selection of unseen positives was not significantly above chance, $t(19) = 2.02$ ($p > 0.05$). This data suggests that subjects did not acquire the underlying invariant rule knowledge, but responded randomly. However, it should be noted that these results approach significance and are in the right direction, thus they may simply lack power (this is also indicated by results in Experiment 8, which has a comparable t-value of 2.98, and comes out significant, see p. 138)

When subjects were debriefed following the test phase, they were told that the 30 sequences in the first part of the experiment all contained a common feature, i.e. the tone G. Furthermore, they were informed of the difference between the paired sequences (i.e. one containing the exemplar G, but not the other). All subjects reported that they had not noticed the common feature in the learning sequences and that they had mostly guessed which sequence in the test pair they had heard previously. Some subjects reported that they sometimes felt they had heard neither of the test pair sequences before. However, they indicated that they attributed this to the general difficulty of the apparent memory task itself, rather than guessing they had been misled. All subjects expressed surprise as to the true nature of the experiment. Overall, this indicates that subjects were not aware of the invariant feature or the true nature of this experiment.

In this instance, subjects did not seem to acquire any implicit or explicit knowledge of the invariant tone G.

Discussion

Subjects in Experiment 7 did not show acquisition of any implicit or explicit knowledge of the underlying invariant tone G. This is surprising, since this was an audio replication of Experiment 1 in the original McGeorge and Burton (1990) study, in which apparent implicit learning of the invariant feature (the digit “3”) was found. The difference in the current experiment was the substitution of the visual material with simple auditory tones. It appears that subjects were unable to extract, implicitly or explicitly, the common tone G from the 30 learning sequences in the current context. It is worthwhile noting that Burns (1999) summarized that even experts (i.e. musicians) are not able to label individual tones correctly, except for a small minority of people who have perfect pitch. Thus, it is a possibility that subjects were unable to distinguish a specific tone (i.e. G) in the current experiment, although they were able to make a relative pitch judgment in the distracter task for each tone sequence. The question arises whether it is possible to utilize a different auditory quality of tones in order to create an auditory invariant rule that can be learned implicitly. Experiment 8 is an attempt to use a different auditory quality by using a harmonic relationship between two tones.

4.4 Experiment 8

Experiment 7 failed to show implicit acquisition of a specific tone in an auditory invariant learning task. At present, it is still not known whether subjects can acquire an auditory invariant feature and if they can, what form this auditory regularity may take. Experiment 8 investigated further whether such learning is possible using a different auditory invariant than the previous experiment. Harmonies between tones may not provide an auditory quality that can be clearly labeled by novices, but it may nevertheless provide a salient auditory property to the untrained listener. All tones in the current series of experiments were taken from the 12-tone scale. The hierarchical structure of the 12-tone scale leads to certain harmonic expectations even by the untrained listener (Krumhansl, 1990). It is a reasonable assumption that a particular harmonic quality between tones from this scale would provide an invariant feature that subjects may acquire implicitly. A Major 4^{th} provides such a salient harmonic quality. By integrating this in the invariant learning paradigm as the invariant regularity, subjects can be provided with an auditory invariant that has the potential for implicit acquisition. Experiment 8 uses this salient feature in its learning sequences. This investigated whether subjects can learn an auditory invariant feature when this is not an individual tone, but a harmonic relationship between two tones.

Although the current series of experiments did not focus on the question of implicit versus explicit knowledge acquisition, the post-task questioning re subjects' potential awareness of the presence of the invariant feature was constrained in order to tap into any verbalizable knowledge, as well as preventing contamination from debriefing.

The main question in the current experiment was whether subjects are able to acquire the harmonic invariant feature (i.e. presence of a 4^{th} between the second and the third tone). Thus, if subjects show above chance selection of positives (those sequences containing a 4^{th} in the second interval) when compared to negatives (those sequences that do not contain a 4^{th} in the second interval) at test, subjects were deemed to have acquired the relevant invariant knowledge.

Method

Subjects

20 subjects were recruited from the student population of the University of Glasgow. All received a small payment in return for their participation. All participants reported having normal hearing.

Materials

All material was generated and presented as in Experiment 7.

There were eight additional tones to Experiment 7's nine, totaling 17 in this experiment. All were drawn from the 12-tone scale. Increasing the number of these in the current experiment was necessary in order to accommodate a sufficient number of 4^{th} intervals, as well as providing a limited set of stimuli. A 4^{th} is defined as a specific interval between two tones. Therefore, it is not the tones themselves that make up the invariant feature here, but the invariant lies in the relationship between two given tones, which may vary themselves. The invariant 4^{th} was further defined to appear in the interval between the second and the third tone (second interval), although there were no restrictions on 4^{ths} appearing in the other intervals. It was necessary to increase the number of tones used here, since a 4^{th} not only stipulates a specific interval between two notes, but this interval is unidirectional, i.e. it always incurs an elevation in pitch in the second note. Thus, each invariant 4^{th} introduced a higher tone, increasing the overall number of stimuli required in the total set. For this reason, a set of 17 tones, including halftones, was utilized and only the first 12 of these were used to create invariant 4^{ths}. The remaining five were not used to create 4^{ths}, but provided the second tones to the 4^{ths} from the first 12 tones. The tones and 4^{ths} (separated by -) used were as follows: 1. C1 (130.8Hz) - F,

2. C1# (138.6Hz) - F#, 3. D1 (146.8Hz) - G, 4. D1# (155.6Hz) - G#,

5. E1 (164.8Hz) - A, 6. F (174.6Hz) - A#, 7. F# (185.0Hz) - B,

8. G (196.0Hz) - C2, 9. G# (207.7Hz) - C#2, 10. A (220.0Hz) - D2,

11. A# (233.1Hz) - D#2, 12. B (246.9Hz) - E2. The remaining five tones not used for 4ths were: 13. C2 (261.6Hz), 14. C#2 (277.2Hz), 15. D2 (293.7Hz),

16. D#2 (311.1Hz) and 17. E2 (329.6Hz).

Tones were 500ms in length. The experimental material consisted of four-tone sequences, drawn randomly from the 17 tones above, with each separated by a 250ms ISI. As before, study sets of these were generated individually for each subject. Firstly, a set of 40 positive items was generated. A positive here is defined as a four-tone sequence that contains a 4^{th} in the second interval. Secondly, a set of ten negative items was generated, with a negative constrained to exclude any 4^{th} in the second interval. Study sets were further restricted so that there were no repetitions of a particular four-tone sequence within a study set.

Design and Procedure

This was a within-subjects design, in which subjects were presented with 30 study items drawn from the set of 40 positive sequences, followed by ten pairs of sequences made-up of the remaining ten positive and ten negative sequences.

The experiment consisted of two stages: a learning phase and a test phase. In the learning phase subjects were presented with 30 positive sequences. Subjects were told that this experiment would test their ability to distinguish tones they heard via headphones. They were informed that they would hear 30 four-tone sequences, one at a time, and that they had to decide which of the four tones had been the highest in pitch in a given sequence. They were instructed to make their responses as in Experiment 7. This, again, served as the distracter task. The learning phase was followed by an unexpected test phase. Subjects were given ten pairs of sequences, presented with a 1500ms pause between each sequence in a pair. Each pair consisted of a positive and a negative sequence with the order of these randomized throughout the ten test trials. Neither of these two sequences had been heard in the learning phase previously. Subjects were falsely told that they had heard one of these two sequences in the previous phase of the experiment. They were asked to indicate which one they had heard before as in Experiment 7. Subjects could replay a pair if they wished to do so, or make their response after the first play. If they were unsure they were asked to guess.

This was a self-paced task and each subject took approximately 10 minutes to complete the entire task.

Once subjects had completed this task, they were asked whether they had noticed anything systematic about the 30 sequences in the first phase. They were then asked whether they had used any kind of strategy to decide which of the pair of sequences they had heard before. Following this, they were fully debriefed about the true nature of the experiment.

Results

The mean number of positives selected in the test phase was 6.25 (out of a possible of 10), with a standard deviation (SD) of 1.83. Comparison with chance performance of 5 (out of the possible 10), using a one-sample t-test, showed that subjects selected unseen positives significantly above chance, $t(19) = 2.98$ ($p < 0.05$). These results show that subjects had a bias towards positive test items. This is an indication that subjects have acquired the underlying invariant rule.

Ten subjects reported that they had not noticed a common feature or anything systematic about the 30 different sequences. Four subjects reported that the learning sequences generally seemed to rise in frequency from beginning to end (e.g. "built up a lot", "always going up"). Four subjects reported that the second tone in the learning phase never seemed to have been the highest, with a further subject reporting that the third tone seemed to have always been the highest. One subject reported that there seemed to have been some repetitions of sequences in the learning phase, but no awareness of any systematic feature was indicated. None of these subjects reported any awareness of the difference between the test pairs, except for one who suggested that they had not heard any of the test sequences previously. This particular subject also indicated, however, that the memory load in this task was very difficult and that they simply had not trusted their own memory and had not suspected any misleading by the experimenter. Out of all the subjects, five seemed to have had some awareness of a rise occurring in the relevant second interval, with a further four reporting a sensation of a rise in frequency throughout. These subjects may have been somewhat aware of a byproduct of the invariant 4^{th}: besides being defined as a five-halftone interval between two tones, it is also always a rise in frequency (i.e. from a lower tone to a higher tone). Although subjects may have used this information, they did not seem to have been aware of applying that knowledge, as can be seen from none of them reporting having noticed any difference in the paired test sequences.

When subjects were debriefed following the test phase, they were told that the 30 sequences in the first part of the experiment all contained a common feature, i.e. a 4^{th} in the second interval. Furthermore, they were informed that the test pair sequences had all been novel. They were also informed of the difference between the paired sequences (i.e. one containing the exemplar 4^{th} in the second interval, but not the other). All subjects expressed surprise at this and the responses they had made.

The pattern of responses found here was, therefore, not based on explicit information of the invariant feature, but was based on implicit knowledge.

Discussion

The results found here suggest that subjects have acquired the auditory invariant rule in the learning phase. This is shown in the selection of positives significantly above chance when compared to selection of negatives at test. However, some subjects seemed to have become generally aware of a frequency rise in the second interval. The question, then, is *what* knowledge subjects acquired and used to make their selection: the actual invariant 4^{th} or a co-variation to the invariant, i.e. a frequency rise between the second and third tone? In the test phase subjects were presented with negatives that could contain a rise, fall or unchanged level in frequency in the second interval. It is possible subjects employed knowledge other than that of the invariant 4^{th} itself in order to make their selection. This would have led them to select the positive sequence containing the invariant without actually having to use any knowledge of it. Co-varying information has been found to be a potential factor of learning effects in previous invariant learning studies (e.g. Churchill and Gilmore, 1998; Cock, Berry and Gaffan, 1994; Newell and Bright, 2002a/2002b; Wright and Burton, 1995; Wright and Whittlesea, 1998) and this possibility is discussed further in Experiment 9. Experiment 9 investigated whether selection of positives in Experiment 8 was based on knowledge of the intended invariant property.

4.5 Experiment 9

Subjects in Experiment 8 displayed a significant bias toward selecting positive sequences at test. This suggests that, contrary to failure of finding such knowledge acquisition with a single-tone invariant in Experiment 7, subjects in Experiment 8 acquired the invariant information when this consisted of a harmonic interval. However, a real question has arisen as to *what* knowledge subjects utilized in order to select the auditory sequence they thought they had heard previously. Cock, Berry and Gaffan (1994) suggested that subjects base their selection at test on a similarity strategy, rather than employing any invariant knowledge. In their study they used the same digits employed in the original McGeorge and Burton (1990) task, utilizing an invariant digit "3" in their study set. Importantly, they manipulated similarity of test to learning items independently of whether the item contained the invariant digit. Their study revealed that the important factor of whether a test item was classified as previously seen was the similarity between a test item and a learning item (the similarity here was based on the occurrence of repetitions within a digit string). Additionally, they found that when similarity was controlled for, there was only a small or no effect of the invariant at test. In context of Cock, Berry and Gaffan's (1994) findings, results in Experiment 8 point to the possibility that subjects may have used a similarity-based strategy at test. The similarity in the auditory experiment would have been based on the consistent presence of a rise in frequency in the second interval in this instance. Since a 4th automatically involves a rise in frequency from one tone to another, this provides additional information other than the invariant feature itself. It is a possibility that subjects used this co-varying information, and not the intended invariant 4th, in order to make their choice at test. Another possibility was suggested by Wright and

Burton (1995). Their study demonstrated that subjects did not have to learn an invariant rule (i.e. presence of the digit "3" in positives) in order to select positive items above chance at test. Their results showed that the effect could largely be explained in terms of rejection of particularly distinctive test items that were predominant in the test negatives (again based on repetitions within digit strings). In light of Wright and Burton's (1995) findings, subjects in Experiment 8 may have used such a rejection strategy of distinct sequences during the test phase. Negative items in the test phase could contain a rise, a fall or an unchanged level in frequency in the second interval. Except in the case of a rise in that position in a negative item, subjects may have rejected the negative item on grounds of a mismatch between a positive (which would always have contained a rise in the second interval) and a negative. This must be considered as clearly distinct from using knowledge of the invariant 4th itself, since subjects, in this case, would have been able to chose the positive by rejecting the negative. Several other studies have suggested use of a rejection strategy in an invariant learning context (e.g. Churchill and Gilmore, 1998; Newell and Bright, 2002a/2002b) pointing to the possibility of such a strategy use in the current study. Wright and Whittlesea (1998) found that subjects become sensitive to such co-varying information, and that they are directed to this knowledge by the distracter task. Although it is beyond the scope of the current experiment to test whether subjects used a similarity-based or rejection strategy, it investigated the possibility that subjects used other knowledge than the invariant rule in order to select positive sequences at test.

Constraining the negative test items to contain non-4th rises only in the second interval removed the possibility of using a similarity-based or rejection strategy. Thus, subjects would have to choose a positive item based on the intended invariant knowledge, otherwise leaving them to select randomly. The negative test items in Experiment 9 were constrained to contain only non-4th rises in the second interval, thus excluding the co-varying information from providing a relevant rule for selection of positives at test. Experiment 9 was a replication of Experiment 8 except for constraining the negative test set as laid out. If subjects acquired the underlying invariant knowledge (i.e. a 4th positioned in the second interval), they should show selection of positives significantly above chance when compared to negatives in a forced-choice recognition test.

Method

Subjects

20 subjects were recruited from the student population of the University of Glasgow. All received a small payment in return for their participation. All participants reported having normal hearing.

Materials

All material was generated and presented as in Experiment 8 and the same 17 tones were used.

The experimental material consisted of four-tone sequences, drawn randomly from the 17 tones above. As before, study sets of these were generated individually for each subject. Firstly, a set of 40 positive items was generated. As in Experiment 8 a positive is defined as a four-tone sequence that contains a 4^{th} in the second interval. Secondly, a set of ten negative items was generated, with a negative constrained to exclude any 4^{th} in the second interval. Additionally, the negatives were constrained so that they always had a non-4^{th} rise in frequency from second to third tone. All study sets were further restricted so that there were no repetitions of a particular four-tone sequence within a study set.

Design and Procedure

This was a within-subjects design, in which subjects were presented with 30 study items drawn from the set of 40 positive sequences, followed by ten pairs of sequences made-up of the remaining ten positive and ten negative sequences.

The procedure is an exact replication of the procedure in Experiment 8. Subjects were presented with 30 positive sequences and instructed to indicate which of the four tones in a given sequence had been highest in pitch. This was the distracter task as before. This was followed by an unexpected test phase, in which subjects were presented with ten pairs of sequences, each containing a negative and a positive sequence in random order. Neither of these two sequences had been heard previously in the learning phase. Subjects were falsely told that they had heard previously in the learning phase. Subjects were falsely told that they had heard one of these two sequences in the previous phase of the experiment. They were asked to decide which one they had heard before. They could replay a pair if they wished, or make their response immediately after first play. If they were unsure they were asked to guess.

This was a self-paced task and each subject took approximately 10 minutes to complete the entire task.

Once subjects had completed this task, they were asked whether they had noticed anything systematic about the 30 sequences in the first phase. They were then asked whether they had used any kind of strategy to decide which of the pair of sequences they had heard before. Following this, they were fully debriefed about the true nature of the experiment.

Results

The mean number of positives selected in the test phase was 5.5 (out of a possible 10), with a standard deviation (SD) of 1.7. A one-sample t-test compared this with chance performance of 5. This showed that selection of unseen positives was not significantly above chance, $t(19) = 1.32$ ($p > 0.05$). These results show that subjects did not exhibit a bias towards positive test items. This indicates that subjects failed to acquire any knowledge of the invariant feature.

Eight subjects reported that they had not noticed any common feature between the learning sequences. Four subjects reported having noticed a pattern in the learning sequences that followed a high-low-high-low order through the four-tone sequences. Three subjects indicated that the second tone had never been the highest and a further two subjects stated that the third tone had mostly been the highest tone in each sequence. One subject reported that there never seemed to have been any repeats at the beginning of a given sequence. The remaining two subjects indicated that the learning sequences seemed to contain something harmonic and seemed to have followed some particular order, but could not be more specific than this. These reports indicate that subjects may have been aware of the presence of the rise in the second interval in all sequences.

When subjects were debriefed following the test phase, they were told that the 30 sequences in the first part of the experiment all contained a common feature, i.e. a 4^{th} in the second interval. Furthermore, they were informed that all the sequences had been novel in the test phase and were told of the difference between the paired sequences (i.e. one containing the exemplar 4^{th}, but not the other). None of the subjects reported having been aware of the difference between each test pair. Overall, this indicates that subjects were not aware of the invariant feature or the true nature of this experiment.

These results indicate that subjects did not acquire any implicit or explicit knowledge of the auditory invariant feature.

Discussion

The results of this experiment indicate that subjects did not acquire any implicit or explicit knowledge of the intended invariant feature itself. Removal of the co-varying information, which would have allowed use of co-varying knowledge, left subjects selecting test sequences at random. These results clearly suggest that subjects in Experiment 8 used knowledge of the co-varying information in order to select positives at test. It is interesting to note that there was some indication that subjects may have acquired some explicit knowledge of the co-varying rise in frequency in the learning sequences. It is, therefore, a possibility that subjects used explicit, not implicit, knowledge in Experiment 8 in order to select positives at tests. The rise of potentially explicit knowledge may have been a byproduct of the distracter task used: the second tone in the learning sequences, by default, was never the highest tone. The consistent frequency rise in the second interval may have been made more noticeable by the nature of the distracter task. It may be the case that the sensitivity to the invariant rule subjects displayed in Experiment 8 was a direct consequence of the nature and demands of the distracter task. This is a possibility strongly argued by Wright and Whittlesea (1998). They demonstrated that subjects were able to discriminate between test sequences that did and did not conform to the underlying invariant rule on the bases of correlated information that went hand-in-hand with the intended invariant knowledge. Additionally, they found that the induction task (i.e. distracter task here) has a major role in directing what characteristics of the underlying rule structure are processed implicitly and, thus, learned.

Overall, it is quickly becoming evident that the auditory material employed here provides different processing demands to visual material in an invariant learning context. The results so far suggest that subjects are able to learn an auditory regularity (i.e. a rise in frequency in a specific interval). However, it appears that auditory qualities such as a specific tone or a specific interval are unavailable for implicit or explicit extraction by subjects. There is a question, however, over the suitability of the 4^{th} placed in the second interval as an invariant feature for implicit learning: grouping processes of the individual stimuli in the four-tone sequences may have run counter to the interval the 4^{th} was placed in. Thus, there is a possibility that subjects group tone sequences, like those used here, in ways that may enhance different parts of a given sequence more than others. In the music perception literature the phenomenon of perceptual grouping of tone sequences is a well-known occurrence. Perceptual grouping may occur in various contexts (such as grouping by temporal proximity, by timbre or by amplitude, Deutsch, 1999). It is worthwhile exploring further the issue of whether subjects can acquire an invariant 4^{th} specifically. This can be achieved by placing the harmonic 4^{th} in a different interval in the four-tone sequences, thus potentially enhancing its salience for implicit acquisition.

Experiment 10, then, expands the investigation of what auditory regularities and contexts may give rise to learning of an auditory invariant.

4.6 Experiment 10

Experiment 10 is a replication of Experiment 9 with the exception of moving the invariant feature, a 4^{th}, to lie in the first interval (previously in the second interval) of the four-tone sequences. This investigates whether placing a harmonic regularity in a different interval than in the previous two experiments has an effect on acquisition of that regularity. This allows further exploration of what kind of auditory invariant features subjects may be able to acquire implicitly. The current experiment constrained the negative test set to contain non-4^{th} rises in the first interval. Therefore, as in Experiment 9, it removed any co-varying information subjects may use to select positives at test.

The main question in the current experiment was whether subjects are able to acquire the invariant feature (i.e. presence of a 4^{th} between the first and second tone). Thus, if subjects show above chance selection of positives (those sequences containing a 4^{th} in the first interval) when compared to negatives (those sequences that do not contain a 4^{th} in the first interval) at test, subjects were deemed to have acquired the relevant invariant knowledge.

Method

Subjects

20 subjects were recruited from the student population of the University of Glasgow. All received a small payment in return for their participation. All participants reported having normal hearing.

Materials

All material was generated and presented as in Experiment 9 and the same 17 tones were used.

The experimental material consisted of four-tone sequences, drawn randomly from the 17 tones above. As before, study sets of these were generated individually for each subject. Firstly, a set of 40 positive items was generated. Here, a positive is defined as a four-tone sequence that contains a 4^{th} in the first interval. Secondly, a set of ten negative items was generated, with a negative constrained to exclude any 4^{th} in the first interval. Additionally, the negatives were constrained so that they always had a rise in frequency from second to third tone. All study sets were further restricted so that there were no repetitions of a particular four-tone sequence within a study set.

Design and Procedure

This was a within-subjects design, in which subjects were presented with 30 study items drawn from the set of 40 positive sequences, followed by ten pairs of sequences made-up of the remaining ten positive and ten negative sequences.

This was a replication of Experiment 9, except for the change of the invariant 4^{th} placed in the first interval. Subjects went through the same learning and test phases as in the previous experiment and were given the distracter task as before. This self-paced experiment took subjects approximately ten minutes to complete.

Subjects were then asked to state whether they had noticed anything systematic about the 30 sequences in the learning phase and whether they had found themselves using any kind of strategy in determining which sequence they had heard previously. They were then fully debriefed as to the true purpose of the experiment.

Results

The mean number of positives selected in the test phase was 5.35 (out of a possible 10), with a standard deviation (SD) of 1.63. A one-sample t-test compared this with chance performance of 5. This showed that selection of unseen positives was not significantly above chance, $t(19) = 0.36$ ($p > 0.05$). This suggests that subjects did not acquire the underlying invariant rule, but they selected test sequences randomly.

Eight subjects reported that the first tone had never been the highest tone in a given sequence, with a further stating that the second tone seemed to have been the highest mostly. A further seven subjects reported not having noticed any regular feature amongst the 30 sequences. One subject said that they had noticed repetitions within the sequences, with another subject reporting having noticed hardly any such repetitions. A further subject indicated the sequences had followed some kind of pattern, but could not be more specific. Lastly, one subject felt there had been something melodic about these sequences. None of these last 11 subjects reported having used any strategy in determining which sequence they had heard before. One subject reported a pattern of low-high-low-high for most sequences and used this knowledge in picking out the relevant sequence from the test pair. It appears that nearly half the subjects noticed that there was always a rise in the first interval. This was a natural byproduct of the invariant 4^{th}. However, this could not be used as the relevant knowledge for deciding which sequence in a test pair subjects thought they had heard previously, since the negative test sequences were constrained to contain rises in the first interval, making them equivalent to the positives in this respect.

When subjects were debriefed following the test phase, they were told that the 30 sequences in the first part of the experiment all contained a common feature, i.e. a 4^{th} in the first interval. Furthermore, they were informed of the novelty of all the sequences in the test pairs and the difference between these (i.e. one containing the exemplar 4^{th} in the first interval, but not the other). All subjects expressed surprise at the actual nature of the experiment and their own responses in light of this. This indicates that subjects were not aware of the invariant feature or the true nature of this experiment.

These results indicate that subjects did not acquire any implicit or explicit knowledge of the auditory invariant feature in this instance.

Discussion

Experiment 10 failed to demonstrate acquisition of the invariant harmonic 4^{th} positioned in the first interval. Additionally, the results did not indicate any explicit knowledge of the invariant feature. Subjects in the current experiment were prevented from using co-varying auditory information to make their selection at test, as had subjects in Experiment 9. The current results add to evidence from Experiment 9 that show that a harmonic relationship or specific interval between two tones (i.e. a 4^{th}) was unavailable for implicit or explicit extraction by subjects. This suggests that such a harmonic relationship is not a suitable auditory regularity for implicit learning in an invariant learning context. So far, the experiments in this study have explored what, if any, kind of auditory regularities subjects may acquire in an invariant learning paradigm. Experiment 8 had shown a bias towards such a regularity at test, suggesting that acquisition of an intended auditory feature is possible. However, both Experiments 9 and 10 indicate that the intended auditory invariant 4^{th} positioned in a specific interval in the learning sequences may not have been the actual information subjects employed to make their selection of positives at test in Experiment 8. The question remains what information subjects actually used in order to select positives at test in Experiment 8. In order to finalize these results and identify *what* knowledge subjects used for selection at test in Experiment 8, a final experiment (Experiment 11) was conducted.

4.7 Experiment 11

The invariant learning paradigm was originally based on the assumption that it provides evidence for implicit abstract rule learning. However, there is now increasing evidence that suggests performance in invariant learning tasks does not rely on subjects implicitly learning the invariant rule they were originally assumed to have acquired (Churchill and Gilmore, 1998; Cock, Berry and Gaffan, 1994; Newell and Bright, 2002a/2002b; Wright and Burton, 1995; Wright and Whittlesea, 1998). The current study adds to this view. Results from Experiment 8 brought up the question whether subjects could have used other information than that of the specified invariant feature in order to make their selection at test. This was a possibility, since the invariant 4th automatically incurred co-varying information besides the invariant rule. Subjects may have based their selection of positives at test in Experiment 8 on this information. Following failure to show above-chance selection of positives in Experiment 9, it is necessary to verify *what* knowledge subjects actually based their selection on in Experiment 8. For this reason, Experiment 11 investigated directly what knowledge subjects used for selection at test in Experiment 8. Subjects in the current experiment were given the same learning set of positive items as had been used in Experiment 8 – each learning sequence contained the auditory invariant 4th in the second interval. For the forced-choice recognition test, a variation on the positive and negative test pairs previously used was employed. Thus, subjects were presented with test pairs that contained the negative items from Experiment 8 paired with negative items from Experiment 9. This left subjects to select between pairs of sequences, neither of which actually contained the invariant 4th. However, negative items from Experiment 8 included a rise, fall or repetition in the second interval, whereas negative items from Experiment 9 were restricted to contain only rises in this position. Thus, at test subjects were forced to make a choice between pairs of items that contained a rise and those that contained a rise, fall or unchanged level in frequency in the second interval. In this case, no knowledge of the invariant rule itself could aid the selection process at test. If, however, subjects used knowledge of the presence of a rise in the second interval, they should select negative items from Experiment 9 (rises only) significantly more often than negative items from Experiment 8 (random), leaving selection of rises only items above chance. The question the current experiment investigated, then, was whether subjects used knowledge of the co-varying regularity in the second interval once they had been exposed to a positive learning set of sequences.

Method

Subjects

20 subjects were recruited from the student population of the University of Glasgow. All received a small payment in return for their participation. All participants reported having normal hearing.

Materials

All material was generated and presented as in the previous experiments. The same 17 tones as in Experiments 8 were used.

The experimental material consisted of the 30 learning sequences used in Experiment 8. All of these contained an invariant 4th in the second interval. A further ten sequences were those used in the negative set in Experiment 8. These were constrained to exclude any invariant 4^{th} in the second interval and contained rises, falls or an unchanged level in frequency at random in this interval. These will be known as *random* negative set. A further ten test sequences were taken directly from Experiment 9's negative set. These were constrained to exclude any invariant 4^{th} in the second interval and were also controlled to incorporate only rises in this interval. These will be known as the *rise-only* negative set. Thus, the two test sets in the current experiment varied from the previous experiments in that none contained the given invariant. All study sets were restricted so that there were no repetitions of a particular four-tone sequence within a study set.

Design and Procedure

This was a within-subjects design, in which subjects were presented with 30 study items containing the invariant, followed by ten pairs of sequences made-up of the remaining ten random and ten rise-only test sequences as stated above.

The experiment consisted of two stages: a learning phase and a test phase. In the learning phase subjects were presented with 30 sequences that contained the invariant. Subjects were told that this experiment would test their ability to distinguish tones they heard via headphones. They were informed that they would hear 30 four-tone sequences, one at a time, and that they had to decide which of the four tones had been the highest in pitch in a given sequence. They were instructed that they had to record their response by pressing 1, 2, 3 or 4 on the computer keyboard (each representing tone 1, tone 2, tone 3 and tone 4 in a sequence respectively). This served as the distracter task, making the true nature of this task unrecognizable to subjects. The learning phase was followed by an unexpected test phase. Subjects were given ten pairs of sequences, presented with a 1500ms pause between a pair. Each pair consisted of a positive and a negative sequence with the order of these randomized throughout the ten test trials. Neither of these two sequences had been heard in the learning phase previously and neither contained the invariant feature in the second interval. Subjects were falsely told that they had heard one of these two sequences in the previous phase of the experiment. They were asked to decide which one they had heard before by pressing '1' on the keyboard if it had been the first one, or '2' if it had been the second one of the pair. Subjects were permitted to replay a pair of sequences one time, or make their response immediately after first play. If they were unsure as to which sequence they had heard previously, they were asked to guess.

This was a self-paced task and each subject took approximately 10 minutes to complete the entire task.

Results

The mean number of rise-only sequences selected in the test phase was 5.75 (out of a possible 10), with a standard deviation (SD) 1.33. A one-sample t-test compared this with chance performance of 5. This showed that selection of sequences from the random set was significantly above chance, $t(19) = 2.46$ ($p < 0.05$).

This is evidence that subjects used knowledge of the rise in frequency in the second interval to aid selection of positives at test.

Six subjects reported not having noticed any regular feature and had not found they were using any strategy in picking out the relevant sequence. A further six stated that the third tone had mostly been the highest, with another three subjects reporting that the second tone had never been the highest. These nine subjects used this information as a strategy when possible in the test phase. Three subjects said that the 30 sequences generally seemed to rise from beginning to end and a further two subjects reported to have largely noticed a pattern of high-low-high-low throughout. These last five subjects reported to have used this in determining which sequence sounded more familiar in the test pair. When subjects were debriefed following the test phase, they were told that the 30 sequences in the first part of the experiment all contained a common feature, i.e. a 4th in the second interval. Furthermore, they were informed of the novelty of all the test pair sequences and the difference between the paired sequences (i.e. one containing rises in all sequences in the second interval, the other having rises/falls/unchanged level in frequency between these tones). Subjects were surprised at the presence of the invariant feature in the learning set. However, as their independent responses above show, it can be argued that subjects became aware of something that could be associated with the fact that the learning sequences all contained a rise from the second to the third tone. Responses indicating that the second tone was never the highest, the third was mostly the highest, the pattern of the sequences was high-low-high-low from tone to tone and even sequences generally rising from beginning to end, can all be associated with some level of awareness of the presence of the incurred rises in the second interval. It is possible, therefore, that subjects used some explicit knowledge in making their selection at test.

Discussion

The results found here indicate that subjects acquired co-varying information with the original invariant feature, i.e. that all positive items contained a rise in frequency in the second interval. By extension, these results indicate *what* knowledge subjects in Experiment 8 actually used to select positives at test. When subjects were prevented from using any knowledge of the incurred rise in the second interval in Experiment 9, they did not show any bias towards selection of positive test items, indicating failure to learn the invariant rule. The current experiment provides a clear answer to the question of *what* knowledge subjects used to select positives in Experiment 8. In this case subjects used information that co-varied with the original auditory feature: a rise in frequency in the second interval. These results add to the view that subjects may employ a strategy other than the invariant rule in an invariant learning task (Churchill and Gilmore, 1998; Cock, Berry and Gaffan, 1994; Newell and Bright, 2002a/2002b;

Wright and Burton, 1995; Wright and Whittlesea, 1998), and extend them to an auditory context.

Although the focus of the current experiment was not whether the acquired knowledge was implicit or explicit, the tentative results from questioning subjects following the experiment indicated that a large number of subjects became aware of something that can be associated with the presence of a rise in the relevant invariant interval in all learning sequences.

Overall, this is indication that subjects are able to acquire an auditory regularity such as a rise in frequency, but that this may have occurred through explicit processes in the current context.

4.8 General Discussion

The purpose of the experiments in this chapter was to explore whether subjects could learn an invariant rule when the stimuli they were exposed to were purely auditory. Furthermore, it investigated what subjects are and are not able to learn in an auditory invariant learning task. Results from this study have shown that, in the current context, subjects are unable to learn an auditory invariant feature that consists of a single tone or a harmonic interval between two tones. Nevertheless, the series of experiments 8, 9 and 11 clearly indicate that subjects can acquire an auditory regularity that consists of a rise in frequency from one tone to another. This reveals, for the first time, that auditory stimuli can provide regularities available for acquisition in an invariant learning context.

A question that always emerges in implicit learning studies is whether the information underlying learning is consciously or unconsciously mediated (e.g. Neal and Hesketh, 1997; Shanks and St. John, 1994). The experiments in this chapter did not attempt to answer this question directly. Nevertheless, tentative results from post task questioning indicate that the auditory invariant regularities were not available for implicit or explicit acquisition. However, as has already been argued, the auditory regularity subjects were shown to have acquired in Experiments 8 and 11 may have been based on explicit processes. Several invariant learning studies have found that at least some knowledge used by subjects to aide selection of positives at test is explicit in nature (Churchill and Gilmore, 1998; Newell and Bright, 2002a/2002b; Wright and Burton, 1995). The current study is in line with this view. Subjects did not only acquire the co-varying regularity, but several subjects were able to express its presence verbally. This is an indication that the acquired knowledge in the current experiments may have been explicit. Churchill and Gilmore (1998) argued that the co-varying feature is not available to verbal report, but provides a selection rule at test, unless the nature of the learning stimuli is changed in order to enhance reporting of explicit knowledge of the co-varying feature. In the current context, the nature of the stimuli appears to have been such that several subjects spontaneously reported explicit knowledge of the co-varying feature. This may have been a confound of the distracter task used. Subjects were asked to indicate which tone in a given four-tone sequence had been the highest. This, most likely, highlighted the fact that the second tone had never been the highest, thus drawing attention to the presence of a rise in frequency in the invariant's position. This is a view supported by Wright and Whittlesea (1998). Their study showed that subjects acquire information about the learning stimuli and, importantly, that this is dictated by the distracter task (termed "induction task" in their study). The distracter task used here was the most suitable found in conjunction with the current auditory stimulus material. In future, it would be useful to find a distracter task that did not draw attention to the presence of the particular co-varying feature found here, thus controlling for potential contamination by explicit knowledge.

Previous invariant learning studies had assumed a residual role for the invariant feature in providing a rule for selection of positive test items (Churchill and Gilmore, 1998; Cock, Berry and Gaffan, 1994; Newell and Bright, 2002a; Wright and Burton, 1995). This was not found to be the case with the auditory sequences used here as clearly indicated by the results from Experiment 8, 9 and 11. This is an indication that auditory stimuli provide different processing demands to visual material in the current context.

The current research would benefit from extending its investigation into auditory material that can be learned in an invariant learning context. It would be of interest to find other auditory qualities, for example different instruments, which may provide a suitable auditory invariant feature.

In conclusion, we now have a clearer idea about what auditory regularities subjects can and cannot learn in an auditory invariant learning context. Thus, this study extends previous findings in the invariant learning literature to other experimental contexts. Most importantly, the current study has shown that auditory stimuli can be used to investigate invariant learning. This makes a useful contribution to invariant learning research, both in exploring issues

particular to auditory invariant learning, as well as to invariant learning in general.

Chapter 5. Summary and Conclusions

This chapter will provide a summary of the results found in the experimental chapters. This is followed by the conclusions that can be drawn from the empirical findings with regards to the theoretical considerations outlined in Chapter 1.

Chapter 2

In Chapter 2 a series of experiments was conducted that employed an SRT task. The overall aim was to establish whether subjects could acquire an auditory sequence of tones by listening, without a motor response tied to the sequence exposure. The findings indicated that subjects were able to learn an auditory sequence when a spatial motor response (keypresses) was tied to the sequence exposure, but no such learning was shown for a non-spatial response (voice response) or by listening alone. Overall, the results from the experiments in Chapter 2 indicated that subjects were unable to acquire an auditory sequence of tones by listening alone, but that learning required a spatial motor response.

Chapter 3

Chapter 3 investigated whether auditory stimuli could give rise to learning without overt responses tied to the sequence exposure in a different SRT task to Chapter 2. This replicated a design by Mayr (1996) utilising auditory material in place of Mayr's (1996) original visual and visuo-spatial stimuli. The results from this study indicated that auditory stimulus sequences were not available for acquisition when the stimuli were not relevant to a motor response. Thus, subjects were able to acquire a verbal auditory sequence, as well as a sequence of voices, when they had to respond to these with a motor response (keypresses). However, no such learning was found when the sequence exposure essentially consisted of listening. This implied that the underlying sequence structure remained unprocessed, unless a motor response was involved. These results added to findings from visual sequence learning studies that had failed to show learning of visual stimuli unless these were presented visuo-spatially or with corresponding motor response (e.g. Mayr, 1996; Willingham, 1999).

Chapter 4

In Chapter 4 a series of experiments investigated learning of auditory invariant features. The overall aim was to establish whether subjects could learn an auditory invariant per se. Overall, the results from the experiments in Chapter 4 provided evidence of what type of auditory features can and cannot be acquired in an invariant learning context. The findings indicated that subjects are able to learn an auditory invariant consisting of a rise in frequency from one tone to another, but no learning of a single tone or a specific harmonic relationship between two tones in this context could be shown. These results extended previous findings from the visual domain to an auditory invariant learning context.

A variety of definitions of implicit learning are currently offered and this can be seen in the different ways in which such learning has been operationlized in the laboratory (e.g. Cleeremans et al, 1998; Frensch, 1998). This diversity has been described as symptomatic of the conceptual and methodological problems of implicit learning studies in general (e.g. Cleeremans et al, 1998). The current research investigated implicit learning using auditory stimuli empirically. In light of the lack of a unitary definition and diversity of methodologies, the methods used here employed some of the main experimental tasks from previous implicit learning research in the visual domain. Thus, the tasks used throughout this thesis reflect the diversity of methodologies found in implicit learning in general. Despite the lack of a unitary definition of implicit learning, the tasks employed reflected the main attributes such learning has been associated with: 1) developing a sensitivity to the structural organisation of the stimulus domain, 2) incidental training conditions and 3) difficulty to express the acquired knowledge verbally.

Internalising the regularities that occur in our external environment plays an important role in our everyday lives. Acquisition of implicit knowledge has been described as an unselective and automatic accumulation of such associated information (e.g. Berry and Broadbent, 1988; Reber, 1989). Furthermore, such learning has been described as a fundamental cognitive process, underlying a variety of complex information acquisition (Reber, 1993). Buchner, Steffens, Erdfelder and Rothkegel (1997) suggested that it is necessary to broaden the scope of implicit learning research to stimulus domains other than the visual, in order to provide empirical evidence to the claim that it is, in fact, a fundamental cognitive process. Auditory material allows widening of the current scope of learning without awareness research. There have been few studies that have used auditory material in implicit learning research and these did not provide consistent results (e.g. Buchner et al, 1997; Perruchet, Bigand and Benoit-Gonin, 1997). The experimental research in this thesis explored whether implicit learning of sound stimuli behaves differently from such learning in the visual domain. In order to do so, the different experimental tasks investigated specific questions that have arisen in visual implicit learning research. In this way, the current research extends previous findings from the visual domain to a wider context.

One of the main overall conclusions that can be drawn from the findings in this study concerns the question of whether implicit learning is a fundamental and general cognitive process (e.g. Reber, 1989). The role of rule abstraction in implicit learning tasks is at the heart of this debate. Early accounts of implicit learning suggested that implicit learning is a powerful and unconscious mechanism that allows developing a sensitivity to a set of stimuli unintentionally. Importantly, this process has been described as capable of abstracting rules that describe the underlying structure of stimulus domains (e.g. Reber, 1967; Reber, 1976). Thus, it has been suggested that if subjects are able to acquire the underlying rule structure this would be indication of more central cognitive processes. In contrast, if learning is bound to more peripheral processes, such learning would consist of learning of stimulus surface features or a sequence of motor responses (e.g. Newell and Bright, 2002b; Perruchet and Pacteau, 1990; Seger, 1998). The evidence from Chapter 2 indicates that subjects are unable to acquire the underlying sequence structure when they simply listened to a regular auditory sequence, but learning occurred when they responded to the same sequence with keypress responses. However, the use of voice responses did not provide clear-cut results and it was suggested that this type of motor response may have lacked sensitivity to capture any sequence learning effects. Reproducing a tone by singing it may not be a natural response, thus making it difficult for subjects to perform. This potential difficulty may have added unaccounted noise to the reaction time data, which may have weakened or even eliminated any effects. Findings from Chapter 3 provided a clear answer to the question whether subjects acquired the underlying rules or whether the learning was bound to more peripheral motor response learning. Subjects in Experiment 6 were shown to learn auditory sequences of voices or colour names when these were relevant for motor responses, but they were unable to learn the same sequences without such motor responses tied to the sequence exposure. This adds to evidence from visual implicit learning studies that have failed to find learning of event sequences when spatial or response selection was not an important factor in processing these (e.g. Mayr, 1996; Willingham, 1999). The findings from the current research show that implicit learning in the current context cannot be described as bound to general cognitive processes, since subjects acquired a set of actions and did not acquire an internal representation of the underlying rules. The RT and awareness data from Chapter 3 add emphasis to this claim. These suggested that the sequence of colour names was particularly salient, since only subjects responding to this sequence acquired explicit knowledge of it. Overall, subjects showed some awareness of the sequence they responded-to, but not the one they did not, suggesting that learning here was bound to explicit processes. The claim that implicit knowledge acquisition is unselective and automatic runs counter to these findings (Berry and Broadbent, 1988; Reber, 1989). It seems what subjects learn are the relevant associations between their actions and a set of stimuli, but not the underlying rules. For this reason, the findings from Chapters 2 and 3 provide evidence that implicit learning cannot be a general cognitive process that involves internal representations, but that learning here involves something that is modality specific, since it is tied to actions alone. It appears that the implicit learning of auditory event sequences behaves in a similar fashion to visual sequence learning. That is, the stimulus domain itself is irrelevant, as long as motor responses are tied to the sequence exposure. Overall, these findings add to evidence from studies that have put into question that implicit learning is solely based on implicit rule abstraction (e.g. Gomez and Schvaneveldt, 1994; Shanks and Johnston, 1998).

In order to broaden the experimental context of Chapters' 2 and 3, Chapter 4 explored implicit learning using a different experimental task than the SRT task.

The findings from Chapter 4 showed that subjects were able to acquire an auditory invariant feature. However, clear differences to previous invariant learning studies were shown when subjects were unable to learn a specific tone in an invariant context. This was in contrast to the original invariant learning experiment by McGeorge and Burton (1990) who found learning of an individual digit (i.e. "3"). Previous invariant learning studies suggested a role for co-varying information, rather than the intended invariant feature, as relevant for subjects' performance (e.g. Newell and Bright, 2002a; Cock et al, 1994). However, until now the exact knowledge that subjects acquired eluded researchers. The current research, for the first time, identified exactly *what* the invariant rule subjects learned was in the current auditory context. Hence, this research has shown what auditory regularities subjects could and could not acquire, extending findings from the visual domain and making an original contribution to invariant learning research in general.

An additional issue concerns levels of awareness in the current context. Subjects in Experiment 11 seemed to have acquired some explicit knowledge along with the unintended auditory invariant feature. This puts into question whether such learning would have proceeded without awareness. Overall, the results provide evidence that translating visual invariant learning tasks to the auditory domain may not replicate results found in a visual context directly. However, auditory stimulus presentation incurs different constraints to that of visual material and it may be this that was responsible for the apparent differences in findings between the two domains.

Considering the findings from both the SRT and invariant learning experiments in this thesis, the evidence runs counter to the idea of implicit rule abstraction in these tasks. Evidence from Chapters 2 and 3 failed to show a role for rule abstraction in a sequence learning context. Results from Chapter 4 showed a potential involvement of explicit processes in the acquisition of an auditory invariant. Thus, an explanation of implicit learning based on rule abstraction or processes excluding conscious awareness is put into question.

There is a rising interest in investigating temporal processing across skills (e.g. Salidis, 2001). This was born out of an interest in music cognition as an example of complex cognitive processes, like language learning and its use. Music perception relies on temporal cognition by nature, as time has a fundamental role in music (e.g. rhythm, length of notes, pauses, etc., Krumhansl, 2000). Salidis (2001) found implicit learning of rhythms in a study investigating learning of temporal regularities. The close relationship between music perception and auditory processes in general may make investigating implicit learning of temporal relations an interesting topic in future. Temporal aspects appear of particular importance to auditory processing. However, temporal qualities may span across different stimulus domains and it may be the relative importance of these processes that will prove to be involved in potential fundamental implicit learning processes. Future investigations need to develop tasks that assess the relative contributions of these types of aspects more closely. This may provide new insights into implicit learning in general and the possibility of some aspects of such learning as fundamental, spanning both the visual and auditory domains.

Hebb (1961) asked whether a sequence of stimuli could leave a permanent trace in the nervous system. He investigated this question using a digit span task. This task presented subjects with randomly ordered lists of digits they had to repeat back to the experimenter. Importantly, every third list of digits was the same, without subjects being made aware of this. Hebb (1961) found that subjects' recall of the repeated lists of digits improved at a greater rate than for random lists. These results suggested that even very briefly retained information could leave a relatively permanent trace in memory. Stadler (1993) pointed to similarities in this pattern of results to findings in implicit learning research in general, and serial learning tasks in particular. Here, the main difference in tasks lies in the continuous repetition of a repeating sequence. Stadler (1993) applied Hebb's (1961) basic design to a visual sequence learning task. Subjects were asked to respond to a sequence of stimuli that were not continuously repeated, but consisted of random sequences interspersed with a recurring repeating sequence. Stadler's (1993) results suggested that subjects do learn such intermittent sequences. It is easy to see a relationship between the simple auditory stimuli employed in this thesis and more complex stimuli of musical content. It would be of interest for future research to consider more complex auditory stimulus material, such as tunes, rather than the item-based material employed throughout this thesis. Using a methodology that incorporates aspects of Hebb's (1961) digit span task may provide a way forward to investigate the possibility of implicit learning of musical structures.

In conclusion, the research in this thesis adds empirical evidence to questions that have arisen in implicit learning research in general. By employing exclusively auditory material the scope of implicit learning research was broadened systematically to stimulus contexts other than the visual domain. Thus, it has been shown that auditory stimuli can be utilized successfully to investigate implicit learning.

References

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.

Aslin, R.N., Saffran, J.R. & Newport, E.L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science*, **9**, 321 – 324.

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

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

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

Bigand, E., Perruchet, P. & Boyer, M. (1998). Implicit learning of an artificial grammar of musical timbres. *Cahiers de Psychologie Cognitive/Current Psychology of Cognition*, **17**(3), 577 – 600.

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

Buchner, A. & Frensch, P.A. (1997). Sequence learning. Phenomena and models. *Psychological Research*, **60**, 1 – 3.

Buchner, A. & Steffens, M.C. (2001). Simultaneous learning of different regularities in sequence learning tasks: limits and characteristics. *Psychological Research*, **65**, 71 – 80.

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

Buchner, A., Steffens, M.C. & Rothkegel, R. (1998). On the role of fragmentary knowledge in a sequence learning task. *Quarterly Journal of Experimental Psychology*, **51A**, 251 – 281.

Burns, E.M. (1999). Intervals, scales and tuning. In D. Deutsch (Ed.) *The psychology of music (2nd Edition)*, USA, Academic Press.

Chandler, S. (1993). Are rules and modules really necessary for explaining language? *Journal of Psycholinguistic Research*, **22**, 593 – 606.

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

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

Cleeremans, A. (1993). *Mechanisms of implicit learning: A connectionist model of sequence processing*. Cambridge, MA: MIT Press/Bradford Books.

Cleeremans, A. (1997). Principles for implicit learning. In D.C. Berry (Ed.) *How implicit is implicit learning?* Oxford: Oxford University Press.

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

Cleeremans, A. & Jimenez, L. (2002). Implicit learning and consciousness: A graded, dynamic perspective. In R.M. French & A. Cleeremans (Eds.) *Implicit learning and consciousness*, Hove, UK, Psychology Press.

Cock, J., Berry, D. & Gaffan, E. (1994). The role of similarity processing in an incidental learning task. *Quarterly 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 – 30.

Destrebecqz, A. & Cleeremans, A. (2001). Can sequence learning be implicit? New evidence with the process dissociation procedure. *Psychonomic Bulletin and Review*, **8**(2), 343 – 350.

Deutsch, D. (1999). Grouping mechanisms in music. In D. Deutsch (Ed.) *The psychology of music (2nd Edition)*, USA, Academic Press.

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

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

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

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

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

Frensch, P.A. (1998). One concept, multiple meanings. In M. Stadler & P. Frensch (Eds.) *Handbook of implicit learning.* Thousand Oaks, CA: Sage.

Frensch, P.A. & Miner, C.S. (1995). The role of working memory in implicit sequence learning. *Zeitschrift fuer Experimentelle Psychologie,* **42**, 545 – 575.

Frensch, P.A. & Rünger, D. (2003). Implicit learning. *Current directions in Psychological Science,* **12**(1), 13 – 18.

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

Goschke, T. (1998). Implicit learning of perceptual and motor sequences: evidence for independent learning systems. In M. Stadler & P. Frensch (Eds.) *Handbook of implicit learning.* Thousand Oaks, CA: Sage.

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.T. (1979). Automatic and effortful processes in memory. *Journal of Experimental Psychology: General,* **108**, 356 – 388.

Hebb, D.O. (1961). Distinctive features of learning in the higher animal. In J.F. Delafresnaye (Ed.), *Bain Mechanisms and Learning* (pp. 37 – 46). New York: Oxford University Press.

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

Howard, J.H., Mutter, S.A. & Howard, D.V. (1992). Serial pattern learning by event observation. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **18(5)**, 109 – 1039.

Hsiao, A.T. & Reber, A.S. (1998). The role of attention in sequence learning: exploring the limits of the cognitive unconscious. In M. Stadler & P. Frensch (Eds.) *Handbook of implicit learning*, Thousand Oaks, CA: Sage.

Huddy, V. & Burton, A.M. (2002). Generate and test: an alternative route to knowledge elicitation in an implicit learning task. *Quarterly Journal of Experimental Psychology*, **55A**(4), 1093 – 1107.

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

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

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

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

Kelly, S.W. & Burton, A.M. (2001). Learning complex sequences by observation? *Psychological Research*, **65**, 15 – 23.

Kelly, S.W., Burton, A.M., Kato, T. & Akamatsu, S. (2001). Incidental learning of real-world regularities. *Psychological Science*, **12**(1), 86 – 89.

Kelly, S.W., Burton, A.M., Riedel, B & Lynch, E. (2003). Sequence learning by observation: evidence for separate mechanisms. *British Journal of Psychology*, **94** (3), 355 - 372.

Knowlton, B.K. & Squire, L.R. (1996). Artificial grammar learning depends on implicit acquisition of both abstract and exemplar-specific information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **22**, 169 – 181.

Krumhansl, C.L. (1990). *Cognitive foundations of musical pitch*, New York, Oxford University Press.

Krumhansl, C.L. (2000). Rhythm and pitch in music cognition. *Psychological Bulletin*, **126**, 159 - 179.

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

Lewicki, P., Czyzewska, M. & Hoffman, H. (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **12**, 135 – 146.

Lieberman, D.A. (2000). *Learning: Behavior and Cognition* (3rd ed.). Pacific Grove, California: Brooks/Cole.

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

Logan, G.D. (1990). Automaticity and memory. *Bulletin of the Psychonomic Society*, **28**(6): 486 - 486.

Manza, L. & Reber, A.S. (1997). Representing artificial grammars: Transfer across stimulus forms and modalities. In D.C. Berry (Ed.) *How implicit is implicit learning?* Oxford, UK: Oxford University Press.

Marcel, A. (1984). Conscious and unconscious perception: An approach to the relation between phenomenal experience and perceptual process. *Cognitive Psychology*, **15**, 238 – 300.

Mathews, R., Buss, R., Stanley, W., Blanchard-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*, **16**(6), 1083 – 1100.

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

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

Merikle, P. M. (1992). Perception without awareness: Critical issues. *American Psychologist*, **47**(6), 792 - 795.

Nattkemper, D. & Prinz, W. (1997). Stimulus and response anticipation in a serial reaction task. *Psychological Research*, **60**, 98 – 112.

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

Newell, B.R. & Bright, J.E.H. (2002a). Evidence against hyperspecificity in implicit invariant learning. *The Quarterly Journal of Experimental Psychology*, **55A**(4),
1109 – 1126.

Newell, B.R. & Bright, J.E.H. (2002b). Well past midnight: Calling time on implicit invariant learning? *European Journal of Cognitive Psychology*, **14**(2), 185 – 205.

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

Perruchet, P., Bigand, E. & Benoit-Gonin, F. (1997). The emergence of explicit knowledge during the early phase of learning in sequential reaction time tasks. *Psychological Research*, **60**, 87 – 97.

Perruchet, P. & Gallego, J. (1997). A subjective unit formation account of implicit learning. In M. Stadler & P. Frensch (Eds.) *Handbook of implicit learning*, Thousand Oaks, CA: Sage.

Perruchet, P. & Pacteau, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General*, **19**, 264 – 275.

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

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

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

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

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

Redington, M. & Chater, N. (1997). Probabilistic and distributional approaches to language acquisition. *Trends in Cognitive Sciences*, **1**, 273 – 281.

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

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

Saffran, J.R., Johnson, E.K., Aslin, R.N. & Newport, E.L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, **70**, 27 – 52.

Saffran, J.R., Newport, E.L, Aslin, R.N., Tunick, R.A. & Barrueco, S. (1997). Incidental language learning: Listening (and learning) out of the corner of our ear. *Psychological Science*, **8**, 101 – 195.

Salidis, J (2001). Nonconscious temporal cognition: Learning rhythms implicitly. *Memory & Cognition*, **29**(8), 1111 - 1119.

Schmidtke, V & Heuer, H. (1997). Task integration as a factor in secondary-task effects on sequence learning. *Psychological Research*, **60**, 53 – 71.

Seger, C.A. (1997). Two forms of sequential implicit learning. *Consciousness and Cognition*, **6**, 108 – 131.

Seger, C.A. (1998). Multiple forms of implicit learning. In M. Stadler & P. Frensch (Eds.) *Handbook of implicit learning*. Thousand Oaks, CA: Sage.

Shanks, D.R. & Johnston, T. (1999). Evaluating the relationship between explicit and implicit knowledge in a sequential reaction time task. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **25**(6), 1435 - 1451.

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

Stadler, M. (1993). Implicit serial learning: Questions inspired by Hebb (1961). *Memory & Cognition*, **21**(6), 819 – 827.

Stadler, M. & Frensch, P. (1994). Whither learning, whither memory? *Behavioral and Brain Sciences*, **17**, 423 – 424.

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

Thorndike, E.L. (1911). *Animal Intelligence: Experimental studies*. New York: Mcmillan.

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

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

Tillmann, B., Bharucha, J.J. & Bigand, E. (2000). Implicit learning of tonality: A self-organizing approach. *Psychological Review*, **107**(4), 885 – 913.

Underwood, G. & Bright, J.E.H. (1996). Cognition with and without awareness. In G. Underwood (Ed.) *Implicit Cognition*, Oxford, UK: Oxford University Press.

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

Warren, R.M. (1993). Perception of acoustic sequences: hearing in complex environments. In S. McAdams & E. Bigand (Eds.) *Thinking in sound. The cognitive psychology of human audition*, Oxford, UK: Oxford University Press.

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

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

Willingham, D.B., Wells, L.A. & Farrel, J.M. (2000). Implicit motor sequence learning is represented in response locations. *Memory and Cognition*, **28(3)**, 366 – 375.

Wright, R. & Burton, A. (1995). Implicit learning of an invariant: Just say no. *The Quarterly Journal of Experimental Psychology*, **48A**(3), 783 – 796.

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

Ziessler, M. (1994). The impact of motor responses on serial pattern learning. *Psychological Research*, **57**, 30 – 41.

APPENDIX A.1

Experiment 1. Direct test of awareness, 24 chunks (A, B, C and D corresponding to each tone from low to high respectively):

Learning Sequence (Blocks 1 to 8 and 10) –	Test Sequence (Block 9) –
ABDACB	ABCADC
ACBADB	ACDABC
ADBCDC	ADCBDB
BADBCD	BACDAB
BCDCAB	BCADCB
BDACBA	BDBACD
CABDAC	CADCBD
CBADBC	CBDBAC
CDCABD	CDABCA
DACBAD	DABCAD
DBCDCA	DBACDA
DCABDA	DCBDBA

APPENDIX A.2

Experiment 1. Rating scale for direct test of awareness:

1 – I am absolutely certain the chunk was new.
2 – I think the chunk was new.
3 – I am guessing the chunk was new.
4 – I am guessing the chunk was old.
5 – I think the chunk was old.
6 – I am absolutely certain the chunk was old.

APPENDIX B

Experiments 2 and 3. Set of maths calculations given to subjects in order to distract attention away from auditory input:

7+2+8=	7-6+8=	4+9-3=
9-5+7=	2x9-2=	3x5+6=
8:2+5=	4x5-3=	8-2+5=
9-3+7=	9:3+3=	7-1+9=
3-5+9=	5-4+3=	5+7-3=
1+7x2=	2+1x3=	4+9+2=
3x5+7=	9-6+2=	8-6-5=
2+5+3=	9+9+5=	3x4+7=
6-3-2=	5-2+9=	7+3x3=
8+7x2=	7x2+6=	2x2-1=
3x3x2=	4-2+7=	1+1x9=
1+4x3=	8:4x9=	7-5+9=
4+9+7=	9+1x7=	3-1+7=
8-2-5=	3+4+5=	1x5x3=
4-5+9=	9-5- 4=	5+9-2=
2x5-2=	7x3-8=	3+7+5=
6:2x3=	6+4-8=	9-6-2=
3x8-4=	9:3+6=	4+9-3=
3x2x3=	5+7+9=	6+6x2=
7+4+9=	3-2-1=	2x3-4=
3+9-2=	8+4+7=	5-2+9=
4x4-6=	8+5-3=	9:3+1=
9+2+9=	6x2+8=	9+9-2=
5x5+7=	4x4+8=	4x4-1=
(9x2):6=	6:2x3+	2+7+4=
(7+5):2=	8+9-7=	8-2-5=
5x3+8=	1x5x6=	5x7+1=
2+2+4=	7-1+7=	9-3x2=
7-1-4=	4-3+4=	5+7x9=
5-2+7=	7-2+8=	3x(4x5)=
8+9+7=	9-1-3=	5-2+9=

(3x8):2=
7x5-7=
(9+7):4=
4+7+3=
6-2+9=
6x3+2=
9+8:4=
7x3+9=
2x4x3=
2x(9+6)=
7-3+6=
4+9-3=
1+7x2=
3x4-7=
8+1x9=
6x2+7=
9x9-5=
5+8+3=
2x(6:3)=
6+7-3=
(9-6)x4=
8-3+4=
9-1-4=
4+2+8=
3x3x2=
9x1-7=
(9+1)x5=
8+9+5=
9x(2+3)=
2+5-1=
(1+1)x8=
6+7-3=
(7+8):3=

9x4x2=
2x(9-5)=
4x4x4=
9:3+8=
2x3x4=
3x(9-4)=
8-3+9=
5x5+9=
9x1-6=
1x1+1=
4+2x4=
2+2x3=
8+8-5=
3-5+7=
4+3-9=
3x(9-6)=
(9+9):3=
7+3+2=
9-3-1=
5x2-7=
8:4x5=
(2+4)x7=
9+4+9=
7+9-3=
7x7+3=
5x9-8=
4x(3+9)=
1+8+6=
4+4+4=
7-2x3=
9x(7-4)=
8x3-2=
5+4+9=

8+3-9=
6+7x2=
3x9+7=
8-4+3=
5x5+6=
2x4-1=
4x6-3=
5+5+3=
9x2+7=
3-5+9=
7+8-3=
4x7-3=
9:3+5=
1x2x3=
9-3-7=
8+5-2=
2x3x4=
4x(8-5)=
(7+5):2=
2-5+7=
8x5-6=
1+4+7=
7-3+8=
7x(9+1)=
8x3+3=
3x3x4=
8-6-1=
9+3-7=
(7+4)x9=
3x(7-4)=
(7+8):3=
4+4x3=
7+8+9=