Incidental sequence learning in humans: Predictions of an associative account

Submitted by Fayme Yeates to the
University of Exeter as a thesis for the degree of
Doctor of Philosophy in Psychology in
August 2014

Abstract

This thesis aims to investigate how well associative learning can account for human sequence learning under incidental conditions. It seems that we can learn complex sequential information about events in our environment, for example language or music, incidentally, without being aware of it. Awareness is, however, a complex issue with arguments for (Dienes, 2012) and against (Shanks, 2005) the existence of implicit learning processes. A dual process account proposes that there exist two different learning systems, one based on conscious, controlled reasoning and rules, and the other based on automatic association formation, which can take place outside of awareness (McLaren, Green, & Mackintosh, 1994). This thesis attempts to use the predictions of an associative account in conjunction with a suitable method for investigating implicit learning: sequence learning (Destrebecqz & Cleeremans, 2003). The research involves a collection of serial reaction time (SRT) tasks whereby participants respond to on-screen stimuli that follow a sequence that they were (intentional learning) or were not (incidental learning) informed of. Following on from the experimental design of Jones and McLaren (2009) this thesis provides evidence that humans differ in their ability to learn different sequential contingencies. After training sequences of trials where the current trial location was twice as likely to be either: the same as (Same rule); or different to (Different rule) the location two trials before this, participants were far better at learning the latter rule. I found that this result was not adequately simulated by the benchmark associative model of sequence learning, the Augmented SRN (Cleeremans & McClelland, 1991), and present a revised model. This model, amongst other attributes, represents all the stimuli experienced by participants and can therefore learn stimulus-response contingencies. These seem to block learning (to some extent) about the Same rule thus providing an associative explanation of the advantage for acquisition of the Different rule. Further predictions regarding the role of additional stimuli alongside sequence learning were then derived from this associative account and tested on human participants. The first of these was that additional stimuli within the task will interact with sequence learning. I found that human participants show increased Same rule learning when additional, concurrently presented stimuli follow the previous element in the sequence. I demonstrate that when participants perform an SRT task where responses are predicted by the colour of a cue, they are able to learn about this relationship in the absence of awareness. Using this cue-response learning I further

investigate cue-competition between sequences and colours under incidental conditions and find evidence that suggests between cue associations may alter the influence of cue competition. These results altogether suggest that stimuli – both simple and sequential – can be learned under incidental conditions. This thesis further proposes that learning about simple and more complex relationships between stimuli interacts according to the predictions of an associative account and provides evidence that contributes to a dual process understanding of human learning.

Acknowledgements

I would first like to thank the contribution of my supervisor Ian McLaren for his enthusiasm, encouragement and support. Ian is able to inspire extraordinary excitement for what is to be discovered and instill the confidence needed to pursue it. I would also like to thank most deeply the contributions of Fergal Jones and Andy Wills, who have been incredible sounding boards, collaborators and confidants without any legal obligation to do so and I have learned much from them both. I am extremely grateful to the University of Exeter for funding my PhD, as well as a fantastic research environment and a host of opportunities. I would like to thank my examiners: Aureliu Lavric and Zoltan Dienes for their insightful comments and suggestions regarding this thesis.

I would also like to thank my colleagues and collaborates in the department at Exeter, who have all contributed directly or indirectly to this thesis. Specific thanks go to Charlotte Forrest, who I thank for taking me under her wing for the first half of my PhD. The final stages of my PhD have largely involved terrifying Maisy Best about what is to come, and I want to thank her specifically for the invaluable care and advice she has provided me with. I want to thank Amy McAndrew for her emotional and academic support and friendship across work, sports, cockroach infested conference hostels and more. I would also like to especially acknowledge Rossy McLaren, Will Bowditch, Ciro Civile, Christine Lees, Alison Metcalfe, and my netball teams, especially the Net Prophets, for their friendship and support, which I treasure.

I would like to thank all of my family, especially my mum Michelle for doing such a fantastic job of instilling me with a sense of wonder about the world, as well as putting up with me and nodding at all the right moments. I would also like to thank my dad Mark for my very early training in computing and science, prevalent themes in this thesis. I would also particularly like to thank my granddad Frank Wallhead who may still be unsure about what I have been in training for all of this time, but who funded my trip to CogSci 2012 in Japan regardless. I not only received some excellent feedback there on my research, I also met Mitchell Travis who remains my best friend, role model and research mentor to this day. I would finally like to thank Mitch from the bottom of my heart for simply everything.

Contents

List of Figures

List of Tables

Declaration

The research reported in this thesis was carried out at the University of Exeter between September 2010 and August 2014 and was supervised by Professor Ian McLaren.

This thesis has not been submitted, in whole or in part, for any other degree, diploma or qualification at any university.

Fayme Yeates
August 2014

Chapter 1. General Introduction

"Learning without thought is labour lost"
Confucius

The central aim of this thesis is to investigate how humans learn sequences incidentally. The first challenge in doing so is demonstrating that humans can, in fact, learn without awareness or intention. Whilst Confucius may suggest that we require thought to learn, it seems obvious that we are able to acquire a number of skills and behaviours without understanding the complex underlying rules that define them (Cleeremans & Dienes, 2008) and this subject as a result has attracted great debate (Mitchell, De Houwer, & Lovibond, 2009; Newell & Shanks, 2014; Shanks & St John, 1994). Sequence learning is perhaps the most popular paradigm in which to study incidental learning (Destrebecqz & Cleeremans, 2003), as participants find it difficult to notice sequential patterns in serial reaction time (SRT) tasks even though they show evidence of learning through improved performance on these sequences after training (Lewicki, Czyzewska, & Hoffman, 1987; Nissen & Bullemer, 1987). How this occurs is not fully understood, as in addition to establishing whether this learning does or does not occur explicitly, it is possible that humans can automatically learn sequential information in a number of ways. Theories regarding incidental sequence learning and implicit learning more generally converge on the notion that we are somehow able to extract abstract, statistical regularities from the environment without being aware of them (Dienes, 2012). It is suggested that an associative learning process similar to that proposed to be found in animals underlies implicit human learning (McLaren, Green, & Mackintosh, 1994). This thesis attempts to investigate the specific predictions of an associative account of sequence learning using computational modelling to simulate and generate predictions regarding how altering various properties of the SRT task affects human sequence learning under incidental conditions.

In the introduction to this thesis I will firstly discuss theories of human learning more generally, covering the debate between single and dual process accounts and how these impact on the study of sequence learning. I will then discuss why the study of sequence learning is a key issue in understanding human cognition, and not only because it lends

itself to demonstrating the existence of automatic learning processes. The ability to learn sequences is a phenomenon worthy of understanding in itself, as sequential order is an intrinsic property of many skills and activities acquired by humans; for example, motor skills, music and language. Specific accounts of how we learn sequences are discussed and their implications are situated within human learning in general. Computational models are also discussed, as they remain a popular way of investigating sequence learning and attempt to formalise and therefore test theories of how these processes occur. Studies of sequence learning are reviewed in terms of the specific issues associated with and development of the paradigm over the last quarter of a century. I will then discuss our current understanding of how humans learn sequences in terms of the variations that have been investigated within a simple SRT task: the influence of time and trial order; as well as the role of stimuli and responses. This leads me to consider the relationship between stimuli and responses within a task, and predictions about cue-competition between elements in a sequential learning SRT from an associative perspective are also discussed.

1.1. Human learning processes

One of the most noteworthy and commonly observable phenomena of the human mind is our ability to learn. We can acquire motor skills, language, concepts and categories to such a complex and extensive degree that these learned processes and information are themselves the subject of intensive study. How our mind understands and is able to interact with the world relies on learning: a process of vital importance to human beings. How do we learn?

If we begin with perhaps the original research method employed by psychologists (see Watson, 1913): introspection, we can think of our own experience of learning to drive; playing an instrument; or learning languages, facts or ideas. We can sit down with the words and music to a song and learn to sing it, and our learning is the product of effortful study and intentional diligence. This, however, doesn't seem to be the only way we can learn, as we may find ourselves singing a song we have heard a few times before and observe that some learning has occurred by accident without even being aware of it. Whilst this idea of being able to learn without awareness is intuitively acceptable, that we can learn both with and without awareness is heavily contested

(Beckers, Miller, De Houwer, & Urushihara, 2006; Brewer, 1974; De Houwer, 2009; Mitchell, De Houwer, & Lovibond, 2009; Shanks, 2010; Shanks & St John, 1994).

It is generally accepted that humans are able to learn by some conscious, rule-based, verbalisable problem solving; whereby mental models or representations of the task are built in some way (De Houwer, 2009; Johnson-Laird, 1983; Matthews, Buss, Stanley, Blanchard-Fields, Cho, & Druhan, 1989). Indeed, the issue is not that these explicit learning processes can be denied; but that these cannot account for all learning observed in humans (McLaren, Forrest, McLaren, Jones, Aitken, & Mackintosh, 2014; Mitchell et al., 2009). A range of authors claim that there is a need to posit a further, separable learning system in humans that is functionally distinct from intentional and conscious learning and it characterized as an automatic process that operates outside of control and intention, that occurs in the absence of awareness of what is being learned (Evans, 2003; Lewicki, 1986; McLaren, Green, & Mackintosh, 1994; McLaren et al., 2014; Reber, 1967; 1989).

These dichotomous learning processes are popularly defined by the terms explicit and implicit (Shanks & St John 1994), which stem from the work of Reber (1967) who coined the term 'implicit learning' to describe the improved performance participants demonstrated on trained artificial grammars without being able to describe any rules or relationships. That there is more of a debate on the processes that underlie learning now than in the early days of psychology (Shanks, 2010) may come as a surprise to the naïve reader, for whom the existence of these two processes may seem introspectively obvious and beyond dispute. The following sections aim to summarise accounts of learning that are still as heavily contested today as they were when Reber (1967) discussed them over forty years ago (Cleeremans & Dienes, 2008).

1.1.1. A dual-process approach

One of the "oldest and most deeply entrenched dual-system theories in the behavioral sciences" is the dual-process account of human learning (Mitchell et al., 2009, p. 183). The appeal of a dual-system approach can been seen across psychology, with sets of two processes for memory (Squire, 1992); control (Forrest, 2012); social cognition (Chaiken & Trope, 1999) to name but a few. The human mind as a whole has itself been proposed to be constructed of these two systems (Kahneman, 2011) and this basic

explicit, controlled, conscious versus implicit, automatic and unconscious binary distinction pervades theorising across both the discipline and as a generally held belief regarding our behaviour.

These two learning systems are described in a number of ways: procedural and declarative (Willingham, Nissen, & Bullemer, 1989); implicit and explicit (Reber, 1967); automatic and controlled (Shiffrin & Schnieder, 1977). Essentially, a dual-process approach to learning suggests that there "exist two qualitatively different types of learning" (McLaren, Green, & Mackintosh, 1994, p. 315). Therefore, whilst the names and processes involved might differ across theories, dual-process accounts argue that a single-system account does not fully explain human learning: as learning can operate automatically outside of control, intention and awareness.

Mowrer (1947) suggested that psychologists should see conditioning and problem-solving as two, functionally separable learning processes. He stressed the nature of the two systems, one (associationism) biologically linked with the autonomic nervous system; the second (hedonism), driven by the central nervous system. McLaren, Green, & Mackintosh (1994) formalised this distinction in contemporary terms and propose two systems: a rule-based system that employs verbalisable hypothesis testing by which to learn information; and an automatic system that learns through the automatic formation of associations (similar to the learning processes believed to occur in animals). Human learning and memory processes have also been explained in terms of both a rule based system (e.g., Simon & Lea, 1974; Nosofsky, Clark, & Shin, 1989) and an instance based system, that simply stores each event experienced (e.g. Medin & Schaffer 1978); however, the more popular explanation of implicit learning is based on the automatic formation of links between mental representations (Mitchell et al., 2009).

1.1.2. A single-process approach

The single-process approach suggests that there is no need to posit an additional, automatic learning system when a single, rule-based or propositional system can account for all instances of human learning. Proponents of this view (e.g. Beckers et al., 2006; De Houwer, 2009; Mitchell et al., 2009) suggest that the automatic, associative links do not occur; and instead all learning is underpinned by qualified mental propositions with truth values, and therefore for learning of a contingency to occur

people must be aware of their relationship for an explicit propositional relation to be learned.

The role of some form of automaticity within mental processes as a whole cannot be denied (Cleeremans, Destrebecqz, & Boyer, 1998) as it is introspectively obvious that we do not have explicit access to the intricate workings of our minds. However, Mitchell et al. (2009) outline clearly that their claims regarding the absence of automaticity refer entirely to learning, as perceptual and memory processes may or may not act in an automatic fashion. Cheng & Novick (1992) suggest that learning occurs through judgements that are based on observed contingencies (positive or negative) between stimuli, but that these must be observed. As a result, one can consider the defining maxim of a single-process account as: awareness of the relationship between events in the environment is the minimum requirement for learning.

1.1.3. Central characteristics of dissociable learning systems

The central issues surrounding the number of learning processes are those of awareness, rationality, and control. Evidence for the automatic component of a dual-process account can be inferred when learning is involuntary, if it occurs without any consciously accessible knowledge, or if it does not produce rational outcomes or behaviour. These definitive characteristics of dissociable learning processes are defined briefly here in the context of wider associative theory, as well as considering the implications for the study of sequence learning.

1.1.3.1. Awareness

Generally speaking, many of the arguments for and against a secondary, associative system centre around awareness. One essential difference between the two accounts is the presence or absence of explicit knowledge regarding a relationship between events. An associative learning system does not require this information to be absent, but suggests that this is not a necessary condition for learning. Therefore, a situation in which learning has occurred, but contingency knowledge is absent is one of the gold standards in support of dual learning processes. In an early review of the associative literature Brewer (1974) suggested that there existed no such demonstration, which is supported by recent critiques of current evidence for learning in the absence of awareness (Mitchell et al., 2009; Shanks & St John, 1994; Shanks, 2010).

However, the case has also been made that a variety of associative and implicit learning studies *do* provide convincing evidence of an absence of awareness and that selective citation is required to conclude that human learning is a single process (Dwyer, Le Pelley, George, Haselgrove and Honey, 2009). The issue of demonstrating an absence of awareness, however, is fraught with methodological issues, which will be discussed further in this section. Nevertheless, for implicit learning to be studied we must strive to achieve conditions whereby participants are unaware of what they are learning about to entertain the possibility that this learning cannot be accounted for by a single, explicit process.

1.1.3.2. Rationality

Another argument used in both criticism and defense of associative learning is rationality. Single-process accounts suggest that one, cognitive, explicit account can produce learning effects through propositions, rules and logical inferences that humans can and do make (De Houwer, 2009). Any instances of learning that do not appear to follow rational thought, therefore, can be suggested to support evidence of an automatic process (Dickinson, 1988; Shanks, 2007; Shanks & Dickinson, 1990). How we define rationality is key, though, and in this sense refers to sub-optimal behaviour as a result of learning (Shanks, 1995). This is unlikely in the context of a sequence learning experiment, as the nature of the task means that participants who learn the sequence will respond quicker and more accurately (Nissen & Bullemer, 1987) and therefore learning has a positive pay-off. Mitchell et al (2009) also point out that whilst propositions have a 'truth' value, this does not itself have to be rational, or indeed true. Consequently, sequence learning and this thesis have relatively little to do with the concept of rationality, which itself can be difficult to define and interpret (Shanks, 1995).

1.1.3.3. Instructions and control

The role of controlled cognitive effort is heavily relied upon in the study of implicit processes, as without the ability to unequivocally demonstrate that participants are unaware, another approach is to define a process by the volitional conditions under which it occurs (Jacoby, 1991). Learning can therefore be defined as intentional (with intention and control) or incidental (without intention and outside of control). Participants who demonstrate that they have learned information when they did not actively attempt to do so are used to provide evidence for implicit learning processes

(Dienes & Berry, 1997). This is the standard manipulation of most implicit learning tasks, which consequently involve some cover story or instructions that avoid reference to learning in order to ensure that participants do not attempt to engage any explicit learning process (Perruchet & Pacton, 2006).

Given a simple instruction, however, participants can produce a large number of learning effects that are supposedly associative in nature. With no previous training, if instructed that there exists a contingency between two events (Cook & Harris, 1937) participants are able to produce the appropriate response with no need for the gradual build-up of associations. Similarly, given associative training schedules that should result in conditioned responding, participants are able to withhold a response if instructed to do so (Colgan, 1970; Lovibond, 2003). It seems that instructions can subsequently produce, alter, reduce and stop apparently automatic processes. These results are used to provide support for a single learning process, as it suggests that learning is not automatic and is in fact under our control (Mitchell et al., 2009). These results, however, do not exclude the possibility that implicit learning did or can occur; as this would imply that humans are able to control the *expression* of automatic mental processes (Jones & McLaren, 2009). Indeed, a study by Wan, Dienes, and Fu (2008) found that participants were able to intentionally choose from two implicitly learned artificial grammars without being aware of them. Therefore, it may be possible that participants can stop themselves from expressing implicitly learned information; or produce explicit learning effects that mimic those that seem automatic. However, these do not provide evidence that the implicit learning process was affected or controlled in any way, as the production of an associatively acquired response may simply be overpowered by explicit knowledge in these tasks.

1.1.4. Conceptual issues in studying dual learning processes

1.1.4.1. Restriction of a binary framework

The case of control, as discussed above, provides an example of how defining dual learning processes as binary opposites (e.g. implicit versus explicit) provides us with a research framework with which to study them. For example, if we propose that explicit learning involves control, then we can expect that people will be able to explicitly choose to *not* learn; whereas implicit learning cannot be controlled and therefore will occur regardless of the intention of a person (Jacoby, 1991). Indeed, the understanding

of implicit or unconscious systems is generally framed in terms of a definition of the explicit or conscious (Reber, 1989). These dissociable, binary characteristics are the fundamental components of a structuralist understanding of the human world, where our understanding of one concept is meaningless without its complement (Hawkes, 2003).

A binary pair of processes, however, is often insufficient to capture what are not two distinct functional systems but a continuum on which two things may seem the polar opposites but simply take up very different points on the same scale (for example, gender, Hird, 2000). Defining a mental function as a set of binary polar opposites may then, provide us with a framework that is introspectively agreeable but which restricts our ability to explore the functional properties of these systems. So it is possible that a dual-process account limits our understanding in an attempt to describe divergent phenomena. Indeed, research on the development of explicit knowledge in sequence learning tasks suggests that with gradual practice participants become aware as a result of the increasing strength of memory representations about the sequence (Shanks, Wilkinson, & Channon, 2003). This is a single process account that suggests that explicit knowledge is not the defining characteristic of one of two learning systems, but instead the product of greater learning within a single learning dimension (Cleeremans, 2006; Shanks, 2005).

1.1.4.2. Falsification of a single process account

Gilbert (1999) suggests, however, that the term ‘dual’ is used instead of ‘two’ as the fundamental aim of psychology in understanding the mind is not to number the amount of processes that it may use, but to infer that a single process is not sufficient itself to account for the phenomena. Based on the ideas of both parsimony and falsifiability, whilst one system that can account for all behaviours is the ideal, the dual process logic proposes that the existence of instances that cannot be accounted for by a single system suggest the need for (at least) dual processes. This is a subtly different use of a dual-process theory as framework and method, as it focuses more on inductive logic with a proof in principle and less on understanding, and is the logic most often employed in the study of implicit learning. As Shanks (2007, p. 297) points out in reference to human learning, a higher order set of problem solving processes are not refuted, there are simply those who suggest there “might be a separate type of thinking (associative) when people make instinctive judgements under conditions of less reflection”.

This is a popular approach to the study of implicit learning, but there are inherent problems with attempting to falsify a single process account of learning. Importantly for the study of incidental learning is the issue of attempting to find evidence for the absence of awareness, which (issues with measurement aside), predicts the null hypothesis (the demonstration of *no* explicit learning). Null-hypothesis statistical testing (NHST) cannot provide us with evidence that something has *not* been explicitly learned, a problem for the conventional statistical method employed across psychology. NHST, however, is not the only statistical method for the interpretation of results and there have been a number of suggestions to circumvent this fundamental issue. Qualitative differences (Jiménez, Vaquero, & Lupiáñez, 2006); dissociations (Perruchet, 1985); state-trace analysis (Bamber, 1979); and Bayesian analyses (Dienes, in press) are all proposed as solutions to this issue, and the use and associated issues with these proposed solutions are discussed in Chapters 2 and 5 of this thesis.

1.1.4.3. The case for ignoring the absolute number of learning processes

Witnauer, Urcelay and Miller (2009) argue that a comparison between the two accounts as opposing theories of learning is flawed, as each is concerned with a different level of analysis. Indeed, the arguments surrounding the number of processes are considered by some to have little value, as Cleeremans and Dienes (2008, p. 401) assert:

> The verbal question of how many learning systems there are is in danger of being vacuous. If God were to tell us how many learning systems there were with a single number (one? two? three?), we would have learned nothing.

Simply attempting to demonstrate implicit learning as a proof of principle, or to falsify a single system account is neither a sound theoretical approach nor the aim of this thesis. Rather than attempting to quantify the number of processes involved per se, this thesis attempts to investigate incidental sequence learning in humans and assumes the possibility that humans may be able to learn automatically. In doing so I aim to better understand whether or not automatic associations can form in humans as observed in animals (McLaren, Green, & Mackintosh, 1994), and how this occurs.

As associative learning in humans is proposed to exist alongside explicit learning within a dual-process framework; any investigation of these processes must consider the single-process account and associated issues of studying a possibly implicit process.

This is important as we must consider the influences of explicit learning, as well as the possibility that this account more parsimoniously accounts for human learning. Therefore, whilst proving the occurrence of automatic, associative processes would be a challenge for any researcher (Shanks, 2010); couching an investigation of associative sequence learning processes within a single versus dual process argument is important.

1.2. Studying implicit learning from an associative perspective

Before moving on to discuss theories of how humans learn sequences, this section will briefly review attempts to study implicit associative learning processes in humans, with the conclusion that one of the most beneficial and methodologically promising is sequence learning in the SRT task. There are a number of literatures that investigate automatic or implicit learning processes in humans, which have considerable overlap in the subject of interest but can diverge widely in terms of theoretical and methodological position. Implicit learning, human associative learning and statistical learning research strands each attempt to determine how we learn, but are generally concerned with a different level of analysis. As a whole, the implicit learning literature has been concerned with methodological issues surrounding the elimination of an explicit explanation and defining the conditions under which implicit learning might occur, with less functional consideration given to the underlying processes and how these might occur. Associative learning research operates at a more detailed level of analysis, which functionally (and algorithmically) attempts to understand learning processes. Statistical learning research falls somewhere in the middle, borrowing concepts, paradigms and language from both pre-existing literatures (Perruchet & Pacton, 2006). Whilst these areas have by no means developed exclusively, the literatures could still offer much to one another in studying automatic or implicit learning in humans.

As a general rule, associative learning studies attempt to further understand how humans come to learn associations between events; but they largely ignore the issues posed by a single-process, explicit account of the data. Implicit learning research, on the other hand, has developed through attempts to uncover the presence of unconscious learning processes and is hence suited to the study of human learning. In contrast, associative learning in humans has developed from the existing framework for animal learning, consequently methodological issues have not been so extensively considered. However, the functional explanations offered by associative learning elegantly explain

observed learning effects and the emergent, seemingly explicit phenomena produced by associative models provide an extensive and detailed account of human learning (Shanks, 2009). I will argue that the study of sequence learning offers an ideal experimental setting within which to investigate associative processes in humans that circumvents the issues faced by other paradigms.

1.2.1. Human associative learning paradigms

Research into human associative learning stems from a tradition of animal learning research, beginning with the famous Pavlov's dogs (1927) and the first demonstration that learning occurred between repeatedly presented, temporally contiguous events in the environment. Classical or Pavlovian conditioning, as it is known, is one of the basic learning effects observed in species as simple as the sea slug *Aplysia* (Carew, Walters, & Kandel, 1981). The basic effect involves training an unconditioned stimulus (US), which evokes some autonomic, unconditioned response (e.g. salivation, eye-blink), alongside a neutral, conditioned stimulus (CS) such as a light, tone or odour. This CS is neutral in that presenting it alone should produce no response, and yet after training where the CS is presented before the US the CS comes to evoke a response when presented alone. CS-US pairings are usually defined by their temporal contiguity so a CS occurs before (trace conditioning) or overlaps the start of the US (delay conditioning).

In humans, a variety of stimuli have been used as a US to provide evidence of conditioning responding, for example: aversive noises (Neumann & Waters, 2006); images (Levey & Martin, 1975); electric shocks (McAndrew, Jones, McLaren, & McLaren, 2012; Vervliet, Vansteenwegen, Baeyens, Hermans, & Eelen, 2005); flavors (Chambers, Mobini, & Yeomans, 2007); odours (Marinkovic, Schell, & Dawson 1989); as well as air puffs to the eye (Perruchet, 1985; Weidermann, Tangen, Lovibond, & Mitchell, 2009). However, participants are not necessarily unaware of the contingencies between CS and US. Whilst authors argue that certain stimuli have an automatic, stimulus-driven impact (Bliss-Moreau & Barrett, 2009), there exists the possibility that conscious expectation of a US following the presentation of the CS could lead to an explicitly produced response to the CS (Mitchell et al., 2009, although see Perruchet, 1985).

Studies of causal reasoning in humans are themselves perhaps the biggest influence on the study of associative learning processes in humans after Dickinson, Shanks and Evenden (1984; Shanks, 1985) observed that when performing these tasks humans demonstrated the sort of associative effects shown in animals. These paradigms involve training participants that certain stimuli lead to certain outcomes, and their judgements of the likelihood of such an outcome given the stimuli are then measured to assess learning. Such research on humans is often conducted using elaborate scenarios, such as the allergist paradigm; where the participant plays the role of a doctor who is supposed to work out the nature of a hypothetical person's food allergy. Participants rate the likelihood that a person is going to have an allergic reaction (the outcome, O) when presented with one or more food items as stimuli (e.g. an apple, A and a banana, B). If participants learn that A leads to O and B does not lead to O, we see their ratings of the likelihood of this event increase and decrease respectively when presented with A and B. Thus, simple discrimination learning and a host of other associative learning effects (e.g. blocking, Le Pelley, Oakshot, & McLaren, 2005; backwards blocking, Shanks, 1985) are demonstrated in such tasks.

However, a number of papers have sought to demonstrate that propositions are used to solve causal reasoning tasks and not an associative system (Beckers, De Houwer, Pineño, & Miller, 2005; Lovibond, 2003). In defense of associative learning, studies with complex designs that produce convincing, non-rational learning effects that seem to occur outside of awareness and follow the predictions of an associative account have been provided (Haselgrove, 2009; Le Pelley, Oakshot, & McLaren, 2005; Karazinov & Boakes, 2007), which are hard to reconcile with conscious, explicit reasoning-based accounts. However, two fundamental issues with such paradigms exist, both of which centre around the use of elaborate cover stories. The first issue is that participants are asked throughout the task to make predictive judgements and therefore whilst the cover story means the task is not presented to the participants as a learning task, participants are explicitly required to attend to contingencies between events; violating the maxims for incidental learning and making learning possibly the product of explicit intentions accounted for by a single propositional process (Vadillo, Orgaz, & Matute, 2008). The second issue is that participants may rely on pre-established causal frameworks within these tasks (Waldmann & Holyoak, 1992). Participants making judgments about the relationship between symptoms and diseases were capable of learning associatively

only when the symptoms predicted the disease, and not when the disease predicted the symptoms. Waldmann and Holyoak (1992) suggest that participants rely on a pre-conceived concept of causality to perform these tasks, suggesting that the process is not automatic nor based entirely on learning simple associations.

1.2.2. Implicit human learning paradigms

As there is a consensus that designing associative experiments where conscious and explicit, propositional processes cannot influence learning is incredibly difficult (Boakes, 2009; Seger, 1994; Shanks, 2007); I will now discuss evidence from studies that attempt to fulfill this criterion with a focus on identifying the most effective paradigm for the study of human associative learning.

Reber (1967, p. 855) coined the term "implicit learning" when he employed an artificial grammar on which participants were trained across seven blocks. Letter strings of six to eight items in length formed the experimental stimuli, which were constructed from a Markovian grammar. These were presented once to participants for five seconds, who were required to immediately reproduce the stimulus in its absence. Reporting the number of errors participants made, whilst identical across the first two blocks, performance of participants in the control group (a random set of letter strings) plateaued, whereas participants who were experiencing letter strings constructed by the artificial grammar continued to improve. Due to being unable to verbally report the rules of the grammar, Reber (1967; 1989) suggested the presence of an unconscious system that could acquire abstract knowledge.

Artificial grammar learning tasks have been used to study implicit learning processes across implicit (Dienes, Broadbent, & Berry, 1991), associative (Dienes, 1992), and statistical (Saffran, Aslin, & Newport, 1996) research areas. The grammatical stimuli involve complex underlying relationships within an abstract structure, which are hard for participants to verbalise and demonstrate explicit knowledge of (Cleeremans & Dienes, 2008). The implicit learning literature also contains experimental designs based on visual search tasks (Chung & Jang, 1999); dynamic system control (Berry & Broadbent, 1984; Broadbent, Fitzgerald, & Broadbent, 1986); hidden covariation detection (Lewicki, 1986); probability learning (Reber & Millward, 1968); as well as sequence learning (Nissen & Bullemer, 1987; Lewicki, Czyzewska, & Hoffman, 1987).

In an attempt to avoid the influence of explicit processes, these paradigms all share or involve one or more of the characteristics outlined in Table 1.1. They share the common feature of complex, hard-to-detect, contingencies between events that Reber (1989) suggests is key in the design of implicit learning studies. If we present participants with simple contingencies, Reber (1989) suggests this will contravene the conditions *necessary* for the observation of implicit learning, as the simple nature of this relationship is readily dealt with by the explicit system. Whilst the implicit processes available are indeed capable of learning such simple associations, increasing complexity of the information to learn will not only reduce the ability of participants to employ an explicit, verbalisable set of rules to learn which may overshadow implicit learning; it will also make effective use of the implicit learning process and give preference to its deployment.

Table 1.1. General characteristics of implicit learning tasks

Shared characteristics	Additional manipulations
High contingency complexity	Cognitive load (e.g. Le Pelley et al., 2005)
Low contingency detectability	Attention (e.g. Curran & Keele, 1993)
Time pressure	Subliminal stimuli (e.g. Weins & Öhman, 2002)
	Low stimulus discriminability (e.g. Stevenson & Boakes, 2004)

1.2.3. Sequence learning tasks

Sequence learning tasks have become the dominant paradigm in the study of implicit learning (Cleeremans & Dienes, 2008) and take the form of a serial reaction time (SRT) task. Participants are usually required to respond to different on-screen stimuli that appear in a certain location with a different key press response as quickly and accurately as possible. The instructions given to participants simply encourage them to be fast and accurate and make no mention of learning or that the stimuli will follow some sequence. Whilst not without their own methodological issues, these studies circumvent the issues with explicit biases based on propositions shown in other implicit learning studies. An example would be tasks using cover stories such as an economic challenge, perhaps in a factory where wages, employee happiness, productiveness, etc. are manipulated by the participants who are instructed to simply alter these variables to produce the best solution; a task on which they show evidence of learning without rule knowledge (Berry & Broadbent, 1984; Broadbent, Fitzgerald, & Broadbent, 1986). However, these

paradigms, along with artificial grammar tasks, may be affected by the preconceptions of participants who "may be guided towards (or diverted from) the discovery of its underlying structure by prior knowledge of related real situations" (Perruchet & Pacton, 2006, p. 237).

By separating the contingency to be learned across time and trials this decreases the opportunity for participants to notice or look for contingencies (Jones and McLaren, 2009) and avoids perceptual issues with training certain concurrent stimulus sets. This criticism is true for artificial grammar tasks (Shanks & St John, 1994) as well as in visual search tasks[1]. Participants in sequence learning tasks are instructed to respond quickly and accurately, which means there is not need to provide participants with an elaborate cover story to mask learning. To some extent, this also provides conditions of cognitive load or divided attention (although manipulations on cognitive load and attention in sequence learning investigate these influences further, e.g. Curran & Keele, 1993; Stadler, 1995), as participants are aware of stimuli and responding to them but the speeded nature of the task and motivation to perform well may reduce the influence of explicit processes.

1.2.4. Conclusions

In this thesis, I therefore use sequence learning as the paradigm within which to study incidental associative human learning for the following reasons. Firstly, because this implicit learning task does not rely on stimuli whose properties carry explicit or perceptual significance that may provide an alternative explanation for any observed learning. Secondly, unlike other associative tasks with elaborate cover stories, this task was designed to be performed incidentally, and therefore offers an opportunity to investigate simple contingency learning from an associative perspective; without the influence of explicit learning processes. Whilst it is by no means the first study of associative incidental sequence learning (Cleeremans & McClelland, 1991; Jones & McLaren, 2009; Lee & Livesey, 2013) and automatic associations should occur regardless of the condition, the literature points to sequence learning as the best method to investigate implicit learning (Destrebecqz & Cleeremans, 2003) and therefore it is the

[1] At the Meeting of the Experimental Psychology Society (EPS) January 10, 2014 David Shanks (Vadillo, Beesley, & Shanks, 2014) presented results from implicit learning studies using visual search tasks (including Beesley & Shanks, 2012), of which Mike Burton offered this critique. This is discussed further in Chapter 6, but similar to Shanks & St John's (1994) criticism of the artificial grammar task, participants may have been learning about some perceptual feature of the visual stimuli that is not captured by the explicit measures on the task.

optimal paradigm to select in order to investigate automatic, associative processes in humans.

1.3. Theories of human sequence learning

In the previous section I argued that sequence learning is the ideal task to investigate associative learning in humans, as it is the best implicit learning task available (Destrebecqz & Cleeremans, 2003) and it is these incidental conditions under which we expect associations to form and therefore associative learning to be observed (McLaren, Green, & Mackintosh, 1994). Further to this, and central to this thesis, is that sequence learning itself is a phenomenon worthy of study in its own right. Sequences of motor-actions, phonemes and words make up our behaviours, speech and language; not to mention how we understand music and learn to play sports or perform any skill. Lashley (1951) noted that serial order was a key issue in understanding human cognition and behaviour, and various theories and computational models have since been developed to account for how humans learn sequences.

It is not in dispute that humans can learn sequences, the literature is concerned primarily with what is learned (Dennis, Howard, & Howard, 2006). This thesis aims to address the next level of analysis in asking the fundamental issue of *how* sequences are learned. Both dual-process theories of human learning and theories of sequence learning converge on an associative perspective, but I will first consider alternative implicit and explicit accounts. I will further discuss computational models of sequence learning, which is a popular research strand as representing time and serial order in models of learning is a complex issue that is much debated within psychology and computing. It is also a major theoretical and methodological component of this thesis.

1.3.1. Rule learning

The most parsimonious account of sequence learning in humans is that we can simply learn rules in order to account for the apparent structure that is experienced when performing tasks. This follows a propositional perspective and requires the common features of a sequence to manifest as a mental representation of their abstract structure (De Houwer, 2009; Mitchell et al., 2009). This has early origins in the study of sequence learning in humans, with the work of Restle (1970) suggesting that

participants explicitly learn sequences by applying tree-like structural rules in order to understand and then acquire these sequences. Without intention to learn, research into the development of explicit knowledge in sequence learning tasks suggests that participants can form propositional rules just from experiencing a task, which can produce a step-increase in performance that correlates with self-reports of awareness (Hoffmann & Koch, 1997; Koch, 2007). These studies, however, also show evidence of gradual performance increases before this occurrence which suggest that early learning was occurring that did not rely on rule-based propositions (Rünger, 2012).

1.3.2. Memory

One defining aspect of all theories of sequence learning is that in some way they must account for the problem that contingencies between sequential elements are separated by time. This nearly always involves the influence of some form of memory (Hsiao & Reber, 2001) and consequently this leads some researchers to suggest that the most parsimonious explanation of sequence learning is simply an instance based account (Shanks & Perruchet, 2002; Shanks, Wilkinson, & Channon, 2003), where the experience of each trial and is stored in memory, and that increased training of certain trial orders would make certain instances gain stronger memory representations, thus producing improved response latencies. As a consequence of these strengthening memories, participants would also be increasingly likely to report and have conscious access to these instances, from which they could derive knowledge about underlying contingencies, sequences or rules (Shanks, 2005). This is not an incompatible approach to the associative perspective put forward in this thesis (Fu, Fu, & Dienes, 2008), as it does not specify how memory representations are formed.

1.3.3. Chunking

A variety of models of sequence learning involve some system that can classify instances according to their wider context within a sequence, for example: that the current trial was preceded by the same stimulus location. This leads to a model that involves a combination of chunks and hierarchic representation derived from the seminal work of Lashley (1951). As a model of sequence learning chunking has had mixed results (Curran, 1995): but it can describe explicit sequence learning effects well (Gordon & Meyer, 1987; Povel & Collard, 1982; Restle & Burnside, 1972); However, chunking itself does not offer a functional explanation of learning processes, as some

method by which chunks are formed and then learned is required. One suggestion for this is through the formation of associations, although other chunking models based on perceptual properties exist (e.g. the Elementary perceiver and memorizer [EPAM], Feigenbaum & Simon, 1962; 1984). When applied as a simple explanation of sequence learning under incidental conditions, chunking has descriptive strength, yet a computational instantiation is yet to demonstrate that simply chunking or grouping together trials into sub-sections sufficiently accounts for human performance on such tasks (Spiegel & McLaren, 2006).

1.3.4. Chaining

One of the oldest theories in representing serial order is that of chaining (Ebbinghaus, 1964) where pairs of sequential elements are associated together, as on each trial the current stimulus is the cue for the next stimulus (Lashley, 1951). Simple chaining models only consider these pairwise associations, which are problematic for longer or more complex sequences that share many elements (Hartley & Houghton, 1996). Models can contain these pairwise associations as part of a common trace formed across the experiment (e.g. the sequential pairwise associative memory [SPAM] model, Wallace & Fountain, 2002), which avoids these issues, however these models are unable to ignore non-predictive elements within a probabilistic sequence, leading to interference and a lack of learning (Spiegel & McLaren, 2006).

1.3.5. The simple recurrent network (SRN)

Jordan (1986) suggested that a model of learning needed to encode for time in some way and produced a network that involved recurrent connections between output and input units within a three-layer connectionist neural network. This network involved input and output layers, with a layer of hidden units in between these that encode an internal representation of the input before this is passed to the output. Hidden layers are suggested to be essential to modelling more than simple contingencies between events (McClelland & Rumelhart, 1986) and are the component of such networks that give them such emergent power (Ellis & Humphreys, 1999). The recurrent connection between output and input meant that output on trial t would be copied back into the model at $t + 1$, therefore, a memory for the last trial would appear within the model when making a prediction about the current trial.

Jordan networks were, ultimately unsuccessful in producing a model of a wide range of sequential phenomena, as Elman (1990) refined the application of a recurrent loop within a neural network with far superior success in accounting for human performance. This network - the SRN is shown in Figure 1.1 and involves a similar structure to a Jordan net, but recurrence occurs with the activations of the hidden units on *t* being fed back into the model at *t* + 1. Thus a memory for the internal representation of the last trial is fed back into the model as input on the next trial. It is a connectionist neural network that is organised into layers of units that constrain the directional flow of activation. A set of input units are activated according to local or distributed external input to represent task stimuli and activation passes forward (hence the term feed-forward, McClelland & Rumelhart, 1986) through multiple connections to a set of output units. The activation of these output units are used to train the network, representing the responses or task outcomes – in sequence learning the SRN is trained to predict the next element in the sequence. The difference between the expected output and actual outputs activation is used to calculate an error term. This is passed backwards (back-propagation) through the model connections, updating the weights according to an error correction rule (Rumelhart, Hinton, & Williams, 1986).

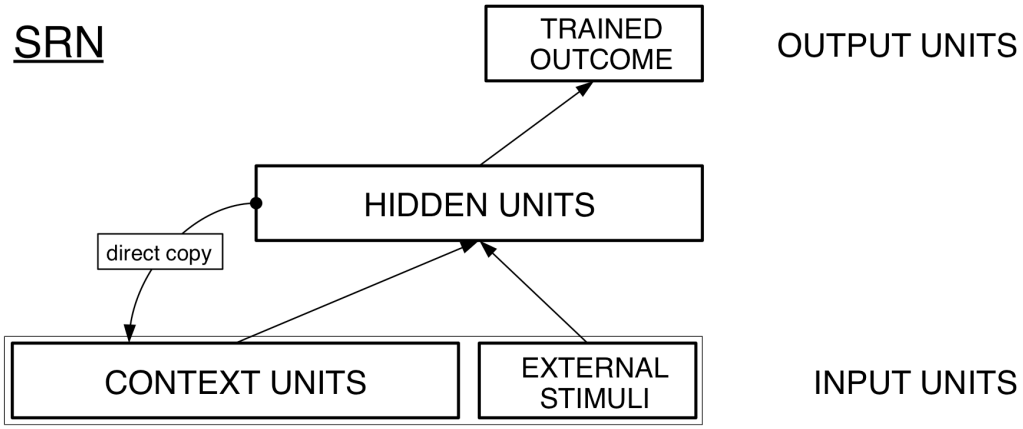

Figure 1.1. The simple recurrent network (SRN, Elman, 1990). Input units feed forward external activation to a hidden layer of units which create a distributed, internal representation of that input. This is fed-forward to output units, which are trained to produce the outcome expected in the task through some supervised error correction. This error is back-propagated: i.e. passed backwards through the model, updating the strength of connections between the units. On each trial the hidden unit activations are copied back into context units, which act then as input alongside external input on the next trial.

The SRN is the most popular model in simulating not only sequence learning but a variety of other forms of implicit learning with great success (Dienes, 1992; Dienes, Altmann, & Gao, 1999), although adaptations to the basic model are required for it to account for certain human learning effects (Beesley & Le Pelley 2010; Cleeremans & McClelland, 1991). The SRN is able to learn complex, probabilistic sequential rules and abstract structures that give it extraordinary explanatory power in accounting for varied and diverse phenomena (Beesley, Jones, & Shanks, 2012).

1.3.6. Other models

There are a number of computational models that can deal with learning structure over time in the associative literature as well as the SRN: the linear associative shift-registrar (Gureckis & Love, 2010); the augmented SRN (AugSRN, Cleeremans & McClelland, 1991); the auto-associative recurrent network (AARN, Dienes, 1992; Maskara & Noetzel, 1993); the temporal recurrent network (TRN, Dominey, 1998); domain-transfer SRN (Dienes, Altman, & Gao, 1999); and the APECS SRN (Jones, Le Pelley, & McLaren, 2002). Further parallel neural models (Hikosaka et al., 1999), synaptic cluster models (Dehaene, Changeux, & Nadal, 1987), hidden markov models (Baum, 1972) and adaptive resonance theory models (ARTMAP, Carpenter, Grossberg, & Reynolds, 1991) all represent time and memory in some way as this is a fundamental challenge for any learning model. Further to this, dual process models (e.g. Keele, Ivry, Mayr, Hazeltine, & Heuer, 2003; Sun, Slusarz, & Terry, 2005) attempt to account not just for implicit sequence learning performance, but further explicit and attentional components of sequence learning that I do not consider here.

Therefore, as the aim of this thesis is the development of understanding of human associative learning using an incidental sequence learning task; the model I will employ as the starting point for my investigations is the Augmented SRN (AugSRN, Cleeremans & McClelland, 1991), which is a version of the SRN that attempts to not only simulate sequence learning, but also to account for sequential effects (which are discussed in detail in section 1.5.1.2). As SRT tasks involve both associative learning of sequences as well as some possible short-term priming of the previous trials, this model incorporates both, and has been successful in simulating the complex pattern of subsequence learning effects demonstrated by Jones and McLaren (2009). The AugSRN is discussed in full detail in Chapter 3, but essentially in adopting this model I intend to

use it to investigate automatic associative processes in humans. The model is therefore chosen based on its prior success in modelling sequence learning and used as a functional complement to associative theory in predicting human behaviour on incidental sequence learning tasks. I do not subsequently propose that a better account of sequence learning itself does not exist; as this thesis examines to what extent the AugSRN can account for human data and thus how effective (this computational model of) associative learning theory is when applied to the incidental learning of sequences.

1.4. Measuring human sequence learning

Now that I have considered how we might account for human sequence learning I will provide a brief overview of the key developments that have occurred in the literature since Nissen & Bullemer (1987) produced evidence of apparently implicit learning on the SRT task, starting with the measurement of explicit knowledge. This is a key concern of this thesis as evidence of truly implicit learning is difficult to obtain. This section converges on the use of both explicit tests of knowledge (despite the issues associated with taking these after the learning event occurs) alongside a computational approach based on associatively predicting human sequence learning in order to find evidence for automatic, associative processes.

1.4.1. The general SRT paradigm

Whilst not all tasks are the same, sequence learning tasks involve participants responding to stimuli (usually on-screen shapes or lights) that appear in certain locations. The number of stimuli and locations differs between experiments, but participants are required to make accurate speeded responses with (usually) spatially compatible key-presses to each stimulus. Therefore, SRT tasks of sequence learning are characterised as visuo-spatial speeded motor response tasks. The first demonstration of implicit learning in SRT tasks is attributed to Nissen and Bullemer (1987) but the task was also used by Lewicki, Czyzewska and Hoffman (1987).

Tasks usually involve three phases: training, test and post-experimental explicit knowledge tests. Training and test phases involve responding to the SRT task, which in training follows a certain sequence that participants are not told about (incidental conditions). Test phases involve responding to sequences of trials that are either random

or follow the sequence on some occasions but not others. After the SRT task has finished participants are then usually asked to perform some test that intends to ascertain the extent of their explicit knowledge about the sequences involved in the task.

Measurements are taken in terms of both response latencies (reaction time, RT) and error rates are taken as an index of incidental performance on the task, as increased speed and less errors on trials that follow the sequence are taken to suggest that participants have learned this sequence and are therefore faster and more accurate (Nissen & Bullemer, 1987). Measurements of explicit learning are discussed in more detail in section 1.4.3 and vary from simply asking participants what they know about the task (Reber, 1967) to more complicated task designs, for example the process dissociation procedure (Jacoby, 1991). Measurement issues in SRT tasks of sequence learning are therefore discussed first, before I move on to examining the sort of sequences used in SRT tasks.

1.4.2. Measuring implicit learning in the SRT task

An SRT task is ideal for measuring implicit learning as it employs an indirect measure of learning in RTs and error rates. Asking participants to predict what will happen next, as in prediction learning tasks (e.g. Reber & Millward, 1968), or indeed how much they know about contingencies or what rules they think may be in play all attract attention to the relationship of the trained stimuli (Shanks & St John, 1994), which is an obvious problem for an incidental task. Such tests can be considered direct, where participants are explicitly asked about their knowledge in some way; on the other hand tests can be indirect, where learning is measured without letting participants know this is being measured (Merikle & Reingold, 1991). RTs and errors can therefore provide an indirect and objective measure of incidental learning, however, there are a number of considerations that must be taken into account.

The first is simply the use of appropriate controls, as speeded responding on a task across training does not provide evidence of learning per se and may represent a non-learning priming effect. SRT tasks of sequence learning must, therefore provide a suitable control group who perform the same task without the presence of sequences (Nissen & Bullemer, 1987) to ensure that improved responding is due to the presence of the sequence itself in the sequence-trained group. Some studies use a within-subject

control of untrained or random sequences (e.g. Jiménez & Vázquez, 2005) or a post-training random test or transfer phase (Shanks & Perruchet, 2002) in order to reduce the number of subjects required for the task. I will demonstrate in section 1.5.1.2 that this is not an adequate design in controlling for other associated issues in SRT tasks.
One important criticism leveled at many studies of sequence learning that otherwise involve a strong experimental design (Anastasopoulou & Harvey, 1999; Shanks & Johnstone, 1999) is the issue of speed-accuracy trade off (Jones & McLaren, 2009). If a group of sequence-trained participants demonstrates reduced response latencies compared to a control group, this may be as a result of faster but less accurate responding in the sequence-trained groups. The presence of a sequence may cause participants to respond differently, but this may not be an improvement over the control group, nor the result of incidental learning. Many studies in the implicit learning literature as a whole, not to mention the sequence learning literature, only report either the reaction times or accuracy of participants. It is, therefore, important to detail both measures of performance on an SRT task to eliminate the possibility that participants trained on sequences are not becoming simply quicker but less accurate, or vice versa, compared to the control groups.

1.4.3. Measuring explicit learning in the SRT task

Measuring the presence or absence of conscious mental processes is a complex area of debate, and similar to all psychological measurements suffers from attempting to measure latent psychological processes through manifest variables (Newell & Dunn, 2008). Generally, tests of learning can be organised into three groups, tests of: acquisition, knowledge and retrieval (Cleeremans et al., 1998). As well as indirect and direct tests as described above (Merikle & Reingold, 1991), tests can also be classified as objective, where a measurement compares learning against chance performance, or subjective, where participants report their belief in their own knowledge (Dienes, 2004). The title of the section is somewhat misleading as measuring any explicit learning is not what researchers wish to do in implicit studies. Whilst researchers wish to demonstrate the absence of explicit learning, in doing so we must ensure that explicit processes are given every chance to appear on test in order to provide convincing evidence they do not, therefore a suitable measure of explicit learning is given far greater attention than the implicit learning measures themselves.

Classically, arguments surrounding tests of explicit knowledge begin with the discussion of verbal self-report, as employed by Reber (1967) and in original sequence learning studies (Nissen and Bullemer, 1987); which are criticised for a lack of sensitivity as they do not measure learning at the detailed level that implicit measures, discussed in the next section do (Shanks & St John, 1994). Forced choice tasks are suggested to provide some test of conscious knowledge at a more sensitive level, which usually involve making recognition judgements (Perruchet & Amorim, 1992), fragment completion (Willingham, Nissen, & Bullemer, 1989), or sequence generation – where participants are asked to predict which trial comes next or produce a sequence (Jiménez et al., 2006). However, these direct measures of explicit knowledge are not free of automatic bottom-up responses when presented with stimuli that have been learned and cannot be assumed to be exclusive of the influence of automatic processes (Merikle & Reingold, 1992).

Subjective measures of explicit knowledge have also been proposed as an index of explicit knowledge, as asking participants whether they are guessing or not can give a criterion of whether participants are engaging in explicit processes (Cheesman & Merikle, 1984; Dienes & Berry, 1997). This is difficult to instantiate during a sequence learning task, as asking questions about the task as it progressed would reveal the nature of learning that was meant to be implicit. Therefore, questions regarding participants introspective performance are insensitive and could be influenced by decay over time or indeed bias (Dienes, 2004). A further suggestion is the zero-correlation coefficient, which suggests that if subjective confidence ratings in performance do not correlate with objective learning measures then there is no evidence of explicit awareness of what has been learned (Dienes & Berry, 1997). However, a corollary of this argument is that subjective and objective measures are independent of one another and are not influenced by the other, which is unlikely to be the case (Jacoby, 1991; Merikle & Reingold, 1992).

Most studies of associative learning attempt to demonstrate the absence of explicit *knowledge* (Cleeremans et al., 1998) rather than acquisition, as this is measured post-learning. Whilst RT and error rates across training provide learning curve data with which to analyse on-line acquisition of learning, explicit tests post-SRT task can always be argued to suffer from memory decay or interference (Shanks & St John, 1994).

Explicit tests that occur after the SRT task therefore may not capture a learning process that occurred through (now inaccessible) explicit propositions (Shanks & St John, 1994).

Consequently, sequence learning tasks offer no definite solution to assess the extent to which participants may or may not have engaged in explicit learning processes, as the nature of implicit serial response tasks ultimately require that explicit checks on learning are conducted post-training. This is due to both keeping participants from uncovering the nature of the task and so as to avoid disruption of the speeded nature of responding and thus the learning of the sequences. Both subjective and objective measures of explicit knowledge were taken in the studies in this thesis after a test (extinction) phase of the experiments, which are consequently subject to criticism for the reasons discussed here. Whilst I acknowledge these criticisms, alongside investigating learning under incidental conditions and predicting human learning performance on the basis of associative theory I hope to provide a convincing account of automatic processes, as I will discuss in the following section.

1.4.4. Other methods to demonstrate incidental learning

Rather than relying on explicit tests entirely, Seth, Dienes, Cleeremans, Overgaard and Pessoa (2008) provide a review of behavioral and biological measures of explicit processes, and conclude that a variety of measures are required to converge on any assumptions about the nature of the processes involved. This section briefly considers additional measures or tests that further attempt to understand what sort of learning occurs in SRT tasks.

1.4.4.1. Process purity

Firstly, it is worth restating here that, although tests may be classified as direct, indirect, subjective, objective, or any other classification; it does not logically follow that they measure any particular mental process (Newell & Dunn, 2008). Reaction times and accuracy on a task, for example, do not exclusively measure performance in the absence of conscious attention and motivation (Merikle & Reingold, 1992). We cannot isolate the study of a mental process such as learning without the influence of explicit thought as well as a whole host of other perceptual and cognitive processes at work, so we cannot assume that any test is process-pure (Jacoby, 1991). Assuming that all measures of learning are therefore capturing other mental processes, which is the best in this case

to use? Indeed, they do not just measure learning, but also the amount of knowledge that is stored, and are therefore linked intrinsically with memory systems. Combined with other task demands, perceptual influences, variables such as motivation and how well participants understand the instructions it is very difficult to compare the results of any two different tasks and be sure of what this difference might mean.

Shanks and St John (1994) propose that this is an intrinsic problem for implicit learning studies that can only be solved under the assumption that measuring learning after the learning event should be used as a marker for what was learned during the event, as to assume that it decays before this point would suggest that the learning is too weak to be of interest. If this is the case, we must measure both implicit *and* explicit knowledge at the same time, with matched sensitivity. The process dissociation procedure (PDP, Jacoby, 1991) attempts to do just this in order disentangle the two processes, originally used in the context of implicit and explicit memory.

In the context of a sequence learning task (Destrebecqz & Cleeremans, 2001; Wilkinson & Shanks, 2004) this involves two tests (whose order may be counterbalanced) where participants are required to: produce responses that follow the trained sequence (inclusion test); and produce responses that do not follow the trained sequence (exclusion test). Implicit knowledge will act to encourage responses consistent with the trained sequences on both tests, whilst explicit knowledge will only positively influence correct sequence generation on the inclusion test. Using this procedure with a response-stimulus interval of 0 seconds (as discussed in Karazinov & Boakes, 2007) participants provided evidence for sequence learning in the exclusion test and therefore of implicit learning. However, Wilkinson and Shanks (2004) failed to replicate the results and a variety of other authors have called this procedure into question (Curran & Hintzman, 1997; Dodson & Johnson, 1996; Graf & Komatsu, 1994).

1.4.4.2. Biological solutions

A proposed solution is to employ a biopsychological or neuropsychological measure as an index of learning. fMRI (Willingham, Salidis, & Gabrieli, 2002); event related potentials (Fu, Bin, Dienes, Fu, & Gao, 2013); eye-tracking (Marcus, Karatekin, & Markiewicz, 2006); PET (Destrebecqz et al., 2005) have all been used to demonstrate some support for implicit learning of sequences. These offer sophisticated and real-time

measures of on-line learning, unlike direct tests. These methods, however, are often costly and give little clear evidence of implicit learning, as we must still infer their index of implicitness from the studies whose status as implicit is the issue under question.

Another solution is to use clinical populations, which may be impaired on learning tasks (e.g. amnesiacs, Clark & Squire, 1998; McGlinchey-Berroth et al., 1997; Squire, 1992), where we can compare performance with normal participants and assess learning in comparison to participants who are unable to learn. Whilst this provides evidence for the sort of neurological structures that may be involved in learning, this has a multitude of issues, namely in comorbidity of deficits and a lack of understanding regarding the nature of the learning impairment, as the brain area may selectively impair perception, attention, acquisition, retrieval or a variety of other associated components of the learning process.

1.4.4.3. Computational considerations

In a comprehensive review of measures of implicit and explicit knowledge Seth et al. (2008) conclude that there is no definitive measure of either implicit or explicit knowledge, and recommend a combination of measures. This thesis takes their view, that measurements should provide results which build on understanding and that an integrated approach based on some theoretical framework offers a way in which to provide greater understanding of these processes. Hence, associative learning theory will be used to make specific, behavioural predictions about how humans learn sequences incidentally; and these will be tested experimentally.

As previously discussed, all tests, implicit and explicit; direct and indirect; biological or behavioural are all problematic as they may not simply measure learning per se, indeed, all psychological measures suffer from being simply manifest variables of the latent variables that we wish to measure (see Newell & Dunn, 2008). Rather than combining behavioral and neurological measures, which both suffer from this manifest issue, a computational approach means that one can directly measure any aspect of the model, then manipulate and quantify the learning processes at work. A computational model of human sequence learning can therefore be used to provide researchers with substantial explanatory power. Indeed, Cleeremans and Dienes (2008, p. 401) suggest that in the

study of human learning processes "what we really need to know are the principles by which a working computational model of human learning could be built". Given that the thesis attempts to investigate associative processes in humans, there seems no better method than to examine the predictions of a computational, associative model and test these on humans.

1.5. Variations within the standard SRT task

There are a number of features that have been varied across the standard SRT sequence learning paradigm that have led to various methodological refinement and also provide insight into how learning of sequences occurs. The obvious variable that one might manipulate within a sequence learning task is the sequence itself, which is discussed first in terms of the sequences used in SRT tasks. Participants demonstrate different learning effects depending on various sequence attributes: whether they are fixed or probabilistic or conditional; their length and the number of sequential elements that they involve. Further to this, participants demonstrate different *response* effects in an SRT task based on various attributes of the task trial order and these sequential effects are discussed in terms of both controlling for and accounting for these effects within a model of human incidental sequence learning. Sequences are not only defined by their trial order, but also by time itself and how this is manipulated within SRT tasks of sequence learning is also discussed. Finally, I will consider the two elements of the SRT task: perceptual characteristics of the stimuli and the nature of the motor responses required, and consider what manipulations of these can tell us about sequence learning and discuss how they are currently represented within models of human sequence learning.

1.5.1 Sequence learning and sequential effects

This section deals with difference between sequence learning and sequential effects. Throughout the thesis I refer to sequence learning as evidence of learning about the trained sequential patterns or contingencies intended by the experimenter. Sequential effects are generally defined as the influence of trial order on speeded responding in an SRT task (Anastasopoulou & Harvey, 1999) and, when considered, are generally taken as an experimental confound to be controlled for in sequence learning tasks (Jones & McLaren, 2009). It is entirely possible that participants automatically learn about many

aspects of the task outside of the sequence of interest, including random elements of trial order. Whilst sequential effects attract theoretical interest and themselves , for the purpose of this thesis they are considered separately to learning about the trained sequence itself for purposes of clarity.

1.5.1.1. Sequence learning

As mentioned in this section introduction, there are a variety of sequences that have been used in the sequence learning literature. These can be divided into roughly two paradigms, the first being those experiments that use one or more fixed, repeating sequences to train participants. Participants in Nissen and Bullemer's (1987) original demonstration of implicit sequence learning, for example, were required to press four different keys for four different stimuli (A, B, C and D) that followed the sequence: D-B-C-A-C-B-D-C-B-A continuously ten times in a block for eight blocks (with no demarcation between each sequence). The second type of sequence learning studies use, instead of this repeating sequence, some sequential structure. Rather than trial locations following a fixed sequence, they can follow a set of underlying rules that results in an abstract structure. Cleeremans and McClelland (1991), for example, employed the same Markovian structure used by Reber (1967) to construct his artificial grammars.

The first, fixed sequential structures are used in most of the early sequence learning literature (Stadler & Neeley, 1997), which lead to a variety of issues. These sequences in some instances do not control for the number of stimulus presentations in each location, therefore participants in Nissen and Bullemer's (1987) study may have been able to respond less to 'D' and 'A' and more to 'B' and 'C' to give them some advantage on the task (DeCoster & O'Mally, 2011). They also involve increased likelihood of certain first order transitions, for example B is followed by C twice, but never by A; as well as there being no possibility that stimuli can repeat. There are therefore two essential problems with fixed sequences: the first being that some elements or structural components within a sequence may be learned, rather than the whole sequence itself. This suggests that explicit knowledge tests may fail the information criterion (Shanks & St John, 1994) as well as demonstrating that participants may not be demonstrating sequence learning at all (Stadler & Neeley, 1997).

The second problem is that certain trial orders may naturally produce response differences, for example participants may be faster to respond to a repetition compared to an alternation (Soetens, Boer, & Huetings, 1985). These sequential effects are a vital point of consideration for sequence learning studies when attempting to demonstrate learning, as well as being worthy of interest in their own right (see section 1.5.1.2). The presence of these sequential effects is an issue as many studies not only use a fixed sequence that may suffer from them but the *same* fixed sequence used by Nissen & Bullemer (1987; Stadler, 1992), or a similar 12-item fixed sequence introduced by Reber and Squire (1998; DeCoster & O'Mally, 2011).

Probabilistic sequential structures do not follow the same, fixed sequence throughout, with a propensity to discourage explicit processes (Cleeremans, 1993); they have an abstract structure (Cleeremans et al., 1998); following the maxim of complex, abstract stimuli relationships suggested by Reber (1989) and hence are preferred in the study of sequence learning (Jones & McLaren, 2009). However, this does not exclude them from sequential effects, as artificial grammars such as the one used by Cleeremans and McClelland (1991) do not control for the number of stimulus presentations in each location nor the number of first order transitions between trials (Anastasopoulou & Harvey, 1999). Reed & Johnson (1994) suggest instead that the SRT task should always use second order conditional (SOC) sequences, where the location of the response stimulus on each trial is uniquely determined the previous two trial locations. These match for the number of locations and first order transitions between stimuli, although these often do not allow repeats to occur, which introduces a sequential effect in itself. Whilst first order transitions can be balanced, there will be an influence of higher order sequential effects (i.e. the effect of trial orders preceding the previous trial; Anastasopoulou & Harvey, 1999; Soetens, Boer, & Huetings, 1985).

If we balanced the number of trial orders precisely in a task then we can propose there would be nothing to learn about, as trial order would be controlled and thus pseudorandom with every possible response location and transition equally likely (Anastasopoulou & Harvey, 1999). Therefore, in order to control for the inevitable influence of sequential effects when a sequence is introduced, we must compare not only performance to a control group who have not experienced these sequential effects, but we must eradicate the influence of these sequential effects in a matched,

pseudorandom test phase where learning in the sequence-trained groups can be demonstrated compared to control participants.

Studies that have been successful in controlling for these sequential influences (Anastasopoulou & Harvey, 1999; Shanks & Johnstone, 1999), as well as presenting a full account of both RT and error data are limited (Jones & McLaren, 2009). I recognise that controlling for sequential effects produces significant methodological constraints in terms of the type of sequences that participants can learn, the need for a between subject control, and the need for a test phase. Following Jones and McLaren (2009) this thesis involves a two-choice SRT task, which enables a simpler control for sequential effects, as the number of stimuli, and therefore of first-order transitions, and therefore higher-order transitions are limited and more easily balanced. This structure also allows for a simple, probabilistic structure to be introduced to relationships between stimuli (Jones & McLaren, 2009; Lee & Livesey, 2013), for example, the likelihood of a first-order repeat or alternation can be changed in the sequence-trained group; and controlled for by the pseudorandom control group where either is equally likely.

1.5.1.2. Sequential effects

Whilst sequential effects have been proposed as something that sequence learning studies are required to control for (Anastasopoulou & Harvey, 1999) they make up a thriving research literature in and of themselves in understanding human perception and performance (Soetens, Melis, & Notebart, 2004). I want to outline how this impacts upon a computational associative account of sequence learning.

Computational models of sequence learning such as the AugSRN (Cleeremans & McClelland, 1991) have attempted to account for these effects, primarily the effect of the previous trial ($t - 1$) on the current trial (t) in order to better account for the variance in human performance on the task. The first-order effect observed by Bertelson (1961, 1963) that participants are faster at responding when the t is in the same location as $t - 1$ (a first-order repeat) was incorporated in the model by adding response-units which introduced short-term priming of the previous response. Jones and McLaren (2009) observed this first-order repeat preference and it was well modelled by the AugSRN, which provided a good explanation of human behaviour on the task. Therefore, the model contains a non-associative (no long-term learning occurs with respect to this

response priming) component in order to account for certain response patterns. This thesis aims to similarly account for human response preferences with a computational model, and so sequential effects as well as sequence learning will be considered.

1.5.1.3. Associative predictions

As mentioned previously, Jones and McLaren (2009) provide a demonstration of sequence learning that is not confounded by either sequential effects or a speed-accuracy trade off, thus making it the methodological starting point for the experimental work presented in this thesis. Additionally, the account that Jones & McLaren (2009) provide of their data is associative, and the AugSRN and human performance under incidental conditions are equivalent, in terms of both sequence learning and sequential effects. The specific details of the paper are discussed at length in the introduction to Chapter 2, but in brief the amount of sequence learning that occurs under incidental conditions differs in the task depending on the particular trial order of certain 'subsequences', which are all taken from a probabilistic structure based on the same underlying rule.

The AugSRN provides evidence that these differences are the result of competition between trial-by-trial associations, which reduces the error term for certain sequences that hence restricts learning. Jones and McLaren (2009) therefore suggest that learning about relationships between previous trials can block learning about sequences on subsequent trials. This might also suggest, however that participants in the task are not able to extract the abstract sequential structure trained on the task and instead learn specific instances. Consequently, the first Chapter in this thesis aims to further investigate these claims regarding competition between trial-by-trial associations within a similar SRT task design, but using a different underlying stochastic structure to the trials experienced by their (Jones and McLaren's) sequence-trained participants to investigate the associative predictions regarding sequence learning.

1.5.2. Time

The study of sequence learning is primarily concerned with variations in trial order, as this is the essential information that is acquired when learning a sequence. The influence of time itself on sequence learning has attracted considerable attention in recent years. SRT sequence learning tasks are generally fast paced, a design feature

borne out of attempting to avoid the explicit development of propositions (i.e. rule induction) that is shared by human associative learning studies (e.g. Karazinov & Boakes, 2007). This suggests that explicit processes are thought to occur given greater time between stimuli within an SRT task, and indeed this was found by Destrebecqz and Cleeremans (2001; 2003) as manipulating response to stimulus interval (RSI) had no effect on incidental learning, but participants demonstrated explicit knowledge when RSIs were increased to 1500 milliseconds (ms). Further authors were unable, however, to replicate these results (Fu et al., 2008; Norman, Price, & Duff, 2006; Wilkinson & Shanks, 2004). Indeed, Frensch and Miner (1994) propose that time has the opposite effect on learning, and suggest that increasing the RSI produces a reduction in learning. Researchers tend however, to agree that RSI has little influence on incidental learning, however fixed or patterned RSIs produce greater incidental sequence learning than random RSIs (Shin, 2008), which suggests that time is encoded and learned about when learning about sequences (Miyawaki, 2006; Rünger, 2012). This is not something that models of sequence learning such as the SRN account for as time is simply represented in a step-wise trial-by-trial fashion (Destrebecqz & Cleeremans, 2003).

In addition to this, time has a differential effect on sequential effects (Soetens, Melis, & Notebart, 2004) as increases in RSI tend to produce a change in first-order response preferences from repeats to alternations. This is again suggested to be due to the influence of explicit expectations, this time akin to a gambler's fallacy heuristic where participants expect an alternation to be more likely (Jarvik, 1951). This suggests that both sequence learning and sequential effects may interact with one another and time to produce the pattern of responding in humans. Jones and McLaren (2009) used a fixed RSI of 500 ms in their task, which is shorter than the 'explicit' RSI manipulation used by Destrebecqz and Cleeremans (2001; 2003) but also longer than similar experiments, such as Cleeremans and McClelland (1991). Jones and McLaren (2009) accordingly increase the learning rate of the AugSRN to better capture the longer RSI, which crudely represents time. The influence of time on sequence learning is addressed by the simulations involved in Chapter 3 of this thesis, which consider the influence of time and order within each trial, as well as between each trial.

1.5.3. Stimuli and responses

One much debated element of sequence learning tasks is to what extent the perceptual and motor components of the task are learned and represented. Whilst Willingham (1999) demonstrated that learning about sequences can be entirely response based; other demonstrations suggest that sequence learning can occur to stimuli that are purely visual in nature (Heyes & Foster, 2002; Howard, Mutter, & Howard, 1992; Marcus, Karatekin, & Markiewicz, 2006) with no responses whatsoever. In a study by Mayr (1996) participants were trained in a task where both the shape of the stimuli and the location of the stimuli followed two, different, independent sequences. Participants showed evidence of being able to learn both sequences when required to respond to shapes and not locations, suggesting that participants were able to learn both response-based sequences regarding shapes alongside stimulus-based sequences regarding locations (Mayr, 1996). It is debated whether sequence learning involves associations between stimuli (S-S learning, Heyes & Foster, 2002; Howard, Mutter, & Howard, 1992); responses (R-R learning, Nattkemper & Prinz, 1997); or the previous response and current stimulus (R-S learning, Ziessler & Nattkemper, 2001). Whilst studies are concerned with which one of these occurs, it is entirely possible that more than one occurs.

Models of sequence learning may consider stimuli and responses a moot point, as when represented locally within a computational model, responses (if correct) can be represented in the same way as stimuli. Associative models of sequence learning generally represent a sequential element as a single event (t) that predicts the next sequential element ($t + 1$). I will argue in this thesis that this does not accurately represent the task conditions experienced by humans in an SRT task. It is entirely possible that an associative system can learn S-S, R-R and R-S associations, as well the often overlooked associations (in this context anyway) between stimuli and responses (S-R learning). Participants are not required to associatively learn that each stimulus in the task requires certain responses, as they are instructed how to respond, which is why models and theories are not concerned with stimulus-response learning; as this is not involved in the learning of sequences. However, there is a perfect contingency between these stimuli (e.g. the light that flashes) and the required response (e.g. the key to be pressed), so we can expect that practice will have some effect on the strength of the

stimulus-response association. Whilst these stimulus-response associations are not a necessary condition for learning about sequences associatively, I will investigate to what extent an associative model predicts how these associations might interact with incidentally learned sequences. Similar to the account of sequential learning effects observed in Jones and McLaren's (2009) study, stimulus-response associations may compete with sequence learning and produce a differential effect.

1.6. Variations on the standard SRT task

There have been a number of variations on the SRT task that introduced dual tasks (Cohen et al., 1990; Curran & Keele, 1993; Jiménez & Mendez, 1999) in order to investigate whether sequence learning requires cognitive effort or selective attention. Whilst these produce interesting results that can inform our understanding of sequence learning, this section will focus on the variations on the standard SRT task that, following from the previous section regarding stimuli and response, associative learning theories may provide specific, testable predictions about. These tasks are those that include *additional* stimuli, which may be expected to interact with the sequence learning in some way. The number of such studies to date is small and I will conclude that they provide no suitable evidence that incidental learning about two sets of contingencies interact. Finding an associatively predicted effect of cue-competition in implicit learning literature as a whole is lacking when the presence of such an effect achieved under incidental conditions would provide strong evidence for the occurrence of automatic, associative processes (Beesley & Shanks, 2012). A central aim of this thesis is therefore to investigate how additional stimuli or cues within a sequence learning task have an effect on the incidental learning of these sequences in order to demonstrate associative learning in humans.

1.6.1. Additional concurrent stimuli

Sequence learning tasks usually involve visuo-spatial stimuli, but some tasks also involve additional stimuli or stimulus elements (e.g. colour or shape). In recent studies by Abrahamse and colleagues (Abrahamse, Lubbe, Verway, Szumska, & Jaskowski, 2012), sequence learning tasks have investigated the claim that sequence learning can be potentiated by concurrent stimuli that also follow the sequence (Robertson & Pascual-Leone, 2001; Robertson, Tormos, Maeda, & Pascual-Leone, 2001). Robertson

and colleagues demonstrated that the learning of sequences of stimulus locations or colours of sequences was stronger when these two sequences were contingent than when only locations or only colours were used alone; however, they did not use the same sequences nor the same tasks at test to assess learning correctly. In doing so Abrahamse et al. (2012) found that participants show no difference between learning about locations alone compared to when colours were congruent with these. This suggests that other stimuli have little effect on sequence learning, and that sequence learning is strongly linked to responding. This merits further investigation, as Abrahamse et al. (2012) investigated location alone sequence learning with a single colour stimulus, whereas the location and congruent colour condition involved four colours and locations. It is possible therefore that colour-location associations in fact interfered with learning about the sequence, or that the colour information was not perceived as separate to the location information and the one colour-response unit was represented as a single compound stimulus.

1.6.2. Additional between-trial stimuli

Nissen and Bullemer (1987) provided evidence that in a dual-task version of their SRT task that an additional counting task that involved tones presented during the RSI disrupted incidental sequence learning. They took this to suggest that reduced attention to the sequence led to reduced learning, indicating that implicit learning requires some element of selective attention. Stadler (1995) provides a different account of these results and suggests that the concurrent tone counting task actually disrupts the sequential order as it is encoded, and that it is not attention but interference from these additional tone stimuli that impairs sequence learning. This, to some extent, may be interpreted associatively, as a representation of a separate element within the RSI may influence the associations made between the sequential stimuli or responses. Indeed, if the responses are the only element of the task implicated in the learning of sequences then additional between-trial stimuli should not have an effect on sequence learning, however, if stimuli also drive sequence learning then we can expect associations between these stimuli to be disrupted.

It is not clear in these dual tasks, however, to what extent incidental processes are interfering with one another, as the secondary task is made explicit to participants and therefore a host of influences including working memory, perceptual and cognitive load

could produce these effects. In this thesis I will aim to investigate further the possibility that sequence learning is affected by additional stimuli presented concurrently within the SRT task. However, instead of examining these stimuli as part of an additional dual-task scenario where additional, explicit influences on processing may account for differences in performance, I will investigate how the associations formed between these additional stimuli and existing elements of the task interact with the learning of sequences.

1.6.3. Additional cues

Further to this, there are a number of studies that investigate sequence learning alongside learning about additional contingencies. Beesley and Shanks (2012) proposed that excellent support for the presence of associative learning processes in humans would be provided if a task could show that learning about two contingencies competed with one another following established learning effects observed in animals, for example blocking (Kamin, 1969). Whilst they found no evidence of cue-competition in humans using a visual search task, I will argue in Chapter 6 of this thesis that such a task may not involve two forms of competing incidental learning, and further examine whether we can find evidence of an interaction between learning about two incidentally acquired contingencies. This would provide strong evidence that humans were learning associatively and is therefore a central aim of this thesis.

In the context of SRT tasks, Cleeremans (1997, p. 74) investigated a novel version of the SRT task where the colour of the stimuli on each trial predicted the next response in the sequence, which was "inspired by work on overshadowing in conditioning experiments with animals". With an aim to investigate whether colour-response learning would overshadow sequence learning, Cleeremans (1997) found that whilst participants learned about colours that this did not interact with sequence learning, providing no evidence of overshadowing and therefore no evidence for associative processes. This conclusion, however, may be flawed for a number of reasons. Firstly, participants were instructed about the presence of colour contingencies and therefore this learning is likely to have been explicit, which could explain the lack of an associative interaction between the two learning processes (Jiménez & Méndez, 1999).

Jiménez and Méndez (1999) adapted Cleeremans' (1997) task and provided evidence from a version of the study where participants did not receive explicit colour-contingency information. They were trained incidentally that colours predicted the next trial alongside sequence learning and no evidence colour-response learning, nor of an interaction between the two was found. Jiménez and Méndez (1999) concluded that there was no evidence of associative overshadowing, however, both of these studies did not provide a control group for sequence and colour contingency learning, which is an investigation that will be conducted within this thesis. Furthermore, the studies did not entertain the possibility that blocking or overshadowing may occur in the other direction: with the colour contingency learning being overshadowed by the sequence learning. Both studies assume that colour learning will be stronger as it is firstly a simple relation to learn, however, this presupposes that associative learning systems have some advantage when learning simple conditional probabilities over more abstract stochastic structures. On the one hand this makes some intuitive sense, but the rich variety of complex and abstract sequences that a simple SRT model have been shown to learn suggest that an associative system is ideally suited for learning these stochastic relationships (Beesley, Jones, & Shanks, 2012). The authors (Cleeremans, 1997; Jiménez & Méndez, 1999) further suggest that overshadowing of sequences is expected by colours due to the temporal contiguity of colours with the next response, whereas sequences occur with greater latencies before the trial to be predicted. Issues of time were discussed previously, and whilst no conclusive answer was reached, it is of course possible that learning increases over time, which was suggested and functionally modelled by McClelland (1979).

1.6.4. Summary

Altogether, these variations on the SRT task provide no conclusive evidence for or against automatic, associative learning processes. The methodologies of the various tasks outlined do not properly control for the number of stimuli across conditions (Abrahamse et al., 2012) or provide control groups for sequence or additional contingency learning for comparison (Jiménez & Méndez, 1999). Subsequently, authors have been unable to conclude that there is evidence for an interaction between stimuli or contingencies and sequence learning; which would provide insight into how incidental learning processes occur. Following Beesley and Shanks (2012) this thesis attempts to investigate cue-competition between contingencies within an SRT as would be expected

if human incidental learning were indeed associative, thus providing strong evidence that automatic learning processes occur through association formation.

1.7. Concluding remarks

Instead of "fighting old battles" (McLaren et al., 2013, p. 194) and merely attempting to demonstrate the presence of implicit learning, this thesis assumes the possible presence of automatic learning processes and aims to investigate whether, under incidental conditions, we can reliably study and computationally model human associative learning. To best study associative processes in humans, a sequence learning approach is adopted which is considered the best task to investigate incidental human learning (Destrebecqz & Cleeremans, 2003). Further to the methodological advantages provided by investigating associative learning using this task, sequence learning is an interesting phenomenon in and of itself.

This thesis aims to investigate human incidental learning at the intersection of computational, implicit and associative approaches. I will examine human learning under incidental conditions and ask whether an associative model can capture the sequential learning observed in participants, as well as the sequential effects that occur in humans. Further to this I will investigate the effect that additional stimuli and contingencies between cues and responses in the SRT task have on incidental sequence learning. Through this, the intention of the thesis is to investigate evidence of the automatic formation of links between events in the environment and to investigate functionally how these associative processes may occur in the human mind.

Jones & McLaren's (2009) results demonstrate that the trial order of the sequences themselves produce effects that, according to associative predictions, can come to interfere with learning about certain sequences. This provides the basis for the first experiment in this thesis, which investigates further the possibility that associations formed between trial-by-trial random orders compete with learning about stochastic structures (trained, probabilistic elements of the trial order) within an incidental sequence learning task. Chapter 3 takes a computational approach in order to investigate how a model of sequence learning (the AugSRN, Cleeremans & McClelland, 1991) accounts for the human data provided in Chapter 2. As outlined in this introduction, the issues of trial order, timing and the role of stimuli and responses are all key variables

within the SRT task, which are addressed in the context of this modelling work. Chapter 4 investigates the role of additional concurrent stimuli within a SRT sequence learning task using predictions of the model outlined in Chapter 3. Further predictions regarding the learning of incidental cue-response contingencies are tested in Chapter 5, which are then incorporated alongside sequence learning in Chapter 6 to investigate whether evidence for cue-competition between two incidental contingencies can be found.

Chapter 2. Sequence learning, subsequence learning and sequential effects

In this chapter I investigated sequence learning under both incidental and intentional conditions in an experiment inspired by the work of Jones and McLaren (2009). In two experiments presented here, participants were trained across two sessions on a two-choice SRT task that followed a stochastic structure determined by one of two rules. Both rules involved the trial before the last, so the current trial could be predicted based on whether the trial two trials previous was a left or a right two thirds of the time. Jones & McLaren (2009) observed a double dissociation between learning of the subsequence XXX under incidental and intentional conditions. As part of learning of a sequential rule, participants demonstrated an absence of learning about XXX under incidental conditions, but exclusively learned this subsequence intentionally. Learning of this subsequence in Experiment 1 presented here suggests that impaired learning of subsequence XXX depends on the structure of sequential contingencies, providing support for an associative account of learning under incidental conditions. Experiment 2 provided evidence of learning under intentional conditions that was extremely similar to Experiment 1, and these results are discussed in the context of methodological issues with dissociations and state-trace analysis (Bamber, 1979). The results of Experiment 1 provide a detailed set of sequence learning and sequential effects that provide a framework within which to investigate computational models of learning in Chapter 3.

2.1. Introduction

2.1.1. Jones and McLaren (2009)

Chapter 1 argued that sequence learning provides an ideal paradigm in which to study implicit or incidental learning in humans, as this involves complex sets of contingencies that are hard to notice. Sequence learning also provides a challenge for traditional models of associative learning, as learning contingencies between events that are separated by time and intervening stimuli is beyond simple associative models that have no means of representing previously experienced stimuli (i.e. some memory). This is an issue for associative explanations of automatic learning in humans, when learning about

sequences and series of events across time. In this section I will outline the procedures and results obtained in the sequence learning task devised by Jones and McLaren (2009), which forms the basic experimental design for Experiments 1 and 2. Jones and McLaren (2009)'s account of the pattern of results, supported by the simulations of the Augmented SRN (AugSRN, Cleeremans & McClelland, 1991), suggest that competing contingencies between events in the sequence produce learning effects for certain trial orders.

Participants were trained over six sessions (two for the intentional participants) of roughly one hour on a simple two-choice serial reaction time (SRT) task that involved pressing one of two keys to one of two on-screen stimuli. These stimuli were simply white circles that always appeared in the same location, either on the left or right hand side of the screen. Under incidental conditions participants were told that the task was simply measuring reaction times (RTs) and accuracy to these stimuli, and no mention was made of the sequence learning that could occur. Under intentional conditions participants were told to look for patterns and sequences and work out what was going on to help them in the task.

As discussed in Chapter 1, this sequence learning was not of repeating strings of trials, but the trials were constructed with a probabilistic structure whereby some trial orders were twice as likely as others. This meant that with only two stimuli, sequential effects could be adequately controlled for within-subjects. The number of right and left stimuli; the number of repeats and alternations experienced between the current trial (t) and the previous trial ($t - 1$); as well as between t -1 and the trial before this ($t - 2$) were equal for each subject, and therefore no difference in the distribution of these sequential features could explain the observed learning. The unavoidable nature of training certain trials orders as more likely however, results in a possible sequential confound, as improved performance in a group trained on certain trial orders could be due to some preference for responding to these trial orders and therefore a (between-subjects) control group is required to control for these sequential effects.

The trial orders trained (i.e. those more likely to occur) were based on an exclusive-or rule, a classic learning problem in connectionist modeling literature (Minsky & Papert, 1969; Rumelhart, Hinton, & Williams, 1985). This rule is based on two items, which

produce either an 'exclusive' case (both items are the same) or an 'or' case (both items are different). Translated into the current task, if the rule occurred on 100% of trials: when one receives two consecutive stimuli on the right side of the screen, or two on the left (the exclusive-case) then the next trial would be a right trial, for example. If the two trials just experienced swapped from right to left, or left to right (the or-case) then the current trial would be a left trial. These contingencies, however, did not occur on 100% of trials and were trained by way of a stochastic structure whereby participants could use this rule to predict the current stimulus location on two thirds of trials during training (although the trials were pseudorandom and the rule predicted the response at 50% during test).

Within this structure, therefore, there was a stochastic, underlying sequential exclusive-or rule to the trial order that participants could learn about. Due to this structure, certain runs of trials to $t - 2$ were more likely. Taking the example above where exclusive-case = right stimulus (R) and or-case = left stimulus (L) then the trial orders: RRR, LLR, RLL and LRL were twice as likely to occur than: RRL, LLL, RLR, LRR. Consequently, as well as investigating the extent to which humans can learn the probabilistic sequential exclusive-or rule, learning could be broken down into these triplets or "subsequences" (Jones & McLaren, 2009, p. 541). As the exclusive-or rule was counterbalanced across participants between right and left stimuli, they are referred to in terms of Xs and Ys rather than Rs or Ls. Jones and McLaren (2009) provide evidence of differential learning of each of these subsequences; and so it follows that an associative model of human learning must account for these subsequence learning effects (for example, more learning of RRR compared to RLR) as well as sequential effects (the effect of the previous trial order – trained or untrained – on responding) and sequence learning (learning of the rule that underlies the probabilistic structure).

2.1.3. Cue competition and subsequence learning

This subsequence learning is investigated further in this chapter, as Jones and McLaren (2009) observed that under intentional conditions, participants only demonstrated learning of the subsequence XXX (and not XYY, YXY, or YYX). The opposite was true of participants who completed the task incidentally, who showed no evidence of learning the subsequence XXX nor XYY, some of YXY and evidence of learning of YYX. The explanation of XXX learning in the intentional case and not the incidental

case was that participants simply noticed this salient string of trials and therefore explicitly performed better on these trials, as evidenced by reports obtained in a structured interview following the task. Learning of this subsequence, Jones and McLaren (2009) suggest, can therefore be used as a marker of whether participants are engaging with rule-based, hypothesis testing.

The apparent dissociation between the incidental learning of subsequences is difficult to reconcile with an explicit rule-based account. An explanation, however, can be provided by the AugSRN, which provides a convincing simulation of human performance on the task. The AugSRN (as described in Chapter 1 and in more technical detail in Chapter 3) learns by way of two sets of connection weights between the units in the model, both with their own learning rates. The 'slow' weights (lower learning rate, no decay) are simply the associative connections that learn the underlying statistical structure of the contingencies in play: these weights learn the exclusive-or sequential contingencies and simulate sequence learning. The 'fast' weights (higher learning rate, experience decay) were introduced by Cleeremans and McClelland (1991) to account for short-term priming effects observed in SRT tasks: transient learning of trial-by-trial mappings that produce sequential effects observed in humans. Jones and McLaren (2009) suggest that it is the competition that occurs between these fast, transient associations and the learning of the underlying statistical contingencies that influences the learning of subsequences.

Jones and McLaren (2009) propose that for the XXX case, the trial-by-trial mappings are the same (X [t – 2] leads to X [t – 1], followed by X [t – 1] leads to X [t]). This leads to an effect akin to blocking: the transient learning of these trial-by-trial contingencies by the fast weights reduces the error term involved in calculating the amount of learning that can occur. This restricts learning of the mapping between t – 1 and t, which is the crucial trial on which the exclusive-or rule predicts a response. As a result, due to the competition between transient trial-by-trial learning (sequential effects) and the more permanent learning of the stochastic structure, simulations of the AugSRN (and humans under incidental conditions) are impaired on learning of this subsequence. This cue competition account of sequential effects and sequence learning not only accounts for the observed subsequence learning differences, but as such provides excellent evidence for automatic cue competition and therefore associative processes in humans.

It may also, however, be possible that participants under incidental learning were using some chunking, or instance based learning, which resulted in the differential learning of some sequences due to some attentional bias towards or away from subsequences with repetitions; or some other motor or perceptual preference for certain subsequences. Indeed, these results could provide strong support for such a theory, suggesting that participants were not learning the stochastic exclusive-or rule at all, but simply learning some of these subsequences as instances. It could also be possible that learning under incidental conditions is limited, and participants were restricted to learning only a small number of subsequences, chunks or instances. This seems unlikely, as there were only four subsequences of three trials to learn in total.

This chapter aimed to investigate Jones & McLaren's (2009) associative cue competition account of subsequence learning in a similar task under both incidental and intentional conditions, matched for training length unlike in the original experiments. The experiments sought to investigate further whether competition from within-subsequence contingencies were responsible for the learning deficit of XXX under incidental conditions, or whether a chunking-based instance account could explain this result. This experimental chapter was further designed in order to generate additional learning effects that a computational model would be required to account for (which are discussed in Chapter 3). The experiments finally attempted to assess whether participants trained on the subsequence XXX have "the potential to be used as a marker as to whether people are engaging error-correcting associative learning or rule-based hypothesis-testing" (Jones & McLaren, 2009, p. 540).

2.2. Experiment 1: Incidental sequence learning

Instead of using a counterbalanced exclusive-or rule for the construction of the sequential contingencies that participants were trained on, this experiment involved a between-subject design comparing learning of one set of sequential contingencies to another. These contingencies followed the same rule structure: one could predict the current trial based on the trial before the last. This differs from the Jones & McLaren (2009) sequential contingency, as the exclusive-or contingency depends on both the previous trial and the trial before the previous trial jointly determining the likelihood of either response on the current trial. The between-subject manipulation was complementary: one group was trained that the current trial was more likely to be the

same as the trial two trials previous (Same), and the other group could predict the current trial as more likely to be the opposite response to the one two trials previous (Different).

Both Same and Different groups experienced the same number of left and right stimuli and the same number of repeats and alternations between responses. The two groups differed only in respect to the construction of their training trial order, which comprised of subsequences of XXX, YYY, XYX and YXY for Same, and XXY, YYX, XYY and YXX for Different (described in more detail below, see section 2.2.1.3.). As the subsequence learning observed under incidental conditions by Jones & McLaren (2009) suggested greater learning of the subsequences YYX and YXY, with some difficulty in learning the subsequences XXX and XYY (those ending in a repeat), this suggests that learning about the Same rule may be reduced compared to the Different rule. However, the Different rule group was trained on the subsequence XYY, which was also learned poorly under incidental conditions (albeit less so than the subsequence XXX).

As mentioned above, participants needed to be sensitive to both of the two previous trials in order to predict the current trial in Jones and McLaren's (2009) study, however this was not the case in Experiment 1. In the Same group (who are trained with XXX), the middle X is not required for a prediction; they are also trained with XYX. Therefore, Experiment 1 alters the stochastic structure of the sequence rule in order to assess the effect of trial-by-trial contingencies on sequence learning. The influence of short-term learning of previous trial order that Jones and McLaren (2009) suggest competes with sequence learning (reducing learning of subsequence XXX) may be altered if the relationship between the sequence elements that leads to these sequential effects is changed. Evidence for this interaction between sequence learning and effects would provide support for some competition between the two producing differential subsequence learning, as suggested by Jones and McLaren (2009). However, if subsequence XXX is not learned under incidental conditions then an instance or chunking based account that suggests that this subsequence is simply harder learn may provide a better explanation of the result.

2.2.1. Method

2.2.1.1. Participants

96 participants (aged between 18 and 29 [*M* = 21.1]; 70 female and 26 male) were recruited from undergraduate students at the University of Exeter and were awarded £10 for participation. Participants provided informed consent prior to taking part in two sessions lasting roughly one hour each. Participants were randomly allocated into one of the four possible conditions, described in more detail in section 2.2.1.3.

2.2.1.2. Materials and Stimuli

The experiment was run on an Apple iMac with Superlab software. Participants were seated roughly 50cm from the screen, which contained two white circle outlines on a black background throughout the task. These white circle outlines were 19 mm in diameter and positioned vertically in line with the screen centre, and 22 mm either to the left or the right of the screen centre horizontally. The response stimulus was a white filled circle (18.5 mm diameter) that was placed within one of the two circle outlines, giving the white circle outline the appearance of lighting up or filling in. Participants were required to press the spatially compatible 'x' and '>' key presses on a QWERTY keyboard to the left or right response stimulus, respectively.

2.2.1.3. Design

The experiment was a two-choice SRT task comprising two sessions of twenty blocks each. Each of these blocks contained 120 trials, with all twenty blocks of the first session and first fifteen blocks of the second session acting as training; and the final five blocks acting as test. Depending on the group that participants were assigned to, across training participants received either blocks containing sequential contingencies (Experimental) or no sequential contingences (Control). There were no sequential contingencies present in any group during the five blocks of test. Participants were randomly allocated either Same or Different sequential contingencies to learn (a dummy variable for the Control group). See Table 1.1 for the between-subject design of Experiment 1.

Pseudorandom block construction. The construction of the trial order of left and right response stimuli experienced in each block followed Jones and McLaren (2009): subsequences were used to construct both training blocks containing sequential

contingencies as well as pseudorandom test blocks. Within a two choice task there are eight possible trial orders that a run of three trials may take (R and L are used here for right and left, respectively): R-R-R; R-R-L; R-L-R; R-L-L; L-L-L; L-L-R; L-R-L; and L-R-R. For pseudorandom blocks (for all Control groups and for the five blocks of test for all groups) all of these eight triplets (5 instances of each: 40 triplets in total) were randomly ordered and concatenated to make a continuous string of 120 right or left trials. The amount of right and left trials in total, as well as the number of repeats and alternations experienced up to two trials before the current trial was therefore controlled for. On any given trial there was no way of using any information from the previous trial(s) to predict the current trial location. Therefore the probability of predicting the correct response on any trial, based on none or any of the previous trials, was 50%.

Table 2.1. Table showing the between-subject design of Experiment 1, with 24 participants in each of the four possible conditions of the 2 x 2 design (Rule: Same or Different; Group: Experimental or Control). Each cell contains the type of sequential contingencies that participants experienced for 35 blocks of training and 5 block of test, where all groups received pseudorandom blocks.

	Same Rule		Different Rule	
	Experimental	Control	Experimental	Control
Training	Same	Pseudorandom	Different	Pseudorandom
Test	Pseudorandom	Pseudorandom	Pseudorandom	Pseudorandom

Training block construction. For training blocks in the experimental groups, a similar triplet procedure was used, but instead of using all eight possible triplets, a subset of half of these were selected (what Jones & McLaren, 2009, p. 541, term "subsequences"). Participants in the Experimental Group were trained on either Same or Different rules, which I will describe in turn. The Same rule involved the triplets RRR; RLR; LLL; and LRL – all of which followed the same pattern: that the first and third item in the triplet is the same. Training blocks for participants in the Same condition were constructed using 10 of each of these four triplets, randomly ordered and concatenated to make 120 with equal likelihood of right or left trials, as well as repeats or alternations.

Unlike the pseudorandom blocks, participants could use the sequential contingencies between the current trial and the one two trials back to help them predict the response stimulus location and therefore their own response. Due to the triplet structure, on every third trial in the block participants were able (in principle) to predict the response stimulus location based on the trial before the previous trial 100% of the time. It is important to note that this structure was not made explicit to participants and trials occurred in series with no demarcation of triplets or any special status attributed to third trials as participants were required to respond to each trial. Every third trial was one of the four trained subsequences. On all other first and second trials in the block of a triplet, this sequential contingency occurred by chance – as the arbitrary ordering of the trained subsequences meant that all possible triplets occurred across a block. For example, RRR followed by LRL and LLL produced the following trial order: RRRLRLLLL. This set of nine trials therefore contains, on a trial by trial basis, the following triplets: Trial 1 and 2 – no triplet (not enough previous trials); Trial 3 – RRR (consistent with Same rule); Trial 4 – RRL (inconsistent with Same rule); Trial 5 – RLR (consistent); Trial 6 – LRL (consistent); Trial 7 – RLL (inconsistent); Trial 8 – LLL (consistent); and Trial 9 – LLL (consistent). Whilst every third trial had a 100% likelihood of the trained subsequences, on the first and second trial of each triplet participants had a 50% probability of being able to predict the current trial based on the trial two trials previous. Overall, this adds up to a partial reinforcement schedule of two out of three trials following the Same rule, a 67% probability of being able to predict the current trial as the same response location as the one that occurred before the last.

Those participants in the Experimental group and the Different condition were trained with the complementary rule, that the current trial could be predicted on 67% of trials as the *opposite* response location to the one that occurred two trials previously. The training blocks for these groups were constructed in the same way as for the Same condition, with ten of each of the subsequence triplets RRL; RLL; LLR; and LRR randomly ordered and concatenated for each block. This produced a 100% sequential contingency following the Different rule on every third trial and chance on other trials, leading to the same 67% sequential contingency as outlined for the Same group above.

During both training and test, it was therefore possible to compare how participants perform in both speed and accuracy of responding to trials that are consistent with those

subsequences that they have been trained on, as well as those subsequences that are inconsistent with the trained sequential contingencies. However, these Consistent and Inconsistent trials did not occur with equal frequency throughout training, hence five pseudorandom test blocks containing an equal likelihood of trials that were consistent or inconsistent with the trained sequential rule allowed for a matched comparison between Experimental and Control groups. Trials in the Control groups were also given dummy Consistent or Inconsistent labels, depending on whether participants were assigned the Same or Different condition. There were no differences between these two Control groups, who received 40 pseudorandom blocks with no sequential contingencies to learn, apart from the arbitrary label of each trial as Consistent or Inconsistent that depended on whether they were labeled as belonging to the Same or Different dummy conditions.

2.2.1.4. Procedure

After obtaining informed consent, participants were instructed to simply respond as quickly and accurately as possible to the stimuli, and that the task was investigating how practice had an effect on peoples' speed and accuracy of responding to simple stimuli. No mention was made of any contingencies, relationships, sequences or learning. Participants were reminded of these instructions at the beginning of the second session.

At the beginning of each block participants were instructed to press any key to start. Each trial began with an inter-trial interval of 500 ms where a black background with two white circle outlines was presented. The response stimulus (the left or right white circle) would then appear on screen until either the participant made a keypress response or the trial timed out after 4000 ms from the presentation of the response stimulus. RTs were measured from the onset of the response stimulus. If participants pressed an incorrect key, or the trial timed out, the computer issued a beep sound.

At the end of each block participants saw a white screen that told them their average RT and the percentage of trials on which they made an error for that block. Participants were not able to move to the next block for 30 seconds, during which a countdown timer was displayed on screen underneath their performance feedback. They were instructed to rest for these 30 seconds, with the next block beginning with a 'press any

key' command so participants could have rested for longer if required. Deviating from Jones and McLaren (2009), no monetary incentive was given for speed or accuracy, which was a decision based entirely on practical constraints. I hoped that participants were incentivised enough to perform quickly in order to minimise the time spent performing the task and that their errors would be tempered by the use of the loud beep tone when making an error, as well as the presentation of error feedback, which, anecdotally, participants seemed very aware of and encouraged them to improve throughout.

A short verbal structured interview was given at the end of the second session, in which participants were asked about what they had noticed in the experiment. They were asked structured questions that led from asking what they thought about the task; to whether they noticed anything interesting about the task; to whether they noticed any patterns or sequences; and finally whether they could describe or guess at any patterns or sequences. Participants were finally debriefed and thanked for their participation.

2.2.2. Results

Average reaction times (RTs) and proportion of errors were taken on each trial for each participant. Trials following an error were excluded from the analysis (Laming, 1979), and errors were counted only from those trials where participants made the incorrect response (e.g., pressed the 'x' key instead of the '>' key) and not for any other incorrect key press or trials that timed out (these and the following trial were excluded from the analysis). Each trial (apart from the first two in each block) was considered in terms of the preceding two trials, and therefore in each block there were eight possible trial types corresponding to each possible triplet. For example, every R trial preceded by two R trials was classified as a RRR trial. Therefore, depending on the rule underlying trained contingencies the eight subsequences can be divided into four consistent and four inconsistent subsequences, preceded by RR, RL, LL, or LR. Consistent sequences for the Same group are RRR, RLR, LLL, and LRL and inconsistent: RRL, RLL, LLR, and LRR; the opposite for the Different rule group. These Consistent and Inconsistent trials were arbitrarily assigned in the case of Control groups.

An average RT and error score was calculated for each of these triplets at test. Across training a weighted average was calculated, as outlined by Jones & McLaren (2009). As

the likelihood of certain trial orders up to three trials prior to the current trial ($t - 3$) was not balanced in the experimental groups this weighted average meant that for each subsequence an average was first taken, for example, of each RRR trial preceded by a R trial (a RRRR trial) and then separately for RRR trials preceded by a L trial (a LRRR trial). The two of these averages were themselves averaged, providing a control for sequential effects up to the third order, with effects any more trials previous to this producing negligible effects of around one millisecond (Jones & McLaren, 2009). Jones & McLaren (2009) found no special status of the third trial (which was 100% consistent with the rule) and so no attention was paid to the position (first, second or third) of the trial within the triplet structure used to construct the sequences. All 118 usable trials in a block (117 trials in a training block because of this weighting procedure) were therefore included in the analysis.

The weighted average procedure meant that 16 possible trial quadruplets were produced, each of which would occur 7.5 times per block on average if equally likely. The reduced likelihood of certain quadruplets of trials coupled with participant errors meant that simply analysing trials in one block was likely to result in missing data. Consequently, these averages were computed across five blocks (the next smallest factor of 35), which we refer to henceforth as Epochs (of which there are seven across training – 35 blocks; and one at test – 5 blocks).

Following this weighted average a difference score was calculated, under the assumption that those subsequences inconsistent with the trained sequential rule were predicted to have higher RTs and errors, hence participants should be slower and less accurate on these trials. Therefore, the average RT and proportion of errors for Consistent triplets were taken from the Inconsistent average RT and proportion of errors for each triplet (e.g. RRL minus RRR; RLL minus RLR; LRR minus LRL; LLR minus LLL for the Same rule group. The opposite calculation was made for Difference rule groups). A higher difference score in participants RTs or errors indicates better performance. This difference score is the dependent measure plotted on the graphs and used to analyse performance on the task. These four Subsequences were therefore referred to as RR, RL, LL and LR when comparing the two groups, referring to the two previous trials experienced before the current trial, which differed depending on the Same or Different rule.

Data was analysed first separately for each sequential rule group, with a mixed analysis of variance (ANOVA) examining the factors Group (Experimental versus Control); Epoch (7 sets of five training blocks); and Subsequence (RR, RL, LL, LR) across training, and Group and Subsequence at test. The data was then compared across the sequence rule groups, with training analysed by an ANOVA comparing: Group; Rule (Same versus Different); Epoch; and Subsequence, with test examined by Group; Rule and Subsequence. Both average RTs and proportion of error scores were analysed and within-subject main effects and interactions are reported with *p* values that correct for a departure from sphericity (Huynh-Feldt) with the unadjusted degrees of freedom.

2.2.2.1. Same rule learning

First we take the Same rule group, who were trained on the subsequences RRR, RLR, LLL and LRL (these were twice as likely to occur as RRL, RLL, LLR and LRR across training).

Evidence of learning. The ANOVA revealed a main effect of Group in both RT difference score, $F(1,46) = 160$, $p < .001$, $MSE = 1757$, $\eta_p^2 = .776$; and proportion of errors, $F(1,46) = 8.22$, $p = .006$, $MSE = .103$, $\eta_p^2 = .152$, across training, suggesting that the Experimental group performed better on trained subsequences than inconsistent subsequences than the Control group, see Figure 2.1. This effect persevered at test in the RT difference score, $F(1,46) = 49.1$, $p < .001$, $MSE = 515$, $\eta_p^2 = .571$; but not quite in the proportion of errors, $F(1,46) = 3.23$, $p = .079$, $MSE = .019$, $\eta_p^2 = .066$. The influence of a speed-accuracy trade off can be ruled out as the direction of the non-significant error difference scores show higher scores for the Experimental group. This suggests that we have good evidence that the Experimental group have learned to perform more quickly and accurately on trials consistent with the trained rule.

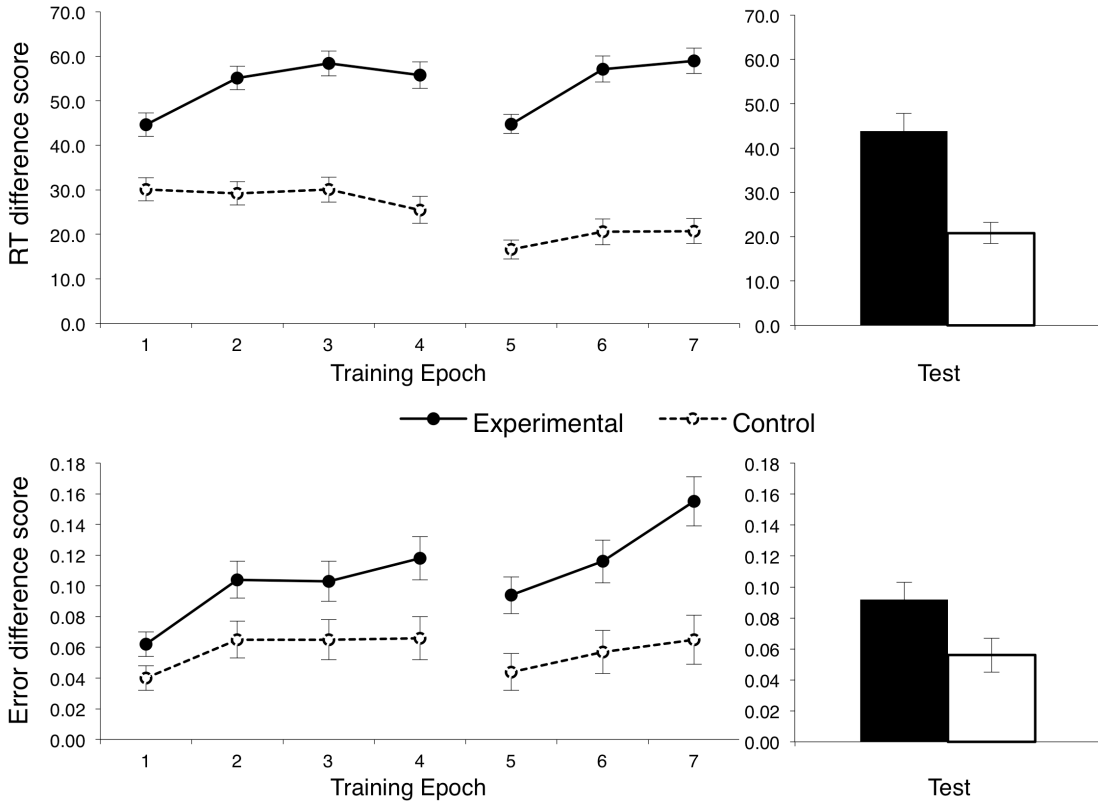

Figure 2.1. RT (top panel) and proportion of error (bottom panel) difference scores across training (left panel) and test (right panel) for Experimental and Control groups trained on Same rule subsequences. Error bars show standard error. Note: the blank between Epochs 4 and 5 represents the break between the two sessions.

There was a significant effect of Epoch, RT difference score, $F(6,276) = 11.8$, $p < .001$, $MSE = 384$, $\eta_p^2 = .204$; and proportion of errors, $F(6,276) = 9.84$, $p < .001$, $MSE = .010$, $\eta_p^2 = .176$. Figure 2.1 demonstrates the increasing trend of interest here, which indicates a general practice effect. The Experimental group performed increasingly better than the Control group as training went on, as demonstrated by a significant interaction between Group and Epoch in both RT difference score, $F(6,276) = 9.57$, $p < .001$, $MSE = 384$, $\eta_p^2 = .172$; and proportion of errors, $F(6,276) = 3.29$, $p = .012$, $MSE = .010$, $\eta_p^2 = .067$. This provides evidence for learning, which is not apparent in the first epoch of the experiment but clearly develops across training.

Subsequence effects. There was an effect of Subsequence on RT difference scores during training, $F(3,138) = 4.92$, $p = .011$, $MSE = 4865$, $\eta_p^2 = .097$; and errors, $F(3,138) = 17.9$, $p < .001$, $MSE = .027$, $\eta_p^2 = .281$. This main effect does not reflect differences in learning of the subsequences (dealt with in the following section), and captures

differences in performance of the subsequences, regardless of whether participants were in Experimental or Control groups. Errors show a preference for responding to the consistent triplets RRR and LLL opposed to RRL and LLR, but the RTs show this is a result of a speed-accuracy trade-off, with participants performing better in terms of speed to the subsequences RLR and LRL compared to RLL and LRR. At test RT difference score is not significant, $F(3,138) = 1.71$, $p = .184$, $MSE = 1079$, $\eta_p^2 = .036$; but proportion of errors show a main effect of Subsequence, $F(3,138) = 4.98$, $p = .009$, $MSE = .012$, $\eta_p^2 = .098$. Subsequences RRR and LLL were performed more accurately compared to RRL and LLR than the other two trained subsequences compared to their inconsistent counterparts.

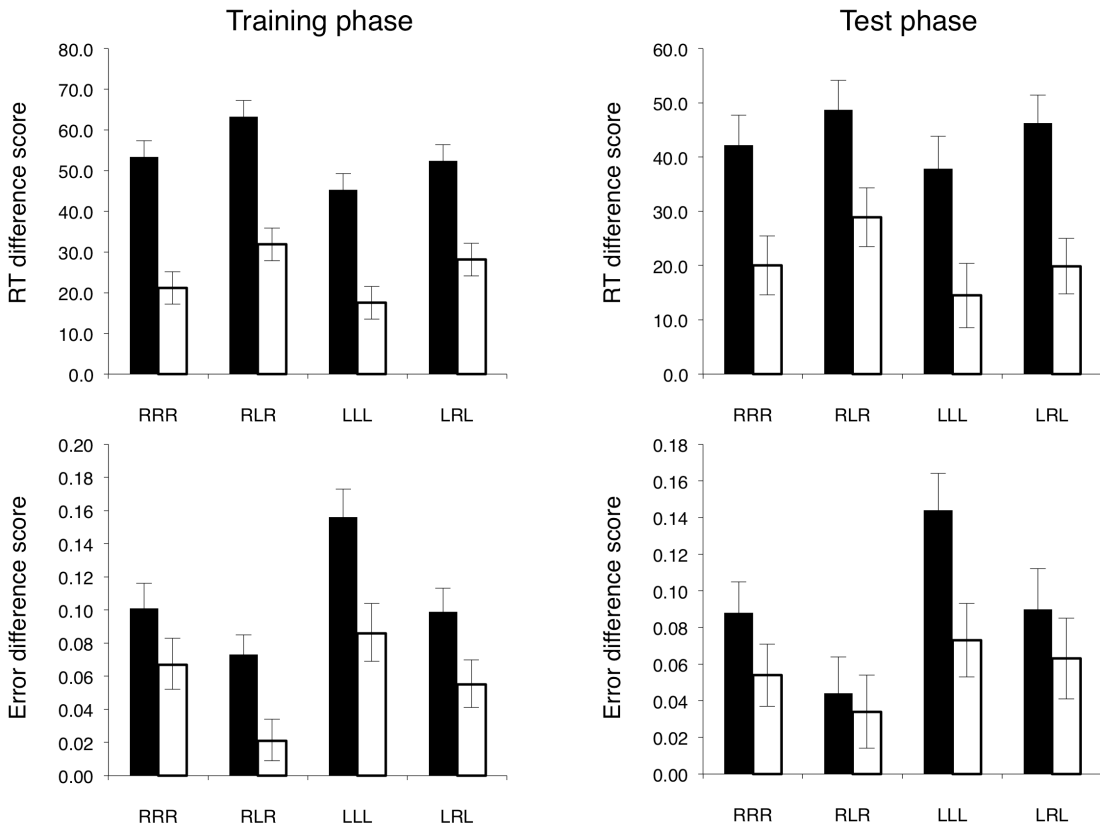

Figure 2.2. RT (top panel) and proportion of error (bottom panel) difference scores across training (left panel) and test (right panel) for the Same rule, Experimental (filled bars) and Control (open bars) groups across the four trained Subsequences. The RRR column, for example, shows performance on 'consistent' RRR trials subtracted from 'inconsistent' RRL trials. Error bars show standard error.

Subsequence learning. One thing each subsequence shares in common is a difference score of greater than zero, which will be discussed and compared across the Rule

conditions in section 2.2.2.3. The main effect of Subsequence suggests only a difference in which subsequences are performed better, regardless of whether participants were trained on these subsequences or not. It is the interaction between Subsequence and Group that can reveal whether any subsequences were learned better than others and thus whether there is any difference in subsequence learning (rather than sequential effects described in the previous section). However, we found no evidence of such an interaction across training, which suggests that there is no evidence that the subsequences were learned differently.

Learning of individual subsequences was analysed separately, taking first the subsequence RRR. Using ANOVA with Epoch and Group across training and Group at test, we find a significant effect of Group in RTs across training, $F(1,46) = 30.6$, $p < .001$; and at test, $F(1,46) = 8.87$, $p < .001$, but not in proportion of errors across training, $F(1,46) = 3.10$, $p = .170$; nor at test, $F(1,46) = 1.67$, $p = .404$. The subsequence LLL demonstrated learning across training RT difference score, $F(1,46) = 26.0$, $p < .001$, and errors, $F(1,46) = 7.47$, $p = .018$; as well as in RT difference scores at test, $F(1,46) = 7.28$, $p = .019$; and error difference at test, $F(1,46) = 5.34$, $p = .046$. Therefore, there is quite good evidence for the learning of subsequences RRR and LLL (or XXX), contrary to the incidental participants in Jones and McLaren (2009). The subsequence RLR also showed evidence of learning across training in both RT difference score, $F(1,46) = 26.1$, $p < .001$, and errors, $F(1,46) = 7.82$, $p = .015$. There was also learning evident at test in the RTs, $F(1,46) = 8.46$, $p = .011$, but not in errors, $F(1,46) = .116$, $p > .9$. The subsequence LRL was also learned, with a Group effect across training in RT difference score, $F(1,46) = 19.7$, $p < .001$ but not errors, $F(1,46) = 4.28$, $p = .088$. This was the same at test, with RT difference scores producing an effect of Group, $F(1,46) = 12.1$, $p = .002$, but not errors, $F(1,46) = .714$, $p = .805$. Therefore, all subsequences in the Same group showed at least some evidence of learning.

Subsequence showed an interaction with Epoch in RT difference score, $F(18,828) = 3.24$, $p = .001$, $MSE = 502$, $\eta_p^2 = .066$; but not proportion of errors, $F(18,828) = 1.17$, $p = .311$, $MSE = .009$, $\eta_p^2 = .025$, see Figure 2.3. This Subsequence and Epoch interaction showed no evidence of an interaction with Group (RT difference score, $F(18,828) = .644$, $p = .806$, $MSE = 502$, $\eta_p^2 = .014$; proportion of errors, $F(18,828) = .962$, $p = .470$, $MSE = .009$, $\eta_p^2 = .020$), suggesting those Subsequence effects observed show no

evidence of being effected by training type or time. Across training the RTs show that the subsequence effect is numerically reduced, with the faster responding to RLR and LRL over inconsistent subsequences compared to RRR and LLL collapsing across the second session.

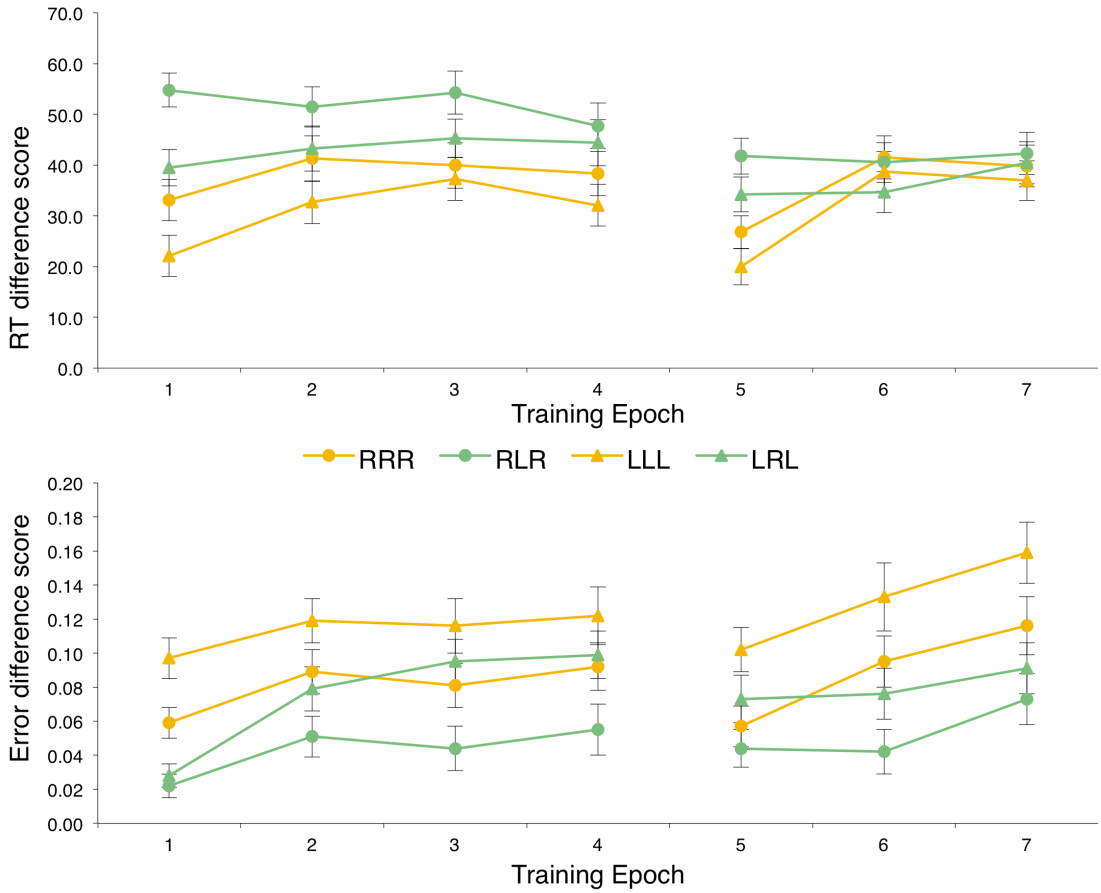

Figure 2.3. RT (top panel) and proportion of error (bottom panel) difference scores across training Epochs for the Same rule groups across the four trained Subsequences. The RRR data points, for example, shows performance on 'consistent' RRR trials subtracted from 'inconsistent' RRL trials. The data are collapsed across all 48 participants regardless of Experimental or Control groups. Error bars show standard error.

Ceiling effects. The observation of good performance on the subsequences consistent with the Same rule regardless of group (difference scores above zero for both Experimental and Control groups) leads neatly to the suggestion that participants may be performing at ceiling and therefore they cannot perform any better on these subsequences, restricting learning. As Jones & McLaren (2009) propose, one can examine this possibility through an analysis of inconsistent trials. If we suggest that on the consistent RRR trials, the Same group are at ceiling, then the inconsistent RRL trials

cannot be at ceiling (in terms of performance) if learning had made these slower and less accurate. Therefore, we can assess learning from these inconsistent trials alone, with slower and less accurate trials in the Experimental group on these inconsistent subsequences providing evidence of learning that cannot be accounted for by a ceiling effect.

Bonferroni corrected comparisons were conducted on inconsistent subsequence RTs and errors in a mixed ANOVA with Epoch, Subsequence and Group across training and Sequence by Group at test. A Group effect (demonstrating learning in the absence of a ceiling effect) was demonstrated across training in the RTs, $F(1,46) = 5.48$, $p = .047$, $MSE = 13508$, $\eta_p^2 = .106$, and errors, $F(1,46) = 10.9$, $p = .004$, $MSE = .031$, $\eta_p^2 = .192$, as well as at test in the errors, $F(1,46) = 17.1$, $p = .001$, $MSE = 2346$, $\eta_p^2 = .271$; but not in RTs, $F(1,46) = .313$, $p > .9$, $MSE = .031$, $\eta_p^2 = .007$, with the direction of the effect showing slower and less accurate responses to inconsistent subsequences in the Experimental group compared to the Control group. This suggests that learning was evident in the absence of a ceiling effect, as learning is not restricted to consistent subsequences that may suffer from a ceiling effect in response speed or accuracy. In summary, all of the subsequences were learned and there was no subsequence interaction with Group, which suggests that there were no differences in how well these subsequences were learned.

2.2.2.2. Different rule learning

The Different rule group, who in the Experimental group were trained on the subsequences RRL, RLL, LLR and LRR (these were twice as likely to occur as RRR, RLR, LLL and LRL across training) was analysed in the same way as the Same group (as described in section 2.2.2.1).

Evidence of learning. There was a very strong main effect of Group for the Different rule participants, in both RT difference score, $F(1,46) = 101$, $p < .001$, $MSE = 4618$, $\eta_p^2 = .686$; and proportion of errors, $F(1,46) = 53.4$, $p < .001$, $MSE = .056$, $\eta_p^2 = .537$, see Figure 2.4. This remained significant and was a strong effect at test, RT difference score, $F(1,46) = 51.7$, $p < .001$, $MSE = 615$, $\eta_p^2 = .529$; and proportion of errors, $F(1,46) = 28.2$, $p < .001$, $MSE = .009$, $\eta_p^2 = .380$, providing good evidence of learning of Different rule subsequences.

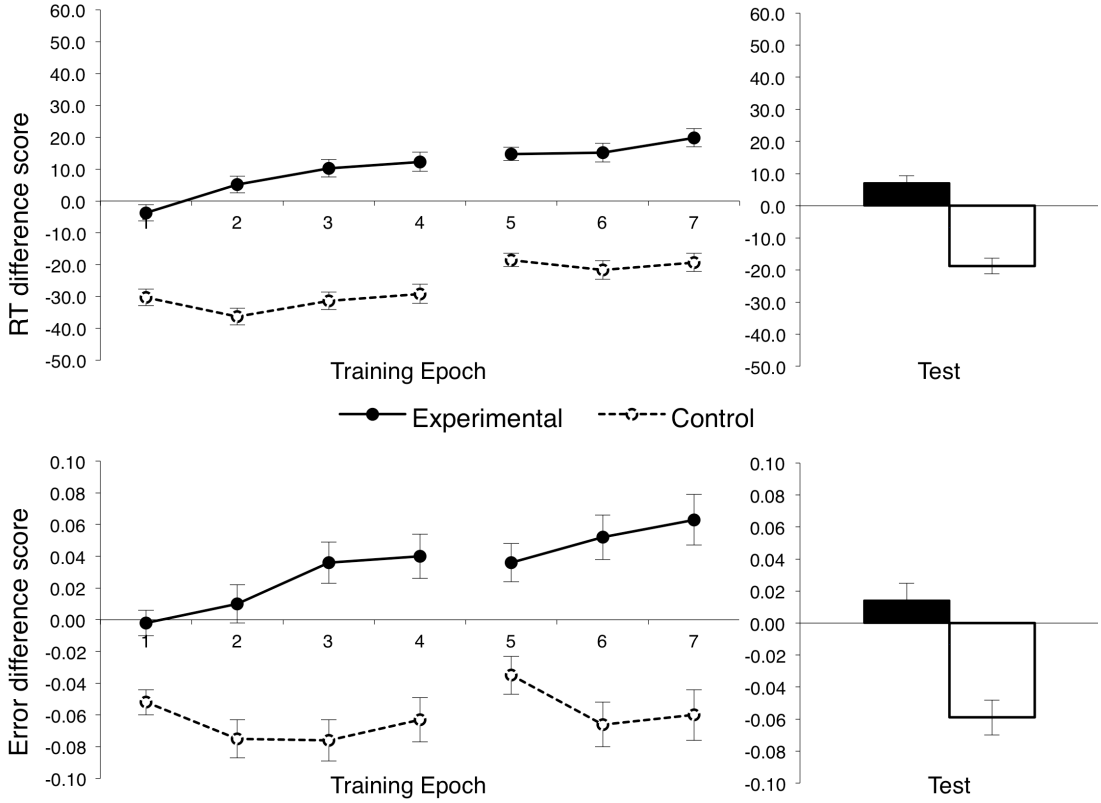

Figure 2.4. RT (top panel) and proportion of error (bottom panel) difference scores across training (left panel) and test (right panel) for Experimental and Control groups trained on Different rule subsequences. Error bars show standard error. Note: the blank between Epochs 4 and 5 represents the break between the two sessions.

As in the Same rule group, there was both an effect of Epoch, RT difference score, $F(6,276) = 32.0$, $p < .001$, $MSE = 326$, $\eta_p^2 = .410$; proportion of errors, $F(6,276) = 7.73$, $p < .001$, $MSE = .005$, $\eta_p^2 = .144$, and an interaction between Epoch and Group, RT difference score, $F(6,276) = 5.55$, $p < .001$, $MSE = 326$, $\eta_p^2 = .108$; and proportion of errors, $F(6,276) = 8.02$, $p < .001$, $MSE = .005$, $\eta_p^2 = .149$, across training. Figure 2.4 demonstrates that the difference between Experimental and Control groups does develop somewhat across training, although this learning is apparent from the first Epoch, suggesting it is acquired rapidly.

Subsequence effects. There was a main effect of subsequence (see Figure 2.5) across training in both RT difference score, $F(3,138) = 9.37$, $p < .001$, $MSE = 6980$, $\eta_p^2 = .169$; and proportion of errors, $F(3,138) = 4.39$, $p = .021$, $MSE = .048$, $\eta_p^2 = .087$, and in proportion of errors at test, $F(3,138) = 7.15$, $p < .001$, $MSE = .009$, $\eta_p^2 = .135$, but not at test in RT difference score, $F(3,138) = 1.25$, $p = .293$, $MSE = 615$, $\eta_p^2 = .026$. This was

expressed as greater accuracy on RLL subsequence compared with the other subsequences, but better RT difference scores to RRL and LLR across training. Crucially, there was no evidence of an interaction between Subsequence and Group across training, RT difference score, $F(3,138) = .727$, $p = .471$, $MSE = 6980$, $\eta_p^2 = .016$; and proportion of errors, $F(3,138) = .197$, $p = .783$, $MSE = .048$, $\eta_p^2 = .004$, nor at test in both RT difference score, $F(3,138) = 2.39$, $p = .096$, $MSE = 615$, $\eta_p^2 = .049$; and proportion of errors, $F(3,138) = .678$, $p = .512$, $MSE = .009$, $\eta_p^2 = .015$. This suggests that there is no evidence that subsequence effects were different due to training experience across subsequences.

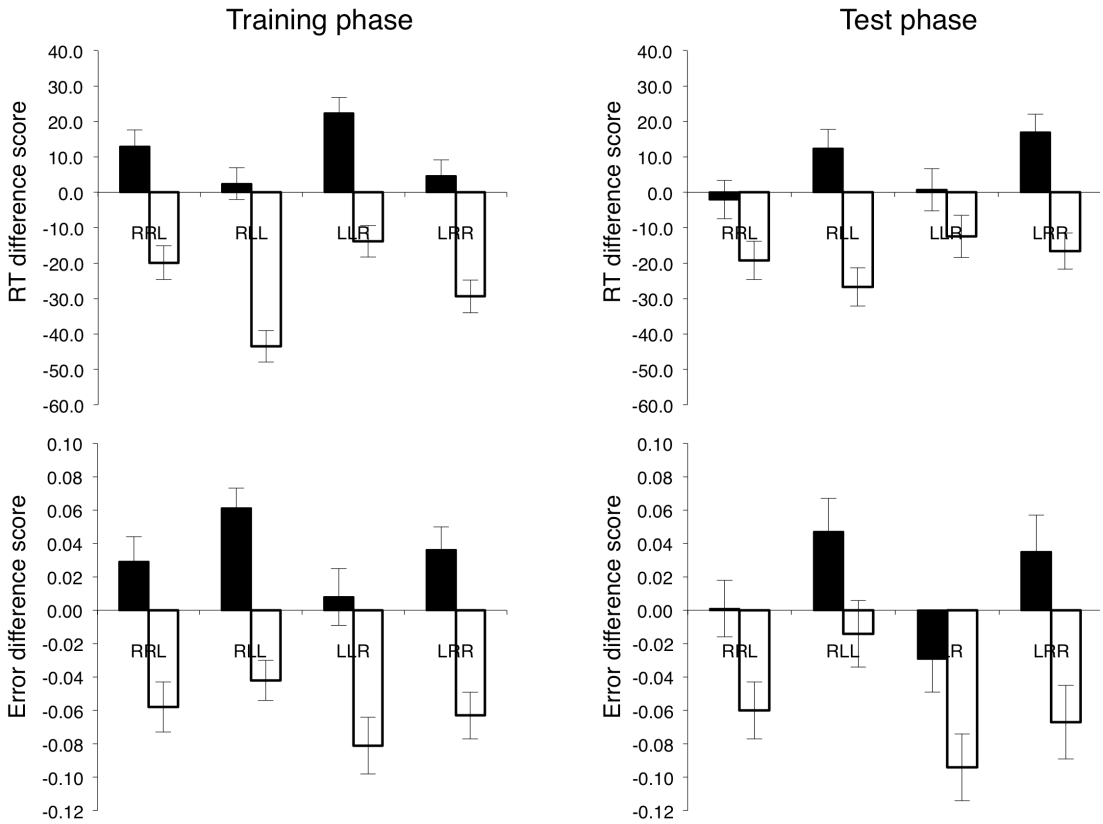

Figure 2.5. RT (top panel) and proportion of error (bottom panel) difference scores across training (left panel) and test (right panel) for the Experimental (filled bars) and Control (open bars) Different rule groups across the four trained Subsequences. The RRL column, for example, shows performance on 'consistent' RRL trials subtracted from 'inconsistent' RRR trials. Error bars show standard error.

Subsequence learning. The learning of each subsequence was assessed through a series of Bonferroni corrected comparisons, taking first the subsequence RRL. Using an ANOVA with Epoch and Group across training and Group at test, we find a significant

effect of Group in RTs across training, $F(1,46) = 18.9$, $p < .001$; and in errors, $F(1,46) = 16.6$, $p < .001$, as well as errors at test, $F(1,46) = 5.95$ $p = .037$; but not at test in RTs, $F(1,46) = 4.78$, $p = .068$. The subsequence LLR demonstrated learning across training RT difference score, $F(1,46) = 26.4$, $p < .001$, and errors, $F(1,46) = 14.4$, $p = .001$; again as well as in errors at test, $F(1,46) = 5.88$, $p = .039$, but not RTs, $F(1,46) = 2.65$, $p = .221$. This is evidence for learning of the subsequences RRL and LLR. The subsequence RLL also showed evidence of learning across training in both RT difference score, $F(1,46) = 50.0$, $p < .001$, and errors, $F(1,46) = 39.2$, $p < .001$. There was also learning evident at test in the RTs, $F(1,46) = 21.6$, $p < .001$; and errors, $F(1,46) = 6.94$, $p < .001$. The subsequence LRR was also learned, with a Group effect across training in RT difference score, $F(1,46) = 21.2$, $p < .001$ and errors, $F(1,46) = 28.1$, $p < .001$. There was also strong learning at test, with RT difference scores producing an effect of Group, $F(1,46) = 23.6$, $p < .001$, with errors also showing evidence of learning at test, $F(1,46) = 17.1$, $p < .001$. Therefore we have strong evidence for the learning of the subsequences RRL and LLR (YYX).

The Subsequence effect itself interacted with Epoch in RT difference score again, $F(18,828) = 3.68$, $p < .001$, $MSE = 617$, $\eta_p^2 = .074$; but not proportion of errors, $F(18,828) = 1.55$, $p = .094$, $MSE = .005$, $\eta_p^2 = .032$, nor Epoch and Group (RT difference score, $F(18,828) = .918$, $p = .529$, $MSE = 326$, $\eta_p^2 = .020$; and proportion of errors, $F(18,828) = 1.19$, $p = .283$, $MSE = .005$, $\eta_p^2 = .025$), suggesting that simply performing the task for increasing time had an effect only on the speed of responding to subsequences, and that differences in Subsequence learning did not change across training. This is shown in Figure 2.6, where we can see that the subsequences speed advantage for RRL and LLR, similar to in the Same group, reduces over the second session.

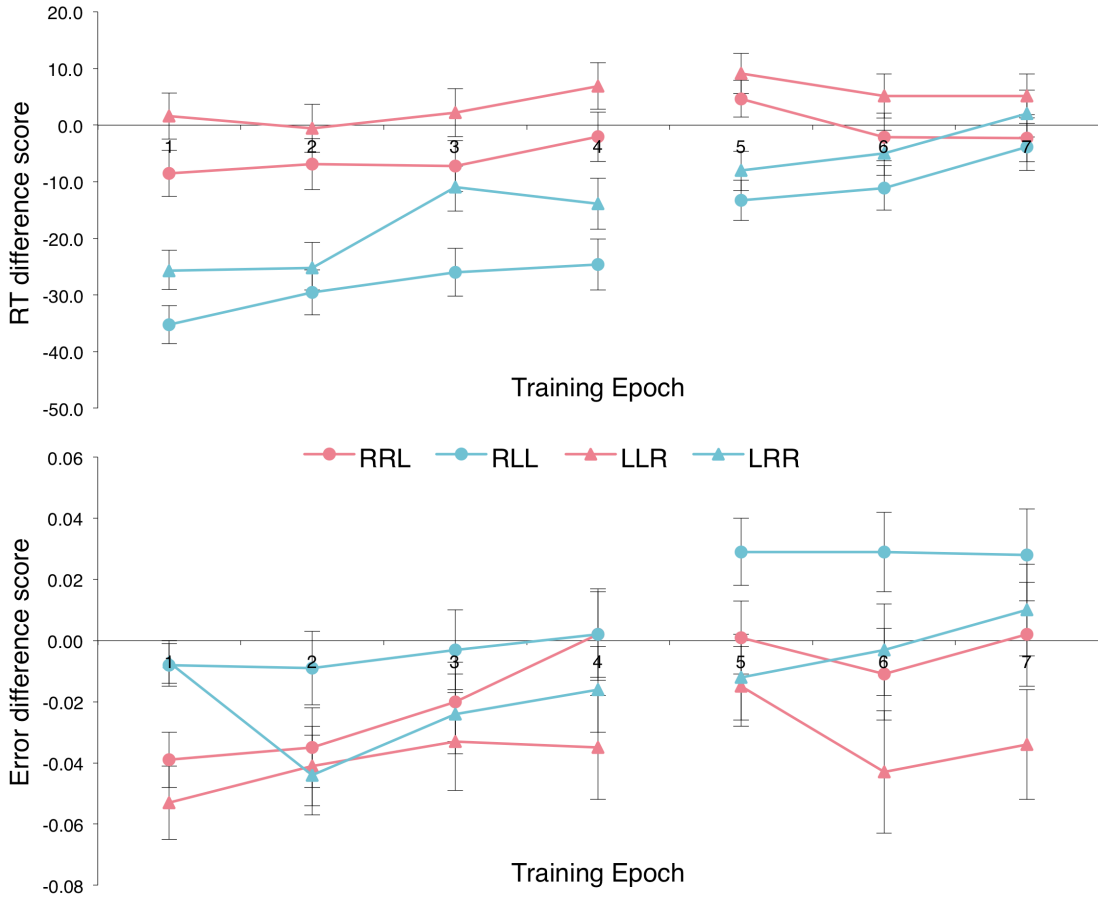

Figure 2.6. RT (top panel) and proportion of error (bottom panel) difference scores across training Epochs for the Same rule groups across the four trained Subsequences. The RRL data points, for example, shows performance on 'consistent' RRL trials subtracted from 'inconsistent' RRR trials. The data are collapsed across all 48 participants regardless of Experimental or Control groups. Error bars show standard error.

2.2.2.3. Same versus Different rule learning

When comparing the data across the different Rules there was a strong main effect of Rule in both RT difference score, $F(1,92) = 469$, $p < .001$, $MSE = 3187$, $\eta_p^2 = .836$; and proportion of errors, $F(1,92) = 78.1$, $p < .001$, $MSE = .079$, $\eta_p^2 = .459$ across training, as well as RT difference score, $F(1,92) = 248$, $p < .001$, $MSE = 565$, $\eta_p^2 = .729$; and proportion of errors, $F(1,92) = 63.6$, $p < .001$, $MSE = .007$, $\eta_p^2 = .409$ at test. The main effect of Rule comprises both Experimental and Control groups, and therefore tells us little about learning, but about the sequential effects underlying performance. The pattern of responding demonstrated by the Control groups is of interest, as this shows a difference score of greater than zero for the Same Control group, suggesting that the Same rule subsequences were overall performed faster and more accurately than the

inconsistent (Different rule subsequences) regardless of training. This main effect therefore probably reflects the overall performance preference for subsequences RRR, RLR, LLL and LRL.

We can see this overall Rule effect in Figure 2.7, which plots the results from both experiments on one graph and highlights the importance of Control groups in controlling for sequential effects. It is immediately apparent that the Control groups are a reflection of one another around a difference score of zero, with better performance on Same rule subsequences compared to Different rule subsequences. Without these Controls it would appear that performance on the Same rule was far better than on the Different rule in the Experimental groups. However, when comparing the groups across their learning through the interaction between Group and Rule across training the converse is apparent, proportion of errors, $F(1,92) = 4.20$, $p = .043$, $MSE = .062$, $\eta_p^2 = .044$ and supported by the numerical trend towards better performance on the Different rule in the RT difference score, $F(1,92) = 3.64$, $p = .059$, $MSE = 3856$, $\eta_p^2 = .038$; and the numerical direction at test, RT difference score, $F(1,92) = .325$, $p = .570$, $MSE = 1233$, $\eta_p^2 = 004$; proportion of errors, $F(1,92) = 2.30$, $p = .133$, $MSE = .007$, $\eta_p^2 = .024$. This provides evidence that the two sequential rules were learned differently, with the Different rule sequences learned better than the Same rule sequences.

This difference between Rules across Experiment and Control groups interacts with Epoch across training, RT difference score, $F(6,552) = 3.16$, $p = .007$, $MSE = 330$, $\eta_p^2 = .033$; but not proportion of errors, $F(6,552) = 1.52$, $p = .184$, $MSE = .007$, $\eta_p^2 = .016$, suggesting that learning of the two different sequential contingencies progressed at a different rate. Indeed, we can see that whilst the learning of the Different rule appears rapidly in the RT difference measure but then improves little across training, that the Same group learns steadily across the training epochs.

Subsequence effects did not show evidence of an interaction with Group and Rule across training: RT difference score, $F(3,276) = .397$, $p = .654$, $MSE = 5938$, $\eta_p^2 = .004$; proportion of errors, $F(3,276) = .529$, $p = .569$, $MSE = .037$, $\eta_p^2 = .006$, nor at test: RT difference score, $F(3,276) = 1.33$, $p = .266$, $MSE = 1095$, $\eta_p^2 = .014$; proportion of errors, $F(3,276) = .979$, $p = .379$, $MSE = .011$, $\eta_p^2 = .011$.

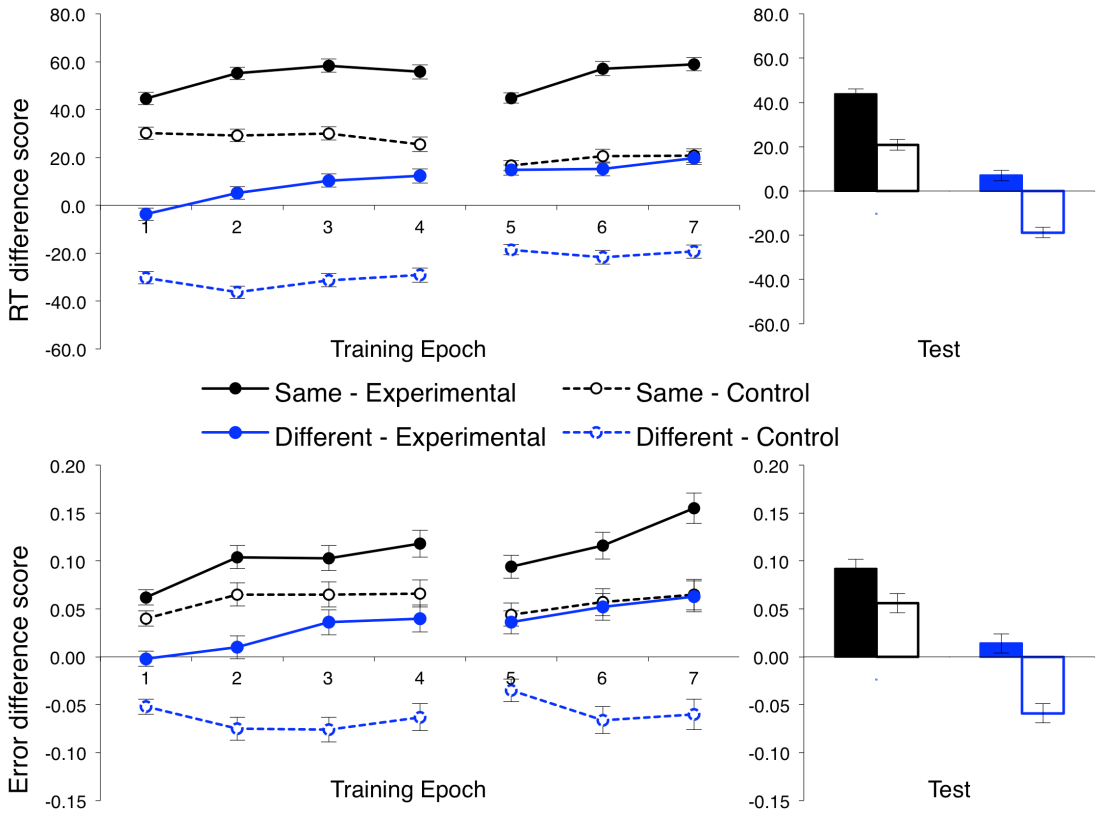

Figure 2.7. RT (top panel) and proportion of error (bottom panel) difference scores across training (left panel) and test (right panel) for Experimental and Control groups trained on Same (black lines) and Different (blue lines) rule subsequences. Error bars show standard error. Note: the blank between Epochs 4 and 5 represents the break between the two sessions.

2.2.2.4 Post-experimental interview

The questions in the post-experimental interview suggest that participants were unable to verbalise any knowledge about sequences, and were all unable to report the rule. Some participants, when asked "Did you notice any patterns or sequences to the responses that you were required to make?" did report noticing long strings of repeats or patterns of alternations, but these were not restricted to the Same Experimental group. Across each of the Experimental and Control groups, participants were just as likely to report noticing repeating trials, but they were also just as likely to report that they thought the sequence was random. A small number of participants reported trying to count or investigate to see whether there was a sequence, even though none had been mentioned, but reported that this made them slower and worse at the task so they quickly gave up on doing this. There were no scaled responses to how much people believed there was a sequence (e.g. Curran, 1997) or how confident they were in these judgments (Dienes & Berry, 1997) therefore no further analysis was carried out.

2.2.3. Discussion

Experiment 1 provided evidence that humans are able to learn both Same and Different rules under incidental conditions. Whilst both groups learned the trained sequential contingencies, the Different rule produced more learning than the Same rule. The Experimental Different group learned more rapidly than the Same group and this learning was more robust, evidenced by the Group interaction apparent across training. The Group effect (the measure of learning) itself survived at test in the Different rule case for errors and RTs, but only in RTs in the Same group.

This supports the predictions of the associative account as the Same group involved the subsequence XXX (both RRR and LLL); which was shown in previous work to be more difficult to learn due to a blocking effect (Jones & McLaren, 2009). Evidence of an Experimental over Control group advantage across RTs in training suggests, however, that participants were at least able to learn the subsequences RRR and LLL in the Same rule group reported here, unlike Jones & McLaren (2009). The associative account predicts better learning of YXY compared to XXX, which was not found here. Further to this, there was evidence for learning of the subsequences RLL and LRR in the Different rule condition, which is contrary to the findings of Jones & McLaren (2009), that subsequences ending in a repetition demonstrated little or no learning under incidental conditions. Indeed, the greatest learning under incidental conditions in Experiment 1 was shown for subsequence XYY (RLL and LRR).

That these subsequence learning effects were not identical to Jones and McLaren (2009) provides evidence that humans are not simply better at learning some forms of subsequence, which might be proposed by some instance or chunking based account of learning. The difference between the two experiments is in the structure of the trained sequences, as in Experiment 1 participants were trained to predict t from $t - 2$, whereas the exclusive-or rule in Jones & McLaren (2009) training participants that t was predicted by the combination of both $t - 1$ and $t - 2$. This suggests that a difficulty in learning subsequences XXX and XYY was not restricted to these subsequences and their trial order, but also to how this related to the structure of the underlying trained contingency.

Support is therefore provided for an associative account of human learning, as altering the relationship between subsequence trial-by-trial order and the trained sequential contingencies produced a different pattern of human learning. It is possible that the exclusive-or rule in Jones and McLaren (2009) may have been simply more difficult to learn. However, Jones and McLaren (2009) trained participants for over twice the amount of trials as Experiment 1, which seems a convincing demonstration of the lack of XXX learning in the exclusive-or case rather than this simply being more difficult to learn; or not having enough time to emerge. If one reduced the subsequences in the Experiment to their structures, for example in the Same group: RRR and LLL both involve three of the same response (or two repeats between first and second order transitions) and RLR and LRL both involve two alternations; then the Same and Different group were only trained on two different subsequence transition structures. Therefore Jones & McLaren's (2009) participants could be argued to have had a more difficult task and given a limited capacity for learning subsequence structures been unable to learn one or two of the subsequences. This suggestion, however, relies on an instance based account that suggests chunks are encoded on the basis of structure and not precise locations (left or right), which somewhat undermines the instance-based premise on which such a theory is built.

It is also possible of course that the nature of the two training contingencies led to a difference in explicit strategy or verbalisable knowledge, as repeating chunks were more likely in the Same group of Experiment 1. In the exclusive-or case (Jones & McLaren, 2009) participants experienced one of either a RRR or LLL trial (as a third trial) on one in four trials in training blocks. Participants in the Same group experienced either a RRR or LLL as a third trial on one in two trials, meaning that they could have noticed these repeating chunks more easily and therefore some verbalisable learning of these subsequences is produced. However, this would provide support for the alterative result, that the Same rule would produce more learning (unless one assumes that a conscious strategy impairs learning, which in the case of Jones & McLaren, 2009 occurred for all subsequences but XXX).

The subsequence effects observed clearly and robustly demonstrated that participants preferred to respond to trials that have been preceded either by two of the same response location (e.g. RRR) as well as two alternations (e.g. RLR). These sequential effects are

unaffected by training or experience with the task and seem to occur either very rapidly or may be preexisting preferences that remain robust across the task. Performance is worse for the complementary Different rule subsequences, RRL and RLL (as well as LLR and LRR). The advantage for simply responding to Same rule subsequences could restrict learning of these subsequences due to ceiling performance, however a Group effect was demonstrated in inconsistent trials alone, suggesting this that cannot explain the Same-Different rule learning difference.

One final thing to note is that the sequential effects observed in the Control groups of Experiment 1 do not match those subsequence sequential effects observed in Jones & McLaren (2009), whose control group performance followed the pattern (from preferred to least preferred): XXX, XYY, YXY, YYX, although XYY and YXY were relatively close around zero on their RT difference scores. The sequential effects observed in the Control groups in Experiment 1 suffer somewhat from speed-accuracy trade-off and in that case are more difficult to interpret, however, difference scores are always above zero for the Control performance on XXX and YXY subsequences, whereas XYY and YYX subsequence performance is always below zero. One explanation for this difference could be the length of training, as one could propose that after five sessions those participants trained with pseudorandom sequences would eventually begin to learn that all subsequences are equally likely, however Jones & McLaren (2009) do not provide evidence of the progression of learning nor sequential effects across training to analyse this possibility. The control groups in both studies were essentially identical but for training length and the presence of a performance reward, which could have influenced the speed and accuracy of responding in Jones and McLaren's (2009) study. Participants were motivated to perform as quickly and accurately as possible, which suggests sequential effects may have been influenced by the speed of responding to stimuli (e.g. Frensch, 2003; Soetens, Boer, & Huetings, 1985).

Altogether, this experiment provides evidence that participants can learn sequences based on a $t - 2$ probabilistic contingency under incidental conditions. Whilst these rules do not differ in contingency nor predictive trial order; they are nevertheless learned differently by participants in this experiment. Participants learn more about sequences that follow a rule whereby t is twice as likely to be different to $t - 2$ (Different rule); over learning that t is twice as likely to be the same as $t - 2$ (Same rule).

The difference in learning seems unlikely to be explained by a rule-based account of learning, and suggests that, as observed by Jones and McLaren (2009), participants may be learning sequential contingencies associatively and therefore trial-by-trial contingencies can interfere with certain subsequence learning.

2.3. Experiment 2: Intentional sequence learning

Experiment 1 demonstrated some learning of the subsequence XXX compared to the lack of learning evident for this subsequence in Jones and McLaren (2009). This may be due the reduced cue competition between transient trial-by-trial contingencies and sequential contingences. As Experiment 1 uses Same and Different rules to train participants, these depend only on trial $t - 2$, and trial $t - 1$ is no longer important for the sequential learning as it was in Jones and McLaren's (2009) study. The subsequence XXX is therefore not a marker for incidental learning, as impaired implicit learning of this subsequence is dependent on the trained sequential rule itself. The second study in this thesis aimed to investigate how participants learned these rules and subsequences under intentional conditions, whereby they were instructed to actively search for and use sequences, patterns or rules that they noticed and were explicitly told that there were contingencies in the task. Participants in Control groups were told the same instructions to control for the effect that the rule-based hypothesis testing may have on the sequential effects themselves. As Jones & McLaren (2009) found that participants were able to notice the subsequence XXX and learn about this explicitly, I expected that participants in the Same group would experience more learning under intentional conditions, possibly reversing the direction of the rule learning interaction.

2.3.1. Method

2.3.1.1. Participants

96 participants (aged between 18 and 35 [$M = 21.1$]; 73 female and 23 male) were recruited from undergraduate students at the University of Exeter and were awarded £10 for participation. Participants provided informed consent prior to taking part in two sessions lasting roughly one hour each. Participants were randomly allocated into one of the four possible conditions.

2.3.1.2. Materials and Stimuli

The materials and stimuli were the same as for Experiment 1 (see section 2.2.1.2).

2.3.1.3. Design

The experiment designed followed that of Experiment 1 exactly (see section 2.2.1.3).

2.3.1.4. Procedure

After obtaining informed consent, participants were instructed to respond as quickly and accurately as possible to the stimuli, but also that the task was made up of sequences of trials that participants could use to improve their performance. Participants were explicitly made aware that whilst we would be recording their speed and accuracy, that the task may contain sequences that might be difficult to notice, but that they should try to discover it and use these sequences to help them make predictions about where the trial would appear next. Participants were reminded of these instructions at the beginning of the second session. Each trial followed the same sequence as outlined in the procedure for Experiment 1 (see section 2.2.1.4).

As in Experiment 1, a short verbal structured interview was given at the end of the second session, in which participants were asked about what sequences they had noticed in the experiment. They were asked the same structured questions as the incidental group, that led from asking what they thought about the task; to whether they noticed anything interesting about the task; to whether they noticed any patterns or sequences; and finally whether they could describe or guess at any patterns or sequences. Participants were debriefed and thanked for their participation at the end.

2.3.2. Results

Weighted averages were calculated and analysed for each subsequence as for the incidental participants in Experiment 1 and separate analyses for Same and Different sequential rules are given first, before a comparison between the two is presented.

2.3.2.1. Same rule learning

An ANOVA investigated Epoch; Group (Experimental or Control); and Subsequence (RRR, RLR, LLL, LRL) across training and Group and Subsequence at test.

Evidence of learning. The main effect of Group (see Figure 2.8) was apparent across training, RT difference score, $F(1,46) = 94.5$, $p < .001$, $MSE = 3256$, $\eta_p^2 = .673$; proportion of errors, $F(1,46) = 27.4$, $p < .001$, $MSE = .054$, $\eta_p^2 = .373$, and in RT difference score at test, $F(1,46) = 19.3$, $p < .013$, $MSE = 1759$, $\eta_p^2 = .295$; and proportion of errors, $F(1,46) = 9.00$, $p = .004$, $MSE = .006$, $\eta_p^2 = .164$. Taken together, this provides evidence that the Intentional participants learned the Same rule sequences, which carried through to test.

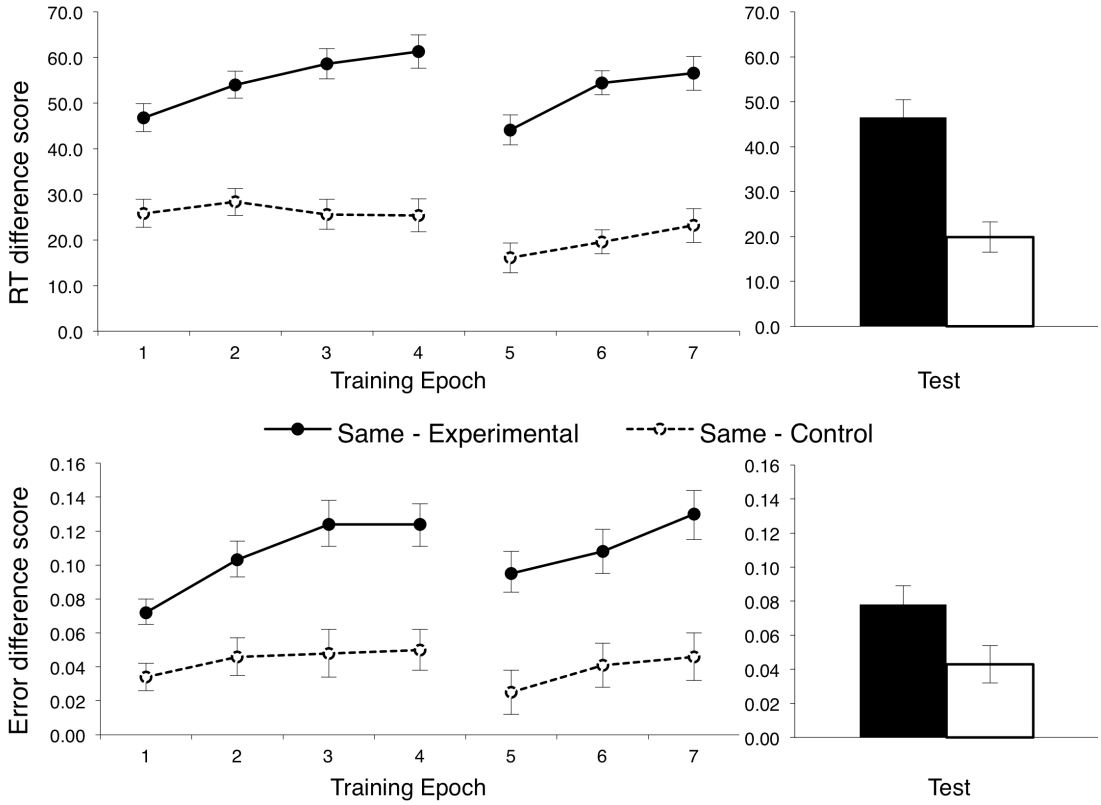

Figure 2.8. RT (top panel) and proportion of error (bottom panel) difference scores across training (left panel) and test (right panel) for Experimental and Control groups trained on Same rule subsequences. Error bars show standard error. Note: the blank between Epochs 4 and 5 represents the break between the two sessions.

There was little evidence that the main effect of Group interacted with Epoch across training in either RT difference score, $F(6,276) = 2.22$, $p = .067$, $MSE = 963$, $\eta_p^2 = .046$; and proportion of errors, $F(6,276) = 1.93$, $p = .092$, $MSE = .007$, $\eta_p^2 = .040$. Whilst the Experimental group had an advantage across training, this was evident from the first Epoch and only numerically increased with more training. Epoch itself had a main effect in both RT difference score, $F(6,276) = 6.02$, $p < .001$, $MSE = 963$, $\eta_p^2 = .116$;

and proportion of errors, $F(6,276) = 6.34$, $p = .015$, $MSE = .046$, $\eta_p^2 = .121$. This suggests that training had no differential effect on Experimental and Control groups, and therefore that the difference between them was relatively stable throughout.

Subsequence effects. The main effect of Subsequence was significant in the RT difference score across training, $F(3,138) = 3.88$, $p = .029$, $MSE = 7267$, $\eta_p^2 = .078$; and in the proportion of errors, $F(3,138) = 16.1$, $p < .001$, $MSE = .020$, $\eta_p^2 = 259$, see Figure 2.9.

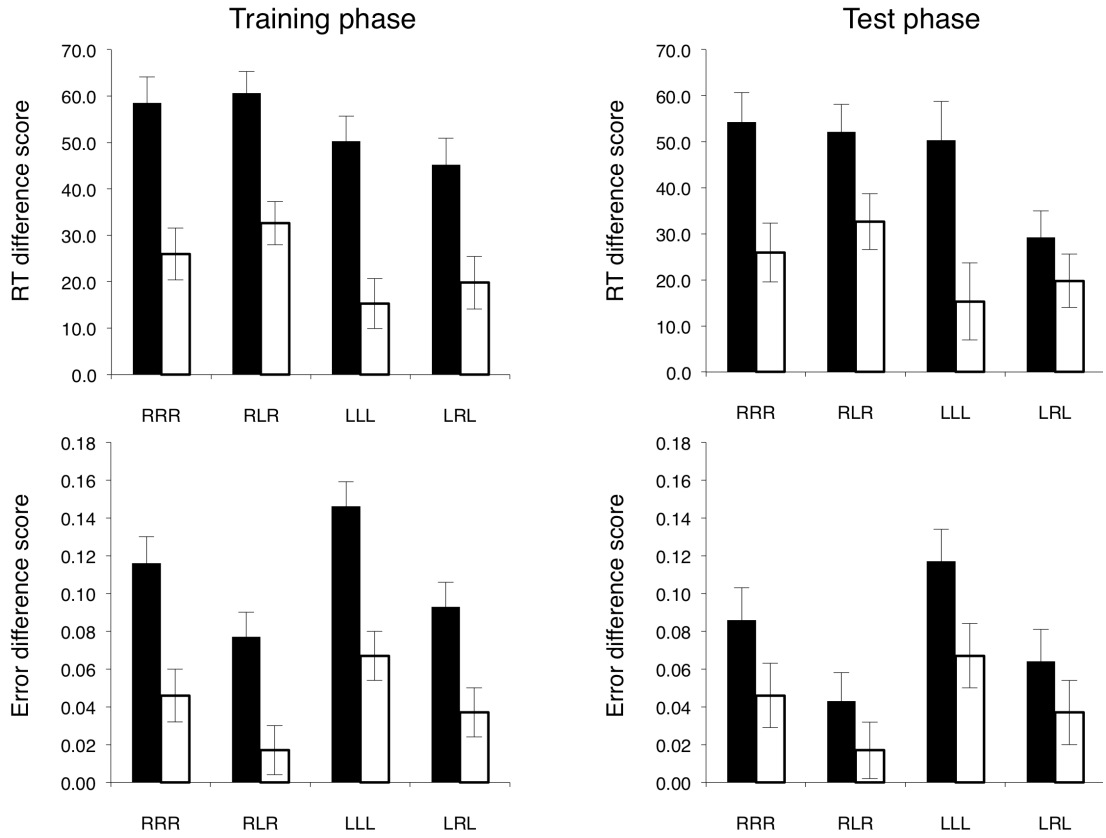

Figure 2.9. RT (top panel) and proportion of error (bottom panel) difference scores across training (left panel) and test (right panel) for the Experimental (filled bars) and Control (open bars) Same rule groups across the four trained Subsequences. The RRR column, for example, shows performance on 'consistent' RRR trials subtracted from 'inconsistent' RRL trials. Error bars show standard error.

There was also a significant main effect in both RT difference score, $F(3,138) = 3.85$, $p = .022$, $MSE = 1559$, $\eta_p^2 = .077$ and proportion of error difference score, $F(3,138) = 11.8$, $p < .001$, $MSE = .005$, $\eta_p^2 = .204$ at test. This Subsequence effect did not show evidence for an interaction with Group at training: RT difference score, $F(3,138) = .363$, $p = .672$, $MSE = 7276$, $\eta_p^2 = .008$; proportion of errors, $F(3,138) = .688$, $p = .508$, MSE

= .020, η_p^2 = .015, or at test: RT difference score, $F(3,138) = .524$, $p = .605$, MSE = 1559, η_p^2 = .011; proportion of errors, $F(3,138) = .178$, $p = .908$, MSE = .005, η_p^2 = .004.

This suggests that the Subsequence effect itself occurred regardless of whether participants were trained or not, and in the RTs across training and test followed the pattern observed in Experiment 1, that subsequences RLR is performed quicker compared to RRR and LRL compared to LLL, however both RRR and RLR produce faster responses than LLL and LRL. Again, as with Experiment 1 this takes the form of a speed-accuracy trade off, as the non-significant pattern of responding to subsequences in the errors across training show a more accurate responding to the RRR and LLL subsequences compared to their inconsistent counterparts than RLR and LRL, with higher accuracy to LRL and LLL subsequences compared to RLR and RRR subsequences.

Intentional participants demonstrated an interaction between Epoch and Subsequence, RT difference score, $F(18,828) = 2.91$, $p = .004$, $MSE = 1255$, $\eta_p^2 = .060$; and proportion of errors, $F(18,828) = 1.88$, $p = .030$, $MSE = .005$, $\eta_p^2 = .039$, see Figure 2.10. These variations across block reflect the opposite pattern to Experiment 1 as the subsequence effects, rather than collapsing across the experiment by the second session, seem to increase across training and may reflect active hypothesis testing. Indeed, greater performance to RRR and LLL begins to emerge in both measures towards the end of training, suggesting greater attention and performance emerges to these subsequences. The interaction did not, however, provide evidence of an interaction with Group, RT difference score, $F(18,828) = .781$, $p = .620$, $MSE = 1255$, $\eta_p^2 = .017$; and proportion of errors, $F(18,828) = .429$, $p = .959$, $MSE = .005$, $\eta_p^2 = .009$, suggesting that any Subsequence variation across Epochs has no influence on learning.

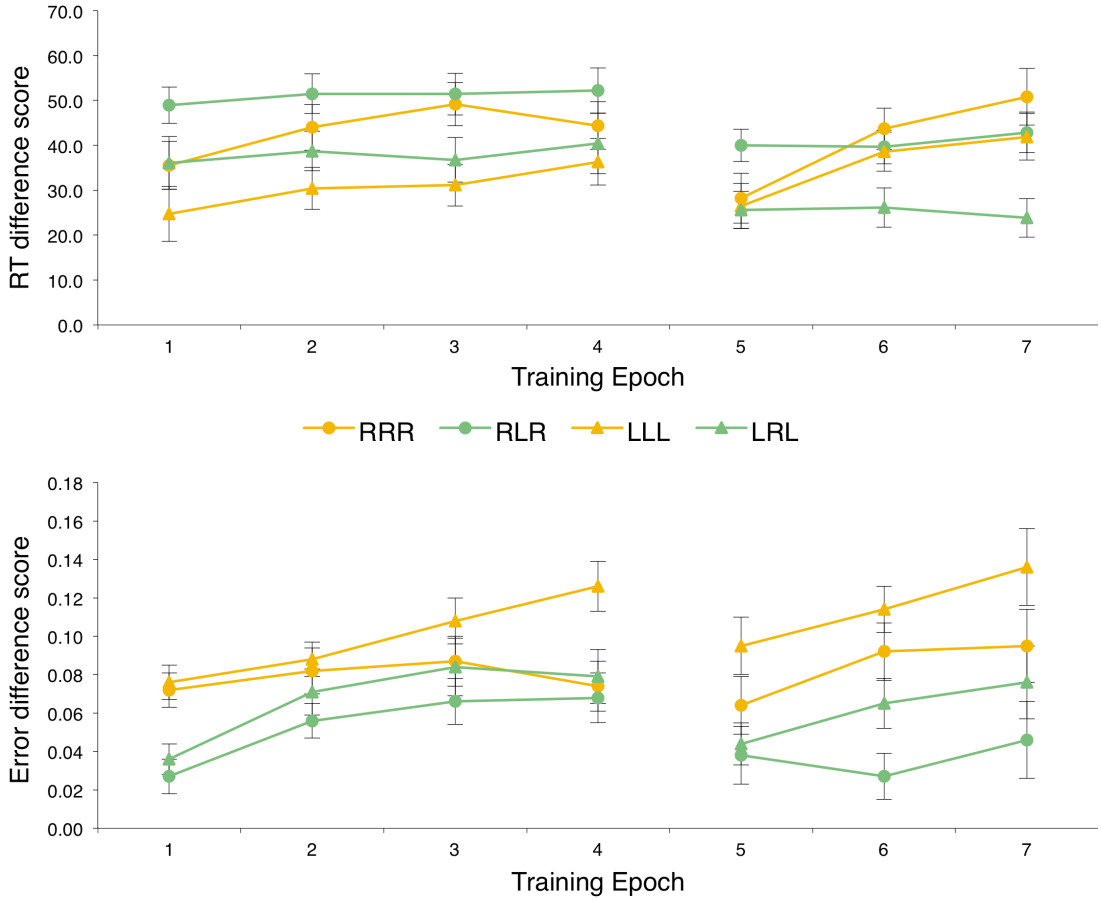

Figure 2.10. RT (top panel) and proportion of error (bottom panel) difference scores across training Epochs for the Same rule groups across the four trained Subsequences. The RRR data points, for example, shows performance on 'consistent' RRR trials subtracted from 'inconsistent' RRL trials. The data are collapsed across all 48 participants regardless of Experimental or Control groups. Error bars show standard error.

Subsequence learning. As for Experiment 1, a series of Bonferroni corrected comparisons were conducted on each subsequence taking first the subsequence RRR. Using an ANOVA with Epoch and Group across training and Location and Group at test, a main effect of Group was found across training in RTs, $F(1,46) = 19.8$, $p < .001$, and proportion of errors, $F(1,46) = 13.8$, $p = .001$. This was supported by learning at test in RT difference scores, $F(1,46) = 6.25$, $p = .032$, but not proportion of errors, $F(1,46) = 1.07$, $p = .614$. The subsequence LLL also demonstrated learning across training in both RTs, $F(1,46) = 19.5$, $p < .001$, and errors, $F(1,46) = 20.2$, $p < .001$; but not quite in test RT difference scores, $F(1,46) = 4.68$, $p = .071$; nor errors, $F(1,46) = 2.45$, $p = .249$. This provides some evidence of learning the subsequences RRR and LLL (XXX). There was, however, also learning of the subsequences RLR and LRL. RLR showed a Group effect

across training RTs, $F(1,46) = 22.4$, $p < .001$, and errors, $F(1,46) = 15.5$, $p = .001$, as well as test RTs, $F(1,46) = 16.6$, $p < .001$, but not errors, $F(1,46) = 5.18$, $p = .055$. The subsequence LRL also showed evidence of learning across both training measures: RTs, $F(1,46) = 15.9$, $p < .001$; errors, $F(1,46) = 14.1$, $p = .001$, as well as in RTs at test, $F(1,46) = 6.46$, $p = .029$, but not test error difference scores, $F(1,46) = 4.89$, $p = .064$. This suggests that participants learned the subsequences RLR and LRL (YXY) as well as the subsequence XXX, in contrast to the intentional participants in Jones and McLaren's (2009) experiment, who learned only XXX.

Ceiling effects. As suggested for the incidental Experiment 1 participants, performance could have been restricted by the subsequence effects experienced in the Same group, as responding to subsequences RRR, RLR, LLL, LRL meant that participants may have been responding at ceiling. RTs and errors measures may have been unable to show a learning effect as participants may have been responding as quickly and accurately as possible to consistent subsequences. Taking the inconsistent trials only (which are expected to be slower and less accurate for the Experimental group, and therefore do not suffer from this ceiling issue) learning was observed in the Same group at test in RTs, $F(1,46) = 5.59$, $p = .045$, $MSE = 4782$, $\eta_p^2 = .108$ and errors, $F(1,46) = 9.00$, $p = .009$, $MSE = .013$, $\eta_p^2 = .164$ (Bonferroni corrected for multiple comparisons). This suggests that participants demonstrate learning regardless of a possible ceiling-effect. However, again in this study there is no Subsequence by Group interaction, as all subsequences are learned and under Intentional conditions participants again do not learn these subsequences differently.

2.3.2.2. Different rule learning

An ANOVA investigated Epoch; Group (Experimental or Control); and Subsequence (RRL, RLL, LLR, LRR) across training and Group and Subsequence at test.

Evidence of learning. The Different rule learning condition showed strong evidence of learning across both training: RT difference score, $F(1,46) = 136$, $p < .001$, $MSE = 4457$, $\eta_p^2 = .747$; proportion of errors, $F(1,46) = 56.4$, $p < .001$, $MSE = .071$, $\eta_p^2 = .551$, and test, RT difference score, $F(1,46) = 61.9$, $p < .001$, $MSE = 708$, $\eta_p^2 = .574$; and proportion of errors, $F(1,46) = 33.2$, $p < .001$, $MSE = .008$, $\eta_p^2 = .419$. This provides strong evidence of learning (see Figure 2.11), with the Experimental group performing

significantly better than the Control group. Similar to the Incidental group in Experiment 1 there was an interaction with Epoch, RT difference score, $F(6,276) = 8.69$, $p < .001$, $MSE = 482$, $\eta_p^2 = .159$; and proportion of errors, $F(6,276) = 8.90$, $p < .001$, $MSE = .007$, $\eta_p^2 = .162$. The learning therefore showed evidence that it developed across the experiment, and the Different rule group shows evidence that this learning was acquired rapidly, with a difference between Experimental and Control groups apparent from the first Epoch.

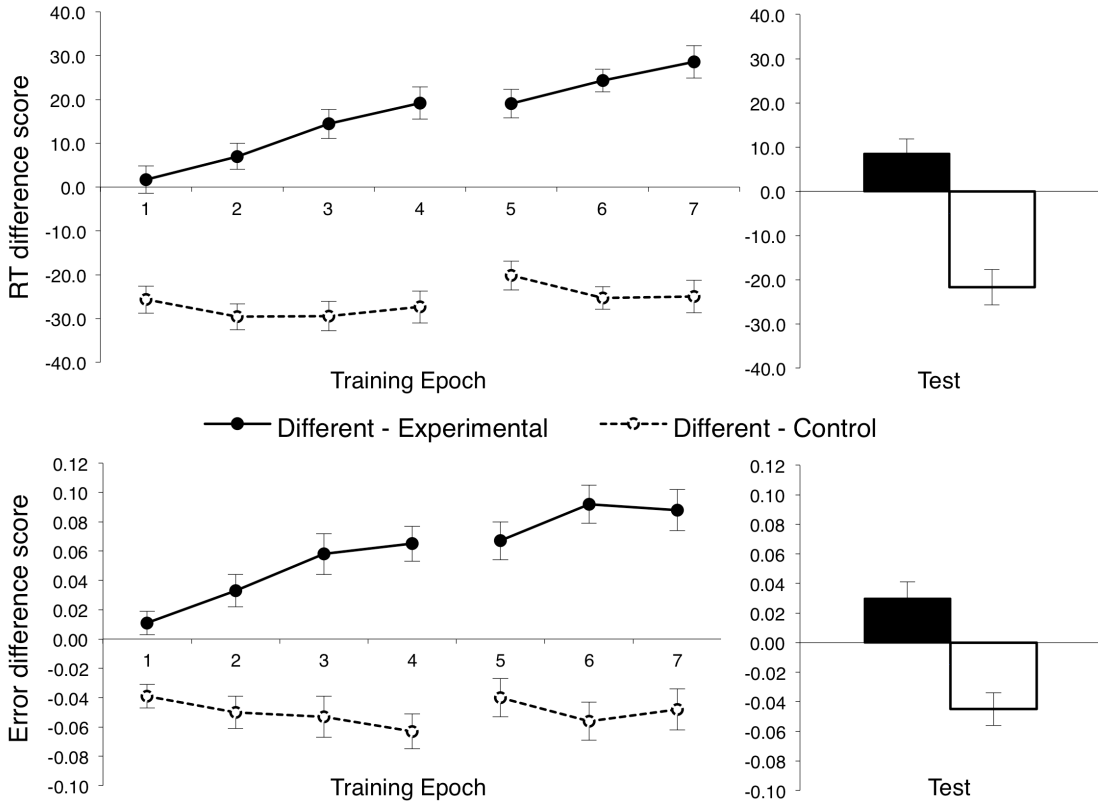

Figure 2.11. RT (top panel) and proportion of error (bottom panel) difference scores across training (left panel) and test (right panel) for Experimental and Control groups trained on Different rule subsequences. Error bars show standard error. Note: the blank between Epochs 4 and 5 represents the break between the two sessions and explains the dip in performance observed between the two blocks.

Subsequence effects. The Intentional Different group showed evidence of a Subsequence main effect across training in errors, $F(3,138) = 6.55$, $p = .004$, $MSE = .045$, $MSE =$, $\eta_p^2 = .125$; but not RTs, $F(3,138) = .405$, $p = .601$, $MSE = 12351$, $\eta_p^2 = .009$, and in both the errors at test, $F(3,138) = 15.5$, $p < .001$, $MSE = .014$, $\eta_p^2 = .252$; and RT difference score, $F(3,138) = 3.69$, $p = .037$, $MSE = 1764$, $\eta_p^2 = .074$, see Figure

2.12. Therefore, regardless of whether participants were trained or not, across training and test they showed a performance benefit for the subsequences RLL and LRR over RRL and LLR in errors, and the opposite effect (a speed-accuracy trade off) in RTs, with faster performance to the subsequences RRL and LLR.

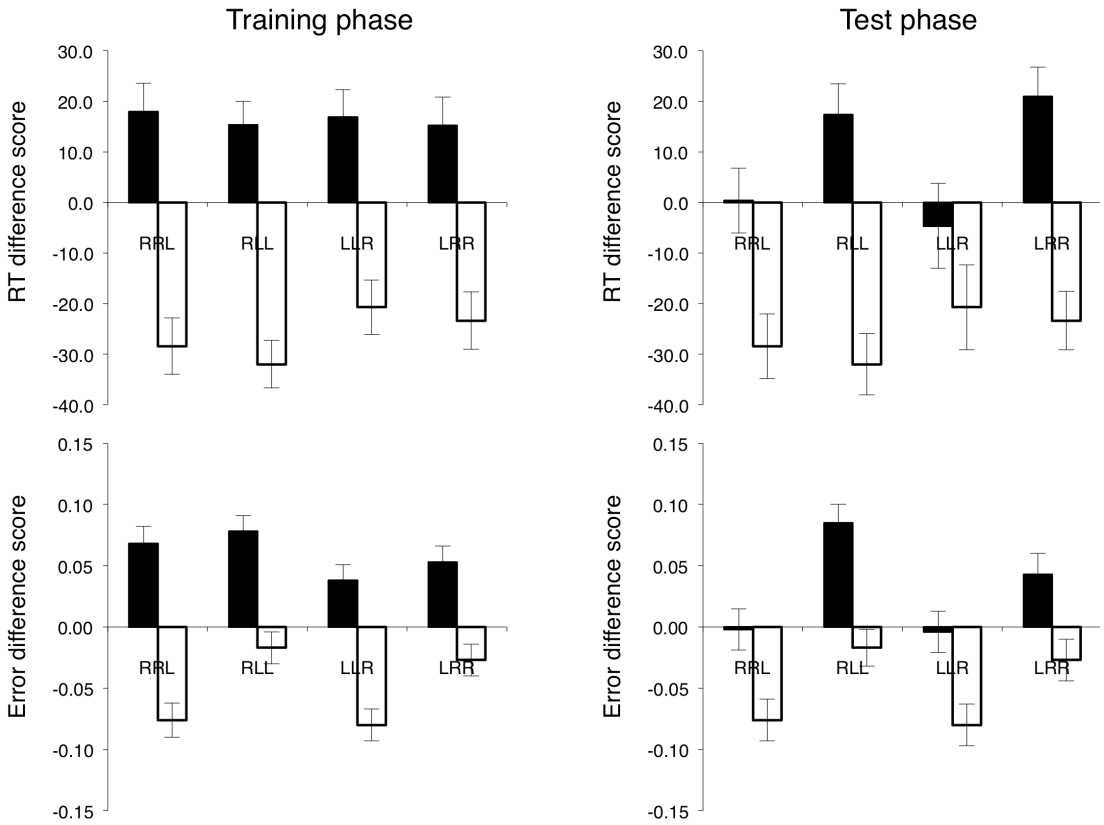

Figure 2.12. RT (top panel) and proportion of error (bottom panel) difference scores across training (left panel) and test (right panel) for the Experimental (filled bars) and Control (open bars) Different rule groups across the four trained Subsequences. The RRL column, for example, shows performance on 'consistent' RRL trials subtracted from 'inconsistent' RRR trials. Error bars show standard error.

There was little evidence for an interaction between Subsequence and Group at training, RT difference score, $F(3,138) = .364$, $p = .626$, $MSE = 12351$, $\eta_p^2 = .008$; proportion of errors, $F(3,138) = 2.52$, $p = .096$, $MSE = .045$, $\eta_p^2 = .052$; although this trend in errors suggests that the relative learning of LRR in Experimental compared to Control groups (and to some extent RLL) is approaching significance and somewhat poorer than for RRL and LLR subsequences. There was very little evidence for an interaction at test: RT difference score, $F(3,138) = .843$, $p = .415$, $MSE = 1764$, $\eta_p^2 = .018$; proportion of errors, $F(3,138) = .736$, $p = .458$, $MSE = .014$, $\eta_p^2 = .016$. The Subsequence main effect

interacted with Epoch (see Figure 2.13) in errors, $F(18,828) = 2.35$, $p = .046$, $MSE = .025$, $\eta_p^2 = .049$; but not RT difference scores, $F(18,828) = 2.00$, $p = .055$, $MSE = 1625$, $\eta_p^2 = .042$; nor with Epoch and Group, RT difference score, $F(18,828) = .633$, $p = .726$, $MSE = 1625$, $\eta_p^2 = .014$; and proportion of errors, $F(18,828) = .679$, $p = .629$, $MSE = .025$, $\eta_p^2 = .015$. This suggests that these effects may not have been based on learning, and again as with the Same group perhaps reflect the changing attention or hypothesis testing of the Intentional participants, regardless of the information available to them, as performance on both RT and error measures increases for the subsequences RLL and LRR across training, regardless of the training that participants received.

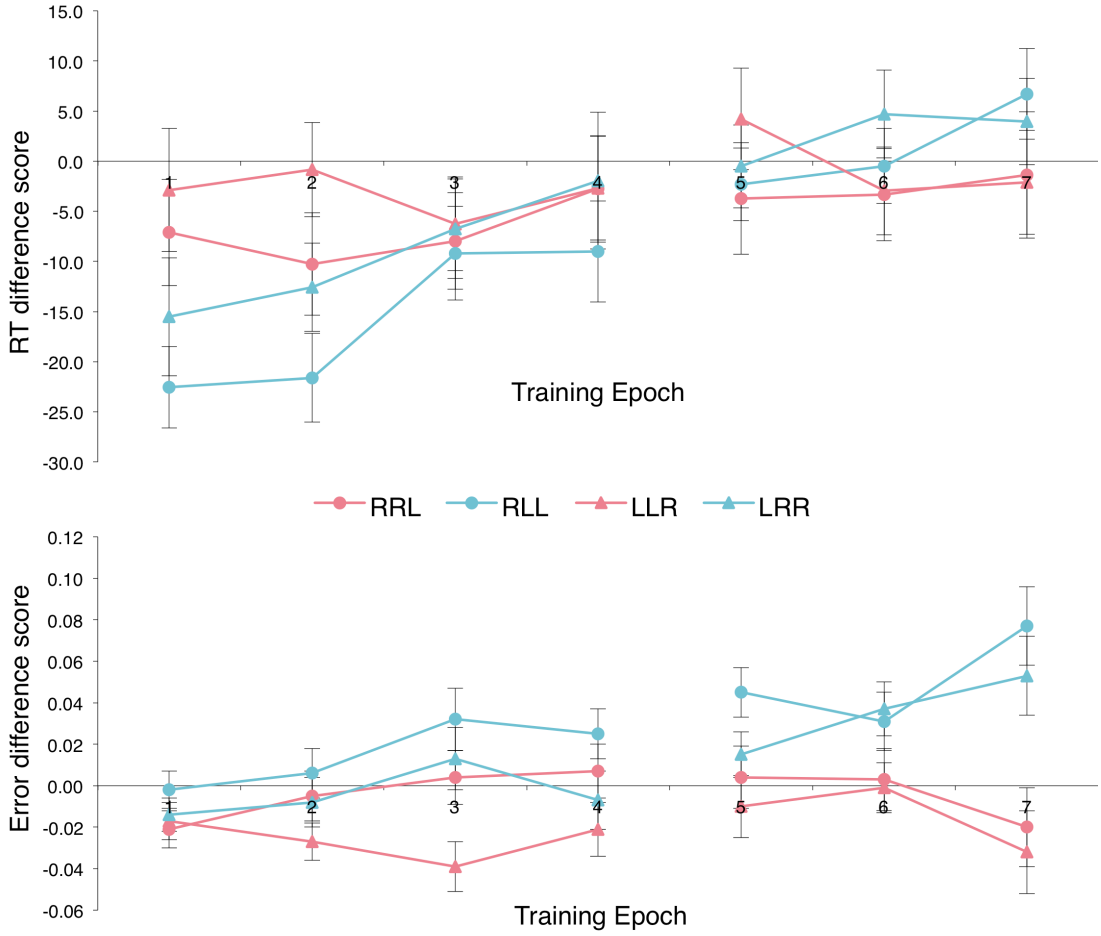

Figure 2.13. RT (top panel) and proportion of error (bottom panel) difference scores across training Epochs for the Different rule groups across the four trained Subsequences. The RRL data points, for example, shows performance on 'consistent' RRL trials subtracted from 'inconsistent' RRR trials. The data are collapsed across all 48 participants regardless of Experimental or Control groups. Error bars show standard error.

Subsequence learning. Again using a series of Bonferroni corrected comparisons the subsequences were taken alone to investigate which of these were successfully learned. An ANOVA with Epoch and Group across training and Location and Group at test demonstrates a main effect of learning across training for the subsequence RRL, RT difference score, $F(1,46) = 30.5$, $p < .001$; proportion of errors, $F(1,46) = 42.3$, $p < .001$, and at test in both RTs, $F(1,46) = 9.58$, $p = .007$, and proportion of errors, $F(1,46) = 15.6$, $p = .001$. The LLR subsequence was also learned well, with a main effect of group across training in both RT, $F(1,46) = 26.1$, $p < .001$, and error difference scores, $F(1,46) = 39.7$, $p < .001$; as well as at test in RTs, $F(1,46) = 6.41$, $p = .030$, and errors, $F(1,46) = 13.6$, $p = .001$. Therefore, there is strong evidence for learning of the subsequence YYX, which is also the case for XYY, as there is evidence of learning across training in both measures for RLL: in RTs, $F(1,46) = 42.7$, $p < .001$, and errors, $F(1,46) = 18.5$, $p < .001$; as well as at test in both RTs, $F(1,46) = 17.6$, $p < .001$, and errors, $F(1,46) = 7.16$, $p < .001$. Subsequence LRR was also learned well, with evidence from training RTs, $F(1,46) = 16.9$, $p < .001$, and errors, $F(1,46) = 15.5$, $p < .001$, as well as test RTs, $F(1,46) = 15.7$, $p < .001$, providing support for learning. Test errors, $F(1,46) = 2.78$, $p = .204$, do not reach significance for the subsequence LRR. Therefore, contrary to the findings of Jones and McLaren (2009), under intentional conditions participants seemed to have no trouble in learning the subsequences RRL, RLL, LLR, and LRR.

2.3.2.3. Same versus Different rule learning

I begin with an ANOVA comparing Epoch; Group; Rule and Subsequence across training and Group; Rule and Subsequence at test. All *p* values reported are corrected for multiple comparisons using a Bonferroni adjustment. There was a significant effect of Rule across training, RT difference score, $F(1,92) = 329$, $p < .001$, $MSE = 3856$, $\eta_p^2 = .781$; proportion of errors, $F(1,92) = 52.8$, $p < .001$, $MSE = .062$, $\eta_p^2 = .365$, and test, RT difference score, $F(1,92) = 123$, $p < .001$, $MSE = 1233$, $\eta_p^2 = .573$; proportion of errors, $F(1,92) = 60.7$, $p < .001$, $MSE = .007$, $\eta_p^2 = .398$. This demonstrates the very strong effect of subsequences, as this main effect occurred regardless of training and reflects the benefit of simply performing subsequences RRR, RLR, LLL and LRL over RRL, RLL, LLR, LRR. Rule also interacts with Group across training, RT difference score, $F(1,92) = 6.46$, $p = .013$, $MSE = 3856$, $\eta_p^2 = .066$; and proportion of errors, $F(1,92) = 4.96$, $p = .028$, $MSE = .062$, $\eta_p^2 = .051$; and at test in proportion of errors, $F(1,92) = 5.63$, $p = .020$, $MSE = .007$, $\eta_p^2 = .058$, but not RT difference score, $F(1,92)$

= .256, p = .614, MSE = 1233, η_p^2 = .003. Figure 2.14 compares the performance of both Rule conditions across training and test, which shows that the Different rule is learned better than the Same rule.

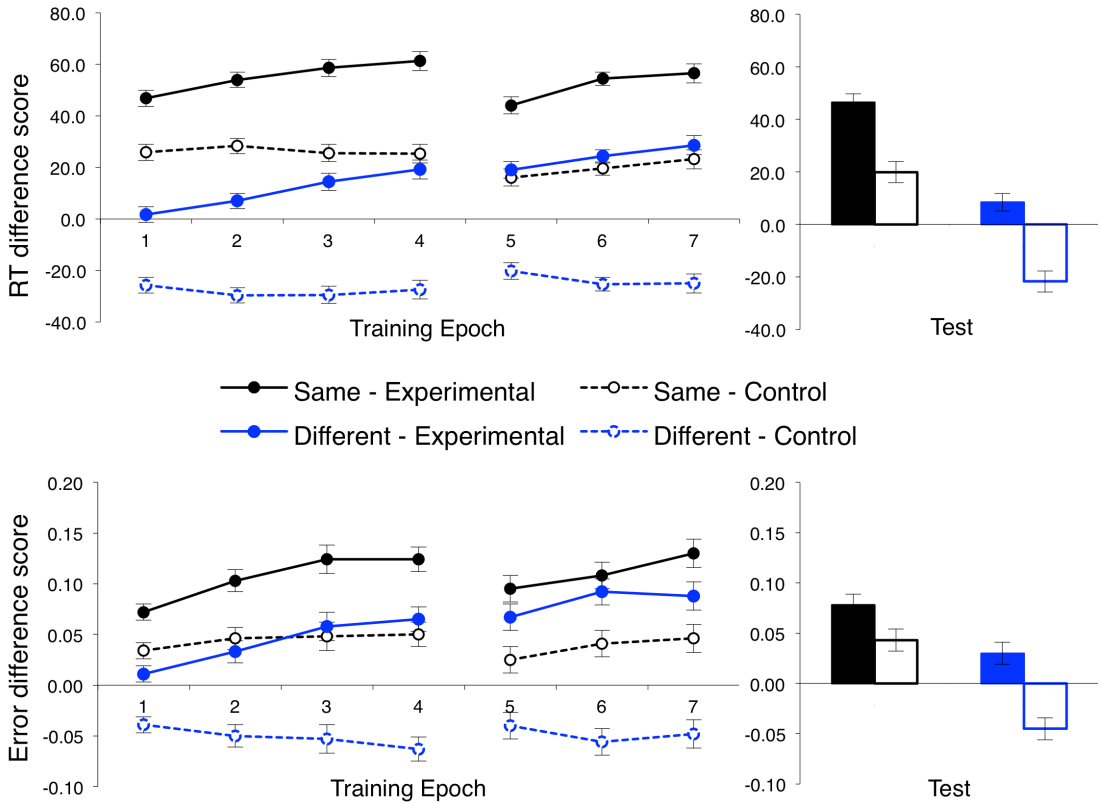

Figure 2.14. RT (top panel) and proportion of error (bottom panel) difference scores across training (left panel) and test (right panel) for Experimental and Control groups trained on Same (black lines) and Different (blue lines) rule subsequences. Error bars show standard error. Note: the blank between Epochs 4 and 5 represents the break between the two sessions.

Whilst the Intentional conditions provided evidence for a difference in learning between the two Rule conditions under Intentional conditions, this is not qualitatively different from the learning observed in Experiment 1. Rule interacts with Subsequence in errors across training, $F(3,276) = 19.6$, $p < .001$, $MSE = .032$, $\eta_p^2 = .175$; but not RTs, $F(3,276) = 2.98$, $p = .065$, $MSE = 9622$, $\eta_p^2 = .031$. There was also an interaction at test in both RT difference score, $F(3,276) = 6.58$, $p = .002$, $MSE = 1595$, $\eta_p^2 = .067$; and proportion of errors, $F(3,276) = 26.4$ $p < .001$, $MSE = .009$, $\eta_p^2 = .223$. This, however, simply reflects the sequential effects observed in Experiment 1, that participants preferred responding to Same rule subsequences and struggled most with Different rule subsequences. These effects showed little evidence of being affected by Group: training,

RT difference score, $F(3,276) = .396$, $p = .625$, $MSE = 9622$, $\eta_p^2 = .004$; and proportion of errors, $F(3,276) = .960$, $p = .377$, $MSE = .032$, $\eta_p^2 = .010$; test, RT difference score, $F(3,276) = .848$, $p = .428$, $MSE = 1595$, $\eta_p^2 = .009$; and proportion of errors, $F(3,276) = .829$, $p = .443$, $MSE = .009$, $\eta_p^2 = .009$, suggesting that there was no evidence for the influence of learning on these preferences. Similarly, the interaction across training between Epoch, Rule and Subsequence, RT difference score, $F(18,1656) = 4.07$, $p < .001$, $MSE = 1094$, $\eta_p^2 = .042$; and proportion of errors, $F(18,1656) = 2.78$, $p = .008$, $MSE = .013$, $\eta_p^2 = .029$, was not itself affected by Group and therefore is not related to learning, RT difference score, $F(18,1656) = .831$, $p = .595$, $MSE = 1094$, $\eta_p^2 = .009$; and proportion of errors, $F(18,1656) = .508$, $p = .826$, $MSE = .013$, $\eta_p^2 = .005$.

2.3.2.4. Post-experimental interview

When asked the same questions as the incidental participants in Experiment 1 participants in Experiment 2 gave surprisingly similar answers, even though they were informed about the presence of sequences. No participants reported being able to work out the rule, and reported that they thought the sequence was random or not with a similar frequency across the Experimental (19 reported that they thought there were no sequences) and Control group (22 participants reported that they thought there were no sequences). Similar to Experiment 1, participants reported noticing strings of repeats and alternations regardless of their group or rule, and reported either giving up on looking for sequences as they thought they did better by not concentrating; or that it was simply too hard to find a pattern and therefore they were led to assume the stimuli were random.

2.3.3. Discussion

Experiment 2 provided evidence of learning of both Same and Different rules under intentional conditions. As with Experiment 1, this learning was greater and more robust for the Different rule group and there was no apparent qualitative difference between Incidental and Intentional conditions. The prediction regarding intentional attention towards and consequently greater learning of the subsequence XXX (RRR and LLL, both Same rule subsequences) suggested that those participants using intentional rules would learn more in the Same rule group did not manifest in greater learning compared to the Different rule group.

That participants could learn RLR and LRL subsequences under intentional conditions (and as well as RRR and LLL) in this experiment when they did not in Jones and McLaren's (2009) study is interesting. Why this difference between participants in Jones & McLaren's (2009) intentional condition and Experiment 2? One explanation could be a ceiling effect as participants may be already responding as quickly and accurately as they can on the RRR and LLL subsequences, however an analysis of the inconsistent trials showed this not to be the case. Both experiments lasted for two sessions, however Jones & McLaren (2009) incorporated a pre-training test phase, where participants experienced pseudorandom control blocks before beginning their training. They were not told that these blocks were pseudo-random and simply instructed as participants were in the case of Experiment 2. Therefore, participants in Jones & McLaren's (2009) intentional group would have been both: trained for fewer blocks in total; and would have experienced a change between pseudorandom blocks and those with sequential contingencies. It seems unlikely that a shorter length of training would produce greater learning of subsequence XXX, but the alternative proposal regarding 10 blocks of pseudorandom pre-training may provide an explanation of this difference, as Jones & McLaren (2009) did not find evidence of differential subsequence learning at test. This suggests that the post-training pseudorandom blocks led to some unlearning or withholding of information regarding subsequence learning. It follows then, that under intentional conditions one would expect a change in trial structure to produce qualitative differences in strategy and performance. If this is the case, it seems possible that experiencing pseudorandom trials before training produces an effect when these trials begin to follow a sequential structure. The sudden appearance after the second Epoch of more strings of repeats would perhaps have attracted more attention than in the case of Experiment 2, where contingencies were in play from the first instance in Experimental groups.

There was no qualitative difference between learning in Different rule groups under intentional and incidental conditions, with all subsequences learned well in both cases. This provides some evidence that the instructional manipulation may not have had a huge effect on the learning processes involved. Jones & McLaren (2009, p. 546) reduced the six-session incidental length to two due to the probability that "participants would find it difficult to maintain hypothesis testing" for this length of time. Jiménez & Méndez (1999) suggest that participants can still demonstrate sequence learning

regardless of whether participants experience divided attention, which may provide some evidence that regardless of whether participants attempt to work out subsequences, that the cognitive strategies they may be using might not impact on the underlying learning processes that seem extraordinarily similar to those observed under incidental conditions.

The possibility that participants simply give up on hypothesis-testing suggests that using intentional participants in order to understand the qualitative differences between the two systems is not methodologically ideal, as participants (in either condition) do not necessarily follow these instructions. Indeed, many participants in the intentional condition reported that they found it easier to 'just do' the task, or that they gave up looking for sequences or patterns as they found it too difficult. Conversely, some participants under incidental conditions reported that they noticed or even looked for subsequences. It has been suggested by a variety of authors that these self-reports of strategies provide a post-hoc way to analyse participants (e.g. Curran & Keele, 1993), however this has a number of issues of control, as well as sensitivity of the information, which may have been the case for participants who did not report doing the task in some other way than instructed. A better, on-line measure of hypothesis testing could be derived from eye-tracking (e.g. Marcus, Karatekin, & Markiewicz, 2006) or using ERP markers (e.g. Fu, Bin, Dienes, Fu, & Gao, 2013) however this requires assumptions regarding the biopsychological measures and would still require an instructional manipulation (which would itself be checked by these assumed markers of hypothesis-testing).

The best evidence provided in Experiment 2 for the occurrence of conscious, rule-based, hypothesis testing is the changing subsequence effects across the experiment, regardless of whether participants were in Experimental or Control groups. This suggests that some shift in performance across the experiment occurs to subsequences, regardless of whether learning is occurring, which I suggest is explained by attentional changes to patterns observed in the data throughout training. That these, possibly intentional, processes do not interact with Group, suggests the evidence of learning that we have in Experiment 2 occurs outside of awareness. Whilst participants may engage with some hypothesis testing, this apparently has no effect on their learning.

Indeed, the results of the post-experimental interview reveal that participants made very little progress in discovering any rules underlying the task, and consequently were inclined to give up on their search. No participants could report the rule, or any element of the rule, which provides a strong indication that these Intentional participants did not possess explicit knowledge about the sequences. Therefore, in summary participants in the Intentional condition did not possess explicit knowledge. They may have been approaching the task differently, and attempting to uncover or search for patterns however it seems as if they were unable to find any patterns or rules as evidenced by the absence of a difference between the two conditions and the lack of explicit sequence knowledge exhibited by the Intentional group even after two hours of experience with these sequences.

2.4. Comparing incidental and intentional learning

Comparing the Incidental (Experiment 1) and Intentional (Experiment 2) experiments in a basic and qualitative way, as suggested by Jones & McLaren (2009) has already been discussed above to some extent. The two experiments do not provide evidence of a learning difference, as all groups learn all subsequences to some extent and the Different rule is learned more than the Same rule, regardless of instructions. Indeed the similarity between the results of the two conditions seems to suggest that participants under both sets of instructions learned in the same way. That this learning was entirely propositional, however, seems unlikely, as Experiment 2 produced evidence of sequential effects that differed across training, suggesting that participants shifted their performance across the task depending on some attention to or focus on different patterns or subsequence elements. Further to this, participants did not follow in either Experiment 1 or 2 the intentional subsequence learning observed in Jones & McLaren (2009) and instead both produced effects better reconciled with those produced by an associative model. How then are we to consider these results?

2.4.1. Comparing incidental and intentional conditions

An ANOVA with a Bonferroni correction for multiple comparisons was conducted to compare learning of the Same rule and, separately, learning of the Different rule across the two instruction conditions. Training data was analysed by Block (7); Group (Experimental vs Control); Subsequence; and Condition (Incidental vs Intentional).

There was no interaction between Condition and Group for the Same rule nor interaction between Condition and any other variable, providing support for the earlier suggestion that participants were performing the task similarly regardless of task instructions. The Different rule was also not learned differently by participants under intentional conditions and no interaction occurred between Condition and any other variable. What these results tell us about functionally separable learning processes is unclear. Jones & McLaren (2009) found that intentional learning followed the opposite pattern to incidental participants, which suggests that a qualitatively different learning processes was involved, as an exemplar or instance was learned by participants in the intentional condition whereas participants learned subsequences on the basis of associative predictions incidentally. All participants in Experiments 1 and 2, however, seem to demonstrate similar patterns of behaviour and therefore can be accounted for by a single process. Consequently, the instructed manipulation did not engage one system versus the other.

2.4.2. Qualitative differences and process purity

A simple 'qualitative differences' account (Jones & McLaren, 2009; Shanks, 2010) of two learning processes has been shown to be more complex than one might assume, as one cannot be sure (with current measurements of human psychological processes, Shanks & St John, 1994) that the learning processes: (a) occur in isolation; (b) occur as a result of the instructed manipulation; or (c) occur in isolation as the result of an instructed manipulation. The dual-process account suggests that verbalisable rule-based learning can control and take precedence over automatic learning (e.g. Reber, 1989), but that the automatic system is processing at all times; forming associations between stimuli regardless of intention (McLaren, Green, & Mackintosh, 1994). This suggests that automatic processes will always occur and that under this assumption we could extract the influence of one from the other (e.g. Jacoby, 1991).

Jones & McLaren (2009) provide evidence that intentional learners only show evidence of XXX learning, which could be due to explicit encoding of this sequence, but they do not consider why the other subsequences have not been learned. There is increasing influence of what some consider non-automatic processes (Shanks, 2010) such as attention, in the explanation of associative learning phenomena (Pearce & Hall, 1980; Mackintosh, 1975). Nissen and Bullemer (1987) demonstrated that attention to other

tasks could reduce incidental sequence learning and whilst the reason for this is debated (Jiménez and Méndez, 1999; Stadler, 1995) this suggests that some explicit influences may interrupt or interact with automatic processes, which is sometimes taken as evidence for the non-implicit nature of these processes (e.g. Shanks, 2010). Therefore, it may not be the case that there are two entirely separate systems and that automatic learning processes encode and learn all information available at all times (Pearce & Bouton, 2003). It may instead be the case that conscious effort can reduce or increase automatic learning by changing the amount of attention to the stimuli (Jiménez & Méndez, 1999) or interrupting the sequence itself (Stadler, 1995) in some way. There may, therefore, be a more complex interaction between the two that we are unable to tease out by simply instructing participants to do so one way or the other. Without understanding of how these two processes may interact we may be unable to dissociate the two (Sun, Slusarz, & Terry, 2005).

2.4.3. Dissociation logic

Many researchers concede the likelihood of an interaction between the two processes, for example Willingham, Nissen, & Bullemer (1989) demonstrate that increased awareness of stimulus regularities improved performance on the 'implicit' SRT task; yet they still suggest that the learning systems involved operate in parallel. These comparisons usually rely on dissociation logic, which as mentioned in the general introduction, has been shown to be essentially flawed (Dunn, 2003). The simple logic of which involves the demonstration that manipulating a given independent variable affects one dependent variable and not another. This single dissociation has been used to provide evidence for implicit learning as a functionally separate system (Jiménez & Méndez, 1999; Reber & Squire, 1994) but greater evidence is provided by authors who provide evidence of a double dissociation, where two independent variables produce complementary single dissociations on the same two dependent variables (Dominey et al., 1998; Jiménez et al., 2006; Jones & McLaren, 2009).

Jones & McLaren (2009) essentially provide evidence of such a double dissociation, with Group (Experimental versus Control) as the independent variable, a single dissociation is demonstrated as learning (the dependent variable) is observed for subsequences YYX and YXY, but not XXX. The opposite is true under intentional conditions (learning of XXX and not YYX or YXY), suggesting the presence of

multiple systems. Loftus (1978) has criticised the use of dissociations with bounded variables, such as accuracy, as floor and ceiling effects can be demonstrated to produce dissociations. Jones & McLaren (2009) avoid this critique with their analysis of inconsistent trials as reported for Experiments 1 and 2 here.

However, Dunn (2003) proposes that the logic of dissociations is fundamentally flawed, as whilst one can infer that a variable has an effect on performance of a given task, one can never infer that a variable has no effect on the performance of another task. In the case of the subsequence learning dissociation proposed by Jones & McLaren (2009) this means that the lack of learning for subsequence XXX under incidental conditions and the other subsequences under intentional conditions cannot be shown to be unaffected entirely by instruction. Indeed, when learning these subsequences as part of a different sequential structure in Experiments 1 and 2 participants do show an effect of incidental learning on the subsequence XXX.

2.4.4. State-trace analysis

As a solution to the study of single/multiple-process accounts of latent psychological processes, Dunn instead proposes the solution of state-trace analysis (Bamber, 1979; Dunn & Kirsner, 1988; Loftus, 1978). Instead of demonstrating single or double dissociations, state-trace analysis instead requires two *dimensions*, representing either one dependent variable measured under two different conditions, or two different dependent variables. In the case of Experiment 1 and 2 this could be the learning of Same and Different rule sequential contingencies. Performance is plotted across the *trace* of the experiment, i.e. across some continuous measure of time or number of blocks to produce the function of interest. In this case, this would correspond to plotting learning across the epochs of training. These plots can then be made for two or more independent variables of interest– these are the *states*. Here an example of a state manipulation would be both an incidental and intentional set of points. The crucial analysis consists of determining whether our two plots are best described as part of one continuous function, or require two distinct functions to capture each trace.

State-trace analysis has been applied to a large number of research areas where the question regarding the number of processes involved in a certain phenomena are disputed, including: category learning (Newell, Dunn, & Kalish, 2010; Newell, 2012),

cognitive development (Mayr, Kleigl, & Krampe, 1996), the face inversion effect (Loftus, Oberg, & Dillon, 2004; Prince & Heathcote, 2009), remember-know judgments (Dunn, 2008; Heathcote, Bora, & Freeman, 2010) and the neuroscience of recognition memory (Staresina, Fell, Dunn, Axmacher, & Henson, 2013). The remember-know judgments fall closest to the current line of investigation, which aim to investigate two (explicit and implicit) memory systems, but the only investigation so far to my knowledge involving two distinct learning processes is the work of Forrest (2012) who found evidence for multiple functions (and therefore for more than one process) underlying task-switching under two different task instructions.

2.4.4.1. State-trace analysis of Experiments 1 and 2

A state-trace analysis was conducted on the data, as described above taking the incidental (Experiment 1) and intentional (Experiment 2) results separately as the two states that we wish to investigate the dimensionality of. The dimensions (Same and Different) rule learning are taken as the dependent variables of interest for the axes, although a further difference score has to be calculated between the Groups as the Experimental or Control groups alone do not produce a dependent measure of learning. A learning score was therefore calculated for each Experimental participant for each Epoch in the first session as the difference between Experimental and average Control performance for that Rule. A by session analysis was used as the second session starts with a performance dip after the break between sessions, which obscures learning as a simple function of training: the trace. Figure 2.15 demonstrates the resulting plots for both RTs and errors across training for the data produced in Experiments 1 and 2.

These state-trace plots demonstrate firstly what the ANOVA showed, namely that the Different rule was learned better regardless of condition (note the larger scale of the Different rule axes compared to the Same), but the plots also seem to show that these functions may be separate on visual inspection (a method used by McCarley & Grant, 2008, in the interpretation of state-trace plots). This provides support that not only have the instructions produced some difference in sequence learning; further to this that there may be two functionally separable learning processes driving this difference. The functions on the error plots do not meet all of the requirements, however, for a full state-trace analysis, as they do not follow monotonic functions and do not demonstrate adequate overlap to infer how many functions are produced. Indeed error data is

suggested as incompatible with state-trace analysis as participants may be near floor on the task (Newell, 2008). The RT difference data are clear, however, and do seem to indicate two separate functions.

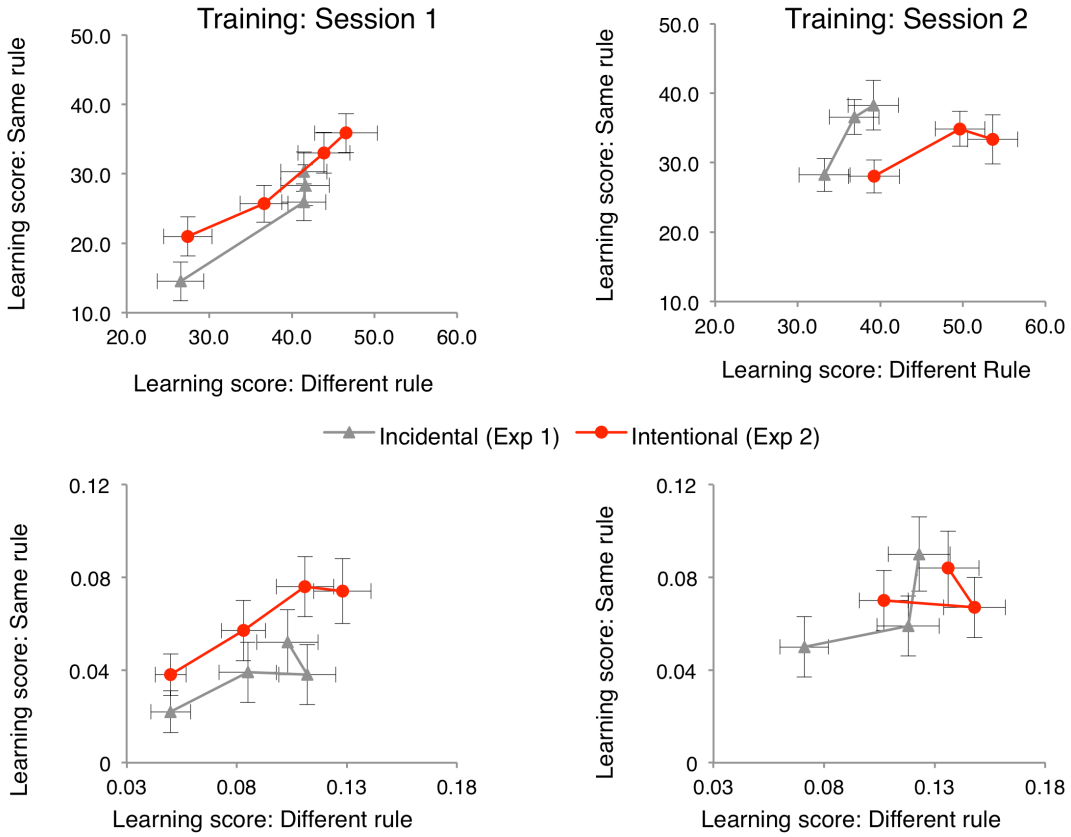

Figure 2.15. State-trace plots of RT (top panel) and error (bottom panel) learning scores across the first four training Epochs (all of Session 1, left panel) and first three training Epochs of Session 2 (right panel) for the two sequence Rules (Same: *y* axis and Different: *x* axis) in both Incidental (Experiment 1, shown in grey) and Intentional (Experiment 2, shown in red) conditions. Error bars show standard error.

There are a number of statistical methods to analyse state-trace plots in the literature (Spearman's Rho [Loftus, Oberg, & Dillon, 2004]; maximum likelihood estimate [Newell & Dunn, 2008] and Bayesian models [Prince, Hawkins, Love, & Heathcote, 2012]), but I follow Forrest (2012) whose work on learning processes is most conceptually related to these data. Using a stepwise multiple regression, Same rule learning (arbitrarily) was predicted from Different rule learning, with Condition added to the model in order to assess whether this variable could significantly increase the fit of the model and provide evidence for multiple functions. In Session 1, adding Condition to the model approached significance, R^2 change = 6.7%, $F(1,5) = 6.28$, p

= .057, increasing the already high R^2 value from 88.1% (R^2_{adj} = 86.1%) to 94.7% (R^2_{adj} = 92.6%). In Session 2, adding Condition to the model does significantly improve the fit of the model, R^2 change = 76.8%, $F(1,4) = 13.4$, $p = .035$, increasing the R^2 value from 5.9% (R^2_{adj} = 1.8%) to 82.8% (R^2_{adj} = 72.3%). This suggests that perhaps there are two distinct learning processes at work within the data. Whilst these Conditions may not have had a qualitative impact on which rule or subsequence was learned, this suggests that learning may have in fact involved more than one process.

2.4.4.2. State-trace analysis of computational simulations

What state-trace analysis can tell us about the dimensionality of the processes involved can, however, be called into question. Using a simple recurrent network (SRN, Elman, 1990, as discussed in Chapter 1 and described in more detail in Chapter 3) to simulate human performance on the task, I produced the same state-trace analysis performed on human data on two sets of model simulations. The task, the rules, number of trials and blocks were identical to Experiments 1 and 2, as described above, however the model was not simulated under different conditions to form the states of interest, but with different learning rates. For a further explanation of the role of learning rates within the SRN, see Chapter 3, but for the purposes of this explanation, by simply increasing the amount of learning on each trial, the simulations of the task using an SRN produced two visibly divergent state-trace functions, see Figure 2.16.

A model with a higher learning rate learns less about Different rule sequences relative to Same rule sequences, this suggests that the state-trace methodology is sensitive to variations within a single process (learning rate), rather than providing evidence of multiple processes. Therefore, that the state-trace plot for incidental and intentional human performance (Figure 2.15) differs does not necessarily imply that two or more processes were at work. This suggests that participants under intentional conditions could have varied in some single process parameter, for example, by analogy to the model simulation they could have learned more on each trial than under incidental conditions. Moreover, it seems that a state-trace plot may claim to circumvent the problems faced when attempting to uncover the number of processes underlying such functions (Loftus, Oberg, & Dillon, 2004) but this simulation result suggests that it may be sensitive to variations within a single process and coupled with the results of Ashby (2014) call the method into question. Importantly for this thesis, despite the difference

between Incidental and Intentional performance on a state-trace plot; this cannot provide conclusive evidence for the presence of functionally separable learning processes, though it is suggestive.

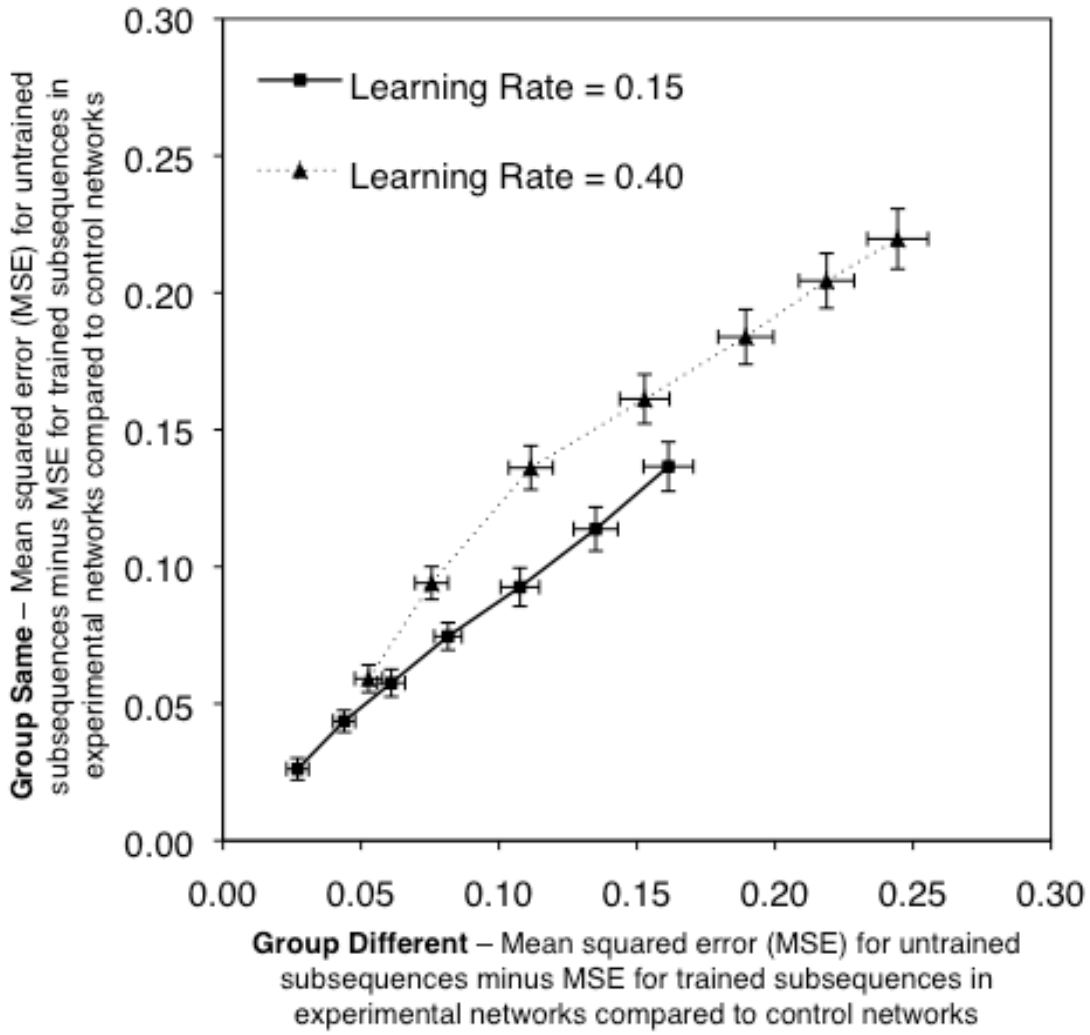

Figure 2.16. State-trace plot of mean squared error (MSE) learning scores across the all seven training Epochs for the two sequence Rules (Same and Different) in both low (0.15) and high (0.40) learning rate simulations of the simple recurrent network (SRN, Elman, 1990). Error bars show standard error.

2.5. General Discussion

Experiment 1 provided evidence of learning under incidental conditions, with the stochastic Different rule learned better than the Same rule. The subsequence learning effects observed in Experiment 1, however, suggest that participants were not impaired at learning subsequence XXX under incidental conditions as shown in Jones and

McLaren (2009). This suggests that rather than participants struggling to learn the subsequence XXX per se, subsequence learning occurs differently depending on the stochastic structure of the trained sequence itself. While the exclusive-or rule trained in Jones and McLaren (2009) required that participants use both of the two trials preceding the current trial to predict the required response; the Same and Different rules in Experiments 1 and 2 only required that participants use the one trial that occurred two trials previous to the current, for which two explanations can be offered.

The first is that participants in Experiment 1 employ a more explicit strategy, that has been shown to increase learning of the subsequence RRR and LLL, which occurred more frequently in the Same rule group than in Jones & McLaren's (2009) exclusive-or rule case. Therefore, participants in Experiment 1 may have noticed the XXX subsequence more and, as a result, learned this subsequence more. It seems unlikely that this is the case, as evidenced by the intentional case (Experiment 2). If indeed the increased occurrence of XXX led to more noticing of this sequence and therefore more learning in Experiment 1, we would also expect to see this (to a greater extent) when participants were instructed to look for and use sequences in Experiment 2. As we do not see greater learning of the Same rule under intentional conditions or an interaction between subsequence learning and conditions it seems unlikely that some verbalisable strategy is accounting for the difference here.

The second explanation of such a result comes from the cue competition account of XXX learning in Jones & McLaren's (2009) model of the results: that the reduction in error term based on the previous two trials following the same mapping from previous to current trial (X leads to X, followed by X leads to X) would block learning of the sequential contingency occurring at *t*. In the exclusive-or case, as both mappings were required to accurately predict *t*, as the first mapping (*t* - 2 to *t* – 1) reduces the error term for the second case (*t* -1 to *t*), then the amount of learning is blocked. In Experiment 1 only the relationship between the first and last items in the subsequence are important for learning the Same rule, therefore, participants may learn that the mapping from *t* – 2 to *t* – 1 has no contingency with *t* and thus learn to disregard this mapping, with the resulting consequence a reduction in the cue competition effect. The effect, however, may not be entirely eradicated as the Same rule is still learned less well, therefore learning the overall structure of the sequential contingencies may not be protected from

the influence of some short-term carry over from the previous trial, even when this has no relationship to the trained contingencies.

This account relies on the short-term priming system applied in the AugSRN (e.g. Cleeremans & McClelland, 1991), which is the subject of Chapter 3. Further work is required to ascertain whether the sequence learning observed with Same and Different rule sequences was indeed produced by this interaction between transient trial-by-trial sequential effects and the nature of the trained sequential contingencies. This highlights the advantage of a computational approach, as the precise mechanisms by which learning can occur can be examined in great detail. Therefore, the next aim of this thesis is to attempt to model the results of Experiment 1, with a very precise set of human learning effects provided by both previous research and the differential subsequence learning effects produced by the Same and Different rule structures.

There was no qualitative difference between Experiment 1 and 2 as defined by a difference in subsequence learning, possibly suggesting that participants did not exclusively perform the task intentionally in this condition and that the instructional manipulation was not effective. Indeed participants were unable to express any knowledge about the sequential rule, even though they had been trained over two long sessions on the task. This suggests that while they may have had volitional control over their intention to learn on the task, that they were still unable to discover a pattern explicitly. It may also be the case that they discovered patterns but did not express this knowledge in the interview or on the task itself (Lee & Livesey, 2013).

It seems that incidental processes remained present in both conditions regardless of intention to learn, as Different rule learning was greater in both experiments and sequence learning was relatively unaffected when participants employed explicit strategies. Whilst it might be possible to take these results as evidence for a single learning process that is not different depending on the volitional conditions under which the task is performed; the results of Jones and McLaren (2009) clearly suggest that this is not the case. Further evidence that both the Incidental and Intentional pattern of learning was driven by automatic, associative processes may be provided if the Different rule learning advantage can be simulated by a model of associative learning, which is the subject of Chapter 3.

The difficulty in comparing these two conditions within a dissociable framework are numerous, and assumptions regarding test and process purity, the lack of interaction between the two processes, and dissociation logic cause problems in trying to establish the processes that underlie learning on such tasks. State-trace analysis (Bamber, 1979) was offered as the best available solution to understanding the processes that underpin performance on a task (Loftus, Oberg, & Dillon, 2004; Newell & Dunn, 2008; Prince et al., 2012), but whilst functions produced by the second session of training (and to some extent the first) provided some evidence for multiple functions, state-trace analysis might reveal variation within some single process and does not necessarily suggest that the instructed manipulation produced a functional difference in the processes involved.

This further suggests that relying on control when dealing with awareness is difficult, as participants are not only subject to the influence of possible automatic processes, but are not necessarily able to or likely to follow these instructions. In a task like this, which is very demanding and quite boring, participants in both Incidental and Intentional conditions reported a whole variety of different motivations, strategies and experiences. The likelihood that all participants in the Intentional condition performed with perfect control and were able to isolate their explicit learning and apply this on each trial is slim, therefore instructing participants does not ensure either the presence or absence of explicit learning; if ultimately participants are unable to uncover what is to be learned. Whether participants could notice and use these Same and Different rule sequences is a question not addressed in this thesis, but forms a future research strand that might involve longer training, the use of hints or even explicitly giving participants the rules and asking them to use them, although this does not always result in learning on the SRT task (Lee & Livesey, 2013).

Further discrepancies between Jones and McLarens' (2009) results and those reported here lie in the sequential effects observed in control groups under both conditions. Jones and McLaren (2009) simply observed a first order repetition advantage, with participants preferring to respond to subsequences XXX and XYY, with control subjects responding to YXY and YYX with a difference score of below zero. However, there was no difference in the control performance of participants in terms of whether they preferred a trial that immediately followed the same response location or not; as there was evidence for a speed accuracy trade-off between faster responses to first order

alternations, and more accurate responding to repeats. Whilst these effects were therefore unclear, subsequences with difference scores above and below zero were clearly RRR, LLL, RLR and LRL (above zero) and RRL, LLR, RLL and LRR (below zero). Therefore, in Experiments 1 and 2 control participants showed a preference for faster and more accurate responding to Same rule subsequences over Different rule subsequences. There should be no difference between the experience of control subjects between Jones and McLaren (2009) and this study, however, the length of training was greater in Jones and McLaren's study, and participants were given monetary rewards for fast and accurate responding. It is not entirely clear why either of these changes would influence sequential effects on control subjects, however Chapter 3 attempts to investigate these issues computationally.

Altogether, while the disparity between the Jones & McLaren (2009) dissociation between intentional and incidental conditions is a convincing demonstration and framework within which to investigate human learning, this is based on the assumptions that: (1) participants followed the incidental and intentional instructions and (2) that these instructed conditions were able to isolate or give preference to particular learning processes. Experiments 1 and 2 provide evidence that an instructed comparison between participants may not produce clear-cut subsequence learning differences, and without an understanding of how the systems operate or a superior method of disentangling these results we are unable to use this research design to better understand human learning.

Chapter 3. Modelling stimulus-response associations and sequence learning

In Chapter 3 I aim to investigate whether the sequence learning effects observed in Experiment 1 can be simulated by the Augmented SRN (AugSRN, Cleeremans & McClelland, 1991), which has previously provided a convincing account of Jones & McLaren's (2009) data by means of competition between transient trial order learning effects and learning of the sequential contingencies. The AugSRN, however, was unable to demonstrate the learning observed in Experiment 1 under the parameters described by either Jones & McLaren (2009) or Cleeremans & McClelland (1991). Further investigations using an SRN (Elman, 1990) and exploring the parameter space both manually and using optimisation procedures provide a convincing demonstration that the SRN or AugSRN is unable simulate the results of Experiment 1. What these models lack, however, is any representation of simple stimulus-response contingencies experienced on each trial, which are not important for sequence learning. When introduced into the AugSRN these produce the learning effects observed in Experiment 1, suggesting that cue competition between simple stimulus-response mappings and sequential information (as well as transient trial-by-trial learning) has an effect on learning of different subsequences.

3.1. Introduction

Chapter 2 presented evidence of differential sequence learning under incidental conditions, as participants trained on the Different rule (that the current trial [t] can be predicted as the opposite location to two trials previous [$t - 2$]) learned more than those trained on the Same rule (that t can be predicted as the same location as $t - 2$), even though the Different rule subsequence are performed worse than Same rule subsequences. As discussed in Chapter 1, the Augmented SRN (AugSRN) is a version of the SRN (Elman, 1990) adapted to account for the trial order effects observed in human serial reaction time (SRT) tasks (Cleeremans & McClelland, 1991). While the SRN is a seminal model regularly employed in the simulation of sequence learning (Beesley, Jones, & Shanks, 2012); the SRN is unable to account for trial-by-trial effects

such as enhanced responding to repeated stimuli, for example. Therefore, the AugSRN is proposed to be the best model of human sequence learning (Jones & McLaren, 2009). This chapter aims to investigate the detailed pattern of the learning and sequential effects demonstrated in Experiment 1, which provide a framework with which to accurately assess a computational model of human learning.

As mentioned in Chapter 2, Experiment 1 showed not only greater performance of Same over Different subsequences in Control groups (sequential effects); but human participants showed evidence of greater Different rule learning compared to Same rule learning. Whilst the Same rule was learned less well, there was still evidence of learning of XXX and strong learning of XYY: although neither of these subsequences were shown to be learned under incidental conditions over greater training in Jones and McLaren's (2009) study. The cue competition effects observed between transient trial-by-trial contingencies could account for these results if participants had learned to disregard the middle trial in a subsequence, as neither group in Experiment 1 required this trial to predict *t.* To ascertain whether the relationship between transient trial order learning effects and more permanent learning of the underlying sequential contingencies within the task could produce these subsequence learning effects, the AugSRN was used to simulate the task.

3.2. The Augmented SRN

The AugSRN is described in detail here, with the construction and components of the model discussed in this section. The first aim of this Chapter was to produce a version of the AugSRN that matched the simulations produced by Jones and McLaren (2009).

3.2.1. Model construction

The model was constructed and simulated using MatLab software and the details of the construction of the model and all of the parameters and aspects of the simulations are described below. The model was constructed using the specifications outlined in previous research and incorporating the parameters given in Jones and McLaren (2009) for the simulation of these results. The method of simulation is also described, which was run in a way that attempted to approximate the human Experiment 1 as closely as possible. Each iteration of the model represents a trial in the human experiment, and

therefore 4800 iterations simulated two 'sessions' worth of one participant's data. The connection weights and the hidden unit activations were then reinitialised, representing the random variation produced across different human participants (Juola & Plunkett, 1998). The model was then run again for 4800 iterations, and this was done for each of the 96 total participants in Experiment 1.

3.2.1.1. Model architecture

The model is constructed of units that are organised into four layers: input, hidden, output and response (shown in Figure 3.1). Each unit in each layer is connected to each unit in the next layer, so each input unit has a connection with each hidden unit, each hidden unit with each output unit, and each output unit with each response unit.

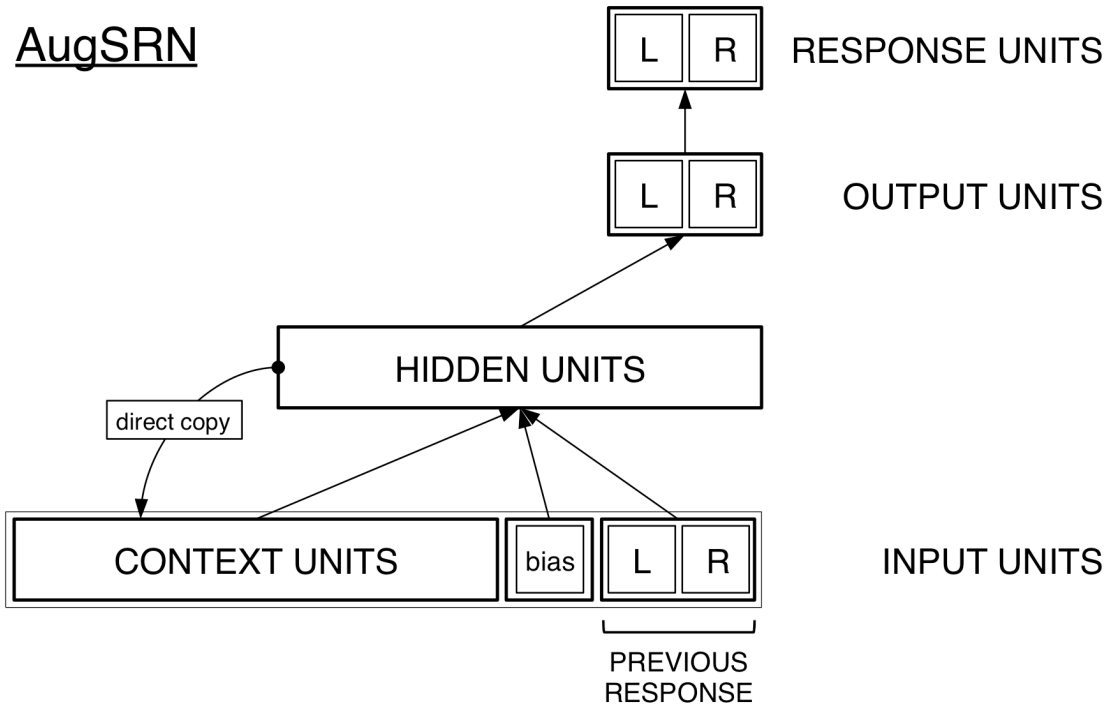

Figure 3.1. Structure of the AugSRN (Cleeremans & McClelland, 1991). Input units (bottom) include both left (L) and right (R) stimuli as well as a copy of the hidden units on the previous trial and a bias (a unit that is always on). Activation flows in the direction of the arrows, with a set of hidden units (middle) passing activation forward once more to output and then response units, again representing L and R stimulus locations and/or responses (top).

Input units. There are three types of input unit, the first (shown bottom right in Figure 3.1) simply representing the left or right stimuli experienced by participants. The second is a hidden bias unit (shown bottom in Figure 3.1), which is a single unit that always has

an activation of 1. The third are context units (shown bottom left in Figure 3.1), which contain a direct copy of the hidden unit activations on the previous trial. There are then, the same number of context units as hidden units (in this case 20); and therefore there were 23 input units in total in the AugSRN.

Hidden units. 20 hidden units (shown in the middle in Figure 3.1) receive activation from the input units, which they feed-forward into the output units. On every trial the hidden unit activations are also copied directly into the context units, ready to feed these internal representations back into the model on the next trial.

Output units. 2 output units (shown second from top in Figure 3.1) form the new layer, one for each left and right stimuli again.

Response units. These form the final layer (shown top in Figure 3.1) in the network and there are again 2 response units, one for each left and right stimuli.

3.2.1.2. Representation and activation of task input and output

Input units and activation. The two input units that discretely represented a left or right stimulus were given an activation of 1 to simulate that stimulus was present; and 0 when it was absent. Therefore, on each iteration of a simulation one of two units would be activated and the other not, depending on the type of trial (right or left) being simulated. As outlined above, activation of the two input units discretely representing stimuli at *t* ($Input_R$ and $Input_L$) was externally manipulated to values of 1 or 0 depending on the trial sequence. The bias unit was always activated to a value of 1, and the context units were given the activation values on each trial exactly equal to the hidden unit activations on the previous trial:

$$Context_{1:20}(t) = Hidden_{1:20}(t-1)$$

(Equation 3.1)

This means that the internal representation of the previous trial was fed back into the model at *t* and is therefore the recurrent element that enables the model to learn contingencies across trials and to represent time.

Hidden units and activation. Each of the 23 input units had two connections to each of the 20 hidden units, which totals 920 connections through which activation flows forward on each trial. The amount of activation that is fed from each of these input units through each of the connections to each hidden unit is determined by the strength of each connection (*w*) between two units and the logistic activation function (Rumelhart, Hinton, & Williams, 1986). Therefore, for a unit *j*, receiving input from one unit *I*, its activation, $Unit_j$ is:

$$Unit_j = \frac{1}{1 + e^{-(Unit_i \times w_{ji})}}$$

(Equation 3.2)

However, as already indicated, for each unit within the AugSRN there are two connections (described in more detail in sections 3.2.1.3 and 3.2.1.4): a 'fast' (*wFast*) and 'slow' (*wSlow*) connection for each unit, so activation for the hypothetical unit *j* as part of an AugSRN would be calculated as:

$$Unit_j = \frac{1}{1 + e^{-((Unit_i \times wFast_{ji}) + (Unit_i \times wSlow_{ji}))}}$$

(Equation 3.3)

Thus, for each unit with multiple connections, these connection weights and activation values can be summed to provide a measure of net input that are converted to activation of the unit in question through the logistic activation function:

$$\begin{aligned} &NetInput_{Hidden_j} \\ &= (Input_R \times wFast_{Hidden_j Input_R}) + (Input_L \times wFast_{Hidden_j Input_L}) \\ &+ (Input_{Bias} \times wFast_{Hidden_j Input_{Bias}}) + (Context_1 \times wFast_{Hidden_j Context_1}) \\ &+ (Context_2 \times wFast_{Hidden_j Context_2}) + \ldots (Context_{20} \times wFast_{Hidden_j Context_{20}}) \\ &+ (Input_R \times wSlow_{Hidden_j Input_R}) + (Input_L \times wSlow_{Hidden_j Input_L}) \\ &+ (Input_{Bias} \times wSlow_{Hidden_j Input_{Bias}}) + (Context_1 \times wSlow_{Hidden_j Context_1}) \\ &+ (Context_2 \times wSlow_{Hidden_j Context_2}) + \ldots (Context_{20} \times wSlow_{Hidden_j Context_{20}}) \end{aligned}$$

(Equation 3.4)

$$NetInput_{Hidden_j} = \sum (Input_i \times wFast_{ji}) + (Input_i \times wSlow_{ji})$$

(Equation 3.5)

$$Hidden_j = \frac{1}{1 + e^{-\sum NetInput_{Hidden_j}}}$$

(Equation 3.6)

Output units and activation. The output units also discretely represented a left or right stimulus, and on each iteration the model was trained to predict the *next* trial. For example, if the first two trials experienced by participants were a right and then a left, the first trial would involve activation of the right input unit ($Input_R$) to 1, and the left input unit ($Input_L$) to 0. The target activation on that first trial was set for the right ($TargetOutput_R$) as 0.1, and for the left (TargetOutput) as 0.9 both activations corresponding to that required on the next trial.

The output units receive input from the hidden units in the same way, with activation for one output unit calculated from the net input of 40 connections (20 fast and 20 slow), and therefore the activation of the right output unit can be calculated as:

$$Output_R = \frac{1}{1 + e^{-\sum NetInput_{Output_R}}}$$

(Equation 3.7)

$$NetInput_{Output_R} = \sum (Hidden_j \times wFast_{Rj}) + (Hidden_j \times wSlow_{Rj})$$

(Equation 3.8)

The difference between this target output and the activation of the output units ($Output_R$ and $Output_L$; discussed in section 3.2.1.2) was taken as an index of how well the model predicts the next trial and therefore how much it has learned and is able to use this to predict the next trial. This was calculated on each trial by taking the average of the squared differences between target and actual output activations, a mean squared error (MSE):

$$Output\ MSE(t) = \frac{(TargetOutput_R(t) - Output_R(t))^2 + (TargetOutput_L(t) - Output_L(t))^2}{2}$$

(Equation 3.9)

The discrete response units may also be used as an index of task performance (as in Cleeremans & McClelland, 1991; and Jones & McLaren, 2009), but the essential comparison between target and actual activation is the same. The response unit MSE can be calculated as:

$$Response\ MSE(t) = \frac{(TargetResponse_R(t) - Response_R(t))^2 + (TargetResponse_L(t) - Response_L(t))^2}{2}$$

(Equation 3.10)

MSE is therefore a measure of prediction strength that can index RT performance on the next trial (Jones & McLaren, 2009), as MSE will decrease with increased prediction strength, as one expects RTs to do with increased prediction strength. Smaller MSE values therefore indicate greater learning (less difference between expected and actual output/responses), and these reflect shorter RTs on the next trial as the MSE measures prediction regarding the next trial, and therefore reflects the speed of responding to the next trial in humans:

$$RT(t) = MSE(t-1)$$

(Equation 3.11)

MSE values were taken from either Output or Response units (depending on the simulation, detailed below) when inputs reflect $t - 1$, which gives RT for trial t for the eight different triplets in each block. Difference scores between inconsistent and consistent subsequences were calculated as outlined in Experiment 1 (see section 2.2.2).

Response unit activation. To capture the influence of the previous trial location or response, two response units were activated depending on the corresponding output unit on that trial, and incorporating a decayed trace (weighted by k, which takes the value of 0.5) of the previous response:

$$Response_R(t) = Output_R(t) + k \times \big(1 - Output_R(t)\big) \times Response_R(t-1)$$

(Equation 3.12)

At this point, a MSE was calculated, using the equation outlined in this section. After this, the response units were set to the appropriate binary response required of the task, thus on the first trial example already given where *t* is right and *t* + 1 is left, $Response_L$ would be set to 1 and $Response_R$ to 0, following Cleeremans & McClelland (1991). Each trial is considered a correct trial as the model approximates RT responses that in human data were only analysed for correct trials. Further decision-making processes that would enable the modeling of error data (e.g. Jones, Wills, & McLaren, 1998; Wills & McLaren, 1997) were not employed as these elements may have interfered with the analysis of the purely associative processes of interest (Jones & McLaren, 2009).

3.2.1.3. Error-correction and back propagation

After the input unit activation on a single model iteration or trial had fed forward activation all the way to the response units, these feed-forward connection weights were updated according to back-propagation of error. This means that, from the output units backwards, a difference between target and actual activations was calculated, forming the basis for the amount of learning that would occur (how much each connection weight would change). This followed the standard back-propagation algorithm for error-correction as developed by Rumelhart, Hinton and Williams (1986), taking the delta rule, whereby the amount of change in connection weight is a function of the difference (error) between output unit activation and the trained target activations. Thus, the error term for the output units was computed as follows:

$$\delta_{Output_R Hidden_j} = (TargetOutput_R - Output_R) \times (1 - Output_R) \times Output_R$$

(Equation 3.13)

The error is then back-propagated to connections between the input and hidden layer using the delta from this first set of error calculations:

$$\begin{aligned} &\delta_{Hidden_j Input_R} \\ &= \delta_{Output_R Hidden_j} \times \left(wSlow_{Output_R Hidden_j} + wFast_{Output_R Hidden_j}\right) \\ &+ \delta_{Output_L Hidden_j} \times \left(wSlow_{Output_L Hidden_j} + wFast_{Output_L Hidden_j}\right) \times (1 \\ &- Hidden_j) \times Hidden_j \end{aligned}$$

(Equation 3.14)

Each 'slow' connection weight between any two units (again the hypothetical *i* and *j*) is updated simply by calculating the change in weight multiplied by both the learning rate (α) parameter (described in section 3.2.1.5) and the activation of the unit feeding activation forward through the connection and adding this to the pre-existing connection weight:

$$wSlow_{ji} = wSlow_{ji} + \left(Unit_i \times \delta_{ji}\right) \times \alpha_{Slow}$$

(Equation 3.15)

For 'fast' learning rates this calculation also involves an element of decay, as the previous connection weight is multiplied by the constant *k* (with a value of 0.5) and hence the learning is more transient as it decays by half at each time step:

$$wFast_{ji} = \left(wFast_{ji} \times k\right) + \left(Unit_i \times \delta_{ji}\right) \times \alpha_{Fast}$$

(Equation 3.16)

3.2.1.4. Learning parameters

Learning rates (α) usually take a value between 0 and 1 and thus reduce the amount of weight change that can occur on each trial: producing gradual learning of contingencies and avoiding radical step-changes on each trial that produce oscillations in performance. These learning rates vary from model to model and across simulations, for instance Jones and McLaren (2009) use a slightly higher learning rate than Cleeremans and McClelland (1991). These issues of parametisation are discussed later in this Chapter (see section 3.3.1), but in general terms the AugSRN has two sets of connection weights between each unit in the model, defined by their learning rates (as well as the presence [fast weights] and absence [slow weights] of decay).

The slow weights enable associative learning through a lower learning rate that encourages strong associations that accrue gradually over time. Fast weights have been suggested by various authors as a secondary component of learning that influences subsequent trials but have little permanent effect (Hinton & Plaut, 1987; McClelland & Rumelhart, 1985). To account for the transient effects observed in SRT tasks, Cleeremans and McClelland (1991) introduced fast weights that had a higher learning

rate than the slow weights, thus producing greater learning. However, these fast weights also decay by half over each trial, meaning that this learning is short lived.

The learning rates employed in simulations are often altered, and are generally accepted to be free parameters that can be altered, depending on the task, stimuli or even to represent individual differences in human performance (McLeod, Plunkett, & Rolls, 1998). Cleeremans and McClelland (1991) used learning rates of 0.15 and 0.2 for slow and fast learning rates, respectively. Jones and McLaren (2009) argued that as their task involved fewer stimuli (two rather than six) and longer response-stimulus interval (RSI: 500 ms rather than 120 ms), this justified an increased learning rate: of 0.4 and 0.533 for slow and fast learning rates, respectively. The proportional difference between the slow and fast learning rates was matched across the two studies. As Experiment 1 is based on Jones and McLaren's (2009) study, the learning rate parameters of 0.4 and 0.533 were the ones used in the simulation of Experiment 1.

3.2.1.5. Simulation procedure

As mentioned previously, the simulation of human performance attempted to match Experiment 1 as closely as possible. 96 simulations of 4800 trials were run, each representing one human participant: 48 Control and 48 Experimental. Half of the simulations were trained on pseudorandom trial orders throughout, as outlined for Control groups in Experiment 1. The other 48 participants were either trained on 4200 trials that followed Same or Different rules described in Chapter 2. The trial construction and ordering was exactly the same as for human participants and data was analysed in the same way, with MSE at output for input at $t - 1$ taken as an index of RT performance on the task, with an inconsistent minus consistent difference score calculated for each block and the four trained subsequences. As for the human participants in Experiment 1, this difference score was calculated for all groups over 118 trials on each test block, even though there was no simulated break or separation between blocks. Weighted averages were calculated for 117 trials across each training block to control for $t - 3$ sequential effects.

At the beginning of each simulation of 4800 trials, to represent a new participant the connection weights between all units were randomly given weights between -0.5 and 0.5 and the hidden units were reset and given activations of 0.5. This provides some

variation and error into the simulations, approximating individual differences between subjects (Juola & Plunkett, 1998).

3.2.2. Simulation 1: Jones & McLaren (2009)

3.2.2.1. Simulation procedure

As a check on the model, the Jones and McLaren (2009) incidental learning experiment was simulated. This involved 80 simulations, 40 each Experimental and Control, with Control groups receiving pseudorandom blocks throughout. There were 10 pre-training blocks (1200 pseudorandom trials) for the Experimental networks, which were then trained for 80 blocks (9600 trials) on sequences constructed using the same procedure as described in Chapter 2, but with an exclusive-or rule rather than a Same of Different rule. Therefore the model was twice as likely to experience: XXX, XYY, YYX, YXY than the complementary, inconsistent subsequences (XXY, XYX, YYY, YXX). Note that Xs and Ys used as R and L was counterbalanced across networks. This was followed by a further 10 post-training test blocks (1200 further pseudorandom trials) and 10 further training blocks (1200 further exclusive-or trials).

3.2.2.2. Results

The results were analysed by an ANOVA on the MSE difference scores taken between inconsistent and consistent subsequences with the factors Epoch (Sets of 5 blocks as for Experiment 1 and 2: 16 across training and 2 across test), Group (Experimental versus Control) and Subsequence (XXX, XYY, YYX, YXY). There was a large main effect of Group across training, $F(1,78) = 592$, $p < .001$, $MSE = .002$, $\eta_p^2 = .884$, and test (post-training), $F(1,78) = 67.1$, $p < .001$, $MSE = .001$, $\eta_p^2 = .463$, with the post-training test results shown in Figure 3.2, alongside a reproduction of these same test results from the AugSRN simulations presented in Jones and McLaren (2009). These results match those found by Jones and McLaren (2009): with an interaction between Subsequence and Group at test, $F(3,234) = 31.4$, $p < .001$, $MSE = .002$, $\eta_p^2 = .287$ that followed the same ordinal pattern as produced in both their human participants and the model. Therefore the AugSRN constructed for this thesis could provide evidence of differential subsequence learning and accurately simulate human performance under incidental conditions as found by Jones and McLaren (2009). Note that the precise sequences experienced by each network, as well as randomization of weights at the start of simulations meant that the two sets of simulations will not match precisely.

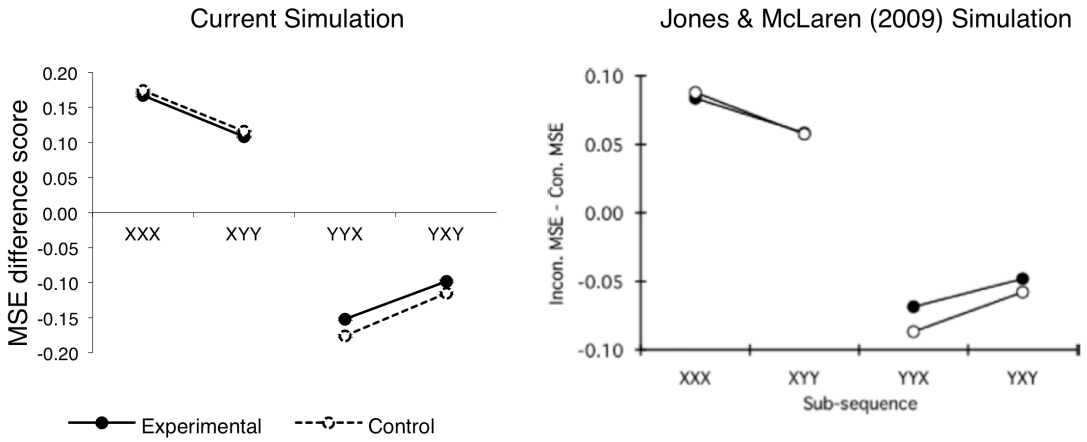

Figure 3.2. MSE difference scores for the post-training phase of AugSRN simulations taken from the current work (left panel) and the work of Jones and McLaren (2009, right panel). Error bars are shown for the current simulation (although obscured by the plot markers as these are very small) and show standard error.

3.3. Simulation 2: Simulating Experiment 1 with the AugSRN

3.3.1. Simulation details

The procedure for simulating Experiment 1 is described in section 3.2.1.5, and matches the experimental design outlined in Chapter 2. Using the AugSRN and the parameters outlined by Jones and McLaren (2009) the model was used to simulate human between-subject performance differences on the incidental sequence learning task. The analysis of each set of simulations was treated in the same way as human RT and error difference scores. ANOVAs investigated learning in each of the two Rule groups (comparing Experimental to Control for Same and Different rules, separately) followed by a Bonferroni corrected analysis of the full model, involving Block, Group, Rule and Subsequence across training and test. The decision was made to average MSE across Blocks rather than Epochs due to the interest in the evolution of learning across time. Whilst analysis at this level of detail was not possible for the human participants as a weighted average across training led to missing values, the AugSRN did not make errors and therefore all trials could be included in the analysis and the course of training could be analysed in greater detail. The model was trained with learning rates of 0.4 and 0.533 for slow and fast weights, respectively and 20 hidden units.

3.3.2. Results

3.3.2.1. Same rule learning

Those networks trained on the Same rule showed a large amount of learning as evidenced by the strong effect of Group: across training, $F(1,47) = 801$, $p < .001$, $MSE = .041$, $\eta_p^2 = .945$; and at test, $F(1,47) = 1133$, $p < .001$, $MSE = .003$, $\eta_p^2 = .960$, shown in black in Figure 3.3. Block had a main effect across training, $F(34,1598) = 19.9$, $p < .001$, $MSE = .007$, $\eta_p^2 = .297$, that interacted with Group, $F(34,1598) = 77.4$, $p < .001$, $MSE = .007$, $\eta_p^2 = .662$ suggesting that learning developed across training. There was also a Block effect at test, $F(4,188) = 142$, $p < .001$, $MSE = .001$, $\eta_p^2 = .751$ that interacted with Group, $F(4,188) = 119$, $p < .001$, $MSE = .001$, $\eta_p^2 = .718$, demonstrating a rapid extinction of learning in the Experimental networks across pseudorandom test trials.

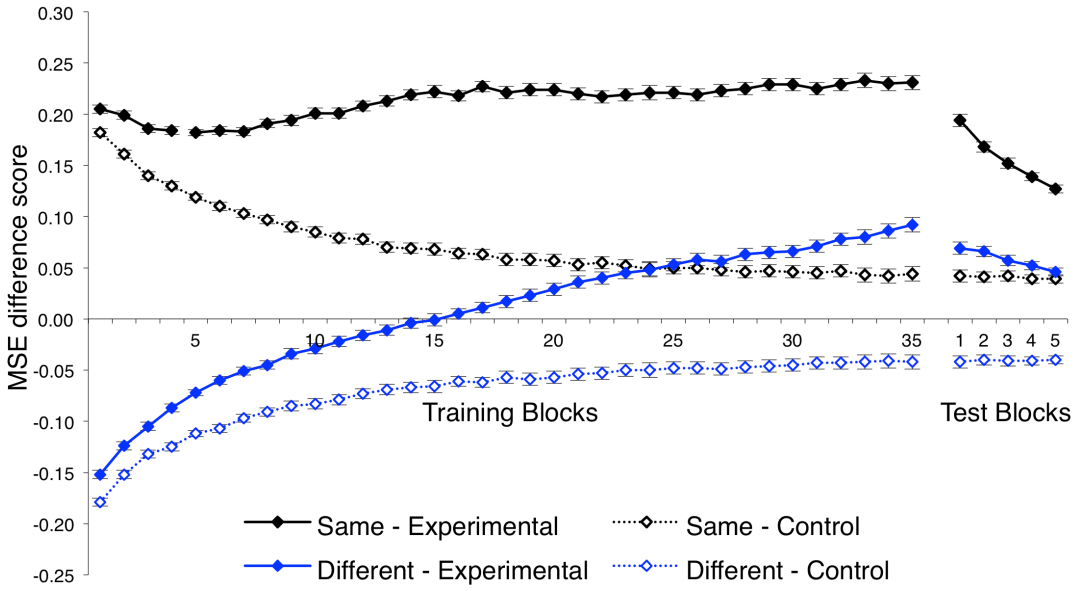

Figure 3.3. AugSRN simulation of Experiment 1 using parameters outlined by Jones and McLaren (2009) for both Same (black) and Different (blue) rules across training blocks and test blocks. Filled diamonds show Experimental networks, open diamonds Control networks. Error bars show standard error.

Subsequence had a large effect across training, $F(3,141) = 13151$, $p < .001$, $MSE = .009$, $\eta_p^2 = .996$; and test, $F(3,141) = 6068$, $p < .001$, $MSE = .001$, $\eta_p^2 = .992$, which is shown in Figure 3.4 in black. This suggests that the model is sensitive to sequential effects, as regardless of training, there was a large preference for responding to subsequences RRR and LLL compared to RRL and LLR over performance to RLR and LRL compared to

RLL and LRR. These sequential effects did not interact with Group at training, $F(3,141) = .510$, $p = .576$, $MSE = .009$, $\eta_p^2 = .011$, nor at test, $F(3,141) = .773$, $p = .501$, $MSE = .001$, $\eta_p^2 = .016$, suggesting that these subsequence effects were not affected by learning, and therefore that the models did not produce evidence of differential subsequence learning within the Same group.

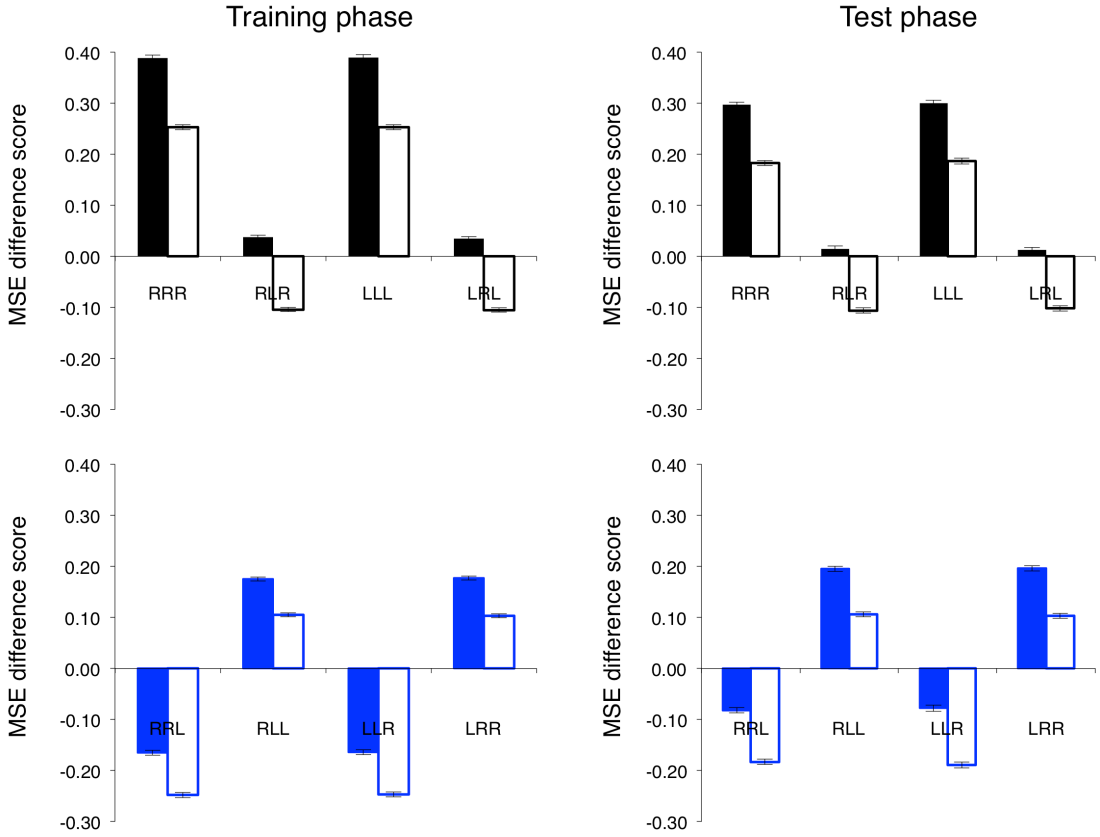

Figure 3.4. AugSRN simulation of MSE difference scores on the different subsequences trained in Experiment 1, using parameters outlined by Jones and McLaren (2009) for both Same (black) and Different (blue) rules across training (left panel) and test (right panel). Experimental networks are shown by filled bars and Control networks by open bars. Error bars show standard error.

Subsequence did interact with Block across training, $F(102,4794) = 118$, $p < .001$, $MSE = .008$, $\eta_p^2 = .715$, and this interaction did itself interact with Group, $F(102,4794) = 2.76$, $p < .001$, $MSE = .005$, $\eta_p^2 = .055$, as shown in Figure 3.5. These effects, while not observed in humans in Experiment 1, demonstrate the gradual learning of the subsequences RLR and LRL (collapsed into XYX) as well RRR and LLL (collapsed into XXX) that develops differently for both sets of networks.

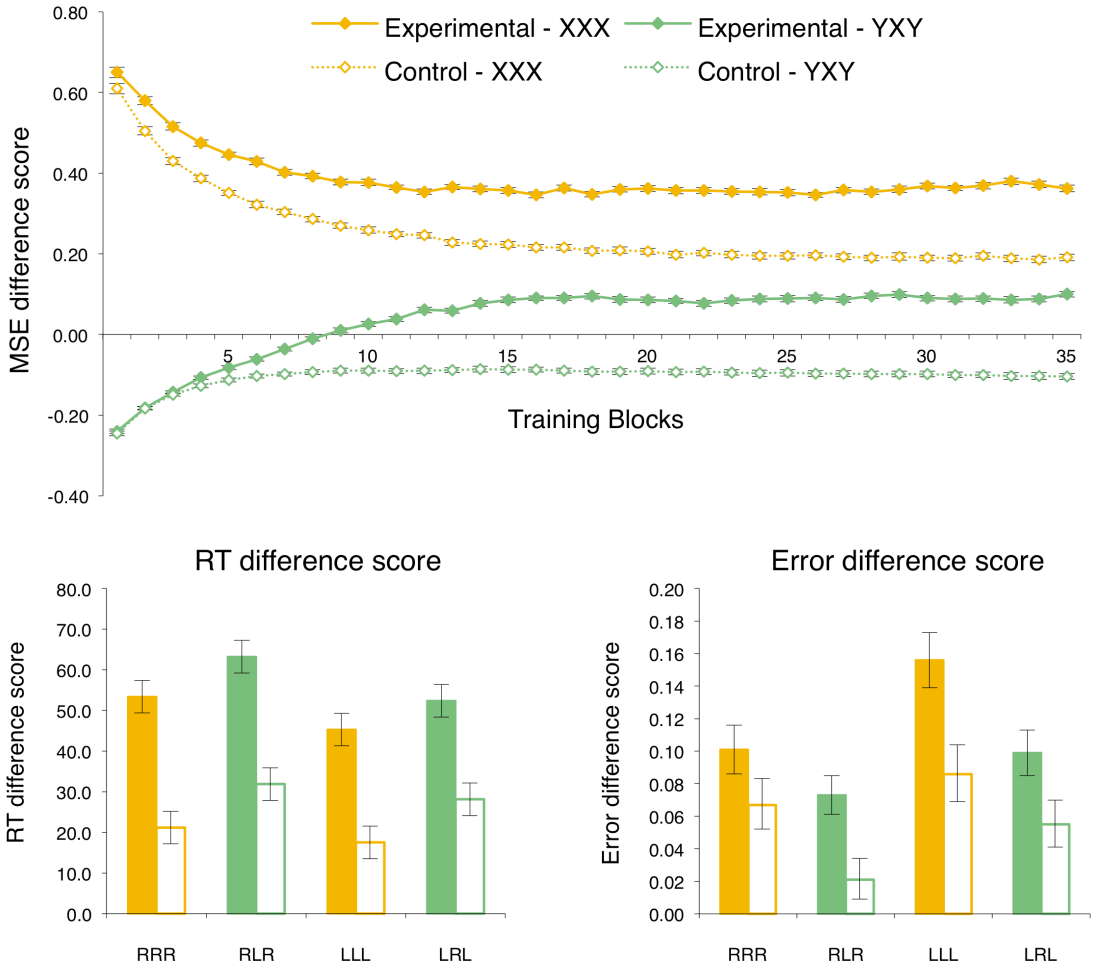

Figure 3.5. AugSRN simulation of MSE difference scores on the different subsequences trained in the Same group as in Experiment 1 (top panel), using parameters outlined by Jones and McLaren (2009) collapsed into XXX (RRR and LLL) and XYX (RLR and LRL). Experimental networks are shown by filled diamonds and Control networks by open diamonds. Bottom panel shows human performance from Experiment 1 across training for Experimental (filled bars) and Control (open bars) groups on all four Same rule subsequences in RT (bottom left panel) and error (bottom right panel) for comparison. Subsequences RRR and LLL (XXX) are shown in yellow bars and RLR and LRL (XYX) are shown in green bars corresponding to the model data. Error bars show standard error.

For the Same rule trained networks, as shown in Figure 3.5 the Control networks approach zero from either side of the *x* axis across training, demonstrating some learning of the *absence* of contingencies within the sequential structure that being to counteract the sequential effects. These reach asymptote, however, and still demonstrate the initial advantage for XXX (over XXY) and disadvantage for XYX (over XYY). These sequential effects do not follow the pattern observed in Experiment 1. Whilst training errors demonstrated an advantage for XXX over XYX, RTs showed the

opposite Sequence effect (although neither of these effects were significant); suggesting that there is no difference between the subsequences. However, the performance on both subsequences were well above a difference score of zero in Experiment 1 in both RT and errors (see Figure 3.5, bottom panel): therefore the AugSRN did not simulate the sequential effects observed in humans.

3.3.2.2. Different rule learning

The networks trained on the Different rule showed an effect of Group across training, $F(1,47) = 134$, $p < .001$, $MSE = .077$, $\eta_p^2 = .741$; and at test, $F(1,47) = 148$, $p < .001$, $MSE = .016$, $\eta_p^2 = .759$, providing evidence that the AugSRN could learn this sequential contingency, as shown in blue in Figure 3.3. There was a significant effect of Block across training, $F(34,1598) = 409$, $p < .001$, $MSE = .017$, $\eta_p^2 = .897$; and test, $F(4,188) = 8.08$, $p < .001$, $MSE = .001$, $\eta_p^2 = .147$, that interacted with Group across training, $F(34,1598) = 44.0$, $p < .001$, $MSE = .017$, $\eta_p^2 = .484$, which demonstrated a gradually increasing difference between Experimental and Control groups with training. There was also an interaction between Block and Group at test, $F(4,188) = 9.75$, $p < .001$, $MSE = .001$, $\eta_p^2 = .172$, which provides evidence of extinction when networks performed on pseudorandom blocks after training.

Subsequence had an effect across training, $F(3,141) = 12641$, $p < .001$, $MSE = .012$, $\eta_p^2 = .996$, and test, $F(3,141) = 3586$, $p < .001$, $MSE = .003$, $\eta_p^2 = .987$ (shown as blue bars in Figure 3.4). This was the inverse set of sequential effects observed for the Same networks, with better performance on subsequences RLL and LLR (compared to RLR and LRL) than on subsequence RRL and LLR (compared to RRR and LLL). This did not interact with Group across training, $F(3,141) = 3.25$, $p = .065$, $MSE = .012$, $\eta_p^2 = .065$, nor test, $F(3,141) = 3.02$, $p = .060$, $MSE = .003$, $\eta_p^2 = .060$, suggesting that the Subsequence effect was not itself affected by learning, nor was there evidence for the differential learning of subsequences within the Different rule. Subsequence did interact with Block across training, $F(102,4794) = 127$, $p < .001$, $MSE = .004$, $\eta_p^2 = .738$, but not test, $F(12,564) = .456$, $p = .787$, $MSE = .008$, $\eta_p^2 = .010$. This is shown in Figure 3.6, which demonstrates the opposite pattern in the Control networks as for the Same rule network (shown in Figure 3.5): the relative advantage and disadvantage for the subsequences XYY and YYX, respectively which converged towards a difference score of zero across training. Human participants in Experiment 1 showed Control

performance on the subsequences in the Different group to be uniformly below zero (see Figure 3.6, bottom panel), and therefore the XYY performance does not accurately simulate the sequential effects observed in humans.

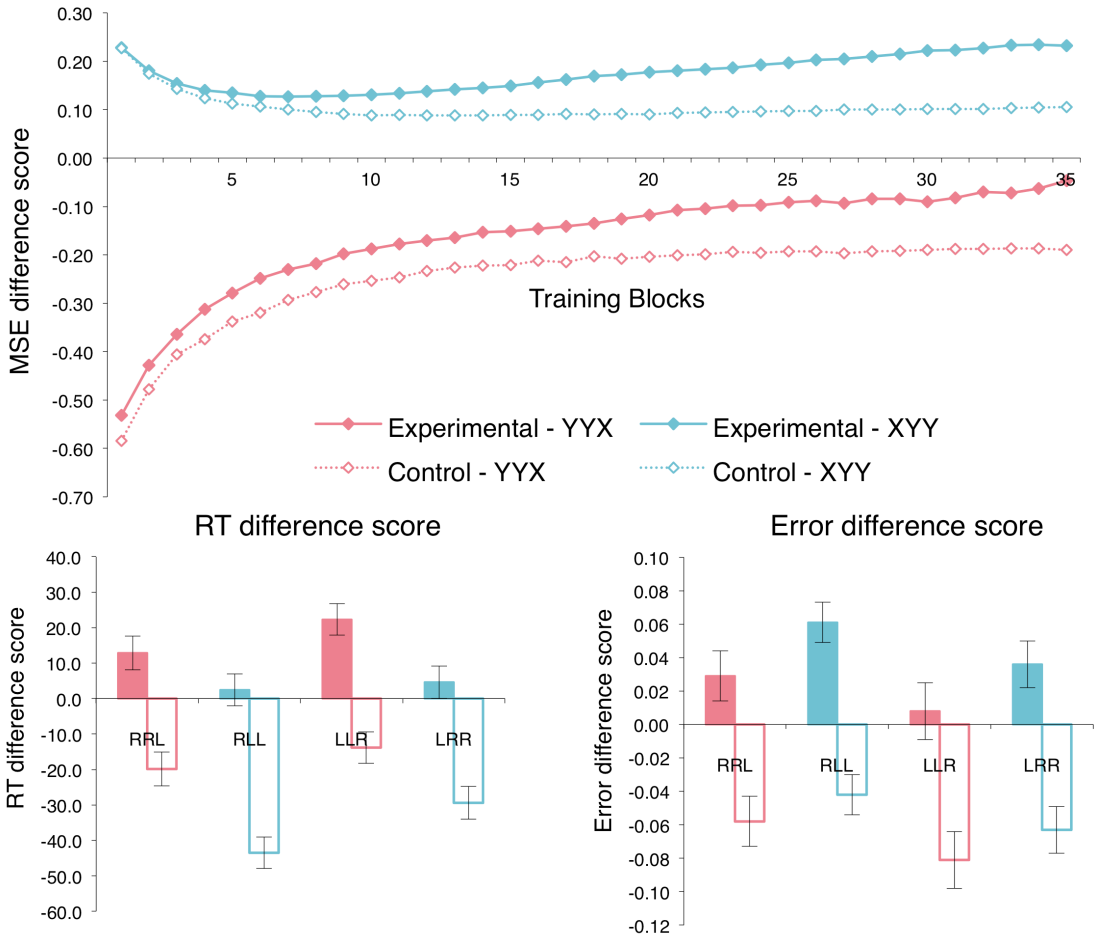

Figure 3.6. AugSRN simulation of MSE difference scores on the different subsequences trained in the Different group as in Experiment 1, using parameters outlined by Jones and McLaren (2009) collapsed into XXY (RRL and LLR) and XYY (RLL and LRR). Experimental networks are shown by filled diamonds and Control networks by open diamonds. Bottom panel shows human performance from Experiment 1 across training for Experimental (filled bars) and Control (open bars) groups on all four Different rule subsequences in RT (bottom left panel) and error (bottom right panel) for comparison. Subsequences RRL and LLR (XYY) are shown in pink bars and RLL and LRR (XYY) are shown in blue bars corresponding to the model data. Error bars show standard error.

3.3.2.3. Same versus Different learning

To assess the between-'subject' comparison of interest, the networks were again compared as the human participants were across training and test with a Bonferroni corrected ANOVA on MSE difference scores with Block, Group, Subsequence and

Rule as the factors of interest. Rule had a large effect across training, $F(1,94) = 1807$, $p < .001$, $MSE = .059$, $\eta_p^2 = .951$, and test, $F(1,94) = 416$, $p < .001$, $MSE = .009$, $\eta_p^2 = .816$; which captures the sequential effects observed in the task, with performance on Same subsequences overall better than for Different subsequences throughout the simulation (see Figure 3.3). This overall difference simulates human participants, regardless of the subsequence effects that did not follow human performance as discussed previously. There was an interaction between Rule and Group across training, $F(1,94) = 54.1$, $p < .001$, $MSE = .059$, $\eta_p^2 = .365$, although not at test, $F(1,94) = 3.70$, $p = .115$, $MSE = .009$, $\eta_p^2 = .038$. This reflects greater learning of the Same rule compared to the Different rule, the opposite sequence learning effect to the one observed in Experiment 1. The non-significant interaction at test also supports numerically greater learning of the Same rule, providing evidence that the simulation does not produce the learning effects observed in humans.

3.3.3. Discussion

The simulation of Experiment 1 produced evidence of learning of both Same and Different rules and of sequential effects, however, the essential learning difference in humans of better Different rule learning compared to the Same rule was not simulated by the AugSRN here. The AugSRN produces the same pattern of sequential effects in Control networks as observed in Jones & McLaren (2009); which as discussed in Chapter 2 was *not* the pattern of sequential effects observed in Experiment 1. Whereas the AugSRN and Jones and McLaren's (2009) participants demonstrated better performance on subsequences ending in a repeat, Experiment 1 found better performance to Same rule subsequences, regardless of whether they ended in a repeat or alternation; suggesting that a higher-order sequential effect was evident that is not captured by the AugSRN.

Running the parameters used by Jones and McLaren (2009), the model learned more about the Same rule than the Different rule, the opposite effect to the one obtained experimentally, providing evidence that the simulation was unable to account for human performance. This may be due to the learning parameters, which deviate from those suggested by Cleeremans and McClelland (1991) and whilst these may have been appropriate when simulating Jones and McLaren's (2009) six-session experiment, they may not appropriately simulate human performance in just two sessions. As the

networks were trained for the same amount of trials as the human participants, this difference between the tasks seems unlikely to merit a change in the rate of learning.

It could be that differences across motivation or participant samples produced a different amount of learning, however, as Jones and McLaren (2009) employed University of Cambridge students and paid them more money for taking part in the experiment, as they were given the opportunity to earn bonuses on top of being paid a minimum of £4 per hour. These differences, some may argue, should not have an effect on incidental learning, which is characterised as operating outside of intention, effortful control or attention (McLaren et al., 2014). However, some authors have suggested that decreased attention has an effect on incidental performance on a incidental learning task (Tanaka, Kiyokawa, Yamada, Dienes, Shigemasu, 2008), or that indeed measures of intelligence, correlate with learning rates (Tomas & Karmiloff-Smith, 2003) or even strength of associative learning (Kaufman, DeYoung, Gray, Brown, & Mackintosh, 2009). These issues, while highly debated, could of course still account for why the AugSRN simulates one set of human performance correctly, but is unable to do so for another set of human participants. To investigate this further, the learning parameters of the model were manipulated in further simulations.

3.4. Different versus Same rule learning

As suggested previously, the learning rates applied to the model are not fixed and can be easily changed. In order to attempt to simulate the human performance observed in Experiment 1, there were first changed to those suggested by Cleeremans and McClelland (1991) when developing the AugSRN.

3.4.1. Simulation 3: Cleeremans & McClelland (1991) parameters

3.4.1.1. Simulation details

Cleeremans and McClelland (1991) suggest smaller learning rate parameters of 0.15 and 0.2 for slow and fast weights, respectively, along with 15 hidden units. These changes were made to the model but all other parameters, model details and procedure remained the same as for Simulation 2.

3.4.1.2. Results

Analysed in the same way as for Experiment 1 and Simulation 2, learning is observed for both rules in the Group effect: Same rule training, $F(1,47) = 155$, $p < .001$, $MSE = .032$, $\eta_p^2 = .768$; and test, $F(1,47) = 135$, $p < .001$, $MSE = .010$, $\eta_p^2 = .742$; Different rule training, $F(1,47) = 82.2$, $p < .001$, $MSE = .016$, $\eta_p^2 = .636$, and test, $F(1,47) = 60.1$, $p < .001$, $MSE = .006$, $\eta_p^2 = .361$. It is clear that there is little difference made by these parameter changes, therefore there is no need to further analyse the detailed pattern of subsequence learning. A significant Bonferroni corrected interaction between Group and Rule across training, $F(1,94) = 24.8$, $p < .001$, $MSE = .024$, $\eta_p^2 = .209$, and test, $F(1,94) = 19.2$, $p < .001$, $MSE = .008$, $\eta_p^2 = .169$, demonstrated that learning was greater for the Same rule. Again providing evidence of the opposite learning effect to that observed in humans.

3.4.1.3. Discussion

The simple parameter change to lower learning rates and less hidden units slightly reduced the overall learning, but had no qualitative impact on the learning effect observed in Simulation 2: that the Same rule was learned better than the Different rule. This suggests that some reduction of attention or learning in Experiment 1 compared to Jones and McLaren's (2009) study does not explain the effects found in my research, and the AugSRN is still unable to simulate Experiment 1.

It could be that the response units in the AugSRN, which prime the previous response and have a substantial effect on the way that the model produces sequential effects, do not reflect human processes. As seen in both Same and Different groups, the model produces a strong preference for subsequences XXX and XYY compared to XXY and XYX, suggesting that the priming of the last response has a strong influence on performance right from the first Block of training. However, MSE difference scores for both Control and Experimental groups on XXX and XYY reduce over the first few Blocks, suggesting that the models are learning to ignore this bias. Indeed, this is observed in the first few blocks in simulations of Jones and McLaren (2009)'s study, seen in Figure 3.7. This overproduced advantage for XXX could be the cause of the Same group learning advantage seen in Simulations 2 and 3. Further to this, as Control networks in the Different group experience a benefit on the subsequence XYY due to this response priming, contrary to the subsequence effect (poor performance) observed

in humans, this may have reduced the amount of learning observed for this subsequence and therefore of the rule as a whole.

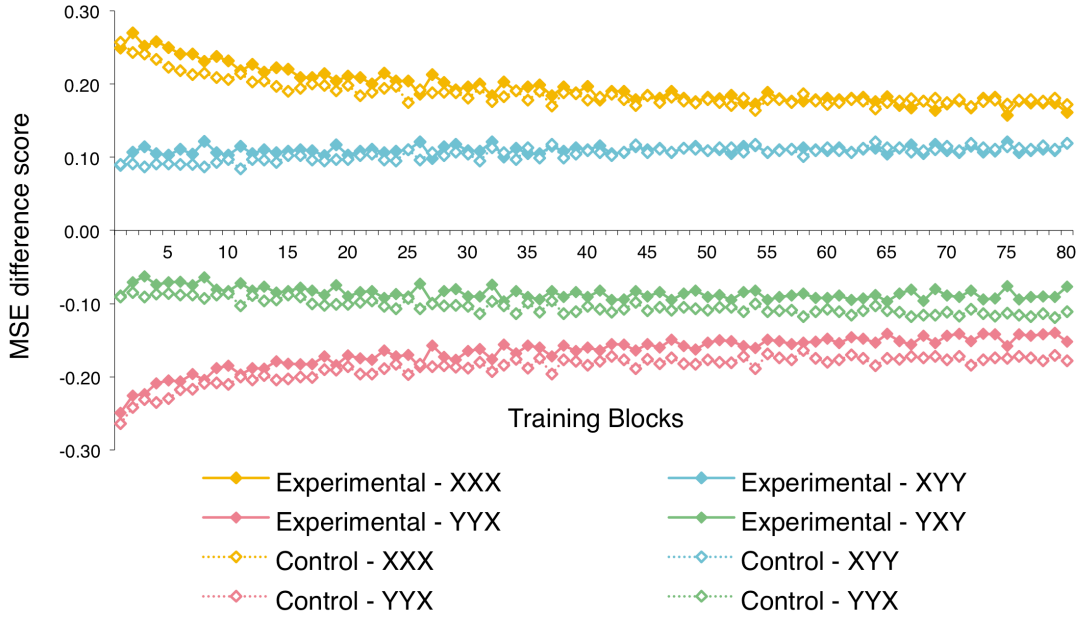

Figure 3.7. AugSRN simulation of MSE difference scores on the different subsequences trained in the Jones and McLaren (2009) exclusive-or task, using parameters outlined by Jones and McLaren (2009). Experimental networks are shown by filled diamonds and Control networks by open diamonds.

3.4.2. Simulation 4: AugSRN without response units

3.4.2.1. Simulation details and results

The response units were removed from the AugSRN, which essentially reduces the model architecture to that of the SRN (Elman, 1990); with input, hidden and output layers only. The models still differ, however, in the existence of two sets of connection weights, fast and slow, for the AugSRN. The SRN only contains one, stable set of weights equivalent to the AugSRN's slow weights that do not experience decay. The learning rate parameters and hidden units were returned to those outlined in Jones and McLaren (2009, slow learning rate: 0.4; fast learning rate: 0.533; hidden units: 20) and the simulations were run according to the procedure outlined for Simulation 2. The MSE was calculated not from the difference between response unit activations and expected responses (as there were no longer response units within the model), but the difference between output unit activations and responses, as outlined in section 3.2.1.2.

The model did not differ from the AugSRN in its preference for Same rule learning over Different rule learning across training and test.

3.4.3. Simulation 5: AugSRN without fast weights

As the removal of the response units made little difference to the pattern of results observed, it is perhaps logical to assume that the response units, which prime responding for a short term are not causing the disparity between human and model performance. The *learning* of transient trial orders, as facilitated by the fast weights in the AugSRN may therefore be the problem. As the response units have one-to-one connections with the output units, and are not involved in error correction or back propagation, they are simply a priming mechanism that gives precedence to the previous response made. The fast weights, with higher learning rates than slow weights, and a half-decay each time step were suggested to account for the short-term learning of contingencies between trials, therefore these do not apply necessarily to response repetitions alone and can produce a short-term influence of experiencing X followed by Y on the subsequent X trial (on which a Y is predicted more likely, Hinton & Plaut, 1987; McClelland & Rumelhart, 1985).

3.4.3.1. Simulation details and results

The AugSRN architecture was reinstated, as outlined in section 3.2.1.1, see Figure 3.1. The new model involved a set of response units, from which the MSE was taken and difference scores calculated. The parameters for learning rates and hidden units were as Simulations 2 and 4 (following Jones & McLaren, 2009) and the simulation procedure the same. The only difference between this model and Simulation 2 was that the fast weights were no longer included in the model. Learning therefore occurred only through one component (learning rate: 0.4) that experienced no decay. The results of Simulation 5 show again that the model, this time without fast weights, was unable to simulate the increased learning of the Different rule observed in Experiment 1. The model was consequently adapted further to investigate whether both the response units and fast weights may be responsible for producing sequential effects that influence Control and Experimental subsequence performance and reduce learning in the Different group.

3.4.4. Simulation 6: SRN

3.4.4.1. Simulation details

Simulation 6 used the SRN to simulate performance on the task, with the architecture depicted in Figure 3.9. The model had no response units and no fast weighted connections. The slow learning rate parameter was matched to Jones and McLaren (2009) and set as 0.4, with 20 hidden units. The MSE was again calculated from the output activation as in Simulation 4. The SRN also produced the same pattern of results as the AugSRN, regardless of whether the response units or fast weights were included in the model. Therefore, simulation of the between-subject difference between human participants on Different rules compared to Same rules was not easily modeled by any standard version of the SRN using the parameters outlined by Jones and McLaren (2009). At least within these learning rate parameters, the model was unable to produce these learning effects and this suggests that either these parameters were incorrect, or that the model has a more fundamental issue. To investigate the learning parameters of the model fully, before discounting the AugSRN as the best model of sequence learning, the parameter space was fully investigated.

There could, of course, be an issue with human participants, and there could be some difference between the human groups in Experiment 1 that the model does not simulate. For instance, one group may simply learn more due to increased attention or motivation as a result of some aspect of the sequences experienced themselves. This is entirely possible, as the Same group contained subsequences XXX and YXY, the experience of which was often anecdotally reported as containing a noticeably high number of strings of repeats and alternations. As participants find these subsequences easier to respond to, they may notice these subsequences explicitly (to some extent) and either: switch to such an explicit system; pay less attention; or be motivated to learn less. While, as mentioned previously, there is substantial debate about whether an incidental, automatic system may or may not be affected by explicit knowledge (Sun, Slusarz, & Terry, 2005) this is a possibility.

Revisiting the human data, we might suspect that if Experimental participants in Same and Different rule conditions were responding differently that there would be evidence of some quantitative difference in speed or accuracy with which participants completed the task, regardless of whether they were responding to consistent or inconsistent

stimuli. Bonferroni corrected comparisons comparing the Same and Different rule Experimental groups in their raw average RTs and proportion of errors to each subsequence (consistent and inconsistent) demonstrate no difference in either RTs across training, $F(1,46) = .005$, $p > .9$. Average RT [ms]: Same, $M = 286.0$, $SE = 4.37$; Different $M = 285.5$, $SE = 4.37$) nor proportion of errors, $F(1,46) = .002$, $p > .9$ (proportion of errors: Same, $M = .076$, $SE = .011$; Different, $M = .075$, $SE = .011$). Therefore, it seems unlikely that either group differs in their strategy, motivation or attention, which would influence responding in some way. It could also be the case that participants were simply not using an associative system to learn the task. The benefit for the Different rule group would be hard to reconcile with this view, as participants find it harder (if anything) to notice or verbalise this rule and therefore a Same rule group advantage might be expected.

3.5. Optimisation of the AugSRN

3.5.1. Optimisation Procedure

To investigate whether the AugSRN was indeed a suitable model of human learning, a search of the parameter space was necessary. Following such a search by hand, it was clear that a non-exhaustive trial-and-error procedure was not going to be successful in producing an AugSRN capable of simulating human learning on this task. The model contains a number of parameters that could be altered, with the absolute values of each as well as the interaction and proportional differences between them providing a multitude of possible conditions in which one could simulate Experiment 1. Therefore, an optimisation procedure was run, which attempted to find the parameters that produced MSE with the best fit to human data.

The MatLab FMINSEARCHBND function was employed in order to attempt to minimise the difference between human and model performance by altering free parameters in the model. These were the: number of hidden units; fast learning rate; slow learning rate; and constant k (fast weight decay). Fast and slow learning rates and constant k were bound between 0 and 1. The number of hidden units was bound between 1 and 9999. The optimisation procedure was given the target of human test performance; the values for errors were chosen as these demonstrated the larger difference. Same and Different Experimental and Control networks were trained in order to minimise the difference between model performance and human performance.

As MSE did not necessarily approximate human error rates, a constant was used for each set of 96 networks which was calculated for each set of parameters as the optimal constant that would transform model MSE values to those equating to human error rates. This constant was calculated by means of an optimisation procedure that attempted to reduce the difference between the MSE and human data. Multiplying the constant by the MSE values produced transformed model MSE for each of the four groups (Same Experimental; Same Control; Different Experimental; Different Control) which were able to be compared to human performance using a further mean squared error calculation (which I will call the Optimisation MSE). This target (the difference between transformed model MSE and human performance) was used to minimise the error in the optimisation procedure.

In order to avoid the fminsearchbnd procedure getting stuck at local minima, the model was run from a variety of starting points. These were selected randomly between the bounds of each parameter. Further to this all four possible model architectures simulated previously (AugSRN; AugSRN without response units; AugSRN without fast weights; SRN) were run. The fminsearchbnd procedure was run several times therefore for each model.

3.5.2. Results

Simulations were classified depending on whether they fulfilled the following criteria, which were identified as the key features of human performance: (1) the correct ordinal pattern of group performance (from highest to lowest: Same Experimental; Same Control; Different Experimental; Different Control); (2) larger learning of Different rule over Same rule; and (3) performance in both Same rule groups and Different Experimental networks with difference score greater than zero, i.e. these three groups all performed above a difference score of 0. A table of the number of models that fulfilled these criteria can be seen in Table 3.1, with 5837 simulations out of 75303 simulating all three correctly.

Table 3.1. Table showing number of simulations (of 96 networks each) that simulated the task as part of the optimisation procedure. Simulations were classified according to whether the pattern of results for each group followed: the correct serial order or not; performance of three of groups above a difference score of zero or not; and whether the Different or Same rule were learned better.

	Correct Serial Order		Incorrect Serial Order	
	Three groups > 0	Three groups not < 0	Three groups > 0	Three groups not < 0
Same > Different	16171	26028	3414	0
Same < Different	5837	13169	10414	0

The results from the human experiment at test were converted into differences for the purposes of visualising the data. Both the sequential effects and sequence learning were measured by such a difference. Firstly, the sequential effects were calculated as the difference between Different and Same rule in performance of Control groups only. Sequence learning was calculated first from the Experimental minus Control learning difference score for each rule, and then a further difference between Same and Different rule learning was taken as an index learning of the sequences. Both differences reflect Different over Same performance or learning, with these scores for human performance shown in Table 3.2, and show that while Different Control group performance is the inverse of Same rule Control effects, that the learning in the Different group is larger than the Same group.

Table 3.2. Results for Experiment 1, showing both RT and error difference score performance for Control groups only for Sequential effects and learning scores (the difference between Experimental and Control difference scores) for the final Epoch of training performance. Sequential effect scores are calculated by taking the difference between the Control group Different rule scores and Same rule scores. Sequence learning scores are taken from the difference between Same rule learning (Experimental minus Control difference scores) and Different rule learning scores (Experimental minus control difference scores).

	Sequential effects			Sequence learning		
	Difference scores			Learning scores		
	Same	Different	Different - Same	Same	Different	Different - Same
RTs	20.8	-18.7	-39.5	23.0	25.7	2.70
Errors	.056	-.059	-.115	.036	.073	.037

These 5937 models are plotted in Figure 3.8, with the Optimisation MSE (the difference between the transformed model MSE difference scores and human error difference scores) plotted on the *y* axis. This is a measure of how close the model is to predicting the human data. The *z* and *x* axes plot two key aspects of the human experimental results, the scores for both the sequence learning effect (greater learning of Different rule than Same rule sequences, *z* axis) and the sequential effects (greater performance on Same rule subsequences than Different rule subsequences, *y* axis).

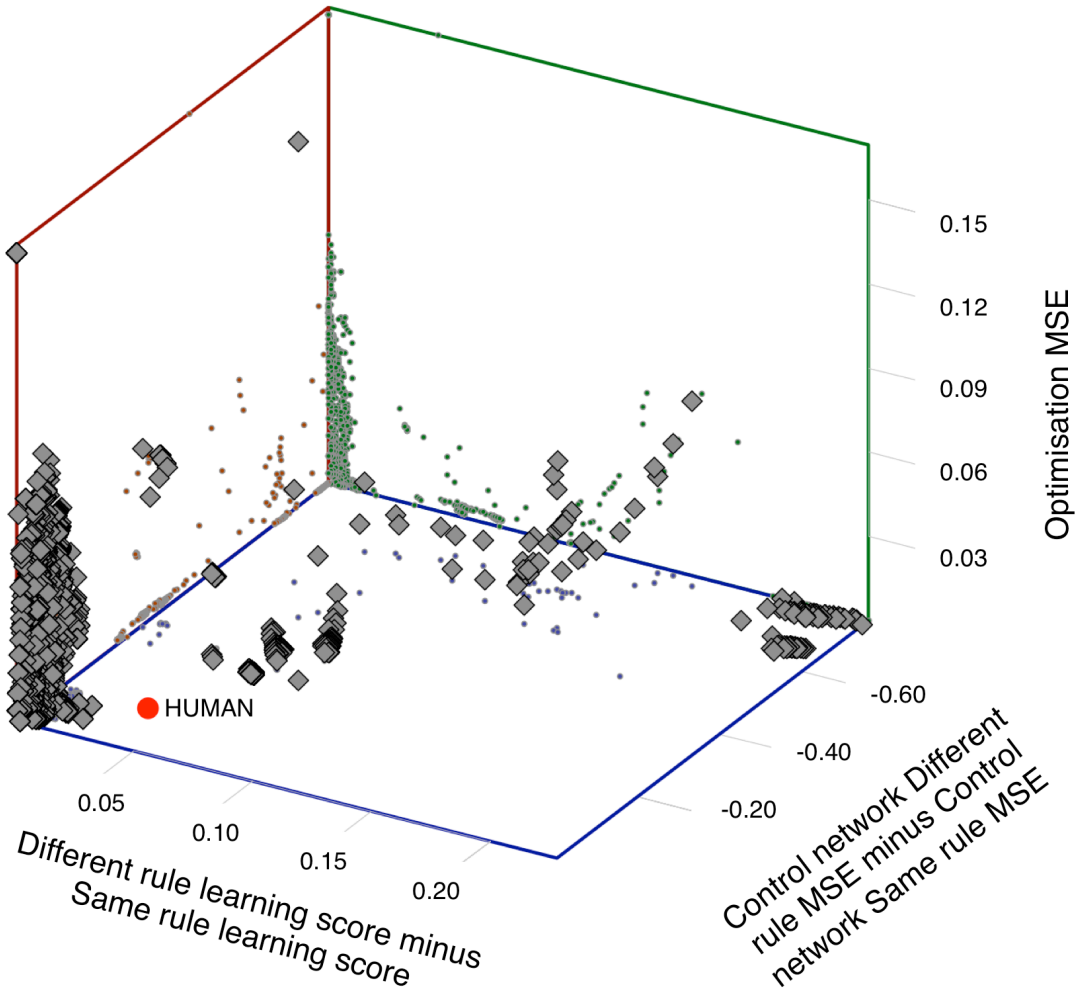

Figure 3.8. Plot of the 5937 models that fulfilled ordinal pattern of results observed in human participants. Mean squared difference between transformed MSE difference scores and human error difference scores is plotted on *y* axis; the sequential effects score (preference in responding to Different over Same rule subsequences) on the *z* axis; the sequence learning score (learning of Different over Same rule subsequences) is plotted on the *x* axis. Note, only the simulations that show greater Different over Same rule learning, as well as greater Same over Different sequential effects. Human performance is shown by the red circle.

As we can observe from Figure 3.8, no models approach the degree of sequence learning found in humans, with no ANOVA of these models producing a significant interaction between Sequence and Group. Therefore, whilst the AugSRN is able to simulate the order of results, the learning difference between Same and Different rules never reaches significance.

3.5.3. Discussion

The result of both a trial-and-error search of the parameter space and extensive optimisation of the free parameters of the model have led to no simulation that adequately captures human performance. Whilst the AugSRN is able to produce simulations that prefer Different rule learning over Same rule learning, these differences are small and not significant. Given the small amount of variance associated with these simulations this provides a convincing demonstration that the SRN or AugSRN is unable to produce the learning effects observed in humans. This suggests that the AugSRN in its current form is not a suitable model of human learning under incidental conditions.

3.6. Stimulus-response associations

3.6.1. Cue competition and subsequence learning

With the AugSRN unable to produce the pattern of responding observed in human participants, the next step was to ask what this model lacks that produced the advantage for Different over Same subsequences. Performance across training and test on errors (which produced a reliable Group by Rule interaction) in Experiment 1 demonstrated learning of subsequences in the following order, from greatest to least: XYY; YYX; YXY; XXX, although this difference was not significant. The explanation offered by Jones and McLaren (2009) for the poor learning of XXX – that trial-by-trial associations compete with sequential contingencies – goes some way to explain the poorer performance of the Same rule group; but this does not explain why YXY is learned less well than XYY and YYX. Similarly, if the transient learning of X -> X reduces the error term (and therefore the amount of learning that occurs) for the second X -> X in the triplet; then we should see the most learning for YYX, as the first mapping (Y -> Y) does not occur in the second instance within a triplet (Y -> X) and therefore the error term will be higher and more learning will occur about this instance.

This is not the case with human participants, as there are no differences between subsequence learning. Whilst the first set of trial-by-trial mappings within a triplet are not important in predicting the Same and Different rule subsequences as they were in the exclusive-or case, an eradication of the transient cue-competition effects proposed by Jones and McLaren (2009) would result in an absence of subsequence learning effects altogether.

It is possible that a vital component of the SRT task and how humans represent the task is missing from the SRN and simulations using such models. The SRN receives input regarding the current trial and produces a prediction, which is taken as an index of responding for the next trial in the sequence. Therefore, in a standard SRN only the current response stimuli or the response made (given that they should be the same thing) are used to make a prediction about the *next* trial, as shown in the top panel of Figure 3.9.

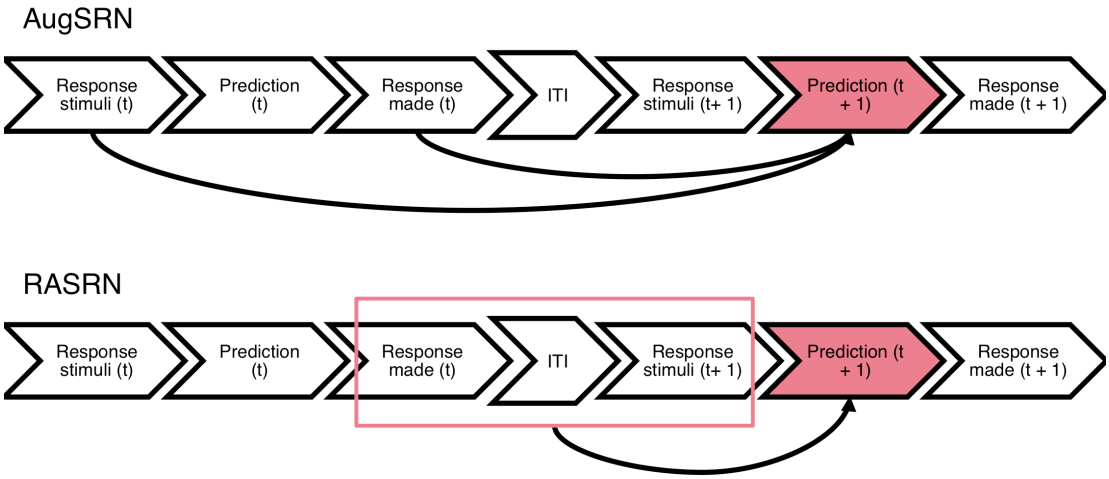

Figure 3.9. Elements in trial sequence of any SRT task at trial *t* and *t* + 1. In the AugSRN and SRN (top panel) either the response stimuli or response made at trial *t* are used as input to predict output (trained to the response made at *t* + 1). In the RASRN (bottom panel) all elements of the trial sequence are used, with the response made at *t* and the response stimuli at *t* + 1 used to predict the response at *t* + 1.

Consequently, the SRN ignores the current stimuli that are presented to humans in an SRT task between *t* and *t* +1 (Destrebecqz & Cleeremans, 2003). Traditionally, as the SRN attempts to model trial-by-trial contingencies and relationships across time, these simple stimulus-response (S-R) mappings on each time step of the model are not

represented, as they are surplus to the requirements of a sequence learning demonstration. Thus, in a standard SRT task, the stimulus that indicates the next response in the sequence is actually *not* represented in a typical SRN simulation that endeavours to predict the next response based on previous trials.

If both the previous response and current stimulus were used to predict the current response required, then these two inputs and their relationship with the previous trials could produce a cue competition effect that would account for human performance on the task. As simple associations form between each trial and their output, this might increase learning about XYY and YYX above the Same rule. That nature of the Different rule sequential contingency is that $t - 2$ is more likely to be in the opposite location to t. Therefore, the stimulus-response associations on $t - 2$ and t are not the same, and therefore learning about the relationship between $t - 2$ and t is not blocked. Learning about $t - 2$ in the Same rule, however, may be blocked by the representation of the stimulus-response association of $t - 2$ which occurs on t.

The model was, as a result, altered to include a better representation of the task given to humans, including input to represent the previous trial in the sequence (the Previous Response) as well as the current on-screen stimuli that participants were required to respond to, see Figure 3.9. To represent the ITI in between these two events, these inputs were given different activation values, with the previous response receiving a higher activation value to represent the increased time that participants had to process this information whilst making a prediction. The current stimuli are only on screen for a short time, with the occurrence of these stimuli prompting an immediate response. Therefore, whilst the current stimulus ($t + 1$) has a perfect relationship with the required response ($t + 1$); it has less time to accrue learning. The context units were given a higher activation value than both previous responses and current stimuli, as this internal representation of the task was in place before the previous response and current stimuli.

3.6.2. Simulation 7: RASRN

3.6.2.1. Simulation details

To better represent the task, the model architecture of the AugSRN was altered to that shown in Figure 3.10. The response units were removed, as these were found to make little difference to the subsequence effects above. As these units accounted for data that

involved both shorter RSIs (Cleeremans & McClelland, 1991) as well as when participants were motivated to respond extremely quickly (Jones & McLaren, 2009) and I observed no bias towards a first-order repeat preference in humans (Experiment 1, this thesis), these units were serving no useful purpose. Another modification was to add two new units, representing the current stimuli, as input with small activation values (of 0 for off and 0.1 for on) to represent the shorter time this input was available when making the prediction regarding the output (or response) to be predicted. The input in the AugSRN (now explicitly representing the previous correct response) was given an activation value of 0.75, with the context units given a higher value still (1.3 times the activation of the hidden units on the previous trial) to represent the time course of each trial. In all other instances the model remained unchanged and a simulation of the task was run with: 20 hidden units; a slow learning rate of 0.2; and fast learning rate of 0.5.

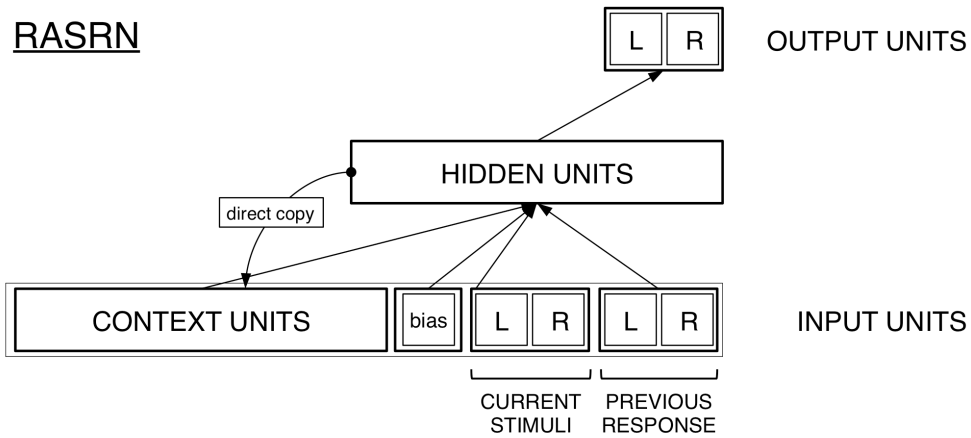

Figure 3.10. Structure of the Revised AugSRN (RASRN). Input units include representations of both left (L) and right (R) previous responses and L and R current on-screen stimuli as well as a copy of the hidden units on the previous trial and a bias (a unit that is always on). Activation flows in the direction of the arrows, with a set of hidden units passing activation forward once more to output and then response units, again representing L and R responses.

3.6.2.2. Results

The revised AugSRN (RASRN) results were analysed in the same way as the previous simulations, and showed a main effect of Group for both Same rule and Different rule networks, see Figure 3.11. Same rule learning: training, $F(1,46) = 755$, $p < .001$, $MSE = .058$, $\eta_p^2 = .943$; and test, $F(1,46) = 562$, $p < .001$, $MSE = .011$, $\eta_p^2 = .916$; Different rule learning: training, $F(1,46) = 220$, $p < .001$, $MSE = .119$, $\eta_p^2 = .827$; and test, $F(1,46) = 80.4$, $p < .001$, $MSE = .040$, $\eta_p^2 = .639$. When compared to one another, Group and

Sequence interacted across training, $F(1,92) = 12.9$, $p = .001$, $MSE = .088$, $\eta_p^2 = .123$; and test, $F(1,92) = .4.48$, $p = .037$, $MSE = .025$, $\eta_p^2 = .046$. The crucial test of the difference between the two groups showed that the Different rule was learned significantly better than the Same rule.

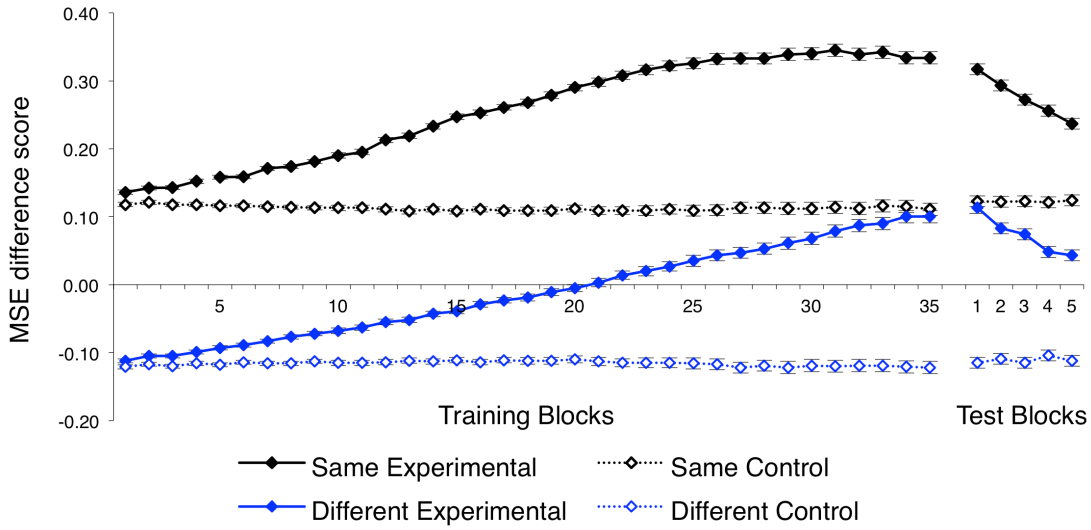

Figure 3.11. RASRN simulation (Simulation 7) of Experiment 1 using a slow learning rate of 0.2; a fast learning rate of 0.5 and 20 hidden units for both Same (black) and Different (blue) rules across training blocks and test blocks. Filled diamonds show Experimental networks, open diamonds Control networks. Error bars show standard error.

The sequential effects and subsequence learning observed in the RASRN was also far closer to that observed in Experiment 1, see Figure 3.12, with performance on XXX and YXY subsequences above zero and for YYX and XYY subsequences below zero (i.e. Same advantage over Different in Controls). The model also followed the ordinal pattern of sequential learning and effects seen in human errors across training and test, with better performance to RRR and LLL over RRL and LLR compared to RLR and LRL over RLL and LRR. Learning was numerically greatest for XYY; followed by YXX; YXY and the least learning was observed for XXX. The Subsequence effects were significant in both Same and Different groups across training: Same, $F(3,138) = 568$, $p < .001$, $MSE = .035$, $\eta_p^2 = .925$; Different, $F(3,138) = 815$, $p < .001$, $MSE = .035$, $\eta_p^2 = .947$; and test: Same, $F(3,138) = 12.6$, $p < .001$, $MSE = .005$, $\eta_p^2 = .215$; and Different; $F(3,138) = 32.1$, $p < .001$, $MSE = .002$, $\eta_p^2 = .411$. The Same group demonstrated an interaction between Subsequence and Group across training, $F(3,138) = 6.74$, $p = .005$, $MSE = .035$, $MSE = .$, $\eta_p^2 = .128$; but not at test, $F(3,138) = 1.07$, p

= .354, *MSE* = .002, η_p^2 = .023; providing some evidence for the differential learning of these subsequences. The Different group showed a significant interaction at test, $F(3,138) = 5.13$, $p = .010$, *MSE* = .005, η_p^2 = .100; but not training, $F(3,138) = 2.86$, $p = .078$, *MSE* = .024, η_p^2 = .059.

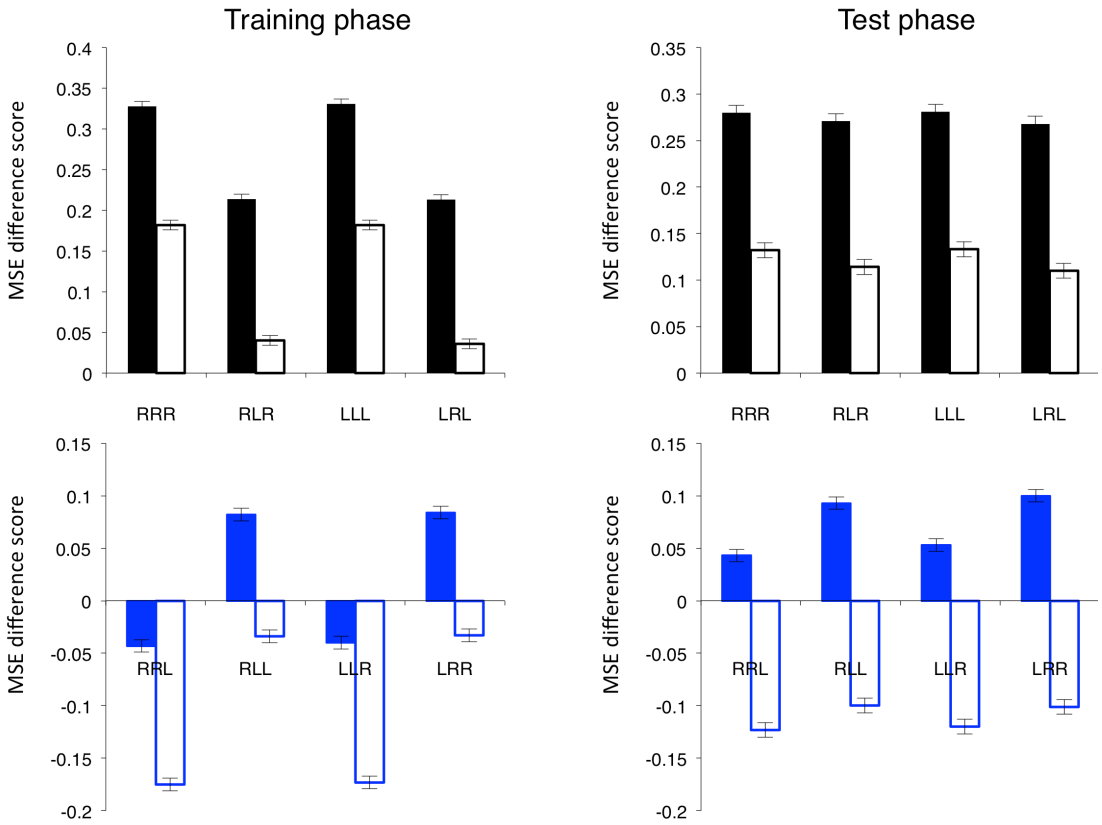

Figure 3.12. RASRN simulation (Simulation 7) of Experiment 1 showing performance of Experiment (filled bars) and Control (open bars) networks on all four subsequences in Same (black) and Different (blue) rules across training (left panel) and test blocks (right panel). Error bars show standard error.

3.6.2.3. Discussion

The RASRN, with the simple addition of units representing the SRT task, was able to simulate the differences that the AugSRN could not. By introducing the current stimuli in the RASRN, the model was able to experience the SR contingencies that participants were exposed to. Therefore, the conditions required to simulate human learning on this task combine competition between sequential contingencies; transient trial-by-trial contingencies, and stimulus-contingencies. Whilst SR contingencies were not previously instantiated in SRN and AugSRN simulations; as they are not required to

produce sequence learning, it seems they are required to accurately simulate human subsequence learning and effects.

There are issues associated with this model still however, specifically including the variability between and within each model. Whilst networks have hidden unit activations and weights reset to represent new participants, the error associated with each network is incredibly small, as well as the variance across blocks. Also the model does not capture the dip in performance seen in human participants between the two sessions. Whilst these issues remain (and are discussed in more detail in Chapter 7) the RASRN demonstrates cue competition between the SR contingencies and sequence learning. This produces less learning of the Same rule, and more learning of the Different rule in line with the sequence learning effects produced by humans. The sequential effects of the RASRN also follow those of humans, with the Same rule sequences showing higher MSE difference scores in control networks, compared to below zero MSE difference scores for Different rule control networks. Therefore, the model provides strong evidence that extending the associative account to encompass these realistic stimulus conditions provides the best explanation of the subsequence learning effects demonstrated in Experiment 1.

3.6.3. Simulation 8: RASRN simulation of Jones and McLaren (2009)

Further issues arose, however, when the RASRN was used to simulate the Jones and McLaren (2009) task as outlined in section 3.2.2. The results are shown alongside Jones and McLaren's in Figure 3.13, where it is clear that the model now demonstrates the sequential effects observed in Experiment 1. Hence, whilst learning of XXX does not occur, and learning of XYY and YXX is still apparent, the learning of YXY has disappeared and the model has reversed the sequential effects observed for XYY and YXY. It is a concern then that the model does not perfectly model both sets of data within these parameters. This may be due to the variability in human performance and as such, using only these parameters, it is not possible to say that the RASRN cannot reproduce the results of Jones and McLaren (2009).

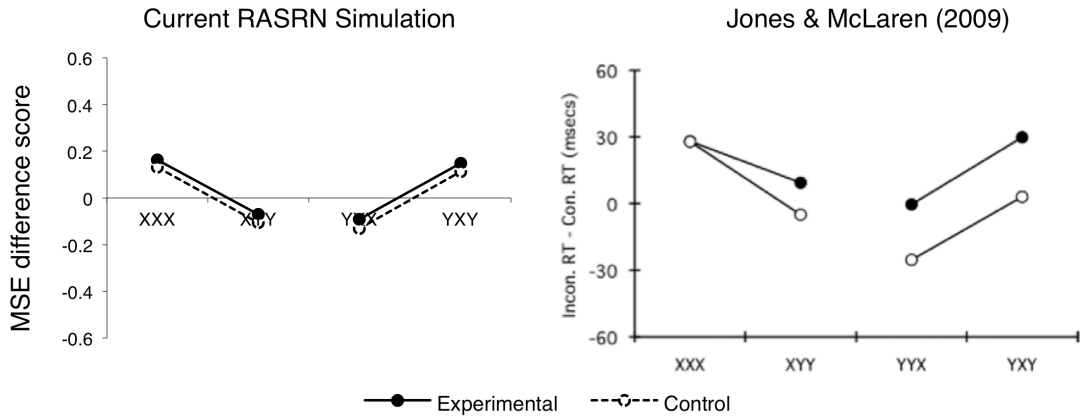

Figure 3.13. MSE difference scores for the post-training test phase of RASRN simulation 8 taken from the current work (left panel) and the human error data taken from Jones and McLaren (2009).

The success of the AugSRN in modeling Jones and McLaren's (2009) results on a similar task, and the RASRN's success in modeling the human data from Experiment 1 do not provide opposing theories about the development of learning under incidental condtions. Both tasks were run independently on different human samples, and whilst I do not dispute that the sequence learning effects were reliable, both remain a single study and only one demonstration of these effects. It would be ill advised to suggest the experimental or computational investigations thus far can reach a definitive conclusion about the exact mechanisms and parameters within which to simulate learning, and I suggest that further work using the sequences in both tasks is needed to establish better how a model could account for both sets of data.

The subsequence effects presented in Experiments 1 and 2 suggest that a far bigger preference may be given to responding to YXY than is indicated by Jones and McLaren (2009) and that participants are faster (but less accurate) on YYX over XYY. As discussed in Chapter 2, the discrepancy between the Control groups in these two studies is the issue here, as both were conducted under the same conditions with control groups trained in similar ways, therefore further studies are required to ascertain the nature of these sequential effects and determine whether the length of training or monetary reward-feedback involved in the Jones and McLaren (2009) study was the cause of these differences, and to ascertain the reliability of these sequential effects.

3.7. General Discussion

Chapter 3 has discussed the construction of the AugSRN and its successful simulation of the Jones and McLaren (2009) human sequence learning results. However, both the same parameters and a variety of alternate versions of the model were unable to simulate the results of Experiment 1, as the model could never produce a simulation that demonstrated significantly greater learning of the Different rule over the Same rule. The model was optimised using a bounded search of the parameter space, with random starting points and a number of possible model architectures used. Still no model could produce a simulation of the human between-subject result that the Same rule was not learned as well as the Different rule.

Considering the subsequence effects observed in Experiment 1 in more detail, the cue competition account offered by Jones and McLaren (2009) did not fit the pattern of results as there was some learning of the subsequence XXX, as well as an advantage for the YXY subsequence in Control groups. This account can, however, explain these results if the task is represented in its entirety. The SRN and AugSRN involve one trial predicting the next, and do not account for the stimulus-response (SR) relationships that occur in the SRT task between the stimulus on-screen and the response required and made. However, humans experience these contingencies and their presence in the task may come to block learning about $t - 2$ when $t - 2$ is the same (Same rule) as t, as these stimuli share the same stimulus-response association.

The RASRN, a version of the AugSRN adapted to include these current stimulus units, replicated the results of Experiment 1 and suggests that this account of the relationship between trial order and current on-screen SR contingencies can account for the incidental learning of humans on this SRT task. This suggests that humans are both sensitive to SR contingencies, and that these contingencies can compete and interact with sequential contingencies and trial-by-trial effects. That humans are sensitive to SR contingencies is, of course, not a novel proposal but the interaction between current on-screen stimuli and trial-by-trial effects and sequential contingencies has not, to the best of my knowledge, been considered nor simulated. Traditional models of associative learning consider SR or stimulus-outcome links on a trial by trial basis, irrespective of any sequential effects or influence of serial trial order. Models of sequence learning do not require SR associations to learn sequences, and therefore do not include them in

general, which does not accurately represent the task conditions (Destrebecqz & Cleeremans, 2003) nor does this account for the possible associations that may occur.

Whilst other models have included representations of the current trial stimuli (Cleeremans, 1993; Destrebecqz & Cleeremans, 2003), these have not included simply providing the model with these inputs (both t and $t - 1$ when predicting t). In these cases, t was introduced to the model as part of a separate, non-recurrent learning system. This poses a fundamental question about how we believe human learning occurs: is it always recurrent, or only sometimes recurrent? It might seem obvious that stimulus-response associations do not require recurrence in order to be learned, but to take this position would be to suggest that certain stimuli are treated differently by our learning system. If we simply presuppose that humans immediately deploy recurrence to stimuli that follow a sequence, but do not do these for other stimuli (e.g. Cleeremans, 1997) this suggests that participants can intrinsically recognise sequential stimuli and have volitional control over whether recurrence occurs. This seems highly unlikely, and therefore there must be some reason to propose that sequential information is learned by a separate, recurrent system over and above the fact that it involves sequential contingencies. Beyond giving a human or model simply one trial, all stimuli, whether they have sequential contingencies or not are presented in some order one after the other.

Separating sequential and non-sequential stimuli in terms of recurrence could be explained if recurrence is taken to represent some characteristic of the sequence that is fed back to a participant that is not true of other stimuli. This may be explained then by some motor response element of sequence learning, as whilst stimuli are responded to, the recurrent loop may represent feedback about the sequence of motor responses made. This may reflect the motor cortex loop with the basal ganglia (Middleton & Strick, 2000), that has been implicated in impaired sequence learning when damaged (Siegert et al., 2006). However, responses alone are not necessary for sequence learning to occur, and participants are able to learn sequences of stimuli that involve no responding (Dennis, Howard, & Howard, 2006). It is generally accepted that the motor component of SRT task learning is dominant (e.g. Bischoff-Grethe, Goedert, Willingham, & Grafton, 2004; Abrahamse et al., 2012). The simulation results of this chapter and the evidence provided by Experiment 1 support Willingham's (1997) suggestion: that both motor and perceptual elements are important.

If recursion were only applicable to motor actions then the application of recurrent models in the study of language (e.g. Elman, 1990; Dienes, 1992) would be compromised. This seems unlikely as language itself is highly recursive in structure (Rohmeir, Dienes, Gao, Fu, 2014). There may be some other *a priori* reason for assuming that only certain, sequential stimuli are learned by a recurrent network; but I suggest that requiring multiple learning systems that separate stimulus-response and trial-by-trial learning is not parsimonious. Further to this, the results of Experiment 1 suggest that learning about these stimulus-response contingencies interacts with sequence learning in a way that provides strong evidence for these stimuli being processed by the same system.

Chapter 4. Concurrent stimuli and sequence learning: Testing a prediction of the RASRN

In this chapter I examine the influence of the two sets of input now included in the RASRN, and how the RASRN predicts that human sequence learning will be affected by variations in these inputs. Chapter 3 demonstrated that both previous responses and current on-screen stimuli were found to be required to produce a simulation of human learning in Chapter 2, and therefore the RASRN was used to predict how human learning would progress when the influence of current stimuli or previous responses were increased. Two new units with activation values that matched the current stimulus units were activated in the same sequential pattern as the current element in the sequence being predicted (Current), the previous element in the sequence (Previous), or were simply random (Random). The sequential contingencies used were those of the Same group in Experiment 1. The RASRN showed that sequence learning was, if anything reduced in the Current condition relative to the Previous condition, suggesting that the increased influence of stimulus-response associations interfered with learning of the Same rule sequential structure. A behavioural experiment was conducted to test this prediction using a modified version of Experiment 1, with all participants performing on Same rule sequences to the same two-choice SRT task. The manipulation of interest was that two new stimuli were introduced: either a high or low tone played concurrently with the appearance of the on-screen response stimuli (Experiment 3); or yellow or purple circle fills within the response stimuli instead of the white stimuli used in Experiment 1 (Experiment 4). In both experiments I found that participants learned more about the Same rule in the Previous condition over the Current condition, following the prediction of the RASRN. Human learning in the Random new stimuli condition, however, did not match up with model predictions, nor did the raw performance scores on consistent and inconsistent subsequences, suggesting that whilst additional stimuli do indeed influence sequence learning, this may not be fully captured by the RASRN. The results of Experiment 3 and 4 suggest that an additional cue may potentiate sequence learning of the Same rule if it provides information about the previous element of the sequence.

4.1. Introduction

The RASRN in Chapter 3 showed evidence of being able to simulate human learning under incidental conditions. The SRN, AugSRN and the parameter space searched between these models were unable to account for the advantage for learning of Different rule sequential contingencies over Same rule contingencies. The AugSRN was able to simulate Jones and McLaren's (2009) data, but this sequential rule relied on contingencies where both the $t - 2$ and $t - 1$ were involved in predicting t. The RASRN represents both the between trial contingencies that participants learn, as well as the stimulus-response contingencies present on every trial. Instead of using the SRN to train the current stimulus predicting the next response, the RASRN was altered so that a representation of the previous (correct) response made as well as the stimuli on screen are both involved in predicting the outcome of a given trial.

The AugSRN does not include a representation of the current stimuli when predicting that trial, which suggests that regardless of the current stimuli that are on-screen or presented to participants; a prediction based on the previous trial will have been made. The RASRN, however, includes such a representation of stimuli and therefore stimulus-response contingencies can develop. Even though these are not required to learn sequences, their influence on sequence learning may be important. In representing both on-screen stimuli and the previous response within the model this changed sequence learning in favour of Different rule sequences. As a result, we have some evidence that by introducing the influence of these two inputs – the current and previous trial – this leads to differences in how well sequences were learned. The RASRN can be used to generate a prediction about the influence of these units on human learning and in doing so guide the design of experiments with which to test its suitability as a model of sequence learning and provide support for the associative account of these processes in humans.

To investigate this further, I wanted to examine how increasing the influence of either the current stimuli (t) or the previous response ($t - 1$) might influence the learning of sequences and subsequences, as well as sequential effects in humans. Following dual-stimulus versions of the SRT sequence learning task (Abrahamse et al., 2012; Robertson & Pascual-Leone, 2001) the model was altered to contain two further input units to represent two hypothetical new stimuli, which were given the same activation value as the current stimulus units with the idea that these would occur at the same time. These

inputs (or stimuli) could then be given corresponding activations to either the current stimulus (Current) or the previous response (Previous) in a sequence, as a between-network (or between-subject) comparison. A random new stimulus condition was also simulated. Therefore, these new stimuli could be congruent (Current) or random (Previous, Random) with respect to the current response, as well as following the trained sequence (Current, Previous) or not (Random).

Concurrent stimuli in the Current group that have a perfect relationship with another stimulus are defined by Abrahamse et al. (2012) as redundant, as there is already a cue (location) that can be used to learn about sequences. Robertson and Pascual-Leone (2001) were perhaps the first authors to present such a redundant dual-stimulus sequence learning task and provided evidence that participants learned more when two correlated and concurrently presented stimuli (locations each with a specific colour) followed a sequence than when just locations or colours were presented. Abrahamse et al. (2012) suggested that Robertson and Pascual-Leone's (2001) result was a consequence of using different sequences across dual and single stimulus conditions. In a replication they found that no improvement in performance occurs in the dual-stimulus group and therefore suggest that these concurrent stimuli are not learned about and have no influence on the task (Abrahamse et al., 2012).

The previous two chapters in this thesis predict instead that stimulus-response associations will (differentially) affect sequence learning, even if they are 'redundant'. As human learning was better simulated by the RASRN when the current stimuli experienced on the trial were introduced, in Chapter 3 I discussed the possibility that stimulus-response associations interfered with learning of the Same rule to a greater extent because of the shared stimulus-response mappings between $t - 2$ and t (the trials that determine the probabilistic sequential structure). In the dual-stimulus task presented by Abrahamse et al. (2012), which used a 12 element second order conditional (SOC) sequence, which involved no first order repeats and only one mapping (2-1-2) that involved two of the same stimuli. It is possible, therefore that stimulus-response associations may not have been able to interfere with such a structure.

Consequently, in this chapter the effect of additional stimuli on sequence learning will be put under greater scrutiny, with the detailed pattern of stimulus-response and sequence learning investigated through the use of subsequences and a simple

probabilistic two-choice structure as used in the previous two chapters in this thesis. Based on the changes to the AugSRN discussed in Chapter 3 as a result of the human pattern of responding found in Chapter 2, the role of current and previous trial information in sequence learning is compared. I intend to use the predictions of the RASRN regarding stimulus-response associations to design and then run a sequence learning SRT task with human participants that involve some additional stimuli with the aim of testing the suitability of the RASRN as an associative explanation of human learning.

4.2. Simulation 9: A prediction of the RASRN

Before conducting any human experimental work, the RASRN was used to simulate possible human performance in order to produce an a priori prediction regarding the effect of additional stimuli that followed the current stimulus (*t*) or previous response (*t* – 1). As mentioned previously, with the aim of matching a possible human experiment, the activation of the two new stimulus units would match the current stimulus activation values (0 and 0.1) and these stimuli would therefore be presented concurrently. Three network simulations were run: one where current stimulus units and new stimulus units were entirely matched (the Current New Stimulus condition), and so each time the left current stimulus unit was activated, so was the corresponding one of the new stimulus units. The second condition (Previous condition) involved the matched activation of each new stimulus unit and the previous response units (although the new stimulus units were activated at 0 and 0.1 rather than 0 and 0.75). Therefore, if the left previous response was activated, so was the corresponding new stimulus unit. A further set of networks were trained on randomly activated new stimuli, with each unit activated an equal amount of times but with no relationship to either current stimuli or previous response activations. Whilst I ran both Different and Same rules on Experimental and Control networks, I will concentrate on the Same rule Experimental condition as this is where we expected (and found) stimulus-response associations to have the biggest effect.

4.2.1. Simulation details

4.2.1.1. Model construction

The model (see Figure 4.1) comprised 2 input units that represented the required response (left or right) on the previous trial (*t* - 1) in a two-choice SRT task which are

shown on the bottom right of Figure 4.1. Two further input units represented the two on-screen response stimuli (left or right) that participants were expected to respond to on the current trial (*t*). The model differed from that described in Chapter 3 with the inclusion of 2 further input units, which represented the new stimuli that I added to the task, representing the new stimuli in this task. These are pictured in Figure 4.1 as a high or low tone (Experiment 3) or two different colours (Experiment 4). A further 20 input units acted as context units and a single bias unit that was always activated to 1 also made up the input layer. 20 hidden units and 2 output units (to represent the response required to trial *t*, left or right) made up the next two layers of the model.

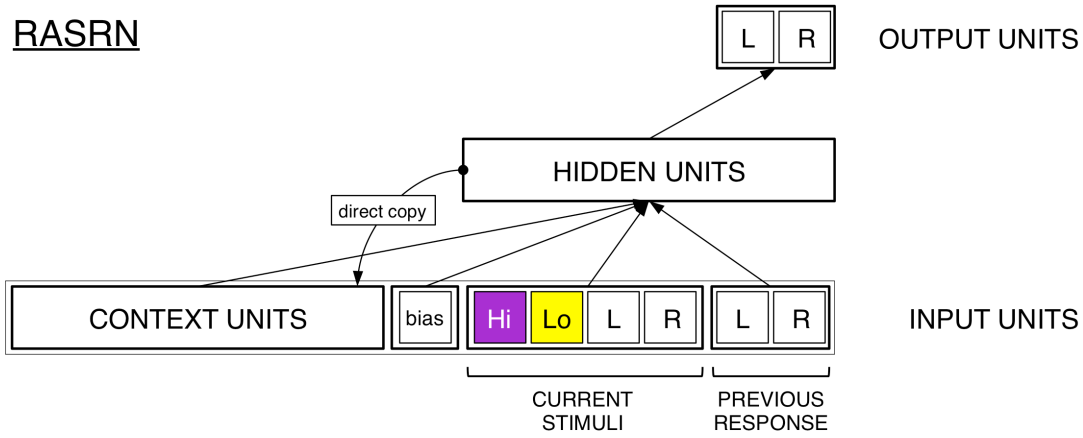

Figure 4.1. RASRN architecture for the simulation of Experiments 3 and 4. The model has input that comprises of the previous response made, as well as the current on-screen stimuli in the task. This involves a left or right on-screen response stimulus as well as either a high (Hi) or low (Lo) tone (Experiment 3) or the colour of the response stimulus itself was either purple or yellow (Experiment 4). Context units and a bias unit make up the remainder of the input, from which activation is fed forwards to a hidden layer of units and again to output units.

All of the simulation details were as described for the RASRN in Simulations 7 and 8 in Chapter 3. Learning rates were 0.2 and 0.5 for slow and fast weights, respectively. The bias unit was always activated to 1, whilst the other input units had different activation values. These were set according to the length of time the inputs have been available, to add some crude approximation to the time course of each trial. Therefore, having just occurred before a response is required, the current stimuli (including both the on-screen response stimuli and the new stimuli) were given an activation of 0 (off) or 0.1 (on). The previous response units were given a higher activation value (0.75 for on, 0 for off) as the input from this response has been around for longer. The context units, which are a copy of the last trial's hidden units, represent the internal representation of the task,

and were updated before the occurrence of a response on $t - 1$, and therefore have the highest activation value, 1.3 times the hidden unit values on the previous trial and 0 for off. The weighted connections between these units and the hidden layer involve activation and error correction as described in Chapter 3 (see section 3.2.1).

4.2.1.2. Design

Following the experiments described in Chapter 2, the training was to last 35 blocks of 120 trials, with 5 blocks of test: totaling 4200 training trials and 600 test trials. Training blocks comprised of Same rule subsequences, constructed in the same way as described in Chapter 2 (see section 2.2.1.3) with the same contingency (two out of three trials followed Same rule across training). Test blocks were made up of pseudorandom trial order.

New stimuli. The key manipulation in this study was in the introduction of two new stimuli, which in the model were represented simply as two new 'current stimulus' input units. The relationship that the stimuli had to the current trial (*t*) differed, according to the group networks were assigned to: Current, Previous or Random (see Table 4.1 for the design). One of the two new stimuli would occur on each trial and each stimulus overall occurred equally within and across all blocks, regardless of group. In the Random group either stimulus was likely on either L or R trials, and therefore there was no relationship between the new stimuli and the response stimuli locations throughout training and at test.

Table 4.1. Design of Simulation 9, with 24 networks run in each of the three conditions.

	New Stimulus	Contingency with response stimulus at:	
		t	*t* - 1
Current	$= Input_{RVL}(t)$	100%	50%
Previous	$= Input_{RVL}(t-1)$	50%	100%
Random	random	50%	50%

The Current group experienced a new stimulus that had a 100% contingency with the current trial: for example, on every R trial new stimulus A occurred, and on every L trial the new stimulus B occurred. This relationship occurred throughout training and test. The Previous group new stimulus had a 100% contingency with the *previous* response stimulus, so if $t - 1$ had been an R trial, trial *t* might be an R or L trial but new

stimulus A would always occur. Similarly new stimulus B would always occur at trial *t* when *t* – 1 was an L trial. In the Previous group case the new stimulus had no relationship whatsoever with the current trial, as A and B occur with equal likelihood on an R or L trial. Therefore, this matches the Random group somewhat in terms of what participants experience on each trial, as this Previous group offers no immediate information to participants about the required response.

The previous response (*t* – 1) is itself not required in the Same rule (predicted from *t* – 2), and it might be predicted to have no positive influence over learning as it simply adds no additional direct information about the current trial. Whilst the Current group might seem intuitively like the new stimuli and offer additional salience or activation to the sequence of trials experienced, the 100% contingency between new stimuli and current response-stimuli could encourage simply stimulus-response learning and may go some way to block learning about sequences, as shown in Chapter 3.

4.2.1.3. Simulation procedure

The RASRN was run over 72 networks, each one with hidden unit activations reset and weights between units randomized between -.5 and .5, to represent the participants that I intended to run on Experiment 3. All of these networks received training on Same rule sequences and pseudorandom blocks at test. I also ran a further 72 Control networks that received pseudorandom training and test blocks which are presented in Figure 4.2 but not analysed here. In all regards, but for the new stimulus units, the simulation followed the RASRN details as in Chapter 3 for the Experimental Group.

4.2.2. Results

4.2.2.1. Same rule learning

MSEs were taken as an index of RT performance on *t* + 1 (see Chapter 3, section 3.2.1) and a difference score calculated, taking MSE on subsequences consistent with the Same rule away from inconsistent MSE. An ANOVA was conducted on the Experimental data only with Block, New Stimulus and Subsequence as factors across training and test separately. Control network results were not analysed here as participants in the following experiments were all trained under Experimental conditions and this analysis attempts to provide the basis for the human experimental design. They are shown in Figure 4.2 for a visual comparison only to demonstrate that the following New Stimulus effects discussed do have an effect on sequence learning

and not just performance. Whilst the Experimental networks show similar performance across the first twenty blocks of the experiment (see Figure 4.2), after this a difference emerges between the groups (main effect of New Stimulus) that is significant across both training, $F(2,69) = 13.6$, $p < .001$, $MSE = .0061$ $\eta_p^2 = .282$, and at test, $F(2,69) = 353$, $p < .001$, $MSE = .017$, $\eta_p^2 = .911$.

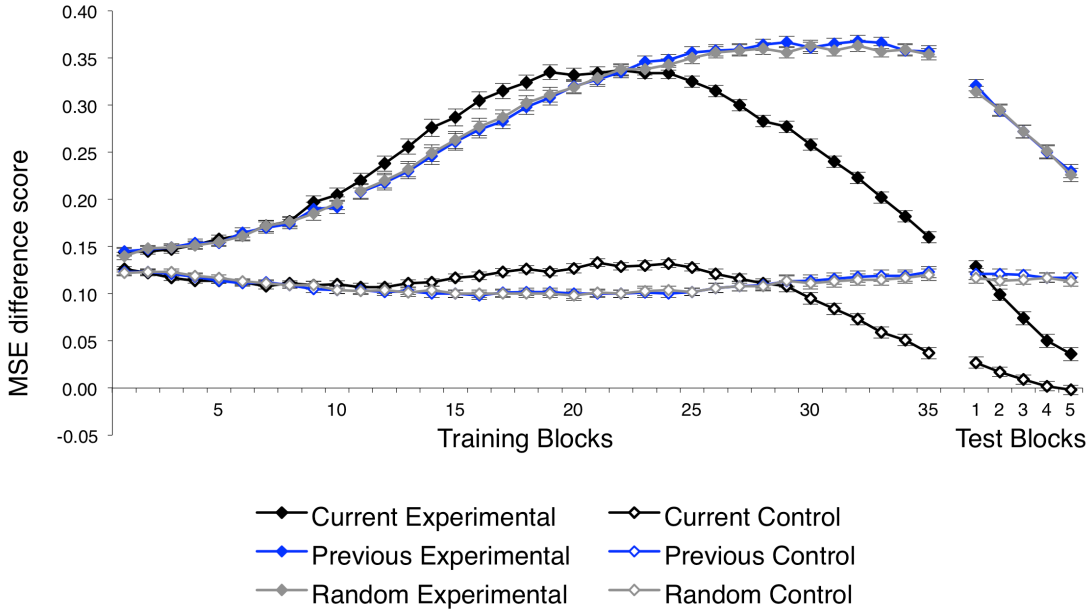

Figure 4.2. MSE difference scores for RASRN simulation of new stimulus task for Experimental (filled diamonds) and Control (open diamonds) networks with Current (black), Previous (blue) and Random (grey) new stimuli conditions across training and test blocks. All networks were trained on Same rule sequences. Note that Control networks are not analysed in this section and are provided as visual data here for comparison. Error bars show standard error.

Analysed further through a series of Bonferroni corrected pairwise comparisons, the Current group perform significantly worse than the Random group across training, $F(1,69) = 19.5$, $p < .001$, $MSE = .001$, $\eta_p^2 = .221$; and test, $F(1,69) = 524$, $p < .001$, $MSE = .001$, $\eta_p^2 = .884$. The Current group also perform worse than the Previous group across training, $F(1,69) = 21.1$, $p < .001$, $MSE = .001$, $\eta_p^2 = .234$; and significantly worse at test, $F(1,69) = 534$, $p < .001$, $MSE = .001$, $\eta_p^2 = .886$.

The RASRN provides evidence that the Current group learns these sequences the worst, with Previous and Random sequence learning at roughly the same level. Before exploring this result further, we can see that the networks showed a large effect of block across training, $F(34,2346) = 673$, $p < .001$, $MSE = .022$, $\eta_p^2 = .907$, which reflects the

learning about the stimulus-response contingencies. In all three conditions the RASRN is able to learn the perfectly predictive relationship between stimulus and response, and this begins to overshadow sequence learning as MSE difference scores trend towards zero. The block effect further interacts with New Stimulus, $F(68,2346) = 58.1$, $p < .001$, $MSE = .022$, $\eta_p^2 = .628$. This demonstrates the emergence of the New Stimulus difference across training, which was stable at test, as while there is an overall Block effect, $F(4,276) = 512$, $p < .001$, $MSE = .001$, $\eta_p^2 = .881$, this does not interact with New Stimulus, $F(8,276) = 1.38$, $p = .231$, $MSE = .001$, $\eta_p^2 = .038$.

To further investigate these relationships the raw MSE for consistent and inconsistent subsequences are presented in Figure 4.3. These show evidence of a trend towards better responding regardless of whether they are consistent or inconsistent, providing evidence for the role of stimulus-response learning occurring throughout the simulation. As we can see, as the Current New Stimulus trained networks experience a higher stimulus-response contingency: this results in greater stimulus-response learning compared to the other two groups and consequently less sequence learning overall.

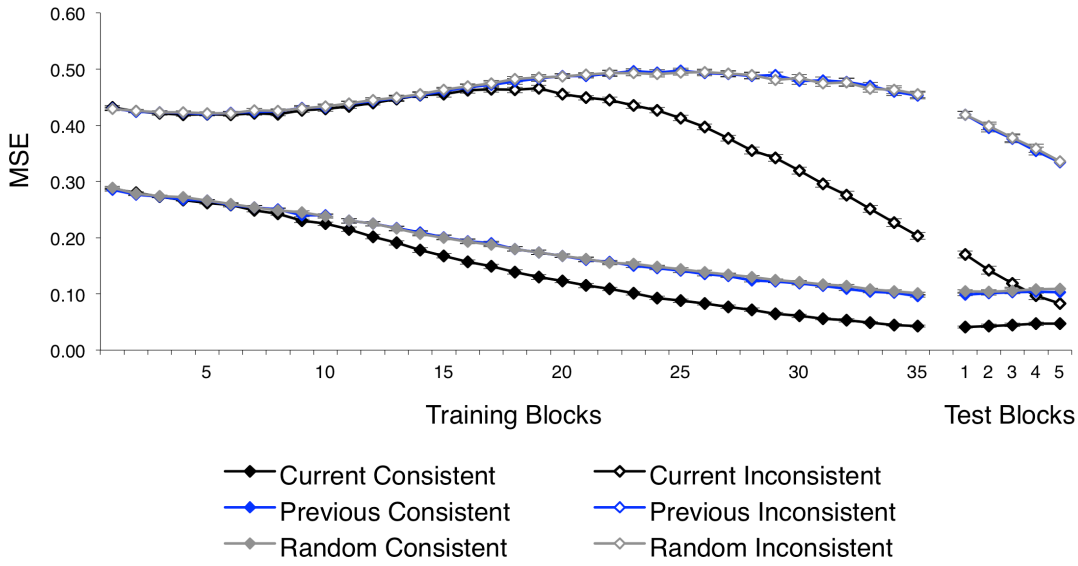

Figure 4.3. MSE for RASRN simulation of new stimulus task for Current (black), Previous (blue) and Random (grey) new stimuli conditions across training and test blocks for consistent (filled diamonds) and inconsistent (open diamonds). All networks were trained on Same rule sequences (Experimental). Error bars (mostly obscured by the markers) show standard error.

As a result, the RASRN predicts that when humans perform the task that the Current group will learn less about sequences due to this cue-competition. Of further interest is

the absence of a difference between the Previous and Random groups, as this suggests that random noise and additional activations that represent the previous stimulus element have the same (or lack of) effect on sequence learning. This suggests that by disrupting the stimulus-response associations formed on each trial the model is able to learn more, as the overshadowing effect of stimulus-response learning on sequence learning is reduced.

4.2.2.2. Subsequence learning and effects

Subsequence also has a large effect on the networks across training, $F(3,207) = 798$, $p < .001$, $MSE = .031$, $\eta_p^2 = .920$, and approached significance at test, $F(3,207) = 2.55$, $p = .082$, $MSE = .002$, $\eta_p^2 = .036$, shown in Figure 4.4. Across training this did not interact with New Stimulus, $F(6,207) = .317$, $p = .817$, $MSE = .031$, $\eta_p^2 = .009$, with networks responding to RRR and LLL better than RLR and LRL. At test, the subsequence effect interacted with New Stimulus, $F(6,207) = 3.83$, $p = .005$, $MSE = .002$, $\eta_p^2 = .100$, as the Current group produced greater learning of subsequences RLR and LRL relative to RRR and LLL, with the other groups showing no difference between subsequences.

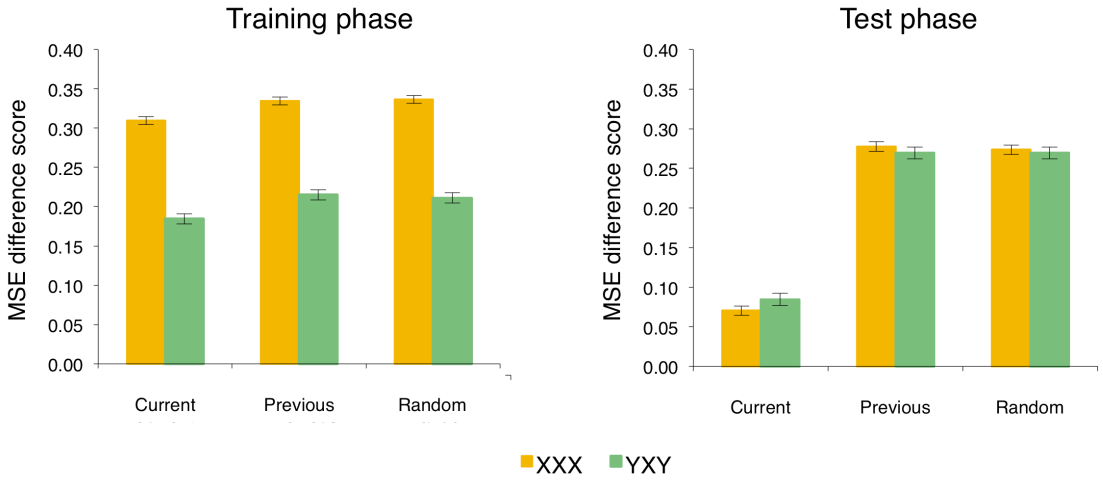

Figure 4.4. MSE difference scores for subsequences XXX (RRR and LLL collapsed, yellow bars) and YXY (RLR and LRL collapsed, green bars) from RASRN simulation of new stimulus task for Current, Prevous and Random new stimuli conditions across training blocks (left panel) and test blocks (right panel). All networks were trained on Same rule sequences. Error bars show standard error.

Whilst all three groups of networks show a strong initial preference for the XXX (RRR and LLL) subsequences, it seems that this disappears by test as networks are able to

learn both the subsequence and the stimulus-response relationships are able to begin to block the influence of trial-by-trial effects on the YXY (RLR and LRL) subsequences. As the Current networks receive more activation of these stimulus-response relationships than Previous and Random networks, this effect is consequently increased.

Block also interacted with Subsequence across both training, $F(102,7038) = 321$, $p < .001$, $MSE = .011$, $\eta_p^2 = .823$, and test, $F(12,828) = 2.01$, $p = .041$, $MSE = .001$, $\eta_p^2 = .028$, see Figure 4.5. This follows from the Subsequence by New Stimulus interaction, as for all groups across training it is easy to see that the subsequences RLR and LRL increase in performance to above that of RRR and LLL, reflecting the increasing influence of stimulus-response associations for all new stimuli, that was shown to be larger then at test for the Current group. However, the interaction between all three variables was not significant across training, $F(204,7038) = .939$, $p = .605$, $MSE = .011$, $\eta_p^2 = .027$, nor test, $F(24,828) = .769$, $p = .729$, $MSE = .001$, $\eta_p^2 = .022$, suggesting that whilst there was an interaction with Subsequence and New Stimulus at test this did not change across the blocks.

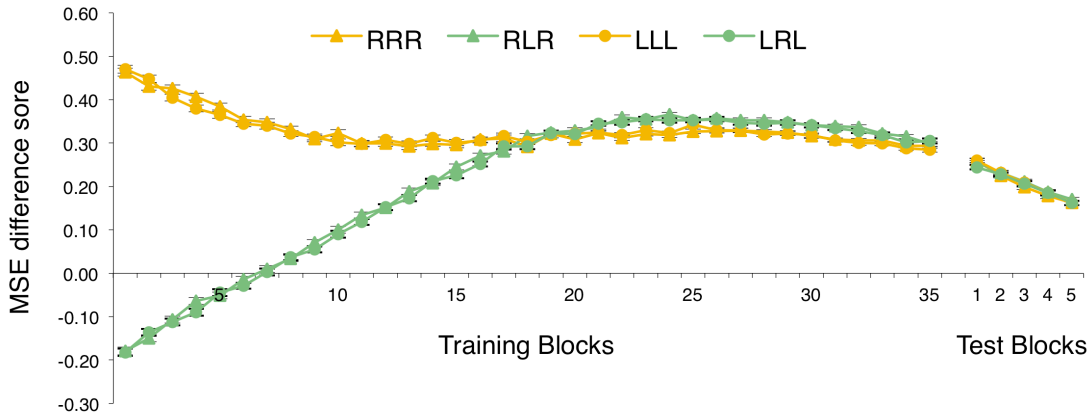

Figure 4.5. MSE difference scores for all four trained subsequences: RRR, LLL (yellow lines), RLR and LRL (green lines) from RASRN simulation of new stimulus task for all Experimental networks, regardless of new stimuli condition across training and test blocks. All networks were trained on Same rule sequences. Error bars show standard error.

4.2.3. Discussion

The RASRN predicts that participants on this task will learn least about sequences in the Current group, with the Previous group and Random group producing the most Same rule sequence learning. This suggests that, if participants are learning using the

mechanisms involved in the RASRN, that increasing the salience of the current on screen stimuli (the Current group) does not improve learning of these sequences relative to the within-subject inconsistent subsequence control (as in Robertson & Pascual-Leone, 2001), it will in fact damage it. An intuitive, performance based account of this task may propose that discriminable stimuli, congruent with the stimuli that you are already learning would be preferred over incongruent, potentially distracting or noise providing stimuli that have no relation to the current response required. However, as the networks in the Current group are given double the amount of activation applicable to training a simple, 100% contingency between units and output, the model starts to learn that these units reliably predict the outcome in such a way that begins to overpower the trained sequential contingency (the Same rule). Therefore, additional concurrent stimuli may be neither redundant nor helpful as the associative account predicts that stimulus-response learning can interfere with sequence learning (as demonstrated in Chapters 2 and 3 of this thesis) and that the Current group will be worse at learning the Same rule sequences.

Interestingly, this effect of additional stimuli is restricted to the Same rule subsequences, as running this task on the Different rule produces no difference between the New Stimuli networks (see Figure 4.6). As the learning of the Different rule is stronger than for the Same rule in the RASRN, this could be protecting the model from the learning of current stimulus-response contingencies. This could also suggest that the learning effect of the New Stimulus is restricted to the subsequences found in the Same group, or that it interacts with this rule in some specific way. This suggestion falls in line with the results from modelling work in Chapter 3, as the introduction of the current stimulus units led to a reduction in Same rule learning, allowing the Different rule learning advantage observed in Experiment 1 to appear. In line with suggestions in the previous chapter, the Current networks trained on the Same rule have a perfect relationship between current stimuli and the required response; as well as the $t - 2$ stimulus location that is used in the learning of the subsequence rule. This provides further evidence that these 'same' relationships between $t - 2$ and t are harder to learn with increased interference from stimulus-response associations.

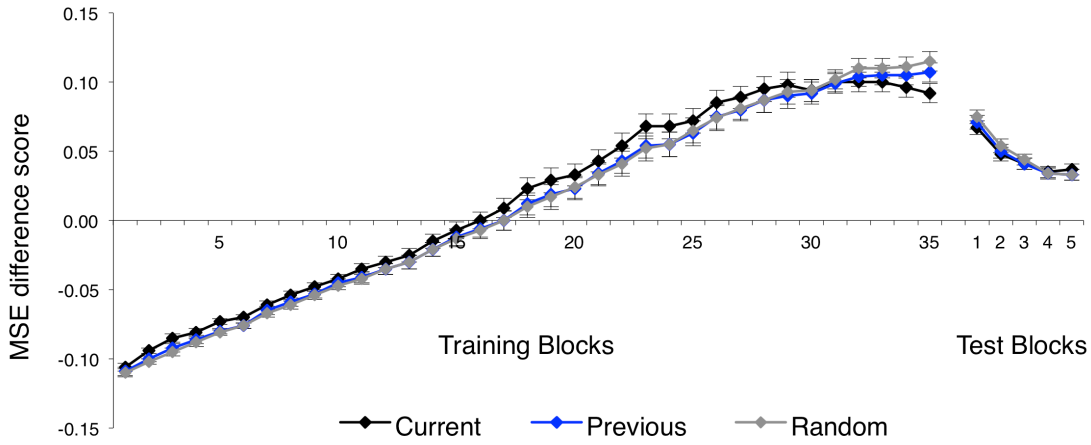

Figure 4.6. MSE difference scores for RASRN simulation of Different rule sequences for Current (black), Previous (blue) and Random (grey) new stimuli conditions across training and test blocks. Error bars show standard error.

4.3. Experiment 3: Tones and Sequences

To investigate the predictions of the RASRN the task outlined for the simulations was run on human participants, with the expectation that those participants experiencing two new stimuli that are entirely congruent with the existing response stimuli in the two-choice SRT task will show worse performance than when these new stimuli correspond with the previous trial or when they are simply random. The new stimuli chosen for the task were two tones, played through headphones for 50 ms with the onset of the on-screen stimuli. The tones were distinguishable by their frequency: a low and high tone.

4.3.1. Method

4.3.1.1. Participants

72 participants (aged between 18 and 48 [$M = 21.7$]; 51 female and 21 male) were recruited from undergraduate students at the University of Exeter and were awarded £10 for participation. Participants provided informed consent prior to taking part in two sessions lasting roughly one hour each. Participants were randomly allocated into one of the three New Stimulus conditions, Current, Previous or Random.

4.3.1.2. Materials and stimuli

The experiment was run on an Apple iMac with PsychToolbox for MatLab software. Participants were seated roughly 50 cm from the screen, which contained two white circle outlines on a black background throughout the task. These white circle outlines

were 19 mm in diameter and positioned vertically in line with the screen centre, and 22 mm either to the left or the right of the screen centre horizontally. The response stimulus was a white filled circle (18.5 mm diameter) that was placed within one of the two circle outlines, giving the white circle outline the appearance of lighting up or filling in. Participants were required to press the spatially compatible 'x' and '>' key presses on a QWERTY keyboard to the left or right response stimulus, respectively. The new tone stimuli were played through headphones and were either high (750 Hz) or low (300 Hz).

4.3.1.3. Design

The experiment was a two-choice SRT task comprising two sessions of twenty blocks each. Each of these blocks contained 120 trials, with all twenty blocks of the first session and first fifteen blocks of the second session acting as training; and the final five blocks acting as test. All participants received training blocks where response stimuli followed Same rule sequential contingencies. This was trained according to a two thirds contingency, as described in Chapter 2 (section 2.2.1.3) with participants able to predict the location of a response stimulus on trial t as the same location as $t - 2$ on two out of three trials. This means that the triplets RRR, RLR, LLL and LRL were twice as likely to occur as RRL, RLL, LLR and LRR. Across the five test blocks all participants experienced a pseudorandom response stimuli trial order, as described in Chapter 2 (see section 2.2.1.3).

Tone stimuli. Regardless of the group participants were assigned to, on every trial a high or low tone played when the response stimuli appeared on screen. Which tone (high or low) occurred depended on the group participants were assigned to. As described above in Simulation 9, the stimuli either occurred with a 100% contingency determined by the current trial (Current); a 100% contingency determined by the previous trial (Previous); or randomly (Random). High and low tones occurred in all groups with equal frequency and in the Random group occurred with equal frequency on both R and L trials, as these were arranged by randomising 30 high and 30 low tones over the 60 R trials in a block, and 30 high and 30 low tones over the 60 L trials in a block. The Random group therefore experienced no contingency between the tones and any response stimulus in the task.

The Current and Previous group tones were constructed according to the response stimulus sequence in place. Across training and test the tones had the same relationship with the response stimuli in these two groups. In the Current group on every trial there was a perfect correlation between the response stimulus and the tone frequency: for example, a high tone occurred on every R trial and a low tone occurred on every L trial. Tone assignment was counterbalanced across participants (half experienced R = High, L = Low; the other half R = Low, L = High).

The Previous group was arranged in much the same way, with a 100% contingency between a response stimulus and tone frequency, but the response stimulus in question was the previous trial. For example, if the previous trial was an R, the tone would always have a high frequency, regardless of whether the current response stimulus was L or R. This was also counterbalanced, with participants experiencing either R(t – 1) = High and L(t – 1) = Low, or R(t – 1) = Low and L(t – 1) = High. The Previous group then experienced no contingency between the response stimulus location and the tone frequency at t; but a 100% contingency between the response stimuli at t – 1 with the tone frequency at t.

4.3.1.4. Procedure

After obtaining informed consent, participants were instructed to simply respond as quickly and accurately as possible to the stimuli, and that the task was investigating how practice had an effect on peoples' speed and accuracy of responding to simple stimuli. Participants were told that at the same time as the stimuli came up on screen, a tone would sound through the headphones. They were instructed to listen to these tones but that they should respond to the on screen stimuli. No mention was made of any contingencies, relationships, sequences or learning. Participants were reminded of these instructions at the beginning of the second session.

At the beginning of each block participants were instructed to press any key to start. Each trial began with an inter-trial interval of 500 ms where a black background with two white circle outlines was presented. The response stimulus (the left or right white circle) would then appear on screen and simultaneously either the high or low frequency tone would sound for 50 ms. The response stimulus would remain on screen until either the participant made a keypress response or the trial timed out after 4000 ms from the

presentation of the response stimulus. RTs were measured from the onset of the response stimulus. If participants pressed an incorrect key, or the trial timed out, the computer issued a low beep sound, which was qualitatively different to the 50msec tone stimulus. The procedure in all other ways directly followed that for Experiment 1, described in section 2.2.1.4.

4.3.2. Results

As for Experiments 1 and 2, average RTs and proportion of errors were calculated for each participant across training blocks and test for each subsequence. RTs were only taken for correct trials that did not follow an error. Errors were calculated as the proportion of trials on which participants made the incorrect key-press response (i.e. they pressed x instead of > for a right response stimulus trial) and not those trials that timed out or involved an incorrect key-press of any other key on the keyboard. The first two trials of each test block and first three trials of each training block were not included in the analysis, as these could not be assigned a subsequence. Training average RTs were weighted across response stimulus location at $t - 3$, and, as this leads to missing values averages for each subsequence, were calculated over two blocks (which constitute an Epoch). This number was used instead of the 5 blocks in Experiment 1 and 2 in order to produce a more sensitive measure of the changing influence of the tones across the experiment, but this meant that block 35 (the final training block) was dropped from the analysis. Difference scores were calculated, taking performance on consistent subsequences (RRR, RLR, LLL, LRL) from inconsistent subsequences (RRL, RLL, LLR, LRR, respectively) and higher scores therefore indicate better performance on trained consistent subsequences over inconsistent subsequences. See section 2.2.2 for more details.

An ANOVA was performed on RT and error difference scores across training and test. Training data compared the seventeen training Epochs; New Stimulus (Current; Previous or Random); and Subsequence (RRR, RLR, LLL and LRL). Test data was analysed for differences between the five Blocks; New Stimulus; and Subsequence. All within-subject main effects and interactions are reported with p values that correct for a departure from sphericity (Huynh-Feldt) with the unadjusted degrees of freedom.

4.3.2.1. Same rule learning

There was a main effect of New Stimulus across training RT difference scores, $F(2,69) = 3.88$, $p = .025$, $MSE = 21248$, $\eta_p^2 = .101$, and errors, $F(2,69) = 5.58$, $p = .006$, $MSE = .611$, $\eta_p^2 = .139$; as well as in errors at test, $F(2,69) = 4.86$, $p = .011$, $MSE = .191$, $\eta_p^2 = .124$; but not RTs, $F(2,69) = 1.44$, $p = .245$, $MSE = 7391$, $\eta_p^2 = .040$, see Figure 4.7. In all cases (training and test, RTs and errors) the Previous group learn numerically more than the other groups, which is supported by Bonferroni corrected pairwise comparisons comparing Previous and Random groups across training RTs, $F(1,69) = 7.08$, $p = .029$, $MSE = 312$, $\eta_p^2 = .093$; errors, $F(1,69) = 11.2$, $p = .004$, $MSE = .009$, $\eta_p^2 = .139$; and test errors, $F(1,69) = 9.59$, $p = .008$, $MSE = .010$, $\eta_p^2 = .122$; but not RTs, $F(1,69) = 2.84$, $p = .289$, $MSE = 370$, $\eta_p^2 = .040$.

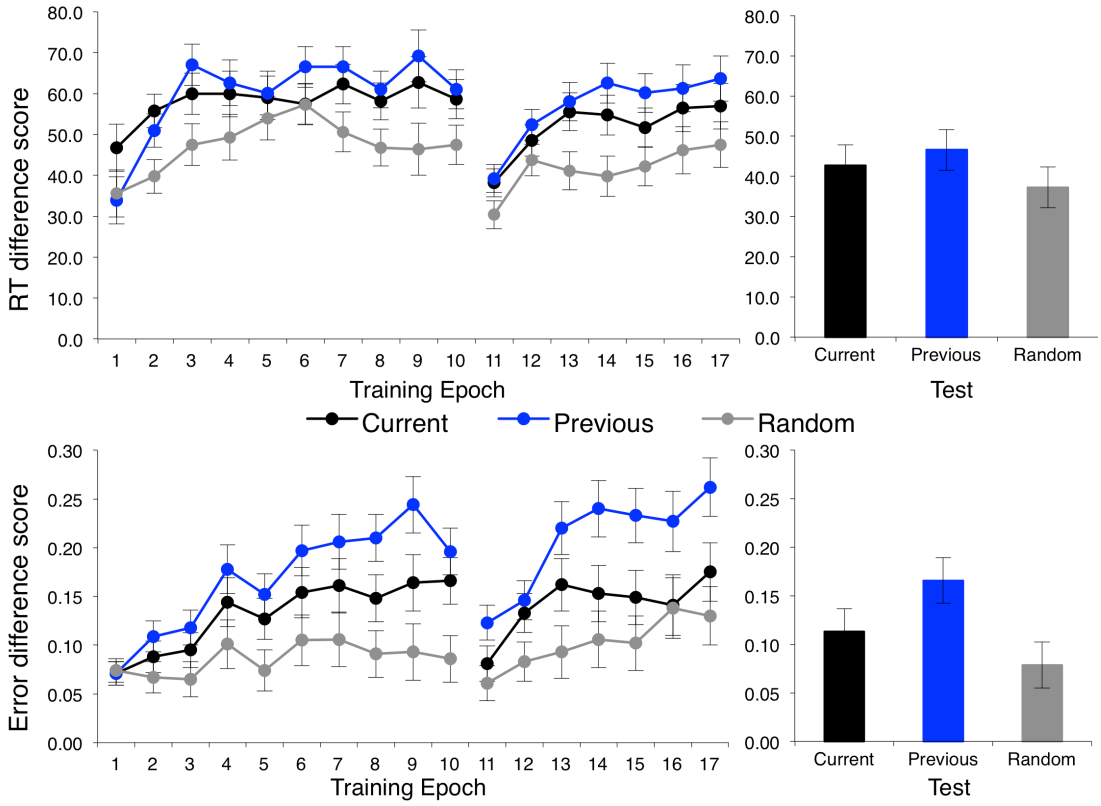

Figure 4.7. RT (top panel) and error (bottom panel) difference scores for participants in Experiment 3 for Current (black), Previous (blue) and Random (grey) new stimuli conditions (tones) across training and test blocks. All participants were trained on Same rule sequences. Error bars show standard error.

Whilst the Previous group was consistently above the Current group, this difference was not significant across training RTs, $F(1,69) = .381$, $p > .9$, $MSE = 312$, $\eta_p^2 = .005$; errors, $F(1,69) = 3.09$, $p = .250$, $MSE = .009$, $\eta_p^2 = .043$; test RTs, $F(1,69) = .489$, $p > .9$, MSE

= 370, η_p^2 = .007; nor errors, $F(1,69) = 3.50$, $p = .197$, $MSE = .010$, $\eta_p^2 = .048$. Similarly, the Current and Random group were not significantly different across training RTs, $F(1,69) = 4.18$, $p = .134$, $MSE = 312$, $\eta_p^2 = .057$; errors, $F(1,69) = 2.51$, $p = .354$, $MSE = .009$, $\eta_p^2 = .035$; test RTs, $F(1,69) = .973$, $p = .982$, $MSE = 370$, $\eta_p^2 = .014$; or errors, $F(1,69) = 1.50$, $p = .674$, $MSE = .010$, $\eta_p^2 = .225$, although the Current group were consistently above the Random group.

Epoch across training had a significant effect in both RTs, $F(16,1104) = 11.4$, $p < .001$, $MSE = 1604$, $\eta_p^2 = .142$, and errors, $F(16,1104) = 14.8$, $p < .001$, $MSE = .035$, $\eta_p^2 = .177$. This provides some evidence of learning across the experiment, although without a control group this development in both RT and error performance cannot be attributed to a speed-accuracy trade off but could be a familiarisation with the task rather than sequence learning. This learning or familiarisation did not show evidence of a significant interaction with New Stimulus in RTs, $F(16,1104) = 1,11$, $p = .327$, $MSE = 1604$, $\eta_p^2 = .031$; but did in the errors, $F(16,1104) = 2.04$, $p = .004$, $MSE = .035$, $\eta_p^2 = .056$, suggesting that learning is occurring more rapidly for the Previous group, as each group are performing with the same accuracy in Epoch one, with the difference between groups emerging across the two sessions of training. The Current group appear to perform higher than the Random group in both RTs and errors for the most part, although a post-hoc Scheffé correction between Current and Random groups just falls outside of significance, $F(2,69) = 2.84$, $p = .065$, $MSE = .013$, $\eta_p^2 = .076$, hence there is no conclusive evidence of greater learning in the Current group.

There is support for a learning difference in the last block of test in the errors, with a significant main effect of New Stimulus, $F(2,69) = 4.84$, $p = .011$, $MSE = .088$, $\eta_p^2 = .123$, which suggests that a learning effect has emerged over training that shows greater learning in the Previous group than the Random group in a Bonferroni corrected comparison, $p = .009$. At test there was a main effect of Block only in the error difference scores, $F(4,276) = 2.57$, $p = .043$, $MSE = .023$, $\eta_p^2 = .036$, not RT difference scores, $F(4,276) = 1.28$, $p = .282$, $MSE = 1710$, $\eta_p^2 = .018$. This reflected an overall trend towards the extinction of learning (which was not contradicted by the non-significant direction in the errors), but this again showed little evidence of an interaction with New Stimulus in either RTs, $F(8,276) = .755$, $p = .618$, $MSE = 1710$, $\eta_p^2 = .021$, or errors, $F(8,276) = .596$, $p = .767$, $MSE = .023$, $\eta_p^2 = .017$.

It is clear when the data is split into inconsistent and consistent subsequence RTs and proportion of errors, see Figure 4.8, that the pattern of results does not follow those produced by the RASRN. In errors, responding to consistent subsequences was the same across the three New Stimulus groups. Whilst the RASRN performed better on inconsistent subsequences in the Current group towards the end of training, the human participants show a worsening in performance in errors across training on inconsistent subsequences. This produces larger difference scores throughout and RT scores show better performance for Current and Previous groups across training on consistent sequences, again not a prediction of the model.

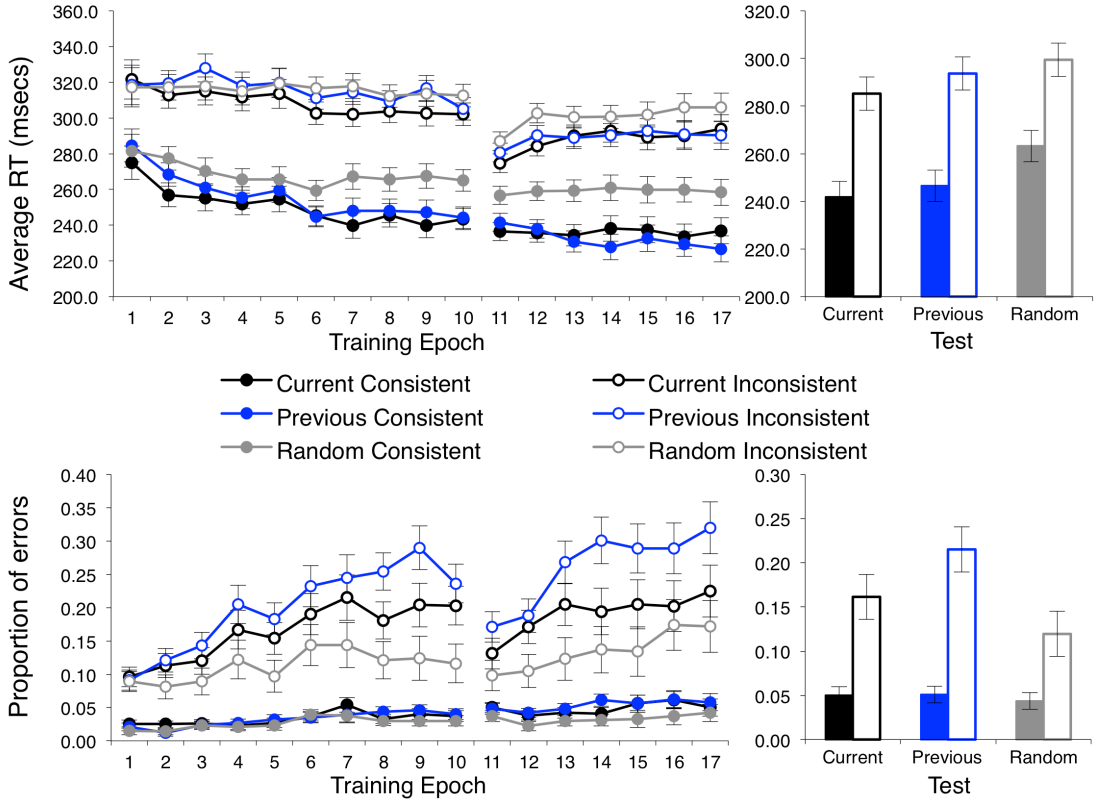

Figure 4.8. Average RTs and proportion of errors for human performance on new stimulus task for Current (black), Previous (blue) and Random (grey) new stimuli conditions across training and test blocks for consistent (filled circles and bars) and inconsistent (open circles and bars). Error bars show standard error.

4.3.2.2. Subsequence effects and learning

Subsequence had an effect across training RTs, $F(3,207) = 9.69$, $p < .001$, $MSE = 15903$, $\eta_p^2 = .123$; training errors, $F(3,207) = 4.35$, $p = .017$, $MSE = .120$, $\eta_p^2 = .059$; test RTs, $F(3,207) = 3.96$, $p = .018$, $MSE = 4998$, $\eta_p^2 = .054$; and test errors, $F(3,207) = 3.93$, p

= .009, *MSE* = .031, η_p^2 = .054, see Figure 4.9. These Subsequence effects demonstrate, as in Chapter 2, a speed-accuracy trade off, with better performance on YXY subsequences (RLR and LRL compared to RLL and LRR, respectively) than XXX subsequences in terms of speed of responding. The opposite is true for accuracy, with better performance on XXX subsequences compared to YXY subsequences. These Subsequence effects showed no evidence of an interaction with the New Stimulus group: relevant interaction for training RTs, $F(6,207) = .407$, $p = .786$, $MSE = 15903$, $\eta_p^2 = .012$; errors, $F(6,207) = .627$, $p = .632$, $MSE = .120$, $\eta_p^2 = .018$; test RTs, $F(6,207) = .713$, $p = .596$, $MSE = 4998$, $\eta_p^2 = .020$; and errors, $F(6,207) = .656$, $p = .685$, $MSE = .031$, $\eta_p^2 = .019$.

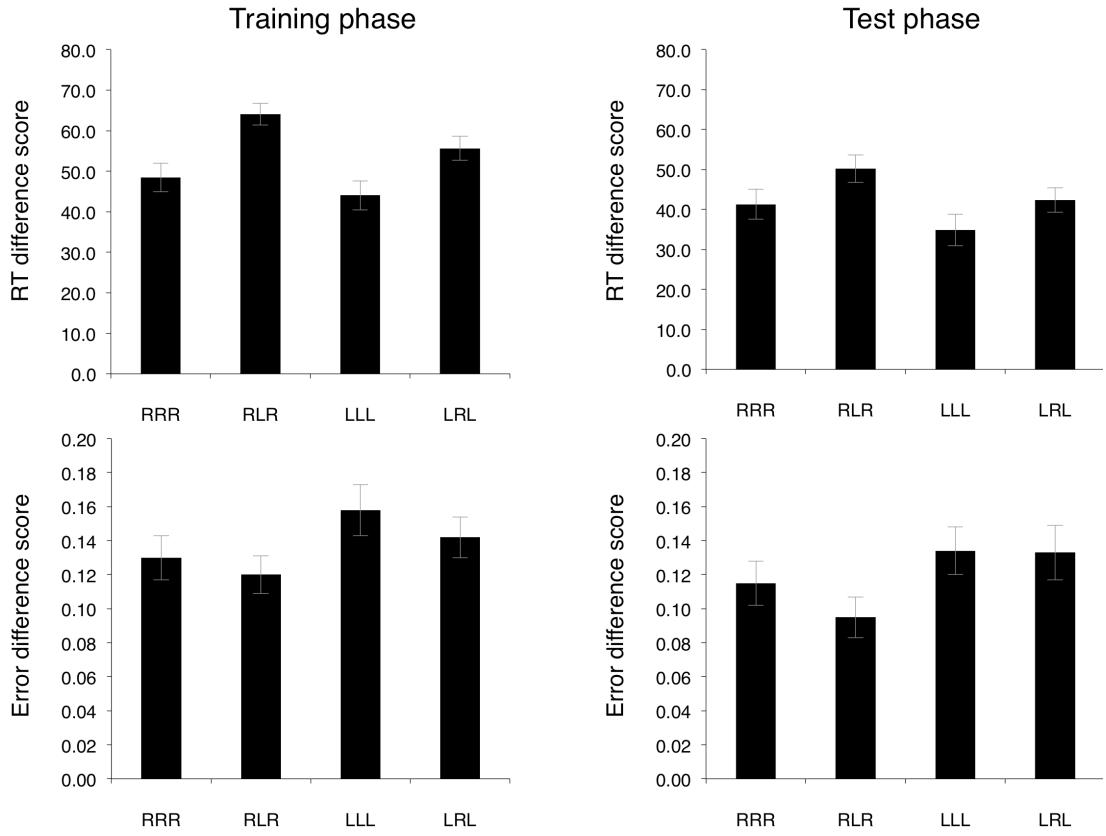

Figure 4.9. RT (top panel) and error (bottom panel) difference scores for subsequences for all participants, regardless of New Stimulus condition, across training blocks (left panel) and test blocks (right panel). Error bars show standard error.

The Subsequence effect did interact with Epoch across training (see Figure 4.10) in the RTs, $F(48,3312) = 2.30$, $p < .001$, $MSE = 2290$, $\eta_p^2 = .032$; but not errors, $F(48,3312) = 1.10$, $p = .322$, $MSE = .026$, $\eta_p^2 = .016$, nor with Block at test in RTs, $F(12,828) = .797$, $p = .623$, $MSE = 1528$, $\eta_p^2 = .011$, or errors, $F(12,828) = .980$, $p = .465$, $MSE = .013$, η_p^2

= .014. This training RT effect also interacted with New Stimulus, $F(96,3312) = 1.40$, $p = .033$, $MSE = 2290$, $\eta_p^2 = .039$, suggesting that whilst performance on YXY remained stable, participants in the Previous group experienced a more rapid learning of subsequence XXX than the other groups in the RTs. This interaction between Block, New Stimulus and Subsequence was also apparent in RTs at test (see Figure 4.10), $F(24,828) = 1.44$, $p = .024$, $MSE = 1528$, $\eta_p^2 = .049$, suggesting that the subsequence XXX suffers from less extinction in the Previous group, as the RT difference score for this subsequence only gradually reduces after Block 2, whereas a more dramatic reduction in this subsequence is seen in Current and Random groups after Bock 1.

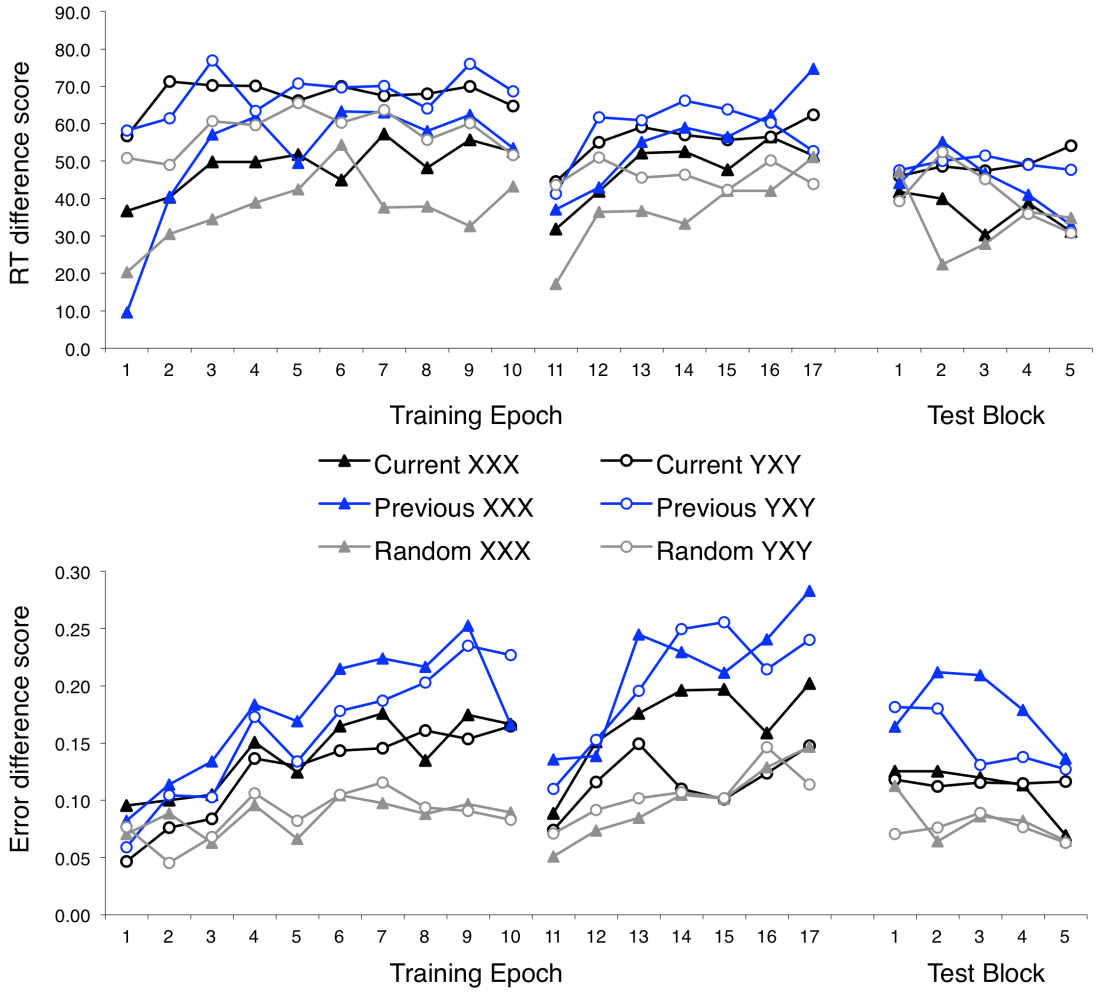

Figure 4.10. RT (top panel) and error (bottom panel) difference scores for Current (black lines); Previous (blue lines); and Random (grey lines) New Stimulus groups on subsequences XXX (RRR and LLL collapsed, filled triangles) and YXY (RLR and LRL collapsed, open circles) across training Epochs.

4.3.3. Discussion

Experiment 3 provides evidence that when learning Same rule subsequences under incidental conditions, participants who heard two tones that corresponded with the on-screen response stimuli learned less than when the tones corresponded with the previous stimuli. This follows an associative account as the Current group learned less than the Previous group, going some way to confirm the predictions of the RASRN. However, both raw RT and error performance on consistent and inconsistent subsequences do not follow the pattern of results observed in the RASRN. Furthermore, the Random group, whose tones had no relationship to current nor previous response stimuli locations, performed (numerically) the worst in learning the Same rule sequential contingencies. It therefore seems unlikely that participants are performing according to the exact predictions of the RASRN.

We can begin our analysis of these results by asking if we see evidence of the RASRN prediction that stimulus-response learning is restricting learning in the Current group in humans? There is no conclusive answer provided by Experiment 3, as the humans do not demonstrate RT and error difference score decreases for any New Stimulus conditions across training as they do for all groups in the RASRN; and more rapidly in the Current networks. Conversely, participants in Experiment 3 demonstrate increasing RT and error difference scores throughout the task, suggesting that the model perhaps involves a stronger stimulus-response association than are formed in humans as this is not enough to overshadow sequence learning in Experiment 3. This does not suggest that stimulus-response associations are not formed, nor that they have no effect; as the Current group respond faster and more accurately than Previous and Random groups, but yet show numerically reduced learning compared to the Previous group.

The Random group provides the strongest evidence of a departure from the predictions of the RASRN, as this group demonstrates the poorest human sequence learning compared to model predictions placing it as equal to the Previous group. This suggests that the influence of these stimulus-response units is perhaps more complex than the RASRN currently predicts; but it is difficult to interpret these differences. The Random group may be worse than the Previous and Current groups as it disrupts the sequence, or it may have no effect on sequence learning and this may provide evidence that the Current and Previous groups both perform better than what we would expect if no dual-stimuli were present. It is also possible that the presence of an adaptive learning rate

may have been adversely effected by the random noise in the Random group. In this case the learning rate may have been reduced, subsequently learning about the sequence reduced. Further modelling using the RASRN with a set of simulations that have no additional stimulus units show precisely the same pattern as the Previous and Random groups, therefore the model predicts that performance in both Previous and Random cases is not affected by these additional stimuli. The only manipulation that affected sequence learning is the enhancement of S-R learning in the Current group; which results in less sequence learning. This is investigated further in Experiment 4 where a group that has no contingencies was tested to enable a comparison between the three existing New Stimulus groups and a no-new stimulus case to attempt to discover in which groups we see more or less Same rule learning than we would expect from the simple single-stimulus SRT task; hence enabling further understanding of the underlying processes and their effect on learning.

A further possibility that may explain the results in Experiment 3 is that the tones produce some qualitative difference between Current and the other two groups, which produces some unwanted effect. Participants in the Current group often said that they used the tone as cues to respond and clearly heard two distinct tones for each side, without fail. Participants in the Previous and Random groups tended to say that they either attempted to ignore the tones, or did not even recognise that there were two different tones, as they thought that these were not part of the task. It seems unlikely that this qualitative difference in and of itself caused the differences on the task, as if this were the case the Random and Previous group performance would match. It could be instead that the influence of the stimulus-response association between each tone and response is reduced in some way in the Current group. This could be because participants pay less attention to the response stimuli as they are simply responding to the tones, or it could be because this condition is simply easier and participants are motivated to learn less. On the other hand, participants may have been provided with a memory of the previous trials, making the sequence learning task simply easier to do and therefore learning may have increased in the Previous group.

An advantage for the Previous and Current groups could also be mediated by the counterbalanced tone mappings in the Current group. On each trial regardless of the tone counterbalancing in the Previous group there was no contingency between a high or low tone and either response stimulus location. As the Current group experienced a

relationship on each trial between the tone and response stimulus location, this group could have been affected by the spatial-musical association of response codes (SMARC, Rusconi, Kwan, Giodano, Umiltà and Butterworth, 2006). It has been shown that people are predisposed to perceive sound in certain locations, with high tones appearing higher in the space or towards the right, and low tones appearing lower and to left. Participants in the Current group were analysed based on the tone counterbalancing, and therefore whether the high and low tones were consistent with the SMARC effect or not were compared in an ANOVA with Epoch (across training) or Block (across test), Subsequence and Counterbalancing (SMARC Compatible or Incompatible). No main effect of Counterbalancing was found across RT training, $F(1,22) = .015$, $p = .904$, $MSE = 18908$, $\eta_p^2 = .001$; training errors, $F(1,22) = .065$, $p = .801$, $MSE = .606$, $\eta_p^2 = .003$ or RT test, $F(1,22) = .001$, $p = .970$, $MSE = 8291$, $\eta_p^2 = .001$, nor errors, $F(1,22) = .885$, $p = .357$, $MSE = .190$, $\eta_p^2 = .039$, therefore it seems that there was no influence of a SMARC effect in the Current group that may have interfered with or influenced the results of Experiment 3.

Whilst the SMARC effect was found not to influence the Current group, the tones may had some additional effect or have been more salient in this group, thus changing perhaps their relative activation level and influence as input. It could also be the case the model was unable to accurately model the data due to its matched temporal representation of the tones (which occurred for less time [50 ms]) than the visual response stimulus (response terminated). As the current stimulus units were all activated at the same level, this difference between on-screen response-stimuli and tones is not represented accurately in these model activations.

Altogether it is difficult to ascertain the effect of the tones on sequence learning with no control group, as the Random group also experienced tones and I am consequently unable to conclude exactly what effect that the tones have. Taken with the possibility that the tones may have been producing unwanted explicit or perceptual confounds to the main learning effect of interest, I decided to investigate these effects without the strong experiential influence of the tones and with a control no-concurrent-stimulus condition for comparison. A replication of the study was therefore designed with New Stimuli that were visual rather than auditory, to investigate whether learning in the Previous group was significantly better than the other groups, or whether some property

of the tones led participants to produce these effects and learning with simply visual stimuli would better follow the predictions of the RASRN.

4.4. Experiment 4: Colours and Sequences

As discussed above, the stimuli chosen to represent the New Stimuli units in Experiment 3 could have introduced several confounds that altered learning as predicted by the RASRN. The salience and length of time that the stimuli were presented for (tones and on-screen white circle response stimuli) were not matched, as well as there being the opportunity to pay attention to one modality or another in the Current group. Participants in this study were trained on the same task, this time for only one session, as the effect in human participants emerged across the first few Epochs of training. The New Stimuli used were colours, rather than tones, with two different colours that took the place of the white circle response stimuli. These colours followed the relationship (or not) with the sequence of response-stimuli locations as described above for Current, Previous and Random groups, as well as a further No-Colour group who were trained with white response stimuli only, which acted as a control for the presence of colours.

4.4.1. Method

4.4.1.1. Participants

96 participants (aged between 18 and 33 [$M = 20.5$]; 71 female and 25 male) were recruited from undergraduate students at the University of Exeter and were awarded £5 or one course credit in exchange for participation. Participants provided informed consent prior to taking part in one session lasting roughly one hour. Participants were randomly allocated into one of the four New Stimulus conditions, Current, Previous, Random or No-Colour.

4.4.1.2. Materials and stimuli

The materials and stimuli used followed Experiment 3 (see section 4.3.1.2), except for the tones and colour of the response stimuli. The tones were no longer involved, and therefore participants did not wear headphones. The response stimulus in this experiement, instead of always being a white filled circle, was in most groups a coloured filled circle (18.5 mm diameter) that was placed within one of the two circle outlines, giving the white circle outline the appearance of lighting up or filling in with a

particular colour. The colours used were purple (RGB: 128,0,255) and yellow (RGB: 255,255,0) for the participants in Current, Previous and Random groups. The No-Colour group experienced white (RGB: 255,255,255) response stimuli on all trials. The relationship between these colours and which location they appeared in is described below. Participants were required to press the spatially compatible 'x' and '>' key presses on a QWERTY keyboard to the left or right response stimulus, respectively.

4.4.1.3. Design

The experiment was a two-choice SRT task comprising of one session of twenty blocks. The first sixteen blocks acted as training; and the final four blocks acted as test. All participants received training blocks of Same rule sequential contingencies, as described earlier. Across the five test blocks all participants experienced a pseudorandom response stimulus trial order.

Response stimuli colour. Participants in the No-Colour group simply experienced white response stimuli on all trials, following exactly the materials and stimuli experienced in Experiment 1. The data was not simply reused from Experiment 1 as those participants were trained across two sessions and tested after 35 blocks of trained; however in the current experiment participants were trained on only one session and tested after sixteen blocks. The Random group experienced response stimuli that were equally likely to be either purple or yellow, with 30 each yellow and purple randomly allocated to the R trials across a block and 30 each yellow and purple response stimuli colour randomly allocated to L trials. There was no contingency between the colour of the response stimuli and the location.

The Current and Previous groups were also organised as described for Experiment 3, with the Current group experiencing a 100% contingency between the response stimulus colour and location, and the Previous group experiencing a 100% contingency between the previous response stimulus location and current colour. Therefore, in the Current group one response stimulus location would fill in purple and the other would fill in yellow in every instance. The Previous group experienced an equal amount of left and right trials that were yellow or purple, but every trial following one response stimulus location would always be a purple trial, and the other a yellow trial. The colours were

counterbalanced across participants so that half experienced R = yellow and L = purple, and the other half R = purple and L = yellow. This continued through training and at test.

4.4.1.4. Procedure

The procedure was as described for Experiment 3 (see section 4.3.1.4), with participants in Current, Previous and Random groups instructed that they were to respond to the side of the screen that the circle appeared and not the colour of the stimulus itself. The No-Colour group were not given this instruction.

4.4.2. Results

Results were collected and RT and error difference scores were analysed using an ANOVA on the eight training Epochs; New Stimulus (Current; Previous; Random; or No-Colour); and Subsequence. At test, an ANOVA compared two Epochs; New Stimulus; and Subsequence. All details regarding the treatment and analysis of data are as for Experiment 3, see section 4.3.2.

4.4.2.1. Same rule learning

There was a main effect of New Stimulus across training in RT difference scores, $F(3,92) = 3.40$, $p = .021$, $MSE = 9665$, $\eta_p^2 = .100$, but not errors, $F(3,92) = .706$, $p = .551$, $MSE = .268$, $\eta_p^2 = .022$, see Figure 4.11. Across test this was the same, there is a main effect across RTs, $F(3,92) = 2.98$, $p = .036$, $MSE = 2887$, $\eta_p^2 = .089$, but not errors, $F(3,92) = .701$, $p = .554$, $MSE = .062$, $\eta_p^2 = .022$. It was clear that the Previous group learned more than the other groups, which was supported by a series of planned contrasts based on the results of Experiment 3, which are all shown in Table 4.2. Previous is compared first to the Current group, and is significantly better in training RTs and test RTs, but not errors across training, nor test. The Previous group learned significantly better than the Random group across training in RTs (again see Table 4.2) and at test. However, again the numerical advantage was not significant in errors across training or test. The difference between Previous and No-Colour groups (see Table 4.2) was significant in RTs across training and test; but not training errors nor test. A series of Bonferroni corrected pairwise comparisons found no other significant differences between New Stimulus conditions.

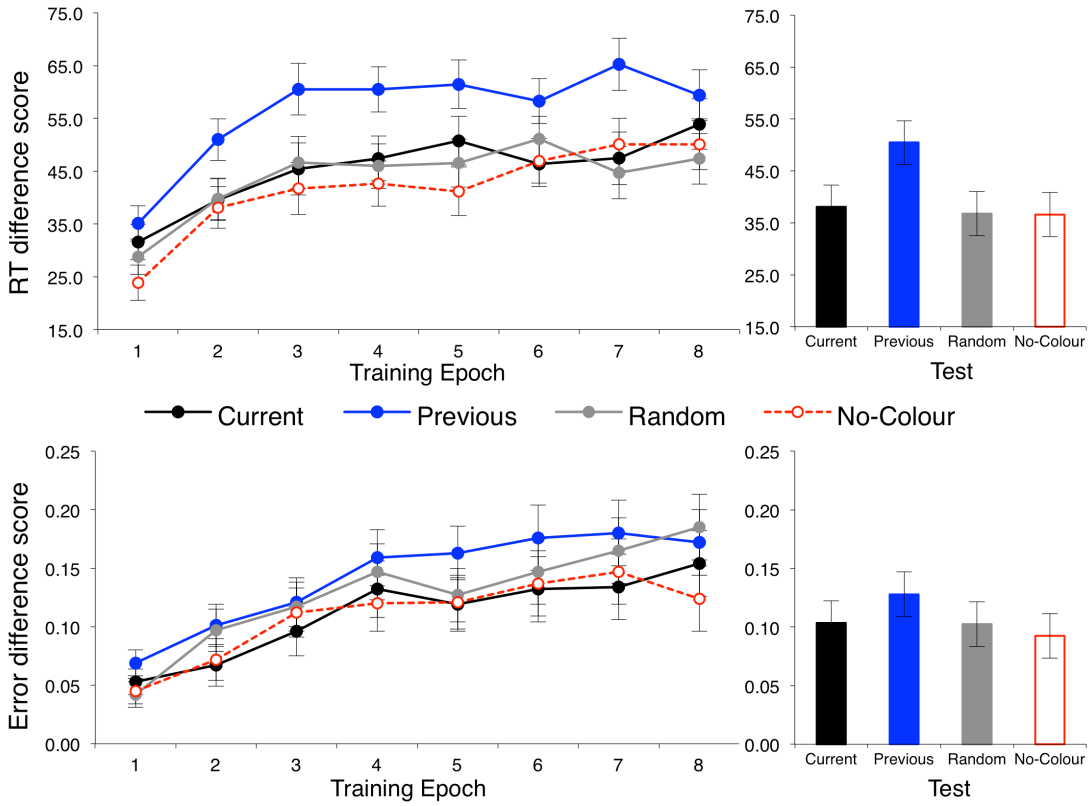

Figure 4.11. RT (top panel) and error (bottom panel) difference scores for participants in Experiment 4 for Current (black), Previous (blue), Random (grey) and No-Colour (red) new stimuli conditions (tones) across training and test blocks. All participants were trained on Same rule sequences. Error bars show standard error.

Table 4.2. Results from ANOVA for planned comparisons regarding the Previous group based on the results of Experiment 3 conducted on both RT and error difference scores for training and test phases of Experiment 4.

	RT difference score: Training			Error difference score: Training		
Previous vs:	F (df = 1,92)	p	η_p^2 (MSE = 302)	F (df = 1,92)	p	η_p^2 (MSE = .008)
Current	4.91	.029	.051	1.46	.231	.016
Random	6.29	.014	.064	.292	.590	.003
No-Colour	8.48	.004	.084	1.55	.217	.017
	RT difference score: Test			Error difference score: Test		
Previous vs:	F (df = 1,92)	p	η_p^2 (MSE = 361)	F (df = 1,92)	p	η_p^2 (MSE = .008)
Current	5.10	.026	.053	.960	.330	.010
Random	6.23	.014	.063	1.01	.319	.011
No-Colour	6.41	.013	.065	1.93	.168	.021

The effect of New Stimulus did not interact with Epoch across training RTs, $F(21,644) = 1.09$, $p = .358$, $MSE = 823$, $\eta_p^2 = .034$; nor errors, $F(21,644) = .514$, $p = .948$, MSE

= .026, η_p^2 = .016. There was also no interaction across test in RTs, $F(3,92) = 1.69$, p = .175, MSE = 548, η_p^2 = .052; nor errors, $F(3,92) = 1.79$, p = .155, MSE = .012, η_p^2 = .055, suggesting that the development of the Previous advantage was not gradual across training, indeed it was relatively stable after the first Epoch. Epoch itself had a main effect across training RTs, $F(7,644) = 29.1$, $p < .001$, MSE = 823, η_p^2 = .240, and errors, $F(7,644) = 25.5$, $p < .001$, MSE = .026, η_p^2 = .217; as well as across test RTs, $F(1,92) = 10.4$, p = .002, MSE = 548, η_p^2 = .101, and errors, $F(1,92) = 7.29$, p = .008, MSE = .012, η_p^2 = .073. This shows the improved performance on consistent subsequences compared to inconsistent across training and the reduction of these difference scores at test, reflecting learning and extinction, respectively.

When we examine the learning effect by dividing the data up into inconsistent and consistent subsequence performance we can see that, like Experiment 3 and unlike the RASRN prediction, that it is the poor accuracy and speed to inconsistent subsequences in the Previous group that produces the learning advantage for this group (see Figure 4.12). Random and No-Colour groups are matched on their consistent and inconsistent performance almost exactly, but both Previous and Current groups experience numerically slower yet more accurate responses to consistent subsequences than the Random and No-Colour groups at test. The model predicts improving performance on Current inconsistent trials, whilst this is not the case the group do demonstrate across training and test the least errors to inconsistent subsequences. This, however, appears to be a speed-accuracy trade-off, with the slowest responding to consistent subsequences. Therefore the pattern of responding replicates that found in Experiment 3, and not the RASRN.

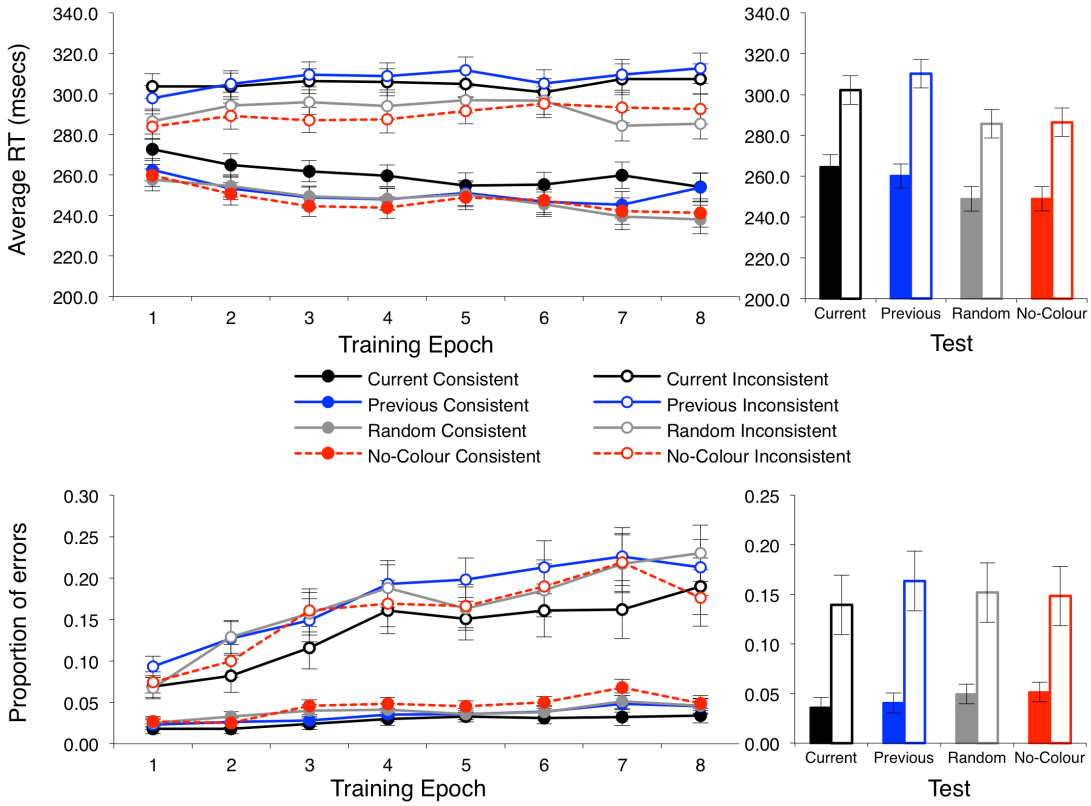

Figure 4.12. Average RTs and proportion of errors for human performance on new stimulus task for Current (black), Previous (blue), Random (grey) and No-Colour (red) new stimuli conditions across training and test blocks for consistent (filled circles and bars) and inconsistent (open circles and bars). Error bars show standard error.

4.4.2.2. Subsequence effects and learning

Participants showed a main effect of Subsequence across training RTs, $F(3,276) = 18.6$, $p < .001$, $MSE = 7389$, $\eta_p^2 = .168$, test RTs, $F(3,276) = 13.1$, $p < .001$, $MSE = 548$, $\eta_p^2 = .125$, and test errors, $F(3,276) = 4.71$, $p = .007$, $MSE = .015$, $\eta_p^2 = .049$; but not training errors, $F(3,276) = 2.35$, $p = .092$, $MSE = .062$, $\eta_p^2 = .025$, see Figure 4.13. We can see that in this case the sequential effects demonstrate faster *and* more accurate responding to RLR and LRL subsequences compared to inconsistent RLL and LRR subsequences over RRR and LLL responding compared with RRL and LLR subsequences. This is entirely in the opposite direction to the subsequence effects found in the RASRN (see Figure 4.4).

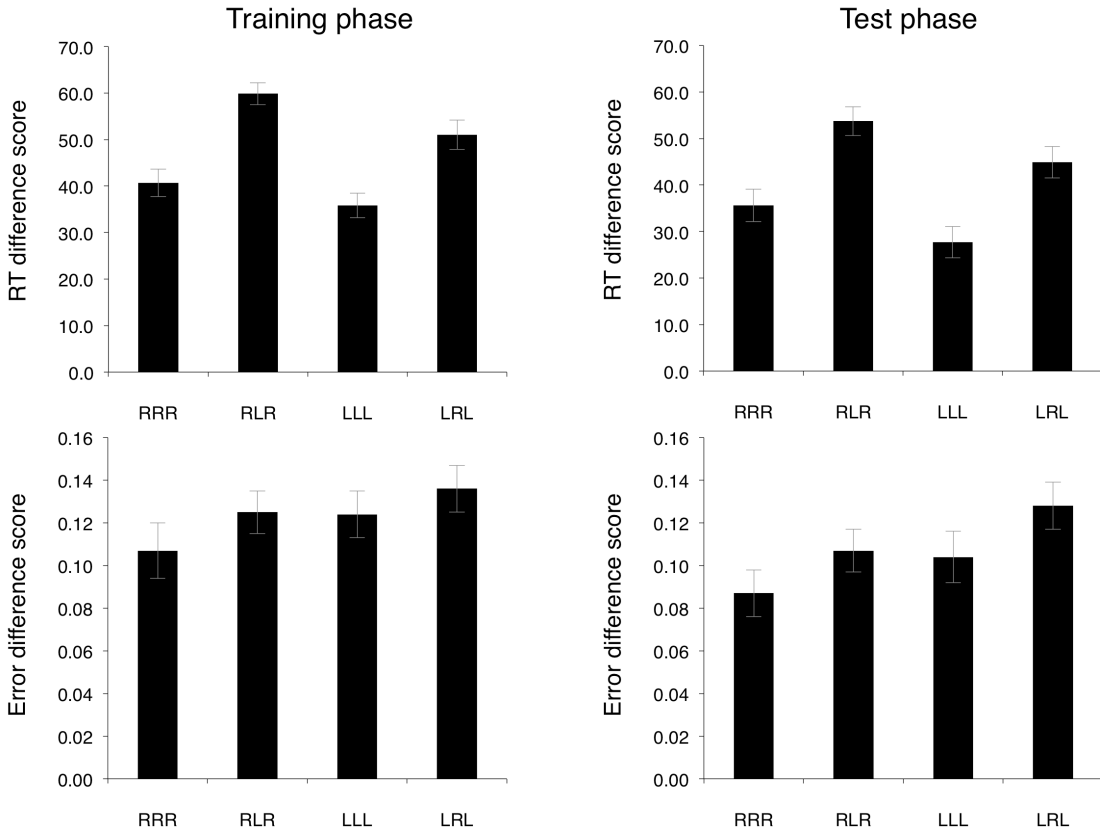

Figure 4.13. RT (top panel) and error (bottom panel) difference scores for subsequences for all participants, regardless of New Stimulus condition, across training blocks (left panel) and test blocks (right panel). Error bars show standard error.

These Subsequence effects did not interact with New Stimulus in any case: training RTs, $F(9,276) = .462$, $p = .830$, $MSE = 7389$, $\eta_p^2 = .015$; training errors, $F(9,276) = .475$, $p = .846$, $MSE = .062$, $\eta_p^2 = .015$; test RTs, $F(9,276) = .437$, $p = .897$, $MSE = 2114$, $\eta_p^2 = .014$; test errors, $F(9,276) = .890$, $p = .516$, $MSE = .015$, $\eta_p^2 = .028$. New Stimulus and Subsequence also had no three-way interaction with Epoch across training: RTs, $F(63,1932) = 1.26$, $p = .098$, $MSE = 807$, $\eta_p^2 = .039$; errors, $F(63,1932) = 1.01$, $p = .464$, $MSE = 824$, $\eta_p^2 = .032$; nor test: RTs, $F(9,276) = 1.11$, $p = .353$, $MSE = 824$, $\eta_p^2 = .035$; errors, $F(9,276) = 1.37$, $p = 200$, $MSE = .008$, $\eta_p^2 = .043$. This suggests that the relationship between the New Stimulus and the trial order did not differentially affect the learning or performance on certain subsequences.

4.4.3. Discussion

The results of Experiment 4 replicate the findings of Experiment 3, that human participants learn more about Same rule sequential contingencies when additional stimuli correspond with the previous response stimulus location. The RASRN predicted

the differential effect of an additional stimulus on the Same rule learning, and correctly predicted that the Current group would be worse than the Previous group. However, in both the case of tones and colours, the Current group also learned the Same rule sequences numerically (though not significantly) better than the Random (and No-Colour) group, which does not follow the prediction of the RASRN. The model predicted that the Current group would actually perform worse than all of the other conditions, with the increased activation of current stimuli producing stronger stimulus-response associations that interfered with sequence learning. The results of Experiments 3 and 4 do not exclude this possibility, but they suggest that the Previous group has a clear advantage over the other New Stimulus conditions.

This suggests that, firstly, tones did not produce some qualitative difference for participants depending on whether they were contingent with on-screen response stimuli or not, or that they were processed at a faster speed than the response stimuli, giving some speed bias to the Current group. That this effect is replicated across two different stimulus types despite their differing characteristics, provides strong support for the improved learning of the Same rule under incidental conditions when additional stimuli reflect the previous response required as being due to the contingent relationship between stimulus and sequence.

This study also provides evidence that No-Colour, Random and Current groups learn sequences at the same rate. This suggests firstly that participants do not seem to be affected by increasing the level of discriminability between stimuli, as in the Current group. Having the purple and yellow stimuli to further separate the representations of left and right response-stimuli did not improve performance on the task compared to a control group, nor did it reduce the learning as predicted by the RASRN. Similarly, adding extra stimuli to the task had no negative effect on learning, as in both Current and Random cases participants were no worse than the No-Colour group. This suggests that these additional stimuli did not distract from sequence learning and therefore this suggests that we have evidence not of any sequence learning disadvantage in any group, but evidence that the Previous groups can learn more about Same rule sequences.

4.5. General Discussion

Whilst adding additional stimuli into the SRT task was shown to have an effect on sequence learning, Experiment 4 suggests that random noise or additional attentional demands have no effect on learning as there was no difference between Random and No-Colour groups in learning difference scores. This provides evidence that the occurrence of these new stimuli did not alter learning just because new stimuli were introduced, reflecting the claims of Abrahamse et al. (2012) regarding redundant stimuli. The Previous group showed significantly more learning of the Same rule sequence than the Random group in Experiments 3 and 4, whereas the Current group was not significantly better than the Random group in Experiment 3 (although it was numerically so) nor either Random or No-Colour groups in Experiment 4. Thus, while I have demonstrated the effect expected as a result of RASRN predictions, that the Current group would be worse than the Previous group, it is not because the Current group was impaired. Instead my results suggest that additional stimuli that have a relationship to the previous elements in a sequence seem to have an effect on learning, even though these stimuli did not provide predictive information about the current trial in themselves.

There are a number of possibilities for this effect, the first being that the Previous condition produces a different level of concentration or attention to the task, as the stimuli (tones or colours) have no contingent relationship with the current response stimuli. This cannot be the case, as the Random condition acts as a control for this manipulation and demonstrates no improvement to learning in either Experiment 3 or 4. Therefore, the Previous group does not simply provide an environment that encourages greater concentration. A further possibility is that non-contingent stimuli produce some increase in error, as there are no contingencies between the two current stimuli (tones/colours and response locations) in Random or Previous groups, however, as the advantage was only observed in the Previous group (and not in Random) this again cannot explain the Previous learning advantage.

A further suggestion is that the pattern of new stimuli (tones or colours) follows the pattern of on-screen stimuli with a lag of one trial, which means that the sequence is experienced twice. This could suggest that participants hear or see a sequence of the same precise order on more than one occasion and this gives them twice the amount of opportunity to encode or remember the subsequence. This suggests that the Previous

group concurrent stimuli are in a sense, non-redundant, as they do offer additional information about the task on each trial, even though this is not directly useful it may be made use of by an associative system. Instead of all stimuli and contingencies going into the same model with some recurrence or memory; this suggests that learning could be in some senses isolated, with one set of stimuli encoded separately to another (e.g. Cleeremans, 1997; Destrebecqz & Cleeremans, 2003). However, this explanation would also require a model that could learn sequences of stimuli separately as the stochastic structure of the Same rule could be extracted from both response stimulus locations and tones or colours, increasing learning of this sequence. This suggests that a model would require separate sets of recurrent networks that sum together (rather than competing) to predict the next trial leading to an increase in the overall representation of the Same rule.

It is suggested that the visual system contains functionally separate areas that encode for spatial features such as location or orientation, while another part of the visual cortex encodes for stimulus properties (e.g. colour), which are dissociated and encode informational separately (Ungerleider & Mishkin, 1982) therefore, it could be entirely possible that the contingencies across stimulus presentations may be restricted to particular dimensions. However, if this was the case then the Current group, which also involves two sets of stimuli or dimensions that both follow the sequential rule should also see an advantage, which they do not. We could instead assume that the Current group's did not represent response stimulus locations and additional concurrent stimuli as two distinct stimuli, as these are contingent and therefore may be bound together and represented configurally at input. If the response stimuli and new stimuli (tones or colours) were not represented locally, but represented as compound stimuli (for example: right-purple, right-yellow; left-purple and left-yellow) then the Current group will receive input from only two units (essentially matching the No-Colour case), while the other groups will receive input from four units. When these four compound representations are used as input to the RASRN, rather than local representations for stimuli locations and colour, this does not produce the differences observed in humans, with the Random group still performing numerically the best on Same rule learning, with no improvement in Previous group responding (see Figure 4.14).

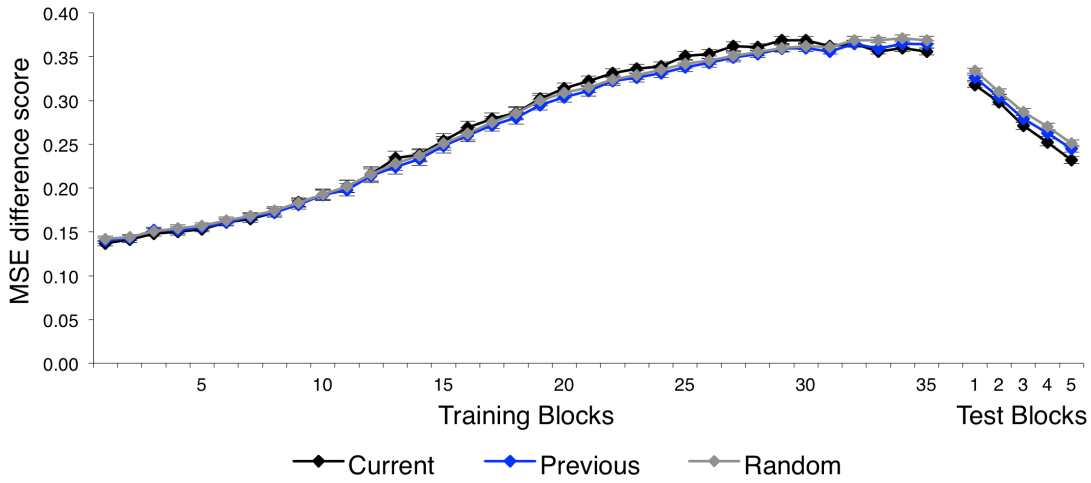

Figure 4.14. MSE difference scores for the RASRN simulations of Current (black), Previous (blue) and Random (grey) conditions for networks trained on Same rule sequences when four units represent compound stimuli (e.g. right-purple) rather than local representations for location and colour, across training and test blocks. Error bars show standard error.

Taken together, it could be possible that Previous groups are given an advantage in sequence learning as two lots of sequence learning sum together; but that representing the Current stimulus conditions as compound stimuli or perhaps the influence of stimulus-stimulus or stimulus-response associations could reduce some advantage that dual-sequence learning occurring may provide. This task does not, however, allow the examination of learning about tones or colours independently of the influence of learning about response-stimuli locations. Whilst the RASRN predicted that the stimuli would simply reduce sequence learning, it is possible that participants then were able to learn about the additional concurrent stimuli and this could in fact increase sequence learning. Without being able to disentangle these two stimulus types it is not possible to understand how much participants may have learned about either stimuli and how these may then interact with one another. A task design is required where dual-stimuli can be separated so that we can investigate how learning about different task elements progresses, and how a relationship between sequences and other stimuli might alter sequence learning.

What can be concluded from these data is that in humans it seems that additional stimuli have an effect on responding, but this was restricted to stimuli that related to the previous element in a sequence of trials experienced. Therefore I have provided evidence that in a sequence learning task the discriminability of stimuli, properties of the stimuli and random variation in some stimuli are not important in sequence learning

itself. However, when these stimuli provide additional information about the sequence (i.e. provide a representation of the previous trial) then participants can use this to significantly improve sequence learning, at least in this case. That this is not observed in the Current group suggests that participants may be able to learn some additional aspect of the sequence in the Previous group, or that stimulus-stimulus or stimulus-response contingencies in the Current group block this effect. However, conclusions regarding the reason for this are tentative, as we are unable to assess whether participants learned anything about the relationship between tones and responding, either with the previous trial or how this impacted on the current trial. Rather than a cue competition effect, then, here we have observed that additional stimuli can produce potentiation of learning. To investigate how additional stimuli interact further with sequence learning, Chapters 5 and 6 adapt the SRT task to include additional colour cue stimuli that are predictive in their own right. In this way I attempt to further investigate how and whether these stimuli are learned about and how they might come to interact with one another.

Chapter 5. Implicit cue-response learning

Chapter 4 introduced additional stimuli to the SRT task potentiated sequence learning when these new stimuli provided information about the previous element in the sequence. To investigate this further, the two-choice SRT task used in Experiments 1 and 2 was adapted to include a separate cue that occurred before the presence of the response stimuli. The aim of this was to create a task that could produce learning about cue-response relationships that could then be placed alongside sequential contingencies to investigate the interaction between the two. In this Chapter this task is piloted without sequential contingencies in order to find evidence of cue-response learning. Participants in Experiment 5 demonstrated learning under incidental conditions about a central colour cue that predicted a response-stimulus location. This learning was simulated across training by the RASRN. A group of participants who completed the task under intentional conditions provided a prior probability for explicit knowledge on the task, and also themselves demonstrated greater learning in Experiment 6. The Intentional performance, when used as a prior in a Bayes factor analysis, provided evidence for the null in Experiment 5: that cue-response learning occurred without awareness of these contingencies. Altogether I provide evidence that participants are able to learn simple cue-response contingencies and this task is therefore suitable for use in a dual-cue version alongside sequential contingencies to investigate the interaction between cue-response and sequence learning in Chapter 6.

5.1. Introduction

Chapter 4 provided evidence that participants were able to learn more about sequences when an additional stimulus on each trial provided information about the previous element in the sequence. As this potentiation of sequence learning was not found when these stimuli were random, this suggests that additional sequential information increased learning. However, as this potentiation was also not seen when the stimuli followed the pattern of the response locations on the current trial, this suggests that a more complex relationship may have been occurring between stimuli and responses. What this relationship was, however, is hard to say as the tones or colours and response-stimuli locations correlated with one another throughout training and test. It was

therefore not possible to assess learning about the relationship the additional stimuli had with responses and isolate this learning from learning about sequences themselves.

In order to understand how the Previous group stimuli were able to interact with and increase sequence learning, we need to be able to separate them in some way, and therefore the additional stimuli would have to be learned about in their own right as well as being related to the sequence itself. If one were able to investigate how simple cue-response associations were formed under incidental conditions then it would be possible to investigate learning about the stimuli and separate this from and understand the interaction with sequence learning. Rather than the stimuli simply correlating with the sequence of pre-existing response locations, stimuli that have their own contingencies with responses could be introduced alongside sequence learning. Therefore, it would be possible to separate learning about these contingencies in a test phase and establish how they interact with one another. These experiments are described in Chapter 6, but first a task design was required that could accommodate sequential contingencies but would demonstrate the effect of simple instrumental cue-response contingencies.

There have been a number of studies that have investigated dual-cue SRT tasks in sequence learning, where not only are there multiple stimuli, but these stimuli are predictive in their own right. Cleeremans (1997) designed a four-choice SRT task where participants showed evidence of incidental sequence learning. The sequence was based on the location of response-stimuli, which could themselves also be one of four colours. These colours also provided predictive information about the location of the next trial and participants learned these cue-response contingencies; but they were also instructed to look for them and explicitly made aware of their presence. This is an obvious issue if I want to examine the influence of automatic, associative learning processes on one another, hence a task is required on which participants demonstrate incidental cue-contingency learning.

In an implicit version of Cleeremans' (1997) task, Jiménez and Méndez (1999) found no learning about cue-response contingencies under incidental conditions, and as such were unable to make any inferences about the effect that cue-response learning may have had on sequence learning (or indeed, vice versa). This chapter attempts to investigate whether people can learn cue-response contingencies incidentally, before

examining this learning within the context of a sequence learning task in the following chapter. However, few tasks are available in the associative literature to base this on, as participants in simple instrumental conditioning experiments often become or are made aware of the contingencies (Perruchet, Cleeremans,, & Destrebecqz, 2006).

5.1.1. Designing a hybrid SRT task with instrumental and sequence learning

As mentioned in Chapter 1, the problem with many studies of simple associative learning in humans is they make it so easy for the propositional system to do all the work. Mr X has avocado and is sick. Mr X has bannana and is not sick. There is no need for an incremental build-up of relationships between contingencies when humans are perfectly capable of deducing these relationships rationally using some conscious reasoning system (Beckers et al., 2006; Mitchell et al., 2009). Whilst authors increase the complexity of the task by increasing the number of stimuli and outcomes (Le Pelley et al., 2005), to some extent any evidence of learning might be driven by explicit knowledge of one or some subset of contingencies within the entire sample. Furthermore, whilst authors claim that elaborate cover stories can in fact mask contingencies that participants make predictions about (Vadillo & Matute, 2010), there remains the problem of using real-world scenarios that attempt to mask the aims of the study but may in fact produce an influence of explicit expectations on learning (Perruchet & Pacton, 2006; Waldmann & Holyoak, 1990; 1992).

The SRT task provides an suitable paradigm to study complex sequential contingencies that participants are unaware of: as they are required to attend to the stimuli as they use these to respond, but the contingencies between previous and current trials are not made explicit. The task instructions therefore do not require that attention is drawn towards trial order or sequential rules that then must be covered up, as task instructions can simply refer to simple key-press responses and performing optimally. The cover story of the experiment is not intended nor required to mask any contingencies that participants can learn to predict. The instructions; conditions under which learning occurs; and the measure of learning all therefore do not require contingencies to be made explicit, which are the principles on which the design of this task was based.

The two colours used in Experiment 4 had an effect on learning that avoided the issues associated with the perceptual differences between tones and visual stimuli. Therefore,

the task was designed to include different colours that would be presented to participants with the aim of training incidental differential instrumental conditioning. The adapted SRT task in Experiment 4 may have demonstrated modified sequence learning as a result of these two colours, but as these were presented concurrently with the response-locations the RASRN would predict small learning of these cue-response contingencies due to the reduced activation at input. Traditionally, when conditioning responses, the onset of the cue or CS presentation occurs before the onset of the US (the response stimuli), as in Perruchet et al. (2006). Therefore the cues introduced to this task were made to occur before the presentation of the response stimuli.

Practically, if the colours were presented before the response stimulus locations filled then they could not be placed within the stimulus location as in Experiment 3, as this would evoke a response. Consequently, the experimental design of Aitken (1996) was used as the basis for the experiment, where participants were required to respond on a two-choice SRT task with response locations at either side of the screen. Participants were instructed to simply respond to stimuli but that before one of these appeared; a shape would appear in the centre of the screen that, unbeknownst to them, provided information about which response location would fill. Aitken (1996) used three shapes: a star; cross, or wedge, with the star or cross as controls that had no relationship with whether that trial would require a right or left response. The wedge stimuli were drawn from a set of categorization stimuli so that, depending on the length of the radius and angle of the wedge, they formed two categories that were each perfectly predictive of a response location (e.g. Category 1 stimuli always occurred before a right response location).

Whilst not intending to study implicit cue-response learning per se, Aitken (1996) provides evidence that participants were able to learn that (a certain category of) cues predicted a response in the absence of awareness (as indexed by a questionnaire). Therefore, the SRT task used in Experiments 1 to 4 was adapted to include a central square stimulus that would fill with a colour before the occurrence of a white circle response stimulus on either side of the screen. The parameters of Aitken's (1996) experiment were followed, as detailed below, with the aim of producing simple contingency learning between colour cues and responding. I was also concerned that these contingencies would be far easier to notice than sequences, and so included direct

tests of explicit knowledge more sensitive than asking participants to verbally describe anything they had noticed.

5.2. Experiment 5: Incidental colour cue-response learning

In this experiment the SRT task was adapted to introduce a cue that occurred before the response-stimuli, similar to the experiment designed by Aitken (1996). These cues were different colours that filled inside of a white square outline in the screen centre. Some of the colours partially (80%) predicted the response location. I chose to include a number of control colours to avoid explicit recognition of contingencies. This initial experiment also attempted to maintain the design continuity between Experiments 1 to 4 as the aim of introducing the cue-response contingencies was to eventually investigate how these would interact with sequence learning. In order to test participants' awareness of the colour contingencies with response-stimuli locations both a post-experimental interview and prediction task were given to participants after the task had finished. The results of which are provided here, but discussed in more detail in section 5.5.

5.2.1. Method

5.2.1.1. Participants

16 participants (aged between 18 and 24 [M = 19.3]; 12 female and 4 male) were recruited from undergraduate students and were awarded £5 in return for participation. Participants provided informed consent prior to taking part in one session lasting roughly one hour.

5.2.1.2. Materials and stimuli

The experiment was run on an Apple iMac with MatLab and Psychtoolbox software. Participants were seated roughly 50 cm from the screen, which contained three white shape outlines on a black background throughout the task. A white outline of a square, 19 mm in height and width, was present in the center of the screen, with two white circle outlines (also 19 mm in diameter) either side of the square. The circles were positioned in line with the square outline vertically either and 22 mm either to the left or the right of the square outline horizontally in line with experiments 1 to 4.

The cue stimulus was a coloured filled square (18.5 mm height and width) that was placed within the white square outline in the centre of the screen, giving it the appearance of lighting up or filling in. This cue could be one of four colours in experimental blocks: red (RGB: 255,0,0); green (RGB: 0,255,0); blue (RGB: 0,0,255); or yellow (RGB: 255,255,0). The response stimulus was a coloured filled circle (18.5 mm diameter) that was placed within one of the two circle outlines, giving the white circle outline the appearance of lighting up or filling in. The colours were the same as for the cue and matched this colour on each trial. Participants were required to press the spatially compatible 'x' and '>' key presses on a QWERTY keyboard to the left or right response stimulus, respectively.

5.2.1.3. Design

The experiment was a two-choice SRT task comprising one session of twenty blocks. These blocks each contained 120 trials, and so the length of training and number of trials differed from Aitken (1996) and these trial numbers were chosen to match Experiments 1 to 4. The first fifteen blocks were training, and the final five acted as test.

Colour contingencies. Colour contingencies with certain responses are shown in Table 5.1. All of the four colours were equally likely within each block and occurred with the same frequency across the experiment. Within the training blocks, two of the four colours positively predicted a certain response (Predictive); and two colours had no positive contingency with a certain response (Non-Predictive), occurring with equal likelihood within and across blocks on right and left response trials. Returning to the Predictive colours, these had an 80% probability of occurring before a certain trial location across training. This followed Posner & Snyder (1975) and meant that the conscious detection of contingencies was made that much more difficult.

The colours themselves (red, yellow, green, blue) were randomly allocated to each experimental Colour for each participant. Colour 1 predicted a left response on 80% of trials, which equates to 24 of the 30 Colour 1 trials per training block (360 across the training). Colour 2 predicted a right response on 80% of trials, and therefore occurred 24 times per block before a right response was required (360 across training blocks). Colours 3 and 4 each occurred 225 total times before a right response and on 225 trials before a left response was required across training. Overall, the number of non-

predictive colours meant that the percentage of trials on which a participant was able to correctly predict the response-stimulus location given the colour on each trial of training was 65%. At test all of these contingencies became 50:50, and each colour was equally likely to predict either response. All blocks included an equal number of right and left response-stimuli and the number of repeats and alternations was controlled with a random sequence of response-stimulus locations that had no relationship to the colours.

Table 5.1. Number and percentage of trials that each of the four Colours across the task in Experiment 5 co-occurred with right or left response stimulus locations.

	Co-occurrence with left response-stimulus		Co-occurrence with right response-stimulus	
	Trials per block	% of total trials for that Colour	Trials per block	% of total trials for that Colour
Predictive				
Colour 1	24	80%	6	20%
Colour 2	6	20%	24	80%
Non-Predictive				
Colour 3	15	50%	15	50%
Colour 4	15	50%	15	50%

5.2.1.4. Procedure

As in previous experiments, participants were instructed to simply respond as quickly and accurately as possible. The only difference in instruction was that they were instructed to fixate on the coloured square, suggesting that it would help them to respond more quickly and more accurately by attracting their attention to the centre of the screen. They were informed that the experiment was intended to measure their reaction times and errors and its aim was to investigate people's ability to respond quickly and accurately to very simple stimuli over the course of an experiment. They were not informed of any relationship between the colours of the stimuli and the responses required and no mention was made of anything to learn about or from. They were told that the colours changed to make the task less dull and to try and retain their attention.

At the beginning of each block participants were instructed to press any key to start. Each trial began with an inter-trial interval of 250 ms where a black background with a white square outline and two white circle outlines was presented. The cue stimulus (a

coloured square in the centre of the screen) would then appear for a variable interval of between 250-500 ms (following Aitken, 1996). The response stimulus (the left or right white circle) would then appear on screen until either the participant made a keypress response or the trial timed out after 4000 ms from the presentation of the response stimulus. RTs were measured from the onset of the response stimulus. If participants pressed an incorrect key, or the trial timed out, the computer issued a beep sound.

Feedback with average RTs and percentage of errors was given at the end of each block and there was a thirty second enforced break, as described for Experiment 1. A short verbal structured interview was given at the end of the session, in which participants were asked about what they had noticed in the experiment regarding the colours. Participants were asked to describe any contingences they may have noticed, and then were asked to explicitly guess which response each colour predicted after being told that two colours were predictive and that one of each of these predicted a left or right stimulus.

After I asked participants whether they had noticed anything about the experiment I asked them to complete a simple prediction task whereby the square in the centre would fill in with a colour and they would have to respond with the keys used in the experiment, although the circle would not fill. They were not put under time pressure to do this and were instructed to simply use their intuition, a guess, or any knowledge they might have about the task to select their response. Participants responded for two blocks of 32 trials with a 250msec RSI, within which each of the four colours were presented eight times. The order of presentation was randomised and there was no feedback given during or at the end of these two short blocks, and participants were informed that when they made a response nothing would happen and neither response stimulus would fill. Participants were finally debriefed and thanked for their participation.

5.2.2. Results

5.2.2.1. Cue-response learning

Inclusion and exclusion criteria for RTs and errors are detailed in Experiment 1, and averages for these were calculated for each of the three Colour Types. These were Predictive Consistent (the 80% of Predictive trials that were consistent with the trained contingency); Predictive Inconsistent (the 20% of Predictive trials that were inconsistent

with the trained contingency); and Non-Predictive (all trials for Colours 3 and 4). An ANOVA was conducted for both average RT and proportion of errors (note: not a difference score) with Colour Type and Block as factors. As before, all within-subject main effects and interactions are reported with a Huynh-Feldt correction to adjust for departures from sphericity, however, the uncorrected degrees of freedom are reported.

Training phase. The main effect of Colour Type was significant across training in RTs, $F(2,30) = 11.2$, $p < .001$, $MSE = 205.1$, $\eta_p^2 = .428$; and errors, $F(2,30) = 9.68$, $p = .001$, $MSE = 002$, $\eta_p^2 = .392$, see Figure 5.8. Participants demonstrated an ordinal pattern of responding in RTs and errors as one would expect to observe if learning had occurred, which is supported by a set of planned contrasts comparing the Colour Types, shown in Table 5.2.

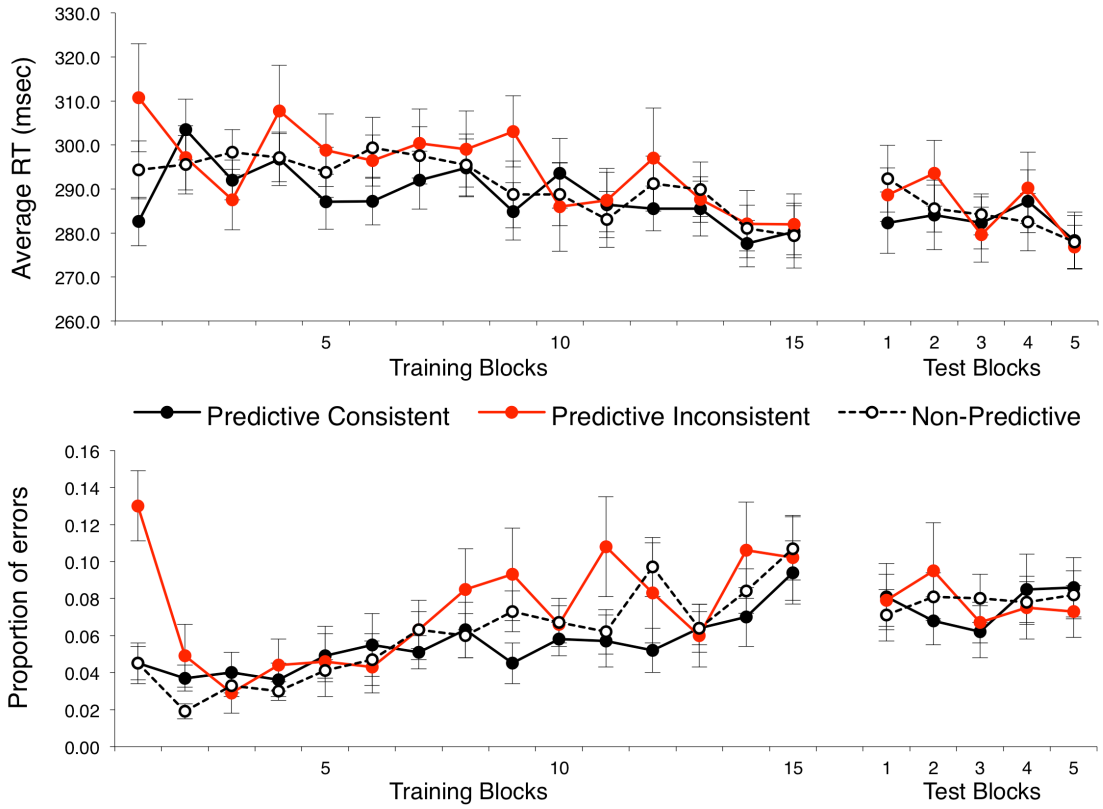

Figure 5.1. Average RTs (top panel) and proportion of errors (bottom panel) for Colour Types across the training and test blocks of Experiment 5. Lines show performance on Predictive Consistent (black filled circles); Predictive Inconsistent (red filled circles); and Non-Predictive (black open circles) colours. Error bars show standard error.

These provide evidence that Predictive Consistent trials are responded to more quickly and accurately than Predictive Inconsistent trials, which are significantly slower than Non-Predictive trials. Predictive Consistent trials are also responded to significantly faster than Non-Predictive trials (again, see Table 5.2) although this trend was not significant in the errors. This provides strong evidence of learning about the contingencies in place between colours and response-stimulus locations.

Table 5.2. Results from ANOVA for planned comparisons on average RT and proportion of errors, comparing the three Colour Types to one another across training.

	Average RT: Training			
	F (df = 1,15)	p	*MSE*	η_p^2
Predictive Consistent vs Predictive Inconsistent	19.0	.001	485.0	.559
Predictive Consistent vs Non-Predictive	5.65	.031	362.8	.274
Predictive Inconsistent vs Non-Predictive	6.70	.021	382.9	.309
	Proportion of errors: Training			
	F (df = 1,15)	p	*MSE*	η_p^2
Predictive Consistent vs Predictive Inconsistent	14.4	.002	.006	.491
Predictive Consistent vs Non-Predictive	1.24	.283	.005	.076
Predictive Inconsistent vs Non-Predictive	12.1	.003	.004	.446

There was also a main effect of Block in RTs, $F(14,210) = 2.66$, $p = .005$, $MSE = 930.9$, $\eta_p^2 = .151$, and errors, $F(14,210) = 5.37$, $p < .001$, $MSE = .002$, $\eta_p^2 = .392$, which is shown in Figure 5.1 as a general trend towards faster and less accurate responding across the task. Block did interact with Colour Type in the errors, $F(28,420) = 1.86$, $p = .020$, $MSE = .004$, $\eta_p^2 = .110$; but not RTs, $F(28,420) = 1.34$, $p = .176$, $MSE = 607$, $\eta_p^2 = .082$, which highlights the increasing errors made to Predictive Inconsistent colours across training; supported by a significant linear interaction contrast comparing Previous Inconsistent and Non-Predictive errors, $F(1,15) = 5.97$, $p = .027$, $MSE = .004$, $\eta_p^2 = .284$; and non-significant trend between Predictive Inconsistent and Consistent errors, $F(1,15) = 4.32$, $p = .055$, $MSE = .005$, $\eta_p^2 = .223$.

Test phase. There was no evidence of a main effect of Colour Type at test in RTs, $F(2,30) = .862$, $p = .433$, $MSE = 201.7$, $\eta_p^2 = .054$, nor errors, $F(2,30) = .077$, $p = .926$, $MSE = .001$, $\eta_p^2 = .005$, although the numerical order of RT results follows the pattern across training and the error pattern only deviates in increased errors to Non-Predictive stimuli (see Figure 59). However, there was a trend towards significance in the first two

blocks of test in RTs, $F(2,30) = 3.16$, $p = .070$, $MSE = 209.6$, $\eta_p^2 = .174$, which resulted from significantly faster responding to Predictive Consistent compared to Predictive Inconsistent colours (Bonferroni corrected), $F(1,15) = 13.2$, $p = .004$, $MSE = 151.3$, $\eta_p^2 = .469$, supported by a non-significant effect in the same direction in the errors, $F(1,15) = 1.09$, $p = .312$, $MSE = .004$, $\eta_p^2 = .068$, see Figure 5.1. The effect size of the RT difference apparent across the first two blocks of test suggested that participants still showed a strong preference to respond quickly (and not less accurately) to Predictive Consistent stimuli, therefore suggesting that some learning remained during extinction. However, this learning did not survive across test and rapidly extinguishes.

5.2.2.2. Direct tests of explicit knowledge

A full consideration of the direct test results is reported in section 5.4.2 alongside results from participants who completed the task under intentional conditions, however, they are reported here for descriptive purposes. On the post-training interview where participants were asked to identify the two predictive colours and report which location they predicted, three participants were able to identify both colours correctly, with two of these participants identifying the correct location. Ten further participants were able to identify one correct colour (five correctly identifying the left colour and five the right); although only half of these participants were able to correctly identify the location (two left; two right). Three participants were unable to report any correct contingencies. This may seem like an alarming number of incidental participants were able to correctly identify the predictive colours, however participants have a one-in-six chance of getting both colours correct by chance and two-thirds chance of getting one colour correct by chance. They are, therefore, just as likely to get both colours correct as to be unable to correctly identify either colour if they are guessing.

On the prediction task there were 32 trials in total, with half of these concerning the Predictive colours and therefore discussed here. Participants scored on average 8.25 correct responses on the task ($M = 8.25$, $SE = .727$), which is shown in Figure 5.2. These correct responses are shown here split for right and left colours, as participants may have performed with eight correct responses out of sixteen, but all of these may have been about one colour and learning driven by knowledge about this single cue-response contingency. These results are analysed further in section 5.4.2.

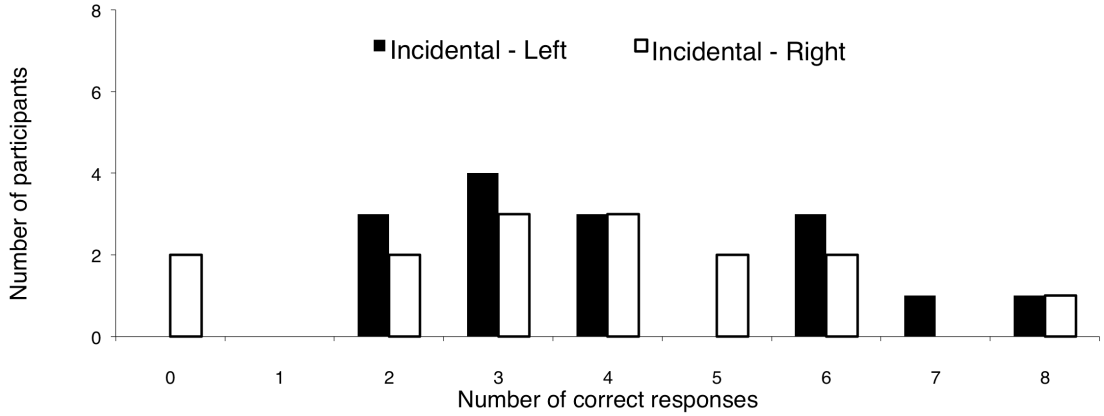

Figure 5.2. Graph showing the number of participants who made each number of correct responses out of the eight total possible for both Predictive Colours 1 (Left) and 2 (Right) on the Experiment 5 Prediction task.

5.2.2.3. Subjective tests of explicit knowledge

Participants were asked whether they had noticed anything about the task and whether, specifically there was any relationship between the colours and responses. Their responses to these questions were sorted into five categories: that they were completely sure that there were contingencies and sure about their identities; that they were confident there were contingencies between colours and responses but could not be confident in saying which colours; that they thought perhaps something may have been going on in the task, but not sure what it was; that they noticed some things but thought it was random; that they did not notice any contingencies and would be surprised to hear there were. Of the 16 participants, four responded that they thought perhaps something was going on in the task, but that they couldn't be sure about it nor identify what the colours or contingencies may have been. The rest of the participants were surprised that there may have been any relationships and all reported that they thought that it was random. Whilst this is a crude measure of subjective explicit knowledge, I found no difference between the two groups on their prediction task score, $t(14) = .783$, $p = .446$, $SED = 1.70$; which goes some way to support the claim that they were subjectively unaware of these contingencies. Whilst the questions did not refer precisely to what degree participants said they were guessing on the task (Cheesman & Merinkle, 1984), they did not correlate with explicit prediction task performance (Chan, 1992), $r(16) = .205$, $p = .446$, although this measure may lack the required sensitivity to accurately conduct this analysis.

5.2.3. Discussion

Altogether Experiment 5 shows evidence of cue-response learning about how colours predicted a response in the SRT task. Participants showed learning across training but only some evidence for learning in the first two blocks of test. This suggests that the colour cue-response stimulus location contingencies were susceptible to rapid extinction. Participants were trained for only 360 Predictive Consistent trials across the experiment for each colour (compared to 90 Inconsistent trials), and alongside 900 Non-Predictive colour trials. This suggests that whilst this was enough training to produce a learning effect across the task itself, that under test conditions where these cue-response associations are removed that participants do not continue to respond with reliably faster and more accurate responses to the Predictive Consistent response location. That this effect reduced rapidly at test was taken into account in Chapter 6, where the length of test was reduced and length of training increased.

Whilst this cue-response learning occurred in the absence of any intention to learn, I wanted to ensure that this cue-response learning developed in the absence of any awareness. Following the suggestions of Z. Dienes (personal communication, August 3, 2012; Dienes, 2014; in press) an Intentional cue-response experiment was run, whereby participants completed the same task but under explicit task instructions. In doing so, the participants in this group would perform the explicit tests of knowledge based on their consciously acquired knowledge.

5.3. Simulation 10: RASRN simulation of cue-response learning

The RASRN using the parameters described in Chapter 3 was again used to predict human performance under incidental conditions, however, this time on a non-sequential learning task. Cleeremans (1997) investigated learning about sequences and cue-response learning and modelled the cue-response component separately to the recurrent network used for sequence learning. In doing so Cleeremans (1997) was able to show evidence that this adapted SRN could learn both sequences and cue-response contingencies, but without any interaction between the two. As I intend to investigate whether cue-response contingencies do, in fact, interact given evidence from Chapter 4, I do not start from the position that this learning occurs as a result of a functionally separable system. Given that the addition of a separate learning system is not

parsimonious I suggest that the RASRN should be able to learn simple cue-response contingencies as well as more complex abstract rules.

5.3.1. Simulation details

The simulation details were mostly as for Chapter 3: Simulation 9 with the slow and fast learning rates set at 0.2 and 0.5, respectively. There were 20 hidden units in the model and therefore 20 context units, as well as 2 output units representing the prediction of the location of *t*. Additionally, there were two input units representing the previous required response ($t - 1$) and two units representing the current on-screen response stimuli (t) there were four additional units to represent the four colours described in Experiment 5. Because these stimuli occurred before the presence of the current on-screen response stimuli but after the previous response, I gave them activation values of 0.4 for on and 0 for off. The previous response (0.75), current stimulus (0.1) and context units (1.3 times the hidden unit activations) remained the same as for Simulation 9. The model was used to simulate the task experienced by human participants in Experiment 5, with 2400 trials run for each network, and 16 total networks run to simulate the 16 participants in Experiment 5. As in the human experiments there were no sequential contingencies in the trial order that networks were trained on, the colour units were activated according to the contingencies outlined in Experiment 5.

5.3.2. Results

The MSEs for each trial were averaged across each Block and Colour Type, as for human average RT and proportion of errors. As with human performance, MSEs were not further manipulated to produce an index of learning, therefore learning can be assessed by the difference between the Colour Types, with lower MSE scores showing less difference between the model's predictions and the next trial location. The results of the simulation are shown in Figure 5.3, and show a significant effect of Colour Type across training, $F(2,30) = 64.4$, $p < .001$, $MSE = .001$, $\eta_p^2 = .811$; and test, $F(2,30) = 199$, $p < .001$, $MSE = .001$, $\eta_p^2 = .930$. It is clear that the networks performed gradually better on Predictive Consistent trials across the experiment and that performance on Predictive Inconsistent trials became steadily worse, supported by a significant Block by Colour Type interaction across training, $F(28,420) = 5.33$, $p < .001$, $MSE = .001$, η_p^2

= .262. This interaction was not present at test, $F(2,30) = .756$, $p = .642$, $MSE = .001$, $\eta_p^2 = .048$, and the networks showed no evidence of extinction.

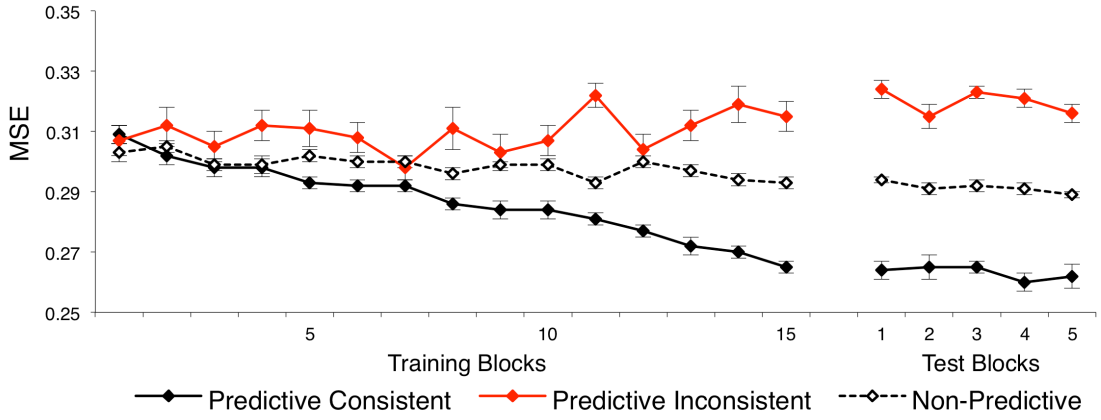

Figure 5.3. MSEs for Colour Types across the training and test blocks of the RASRN simulation. Lines show performance on Predictive Consistent (black filled circles); Predictive Inconsistent (red filled circles); and Non-Predictive (black open circles) trials. Error bars show standard error.

5.3.3. Discussion

The RASRN can learn these cue-response contingencies as the humans do across training. However, the networks show strong performance in the absence of extinction across test which was not observed in human participants. Whilst the RASRN can again account for human learning, this time of cue-response contingencies, it does not fully model the pattern of human behaviour. As discussed in the previous chapter, the RASRN seems to be missing some elements or parameters that would help it account for human behaviour, but these will be discussed further in Chapter 6. For the purposes of simulating Experiment 5, the RASRN produces an adequate simulation of human performance.

5.4. Experiment 6: Intentional cue-response learning

Primarily, the aim of the Intentional version of the task was to provide direct knowledge of test performance priors in order to compare Incidental performance to the plausible performance that we would expect from participants with explicit knowledge of cue-response contingencies in order to assess evidence for the null (Dienes, 2011; 2014). In doing so, the participants were run on a matched experiment, as even though explicit knowledge was expected to develop quickly in the task. It was my intention that

participants should experience the same task demands in terms of length and number of trials, as well as the pseudorandom test phase, in order to provide a complementary condition with which to assess the absence of explicit knowledge in the Incidental learning condition. Therefore, the tasks differed only in their instruction and all other details were matched entirely.

5.4.1. Method

5.4.1.1. Participants

16 participants (aged between 18 and 26 [M = 19.4]; 15 female and 1 male) were recruited from undergraduate students and were awarded £5 in return for participation. Participants provided informed consent prior to taking part in one session lasting roughly one hour.

5.4.1.2. Materials and Stimuli

Materials were the same as for Experiment 5 and 6, except that participants were provided with a piece of paper at the start of the task with instructions about the task on, shown in Figure 5.4.

5.4.1.3. Design

The design of the experiment followed Experiment 5.

5.4.1.4. Procedure

The procedure followed Experiment 5 and participants were again instructed that the experiment intended to measure speed and accuracy of responding and that the colour acted as a fixation point. They were also told that they would be provided with a clue as to how to use the colours themselves to improve their performance (as shown in Figure 5.4). Participants were additionally required to write down which colours they thought were Predictive and which colours they thought were Non-Predictive at the end of each block on the instruction sheet in Figure 5.4.

There are four different colours in this experiment: red, yellow, green and blue. You will see these all and equal number of times in each block.

ONE x colour will mostly be associated with a RIGHT response
ONE x colour will mostly be associated with a LEFT response
TWO x colours have no relationship with a response: they are equally likely to occur with a right or left response and are therefore NON-PREDICTIVE

The colours are randomly allocated to either predicting RIGHT or LEFT or to being NON-PREDICTIVE at the start of the experiment. These allocations do not change throughout the experiment, so the RIGHT colour will predict a right trial throughout and the LEFT colour will predict a left trial throughout. The NON-PREDICTIVE colours will never predict either trial.

Please write either RIGHT, LEFT or NON-PRED at the end of each block depending on what you think the relationship might be. Again, this is not because these relationships change, we are interested in when you learn this information, so please do not go back and 'correct' your answers to earlier blocks.

Block				
1				
2				
3				

Figure 5.4. Instructions available to participants in the Intentional condition describing the nature of the relationships in the task and additional task requirement to note down what they thought the contingencies were after each block. Spaces to write which colour was which for Blocks 4-20 not shown.

5.4.2. Results

Results were analysed as for Experiment 5 with an ANOVA on RT and proportion of errors across Block, Subsequence and Colour Type.

5.4.2.1. Cue-response learning

Training phase. There was a main effect of Colour Type across training RTs, $F(2,30) = 45.9$, $p < .001$, $MSE = 1668$, $\eta_p^2 = .754$; and errors, $F(2,30) = 18.8$, $p < .001$, $MSE = .022$, $\eta_p^2 = .557$, which are shown in Figure 5.5. This clearly demonstrates the same ordinal pattern as we would expect from participants if they had learned about colour cue-response contingencies, supported by a set of Bonferonni corrected comparisons, shown in Table 5.3.

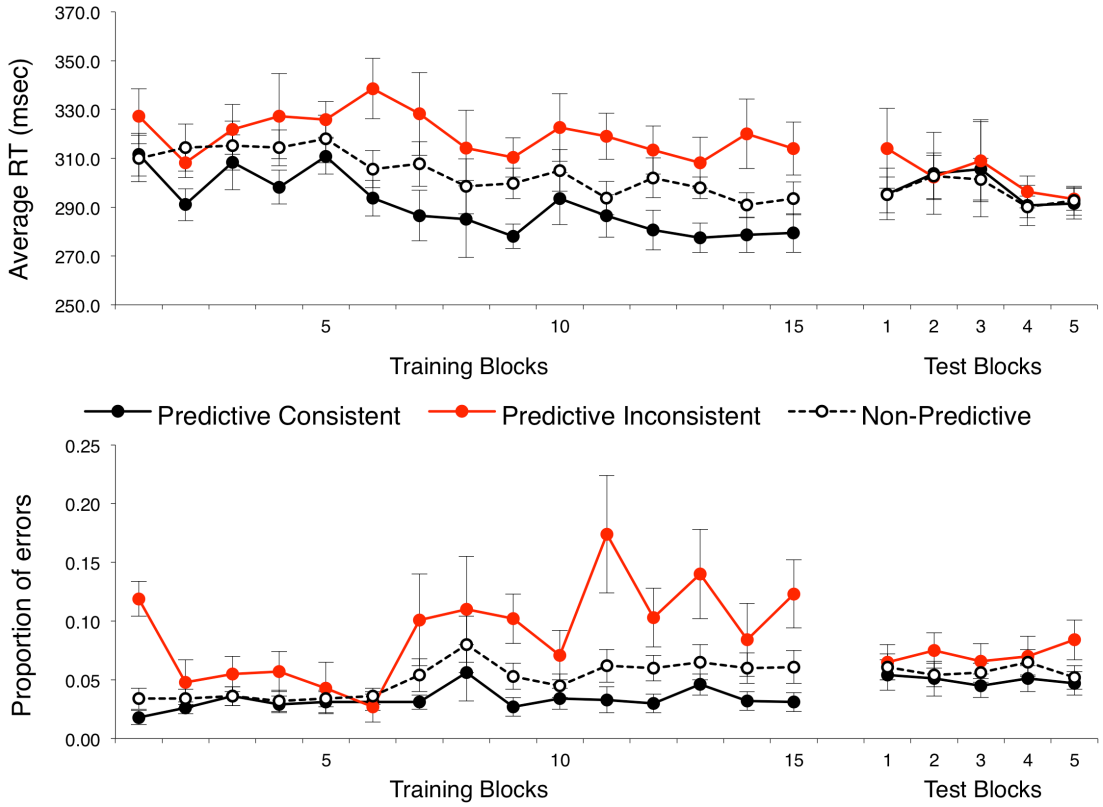

Figure 5.5. Average RTs (top panel) and proportion of errors (bottom panel) for Colour Types across the training and test blocks of Experiment 5. Lines show performance on Predictive Consistent (black filled circles); Predictive Inconsistent (red filled circles); Non-Predictive (black open circles) trials. Error bars show standard error.

These demonstrate significant differences across all pairwise comparisons, with faster and more accurate responding to Predictive Consistent colours, with slower and less accurate responding to Predictive Inconsistent colours, both of which were different to the Non-Predictive colours. Therefore, we have convincing evidence of learning across training. There was a main effect of Block across training (see Figure 5.5) which showed decreasing RTs, $F(14,210) = 2.04$, $p = .044$, $MSE = 2757$, $\eta_p^2 = .119$; and increasing errors, $F(14,210) = 2.66$, $p = .045$, $MSE = .023$, $\eta_p^2 = .150$, across the task. This interacted with Colour Type in errors, $F(28,420) = 2.26$, $p = .028$, $MSE = .011$, $\eta_p^2 = .131$, which demonstrates the development of learning across training.

Test phase. There was a main effect of Colour Type at test in the proportion of errors, $F(2,30) = 4.49$, $p = .023$, $MSE = .003$, $\eta_p^2 = .231$; but not in RTs, $F(2,30) = 1.68$, $p = .207$, $MSE = 698$, $\eta_p^2 = .101$, see Figure 5.5. The Predictive Consistent accuracy advantage was still significantly better than for Predictive Inconsistent colours, $F(1,15)$

= 7.36, p = .048, MSE = .006, η_p^2 = .329, suggesting that some learning in the Intentional group was able to survive extinction. The first two blocks of test did not provide a significant difference in Colour Type in RTs, $F(2,30) = 1.59$, p = .220, MSE = 543, η_p^2 = .096; nor errors, $F(2,30) = 1.63$, p = .217, MSE = .003, η_p^2 = .098, suggesting that rapid extinction also affected participants under Intentional conditions.

Table 5.3. Results from ANOVA for planned comparisons on average RT and proportion of errors, comparing the three Colour Types to one another across training.

	Average RT: Training			
	F (df = 1,15)	p	MSE	η_p^2
Predictive Consistent vs Predictive Inconsistent	52.4	< .001	3933	.777
Predictive Consistent vs Non-Predictive	31.4	< .001	1454	.677
Predictive Inconsistent vs Non-Predictive	42.5	< .001	1359	.739
	Proportion of errors: Training			
	F (df = 1,15)	p	MSE	η_p^2
Predictive Consistent vs Predictive Inconsistent	19.2	.002	.042	.561
Predictive Consistent vs Non-Predictive	17.0	.003	.004	.532
Predictive Inconsistent vs Non-Predictive	18.5	.002	.021	.552

5.4.2.2. Direct tests of explicit knowledge

All of the participants correctly identified both colours and their locations in the post-training interview. Performance was therefore perfect without fail on colour identification, suggesting that under Intentional conditions it was possible for all participants to have full explicit knowledge of the contingencies in the task. Participants also performed well on the prediction task, with the results shown in Figure 5.6. Participants scored 14.4 on average on the prediction task (M = 14.4, SE = .288) which suggests that participants are able to perform well on this task when they have explicit contingency knowledge. It is interesting to note, however, that they were unable to respond correctly on all of the trials even though they reportedly knew which of these responses the colour predicted. Whilst all of the participants were able to make the correct response on more than half of the trials for each predictive colour, left and right, the task seems to suggest that participants may not be entirely sure, or that the influence of the experience of Predictive Inconsistent colours may influence performance on this task (Merikle & Reingold, 1992), or that participants did not give the task their full and effortful attention.

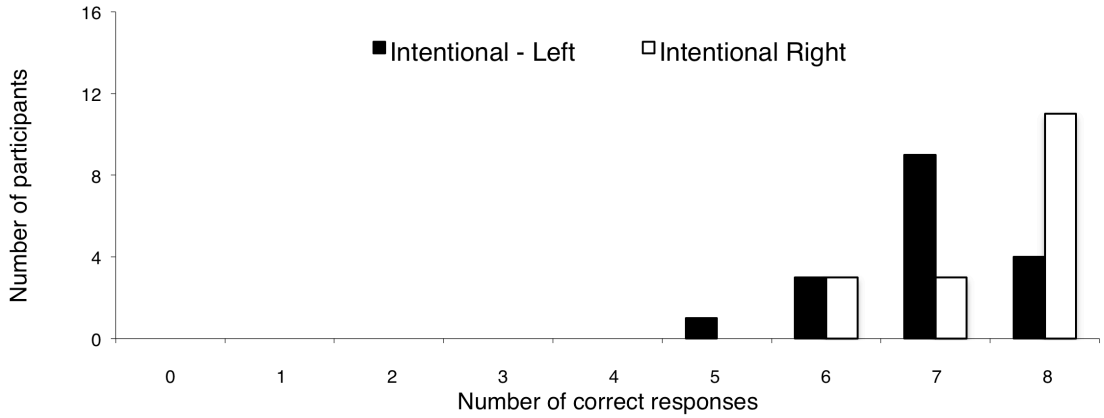

Figure 5.6. Graph showing the number of participants who made each number of correct responses out of the eight total possible for both Predictive Colours 1 (Left) and 2 (Right) on the Experiment 6 Prediction task.

5.4.2.3. Subjective tests of explicit knowledge

The direct tests of explicit knowledge are supported by the subjective questions regarding how confident participants were about the presence of contingencies in the task. Participants all reported that they were sure that contingencies existed and that they were confident in their identification of the predictive colours and locations.

5.4.3. Discussion

Participants who performed the task under Intentional conditions clearly, correctly and confidently showed evidence of explicit knowledge about these contingencies. This was the intended outcome of the instructional manipulation and these results can therefore be used in a Bayes factor analysis to attempt to find evidence for the absence of explicit awareness in the Incidental group in the following section. Interestingly, unlike in Experiment 2, the Intentional instructions seem to have had a strong effect on how much learning was demonstrated by participants, however without a direct comparison we are unable to tell whether this is the case. Participants under Intentional conditions seemed, like the Incidental group, to suffer from extinction during the test phase in their RT and error performance, although a difference between Predictive Consistent and Inconsistent errors remained. Therefore, as well as investigating the Incidental group's explicit knowledge I also compared the indirect measures of learning across the task between the two groups.

5.5. Evidence for implicit learning

Establishing whether participants were learning colour cue-response contingencies in an implicit manner on this task is important. If participants were showing evidence of learning under incidental conditions that was driven by explicit knowledge then attempting to examine any interaction between cue-response and sequence learning may involve two separate processes. Therefore both the indirect RT and error performance across the task and the post-training explicit knowledge tests were compared between incidental and intentional participants. Rather than relying on this difference to produce any conclusive result regarding learning (as in Chapter 2) I decided to use a Bayes factor analysis to find evidence for the absence of explicit knowledge. Using the procedure outlined by Dienes (2014) I used the explicit performance of the Intentional group as a prior probability of aware responding to both post-training direct tests.

5.5.1. Incidental versus Intentional SRT task performance

Intentional learners were compared to incidental learners using an ANOVA on Block, Colour Type and Condition (Incidental versus Intentional). Bonferroni corrections were applied to significance values in order to correct for multiple comparisons. There was no main effect of Condition across training or test, but Figure 5.7 clearly shows the interaction between Colour Type and Condition across training (left panel) in both RTs, $F(2,60) = 24.1$, $p < .001$, $MSE = 1329$, $\eta_p^2 = .446$; and errors, $F(2,60) = 6.80$, $p = .018$, $MSE = .011$, $\eta_p^2 = .185$, with the pattern of learning exaggerated for the Intentional group. This suggests that participants are able to learn more when they actively search for colour-response contingencies. At test all groups show a flattening of the Colour Type effect, which does not show evidence of an interaction with Condition in RTs, $F(2,60) = .710$, $p = .978$, $MSE = 428$, $\eta_p^2 = .023$; nor errors, $F(2,60) = 2.59$, $p = .171$, $MSE = .002$, $\eta_p^2 = .079$, suggesting that both groups suffer from extinction at test.

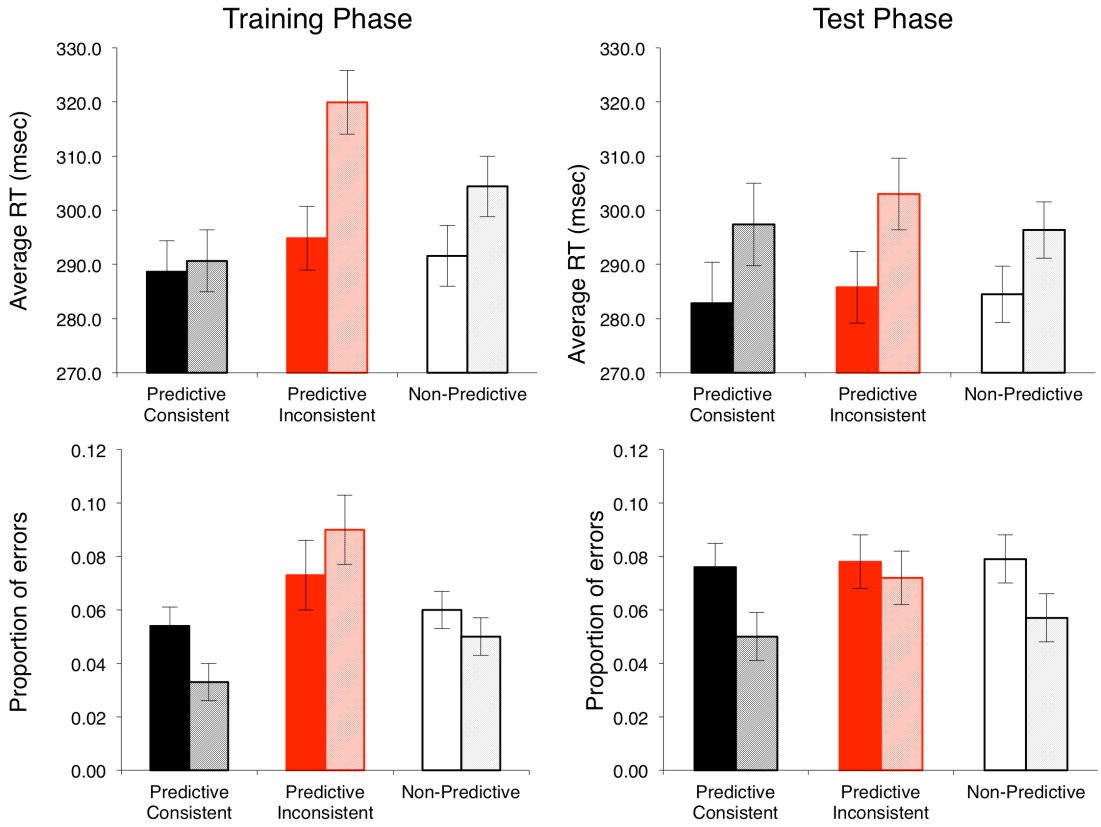

Figure 5.7. Average performance in average RTs (top panel) and proportion of errors made (bottom panel) across for Incidental (solid bars); and Intentional (lighter shaded bars) on Predictive Consistent (black bars); Predictive Inconsistent (red bars); and Non-Predictive (black open/grey bars) Colour Types. Error bars show standard error.

A series of Bonferroni corrected comparisons were conducted on the interaction between Condition and the different levels of Colour Type, which are shown in Table 5.4. These show that across training there was a reliable difference between how well the colour-response contingencies were learned, as the Predictive Consistent – Predictive Inconsistent difference was significantly larger for the Intentional group in both RTs and errors. This was also true of the difference between Predictive Inconsistent disadvantage compared to the Non-Predictive colours, as RTs were far slower and errors more frequent in the Intentional group. Whilst the interaction was significant only in RTs for the difference between Predictive Consistent and Non-Predictive colours; both this and the numerical direction of the error difference provide support to the other training interaction contrasts – that cue-response learning was greater in the Intentional group across training. These effects were, however, eradicated at test, with no difference observed between either of the two Conditions across the Colour Types.

Table 5.4. Results from ANOVA for Bonferroni corrected comparisons on average RT and proportion of errors, comparing the interaction between Condition and the three Colour Types to one another across training

	Average RT: Training			
	F (df = 1,30)	*p*	*MSE*	η_p^2
Predictive Consistent vs Predictive Inconsistent	29.0	< .001	2209	.492
Predictive Consistent vs Non-Predictive	15.6	< .001	909	.342
Predictive Inconsistent vs Non-Predictive	20.7	< .001	871	.408
	Proportion of errors: Training			
	F (df = 1,30)	*p*	*MSE*	η_p^2
Predictive Consistent vs Predictive Inconsistent	7.35	.022	.024	.197
Predictive Consistent vs Non-Predictive	3.49	.141	.004	.104
Predictive Inconsistent vs Non-Predictive	6.88	.027	.013	.187

5.5.2. Evidence for the absence of awareness

5.5.2.1. Post-experiment interview

The results of the post-experiment interview are shown in Table 5.5 and show that all participants in the Intentional condition were able to accurately identify not only the two colours but also their locations. Of the Incidental condition, only three participants were able identify both colours, two of these getting both locations correct. The sample mean for the Incidental group was exactly at chance level of performance on the prediction task (M = 1.00, SE = .153) whilst the Intentional condition performed perfectly on these colour identification questions (M = 2.00, SE = .000). Using the procedure outlined by Dienes (2014) a Bayes factor was calculated, by establishing the prior probability that participants had explicit knowledge about the task taken from the difference between Intentional participants' performance and chance (1.00). As we expect participants in the Incidental experiment to produce less explicit knowledge, it is plausible that any value between chance and the performance of Intentional could occur, and hence a uniform distribution was chosen. Using 0 then as the sample mean (the difference between Incidental performance and chance), a uniform distribution was used to calculate a Bayes factor from chance to the prior probability provided by Intentional performance (Intentional average score of 2 minus expected score by chance of 1 = 1). This produces a Bayes factor of 0.20; which provides evidence for the null and suggests that participants in the Incidental group do not respond to these questions with explicit knowledge about colour cue-response contingencies.

Table 5.5. Table showing the number of participants in Incidental and Intentional conditions who correctly guessed zero, one or both colours as Predictive. Number of participants who guess the correct colour regardless of location are shown, with the number of participants who also guessed the correct location in brackets.

	Neither Correct	Left Correct	Right Correct	Both Correct
Incidental	3	5 (2)	5 (3)	3 (2)
Intentional	0	0 (0)	0 (0)	16 (16)

5.5.2.2. Prediction task

The prediction task data are shown for both groups again in Figure 5.8 for comparison, which clearly show that participants performed with higher accuracy in Intentional groups compared to Incidental groups. Intentional mean performance (M = 14.4, SE = .288) was used as a prior, participants in this group scored 6.44 higher than chance (8 correct), and this was used to plot the normal distribution of plausible explicit performance with a standard deviation of half the mean, 3.22.

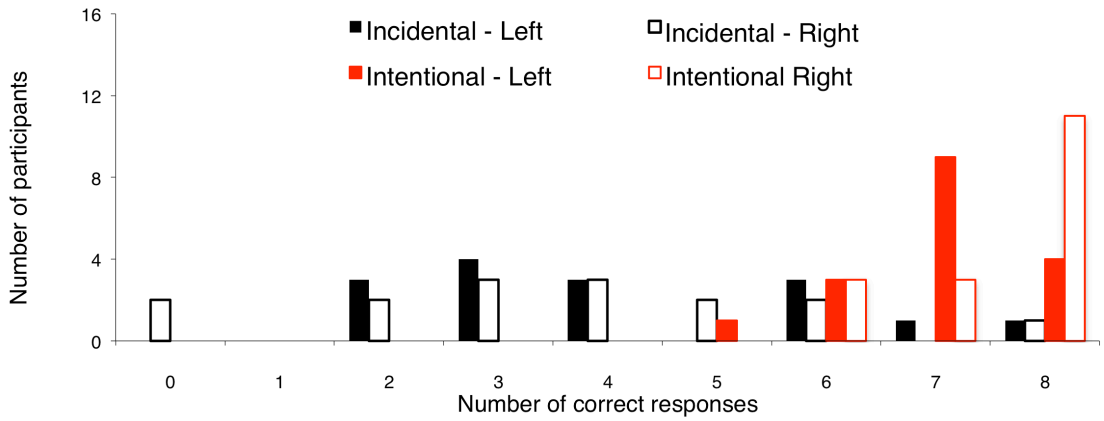

Figure 5.8. Graph showing the number of participants who made each number of correct responses out of the eight total possible for both Predictive Colours 1 (Left, filled bars) and 2 (Right, open bars) on the Prediction task for Incidental (black bars) and Intentional (red bars) groups.

Participants in the Incidental condition scored just 0.25 over chance on the prediction task (M = 8.25, SE = .727), giving a Bayes factor for a normal distribution of .04, providing evidence for the null. Rather than assuming that the likelihood follows a normal distribution with explicit performance more likely, it is possible only a number of participants have explicit knowledge and therefore the full range from 0 (chance) to

6.44 (Intentional prediction task performance above chance) were used in a uniform distribution, giving a Bayes factor of .19, still providing good evidence that participants did not show explicit knowledge of the colour-response contingencies.

It is also possible, however, that the Incidental group learned far less than the Intentional group (see Figure 5.7). To correct for this possibility the estimate for incidental performance can be scaled down to approximate what we might expect if learning had been less. Using the training data (both RT and proportion of errors) I calculated the difference between all three pairs of the Colour Types for Incidental and Intentional groups. I then calculated the proportional difference between these differences. For example, the advantage for the Incidental group of Predictive Consistent versus Predictive Inconsistent colours was 6.19 ms, whereas it was over four times bigger in the Intentional group, 29.3 ms: proportional difference of 4.73 times more learning. I calculated the average of these across RT and errors and found that learning was just under four times greater in the Intentional group, $M = 3.85$, $SD = .967$. Using this proportional difference I scaled the Intentional performance (6.44) to expect 1.67 over chance for the Incidental group, with a uniform distribution. This gives a Bayes factor or 0.70, which is an inconclusive result. Therefore more participants are required to establish confidently that the poor prediction task performance in the Incidental group reflects a lack of explicit knowledge, or simply reduced knowledge or learning about the task altogether.

5.5.3. Discussion

Crucially, these analyses provide evidence that participants were able to learn colour-response contingencies under incidental conditions in the absence of explicit knowledge. Participants under Incidental conditions showed evidence for the null in both post-experiment interview questions and a prediction task, with the results of the Intentional participants acting as the prior probabilities for explicit knowledge.

Participants in the Intentional condition learnt significantly more than those in the Incidental condition, although they also suffered from extinction. The increased colour-response learning in the Intentional group compared to the Incidental provides an interesting insight into the importance of conscious expectancy when learning such contingencies. Explicit knowledge in sequence learning tasks (e.g. Experiment 2; Jones

& McLaren, 2009) does not necessarily improve learning, but in the case of Experiment 6 this seems to be the case. Participants were more cautious in the Intentional condition as they were slower and more accurate when performing the task. This goes some way, perhaps, to support Reber's (1989) claim that explicit knowledge is suited to contingencies whose relationships are simple, deterministic and easy to verbalise.

5.6. General Discussion

Participants were able to learn cue-response contingencies under both Incidental and Intentional conditions and showed evidence of faster and more accurate responding to stimuli consistent with the trained colour-response location than the opposite location. The first thing to note about these results is that the Intentional group clearly demonstrated explicit knowledge about the colour-response contingencies that was consciously accessible. Participants were able to produce confident judgments about the presence of these contingencies as all participants were sure of the contingencies and correctly reported the colour and locations that they predicted. The incidental participants were not sure of these contingencies and performed at chance on guessing the colours and a prediction task, which both provided evidence for the null (no explicit knowledge) using the Intentional group as the prior probability (Dienes, 2011; 2014).

This suggests that participants are able to learn cue-response contingencies under Incidental conditions, but this seems to be somewhat weaker than the sequence learning observed in previous chapters. Whilst no direct comparison was made, as the cue-response learning occurred in the absence of any sequences to learn, the effect was quickly extinguished at test suggesting that this incidental learning may not have been very robust. This may explain the results of Jiménez and Méndez (1999) who found no colour-response learning in their SRT task alongside sequence learning. Cleeremans (1997) suggested that colours should be easy to learn about, as these response contingencies are simple and do not require complex abstract structures to be learned like sequential contingencies do. This is the reason that both sets of authors provide for expecting cue-response contingencies to block or overshadow sequence learning.

The cue-response results from Experiment 5 suggest that, when isolated from sequence learning, human cue-response learning under incidental conditions is not simply stronger because it has a simpler structure. Jiménez and Méndez (1999) did not provide

evidence that participants could learn these cue-response contingencies; without the influence of sequences. This chapter provides some evidence that cue-response learning may be less robust than sequence learning and that a demonstration that participants can do this without sequences is needed within such a dual-cue task. Cleeremans (1997), Jiménez and Méndez (1999) did not consider the possibility that sequential learning may be ideally learned under incidental conditions because of the abstract and complex nature of the relationships that are learned (Reber, 1989). When matched for contingency, number of trials and number of instances: will humans be able to learn sequential contingencies far better under incidental conditions than they can cue-response contingencies? These issues will be dealt with further in the following chapter.

Whether this Chapter provides suitable evidence of a lack of explicit knowledge is an important question, and I conclude that whilst a number of definitions of implicitness are addressed, it would still not convince the most determined of single process champions. Firstly, the explicit knowledge tests occur after the task and therefore after extinction. Any small amount of explicit knowledge that the Incidental group may have had could have been eradicated by this period where cue-response contingencies were absent. Explicit knowledge does survive this extinction in the Intentional group – who show far reduced learning at test than at training that is no different from the Incidental group. Therefore, the Intentional learners showed evidence of extinction but their explicit knowledge remained. However, as they also demonstrated some knowledge on the indirect measures of learning at test while the Incidental group did not, it could be suggested that there is less learning in the Incidental group and therefore easier to extinguish, hence it did not survive to test, either indirect or direct tests.

Whilst the subjective measures of knowledge were not entirely sensitive, these go some way to address these issues. Further to being apparently unable to perform higher than chance on direct test of explicit knowledge, participants were also subjectively unaware of these contingencies. All of the participants in the Intentional group were completely sure that there were contingencies in the task, correlating perfectly with their perfect performance in colour identification. Of the four participants under Incidental conditions who reported that they thought something may have been going on in the task, there was no difference in their prediction task score to the rest of the group. This

suggests that if this measure captured subjective awareness that this did not correlate with performance, thus providing further evidence for the absence of awareness.

However, it could be the case that participants in the Incidental condition were reluctant to express their knowledge with confidence as they were nervous about being incorrect; whereas Intentional participants knew that there were contingencies present and therefore that they were correct and thus had no trouble in confidently displaying their knowledge. Altogether the sensitivity of the measure of subjective awareness could have been improved in order to provide a more detailed account of whether participants satisfied the zero-correlation (Chan, 1992) or guessing (Cheesman & Merinkle, 1984) criterion.

In summary, Experiment 5 provides good evidence that cue-response contingencies can be learned and that they can be learned incidentally. From the perspective of a volitional (Jacoby, 1991), subjective (Dienes & Berry, 1997) or knowledge based (Shanks, 2005) account of conscious knowledge, this task converges on evidence that participants were not aware of cue-response contingencies in the Incidental group. Whilst this learning suffered from quite rapid extinction and was not as large as if the task was performed with explicit contingency knowledge, it is possible for participants to learn these contingencies. The task is therefore usable in attempts to investigate the effect of additional stimuli that are themselves predictive of a response on sequence learning, as discussed in Chapter 6. The RASRN was able to provide a simulation of these incidental learning results, which will be tested alongside the presence of sequential contingencies in Chapter 6.

Chapter 6. Cue-competition and sequence learning

This chapter continues to examine the effect of additional stimuli on sequence learning. The experiments outlined in Chapter 4 used colour and tone stimuli, whose relationship with the sequence had a potentiating effect on sequence learning. These relationships meant that it was not possible to isolate learning about sequences from learning about additional stimuli, and therefore how additional stimuli improved sequence learning. Chapter 5 introduced a task that involved a cue in the centre of the screen that was sometimes predictive of a response location, which was used in this chapter to investigate the interaction between cue-response and sequence learning. In this chapter I used a between-subjects design, comparing learning of participants experiencing either one or both of these contingencies. McLaren et al. (2013) observed that when both contingencies were in play participants learned sequences but not about colours, evidence of cue-competition and an overshadowing effect. Experiment 7, however, provides evidence that when the colour contingencies are positively correlated with the trial order, this overshadowing effect can be counteracted and the Dual group provides good evidence of colour-response learning at test. Experiments 8 and 9 further investigated the role of cue-sequence relationships and cue competition by comparing two groups who experienced both colour-response and sequential contingencies: in one case these cues were correlated, in the other they were not, however conclusive evidence of either cue-competition or facilitation effects was not obtained as evidence of any colour learning in any group was weak. The RASRN predicted that when sequences and cues were correlated, the Dual group would show cue-response potentiation; but when uncorrelated the Dual group would show overshadowing of sequence learning. Whilst these results are not established by the human data reported in this chapter, they are consistent with it, and this provides an indication that the RASRN can, to some extent, simulate the relationship between simple contingency learning and sequence learning and predict whether overshadowing or the opposite effect might be expected. The results provide evidence that the absence of cue-competition does not simply provide support for a propositional account of learning, and that automatic associative approaches can both produce evidence of cue-competition (McLaren et al., 2013) as well as a suitable account of why it does not always occur.

6.1 Introduction

Chapter 4 demonstrated that sequence learning was potentiated by concurrently presenting stimuli that had a relationship to the previous trials experienced (Previous condition), even though they provided no predictive information about the current trial itself. This effect was predicted by the RASRN, as concurrent stimuli that correlated with the response stimuli locations showed increased S-R learning that blocked sequence learning (the Current group). It appeared that the Previous condition increased learning of the Same rule sequence, perhaps through activating the previous sequence element and therefore the $t - 2$ response-location required to predict the current trial (t). However, the mechanism that produced this learning effect could not be isolated, as the concurrent stimuli (tones or colours) and response-stimuli (locations) followed the same sequence. There was no way to separately assess the learning occurring about the tones or colour stimuli and the location stimuli, and therefore I was unable to assess how these may have interacted to cause the sequence learning effect observed in Chapter 4. There was no condition tested in Experiment 3 and 4 in which the concurrent stimuli (tones or colours) had a predictive relationship with the required response that was not related to the sequence. The aim of this chapter was to attempt to isolate stimulus-response learning, and investigate further how a trained relationship between stimuli and responses would interact with sequence learning.

Finding evidence that incidentally learned contingencies follow associative learning effects would provide a strong case for human associative learning (Beesley & Shanks, 2012). It seems that there is a limited capacity for learning relationships between events, as when two contingencies are trained simultaneously (AB+), learning is less than if they were trained separately (A+ and B+). This overshadowing effect (Kamin, 1969) is an example of cue-competition, when multiple cues come to predict the same outcome. Blocking is another example, where one cue is trained to predict the outcome (A+) and after training on this contingency another cue is trained in compound (AB+), which results in far less learning about the additional cue (B) compared to the pre-trained cue (A). There have been a number of studies that have demonstrated cue-competition effects in human contingency learning studies (Dickinson, Shanks, & Evenden, 1984; Le Pelley, Beesley & Suret, 2007); however, as discussed in previous chapters, these tasks are often confounded by the presence of explicit contingency knowledge (Beckers et al., 2006; Mitchell et al., 2009). There are a number of studies that aim to avoid these issues by attempting to mask learning with cover stories, for example Vadillo, Orgaz

and Matute (2009) who had participants perform an overshadowing experiment based around a complex refugee saving task where road mines were partially predicted by cues contained within a spy radio on screen. Whilst the object of this task was not to learn, participants were still required to pay attention to cues and use them, as well as the task involving a pre-established causal framework (Waldmann & Holyoak, 1992).

Beesley and Shanks (2012) argue that the solution to these issues lies in implicit learning tasks and choose contextual cuing within a visual search task for their methodology. Using a task where participants are required to identify the orientation of a single target letter (e.g. a 'T') amongst a display of distractors (e.g. rotated 'L's). Unknown to participants, some of these display patterns were predictive of the location (not the response cue of orientation) of the target letter. Participants were pre-trained on certain distractor patterns (A+) and further patterns were added to the stimulus array (AB+), see Figure 6.1 for a schematic representation of the design.

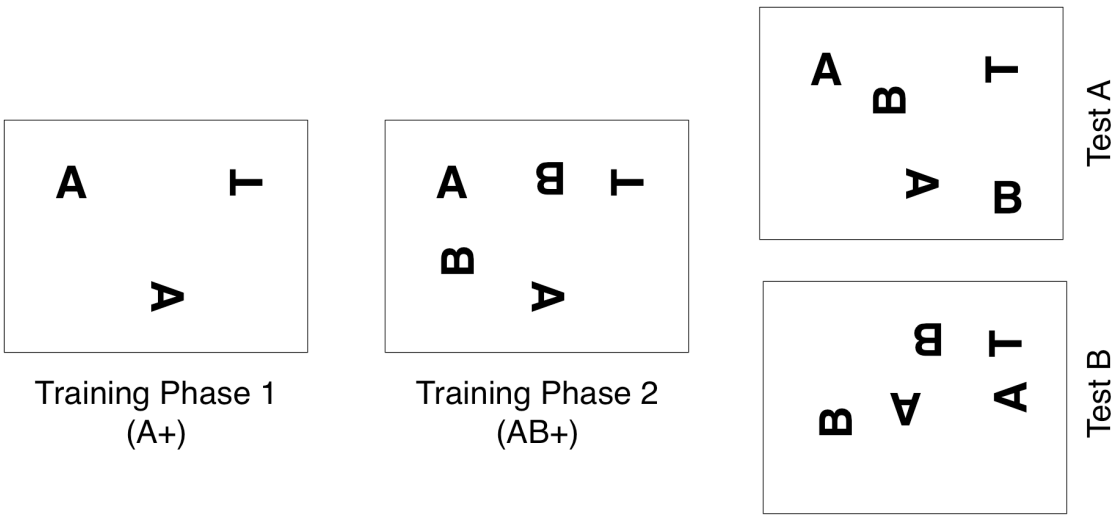

Figure 6.1. Representation of the blocking designed used by Beesley and Shanks (2012) in a contextual cuing visual search task. The target stimulus (T) is predicted in training phase 1 (left panel) by the pattern of the distractor letters ('L's) here shown as pattern A. The second training phase (central panel) involved training with both pattern A and a further pattern of distractor stimuli, shown by Bs. Note, these were the same letters in the task (e.g. 'L's). At test, participants were exposed to both A and B patterns separately, alongside random locations of the B or A distractor letters.

Participants showed learning about each pattern, but no evidence of blocking to the B patterns that were trained in compound with the pre-trained A patterns. Indeed, the

opposite effect (an increase in learning was observed to the B cue compared to a control set of predictors), was observed which suggests that potentiation of learning may have occurred. Beesley and Shanks (2012) suggest this provides evidence that humans do not exhibit blocking under incidental conditions, which supports a propositional account of cue-competition effects in humans (e.g. Mitchell et al., 2009).

As mentioned in Chapter 1, there is an issue with this visual search task as participants may not have learned an association between the entire A pattern and the outcome. Instead they may have learned about a certain set of visual features within the distractor array that are entirely different when presented with the compound training in phase 2[1]. This is similar to Shanks and St John's (1994) criticism of artificial grammar tasks, on which participants may have been learning that words containing high frequencies of certain letters or groups of letters (e.g. a string of 'T's) was actually driving improved performance in classifying the words as consistent with the grammar rules or not. Therefore these tasks fail the information criterion, as the learning may not be an association between pattern A and the outcome; but it may be that the entire visual array, or sections of it (for example, a clump of Ls in one corner) are associated with the outcome. This further suggests that pattern AB may not show cue-competition effects as these may be encoded visually as a whole rather than a separate set of elements.

As reported in the previous chapter, a number of studies with sequence learning and additional cues were unable to show evidence of cue-competition (Cleeremans, 2007; Jiménez & Méndez, 1999) however these studies failed to provide evidence of incidental learning about the cue in its own right. As Chapter 5 demonstrated that participants are able to learn about contingencies between a colour cue and the location of the required response on that trial, this chapter aims to investigate whether colour or sequence learning compete with one another. Previous studies assumed that simple contingency learning would overshadow sequences, whereas this thesis suggests that either is possible. As Chapter 5 demonstrated rapid extinction of incidental cue-response contingency learning, whereas Chapters 2 and 4 demonstrated strong sequence learning that did not suffer from extinction, when these tasks are put together it seems, on this basis, more likely that sequence learning would overshadow learning about colour cue-response contingencies.

[1] Criticism offered by Mike Burton at the London meeting of the Experimental Psychology Society, January 10, 2014

Indeed, in a recent study using the paradigms reported thus far in this thesis, McLaren et al. (2013) used a dual-cue task using both the incidental sequence learning methodology (Jones & McLaren, 2009) concurrently with the addition of colour-response contingencies similar to Experiment 5. Therefore on each trial there was a contingency between the central colour cue and the required response as well as a contingency between the previous trials and the current response stimulus location. McLaren et al. (2013) trained a group that learned both sequential and colour-response contingencies (Dual group) in comparison to sequence-only and colour-only groups: who received random colours or pseudo-random sequences, respectively. These participants showed that sequence learning in Dual and Sequence groups were no different to one another, therefore sequence learning was not affected by the presence of colour-response contingencies. It was the colour cue-response learning that was affected, with little or no evidence of colour learning in the dual group, suggesting that the presence of sequences had overshadowed learning about colour-response contingencies. Therefore it is possible to observe cue-competition effects under incidental conditions. These results could, furthermore, explain the difficultly that Jiménez and Méndez (1999) had in showing evidence of cue-response learning alongside a sequence learning task.

This McLaren et al (2013) study firstly provides evidence of incidental cue-competition, as colour cue-response learning was overshadowed by sequence learning. As the presence of additional stimuli that related to the previous element in the sequence potentiated sequence learning in Chapter 4, I was interested to investigate how the relationship between cue-response and sequence learning may act to *increase* learning. Beesley and Shanks (2012) suggest that this is a possibility that may explain their results, as there may be some within-compound or configural representation that might have increased learning about the 'blocked' B pattern (e.g. Urcelay & Miller, 2009). To some extent there may be a possible interaction then between competition for learning about an outcome and learning about associations between cues. Therefore, the aim of the studies reported in this chapter is to systematically investigate cue competition under incidental learning in humans. Two separate stimulus types were used to examine these effects (colour-response contingencies and exclusive-or sequential contingencies of response stimuli locations), which were spatially separated.

6.2 Experiment 7: Related cue-response and sequence learning

A new version of the SRT task was designed that involved the two-choice SRT sequential rules employed by Jones and McLaren (2009), and compared learning of these contingencies to learning about simple cue-response associations between a central colour cue that appeared before the response-stimuli, as outlined in Chapter 5. Unlike in the work by McLaren et al. (2013) this task aimed to investigate whether cue-competition effects would occur if sequences and colours were themselves related. So whilst both sequential information (the previous two trials, $t - 2$ and $t - 1$) and the colour of the cue were predictive of the response location, they were also predictive of one another with the same contingency – would cue-competition still be observed?

This meant using colours that each had a lower contingency (66%) with a response, rather than 80% in Experiment 5, to match the contingency of sequences and colours. By increasing the number of training blocks to eighteen as well as reducing the number of test blocks to two, the intention was to increase colour learning and attempt to reduce the impact of extinction at test seen in Chapter 5. Therefore, after providing evidence that participants can learn both sequence and colour information separately in conceptually similar tasks under incidental conditions (Chapters 2 and 5); the first experiment in this cue competition series aimed to investigate whether experiencing both colour and sequential contingencies that are related to one another in the same task would alter learning of either of these relationships.

6.2.1 Method

6.2.1.1 Participants

48 participants (aged between 18 and 30 [$M = 20.6$]; 38 female and 10 male) were recruited from first year psychology undergraduate students at the University of Exeter were awarded either one course credit (N = 8) or £5 (N = 40) in return for participation. Participants provided informed consent prior to taking part in one session lasting roughly one hour. Participants were allocated into one of three between subject groups: Dual (positive colour and sequence contingencies); Sequence (no colour and positive sequence contingencies) or Colour (positive colour and no sequence contingencies).

6.2.1.2 Materials and Stimuli

The experiment was run with all stimuli and materials as in the Incidental condition of Experiment 8, see section 5.5.1.2) with a stimulus display of two white circle outlines with a square outline in the screen centre. Both the central square (cue-stimulus) and the circles to the left and right (response-stimuli) filled with one of four colours: red; yellow; green; or blue.

6.2.1.2 Design

The experiment was a two-choice SRT task comprising of one session of twenty blocks. These blocks each contained 120 trials, with the first 18 blocks acting as training and the final two blocks acting as test. Participants received either training on blocks containing sequential contingencies (Dual, Sequence) or no sequential contingences (Colour); and either colour contingencies (Dual, Colour) or no colour contingencies (Sequence). There were neither colour nor sequential contingencies present in any group during the two blocks of test, so a colour was equally likely to occur with either response location, and trial order was pseudorandom.

Sequence construction: Experimental blocks. Those blocks involving sequential contingencies were constructed from 40 subsequence triplets of right and left stimuli in a similar way for Experiments 1 to 4, as outlined in section 2.2.1.3. This experiment, however, followed an exclusive-or rule used in Jones and McLaren's (2009) study of sequence learning. This rule states that if the previous two trials are the same (the exclusive case [e.g. right & right, or left & left]) that the current trial will be one response (e.g. right), whereas if the previous two trials are different (the or case [e.g. right & left, or left & right]) then the current trial will be the other response (e.g. left). Right and left response stimuli were counterbalanced across participants and are forthwith referred to in terms of Xs and Ys. As in my earlier experiments, of the four possible subsequence triplets that follow this rule (XXX, XYY, YYX, YXY), ten of each subsequence were arbitrarily randomised and concatenated to form the 120 trials for a training block with sequential contingencies. The contingency of the exclusive-or rule was therefore 100% on every third trial, as each third trial in the subsequence triplet used to construct the trials followed the rule and the overall contingency was 2/3. This construction method resulted in a balanced number of X and Ys in each block, as well as controlling for sequential effects up to two trials before the current trial (t - 2).

Sequence construction: Control blocks. Those blocks involving no sequential contingencies (Colour group and all groups at test) were pseudorandom and constructed in the same way as for pseudorandom blocks in my earlier experiments (see section 2.2.2.3). The contingency experienced between the previous two trials (or indeed any trial) predicting the current trial using an exclusive-or rule (or any other combination of preceding trials) was therefore 50% and participants should have no sequential information to learn about. These blocks acted as a control for the sequence learning possible in the Dual and Sequence groups. The amount of right and left response stimuli, as well as the number of repeats and alternations were balanced and controlled for.

Colour contingency construction. Once the sequence of X and Ys was constructed, the colour of the cue stimulus was arranged across the trials within each block, with the procedure for this depending on the group that participants were assigned to. All groups experienced all four colours equally across and within training and test blocks, with the colours (red, blue, green, yellow) themselves randomly allocated to Colours 1 to 4. For those groups that received a colour contingency (Dual and Colour), during training Colours 1 and 2 (Colour) or all colours (Dual) were Predictive and occurred with a 67% contingency with an X or Y trial respectively. Colours 3 and 4 were Non-Predictive for the Colour group (and all colours for the Sequence group) and therefore occurred with equal likelihood on either X or Y trials.

Colour contingency construction: Dual group training blocks. Across training blocks, the Dual group was allocated colours according to the subsequence triplets that made up the right and left responses. All colours were assigned to a subsequence 'pair': the first two trials of a triplet (e.g. XY in XYY). Regardless of the position within the block, any trials that are preceded by the first two trials of each of the four trained subsequences were assigned a unique colour. For example, Colour 1 would always occur after two XX trials. As the third trial followed the exclusive-or rule on 67% of trials, this results in a 67% contingency between Colour 1 and X as a result of the contingency of the exclusive-or rule already in place. The contingency between the first pair of trials within the subsequence themselves and a colour was 100% with the previous trial pair. Colour 2 occurred after the subsequence pair XY; Colour 3 after subsequence pair YX and Colour 4 after subsequence pair YY. There were no control colours in this group and all colours therefore had a 67% contingency between cue stimulus colour and response stimulus location across training.

Colour contingency construction: Colour group training blocks. The colours that occurred during training blocks in the Colour group followed the same structure as Experiment 5, with Colour 1 predicting an X trial on 67% of those trials and 67% of Colour 2 trials predicting a Y response. As in Experiment 5, and contrary to the Dual group, these colours were simply assigned according to these contingencies randomly across a block with no further constraint or relationship to the subsequences or response location.

Colour contingency construction: Control blocks. Across training in the Sequence group and at test for all groups, there were no colour contingencies. All four colours occurred with equal likelihood on either X or Y trials. These were allocated randomly across these trials with no further constraint or relationship to the subsequence or responses, meaning the colours had a 50% chance of preceding either response location.

6.2.1.4 Procedure

The procedure followed previous experiments: all participants were instructed to simply respond as quickly and accurately as possible. They were instructed to fixate on the coloured square, suggesting that it would help them to respond quicker and more accurately by attracting their attention to the centre of the screen. They were informed that the experiment intended to measure their reaction times and errors and its aim was to investigate people's ability to respond quickly and accurately to very simple stimuli over the course of an experiment. They were not informed of any relationship between the colours or sequence of the stimuli and the responses required and no mention was made of anything to learn about or from. They were told that the colours changed to make the task less dull and to try and retain their attention.

At the beginning of each block participants were instructed to press any key to start. Each trial began with an inter-trial interval of 250 ms where a black background with a white square outline and two white circle outlines was presented. The cue stimulus (a coloured square in the centre of the screen) would then appear for a variable interval of between 250-500 ms. The response stimulus (the left or right coloured circle) would then appear on screen until either the participant made a keypress response or the trial timed out after 4000 ms from the presentation of the response stimulus. The colour of the cue and response stimuli were the same on each trial, so if the cue-stimulus was a red square then either response stimulus would be a red circle. RTs were measured from

the onset of the response stimulus. If participants pressed an incorrect key, or the trial timed out, the computer issued a beep sound.

Feedback was given as in all previous experiments at the end of each block and at the end of the twenty blocks a short verbal structured interview was given, in which participants were asked about what they had noticed in the experiment regarding both sequences and colours. Participants were asked to describe any contingences they may have noticed, and were required to identify which response each colour predicted. Participants were finally debriefed and thanked for their participation.

6.2.2 Results

RTs and responses were recorded for each participant, with exclusion and inclusion criteria as described previously. Average RT and proportion of errors were calculated for each of the eight subsequences and separately for each Colour Type. All four colours were analysed for Dual and Sequence groups with Predictive Consistent and Predictive Inconsistent trials, however only the two Predictive colours were analysed in the Colour group. Non-Predictive colours were not included in the analysis as there was no equivalent in the Dual group, which exclusively involved Predictive trials. All colours in the Sequence group were Non-Predictive, but two (Colours 1 and 3) were assigned X as 'Consistent' and Y as 'Inconsistent', with the opposite dummy label assigned to responses for Colours 2 and 4. Sequence learning and colour-response learning are analysed separately in the following sections.

Sequence learning: Difference scores. As in Experiments 1 to 4, difference scores were calculated from RTs and errors as an index of performance on consistent subsequences (XXX, XYY, YXY, YYX) taken from the corresponding inconsistent subsequence (XXY, XYX, YXX, YYY), with higher scores reflecting better performance on subsequences consistent with the exclusive-or rule.

Colour-response learning: Difference scores. Difference scores were also calculated using the same principles as for sequence learning, with Predictive Consistent average RTs and proportion of errors taken from Predictive Inconsistent trials to provide a difference measure that with greater values reflects better performance on trials where

the Predictive colour resulted in the trained contingent response over the response Inconsistent with the trained relationship.

6.2.2.1 Sequence learning

An analysis of variance was conducted on both RT and error difference scores across training and test with the factors Block (Training: 18; Test: 2) x Subsequence (4) x Group (3).

Training phase. There was a large significant main effect of Group, RT difference score, $F(2,45) = 23.5$, $p < .001$, $MSE = 2679$, $\eta_p^2 = .511$; proportion of error difference score, $F(2,45) = 13.2$, $p < .001$, $MSE = .014$, $\eta_p^2 = .358$, see Figure 6.2.

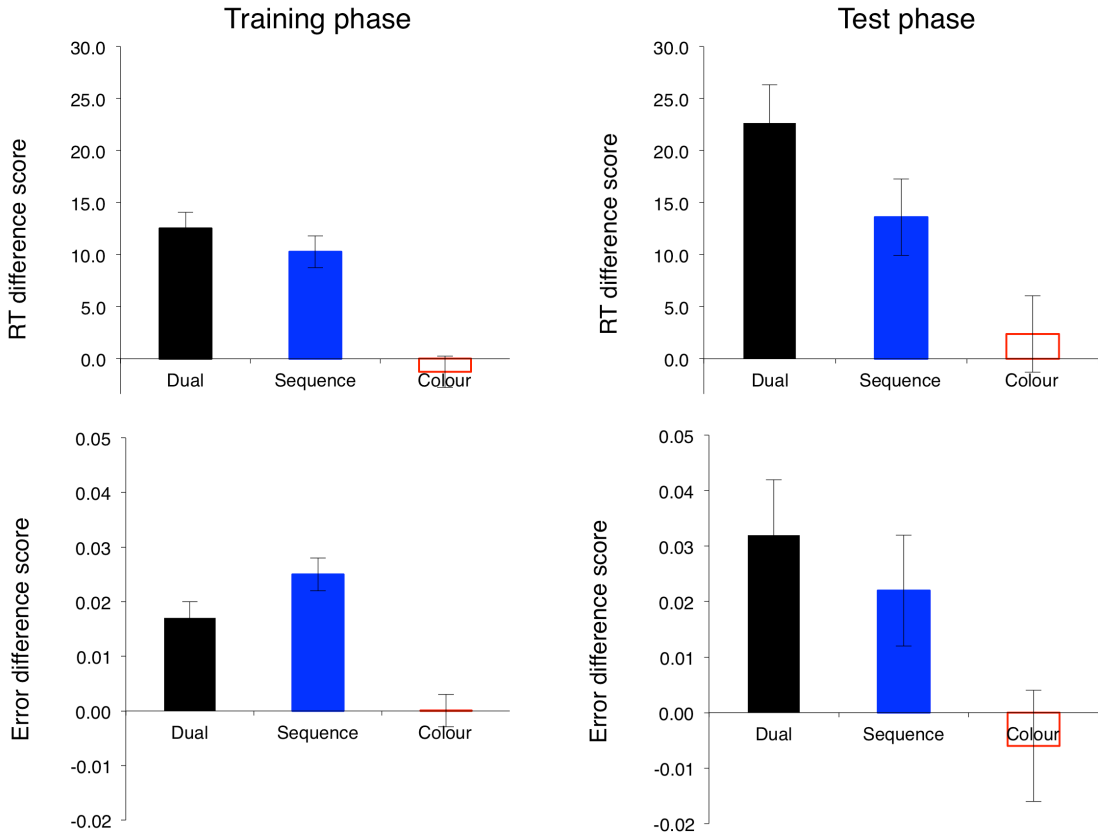

Figure 6.2. RT (top panel) and error (bottom panel) sequence learning difference scores across training (left panel) and test (right panel) for the three different groups. Error bars show standard errors.

A series of planned contrasts compared the groups, as shown in Table 6.1. Using a series of planned comparisons, the Dual group performed significantly better than the Colour group, demonstrating strong evidence of sequence learning in the Dual group

(see Table 6.1). The Sequence group showed no difference to the Dual group also demonstrating strong evidence of sequence learning compared to the Colour group (see Table 6.1). Evidence of sequence learning was therefore provided in both Dual and Sequence groups. There was no evidence of a main effect of Block, nor for the interaction of Block with Group in the RT difference score, which suggests that there was no evidence for a change in sequence learning across the experiment.

Table 6.1. Results from ANOVA for planned comparisons on average RT and proportion of errors sequence difference scores, comparing the three Groups to one another across training in Experiment 7.

	RT difference score: Training			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual vs Colour (control)	40.8	< .001	37.2	.476
Sequence vs Colour (control)	28.6	< .001	37.2	.389
Dual vs Sequence	1.08	.304	37.2	.023
	Error difference score: Training			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual vs Colour (control)	12.2	.001	.001	.213
Sequence vs Colour (control)	25.1	< .001	.001	.358
Dual vs Sequence	2.30	.130	.001	.049

Test phase. The variable Group demonstrates a main effect at test, RT difference score, $F(2,45) = 7.67$, $p = .001$, $MSE = 1724$, $\eta_p^2 = .254$; proportion of error difference score, $F(2,45) = 3.74$, $p = .031$, $MSE = .013$, $\eta_p^2 = .142$, see Figure 6.2. This provides strong evidence of sequence learning that was not eradicated by extinction. A series of planned contrasts (see Table 6.2) show that the Dual group was quicker and more accurate in responding to trained subsequences at test compared to the Colour group. Again, there was no difference between the Dual and Sequence groups; as the Sequence group also showed evidence of reliably faster responding and numerically an accuracy advantage for trained subsequences compared to the colour group, see Table 6.2.

Table 6.2. Results from ANOVA for planned comparisons on average RT and proportion of errors sequence difference scores, comparing the three Groups to one another across the two blocks of test in Experiment 7.

	RT difference score: Test			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual vs Colour (control)	15.3	< .001	216	.254
Sequence vs Colour (control)	4.68	.036	216	.094
Dual vs Sequence	3.05	.088	216	.063
	Error difference score: Test			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual vs Colour (control)	6.94	.012	.002	.134
Sequence vs Colour (control)	3.81	.057	.002	.078
Dual vs Sequence	.466	.498	.002	.010

The two blocks of test did not have a significant effect on responding, nor with Group; Subsequence; or the three way interaction, suggesting that extinction did not occur across the two blocks.

6.2.2.2 Colour learning

To investigate how participants learned about colours an ANOVA was conducted on both RT and error difference scores across training and test with the factors Block (Training: 18; Test: 2) x Group (3).

Training phase. Across training there was a large effect of Group, RT difference score, $F(2,45) = 21.5$, $p < .001$, $MSE = 2063$, $\eta_p^2 = .489$; proportion of error difference score, $F(2,45) = 4.60$, $p = .015$, $MSE = .015$, $\eta_p^2 = .170$, see Figure 6.5. This can be unpacked by planned contrasts (see Table 6.3) with faster and more accurate responding to consistent colours compared to inconsistent colours in the Dual group compared to the Sequence group. Evidence of learning is also apparent in the Colour group, with significantly faster and numerically more accurate responses to trained colours in the Colour group compared to the Sequence group; see Table 6.3.

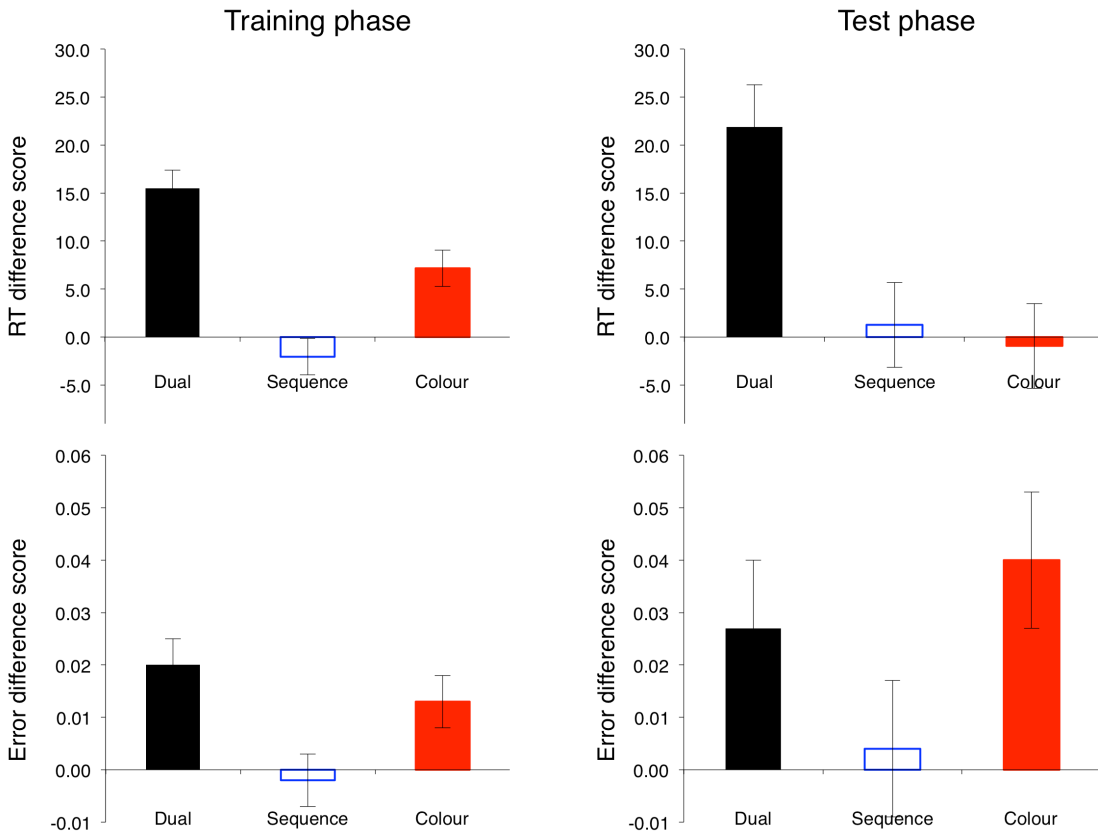

Figure 6.3. RT (top panel) and error (bottom panel) colour learning difference scores for the three Groups (Dual: black filled bars, Sequence: blue open bars; Colour: red filled bars) across training (left panel) and at test (right panel) for Experiment 9. Error bars show standard errors.

There was also a significant difference between Colour and Dual groups in RT difference score but not in proportion of error difference score (see Table 6.3). It is impossible, however, to ascertain whether this difference is caused by some potentiating effect of sequence learning in the Dual group, as Consistent colours and sequences are perfectly confounded and across training as the two are correlated measuring colour differences scores captures sequence learning in the Dual group. There was no effect of Block, nor was there evidence that this interacted with Group.

Table 6.3. Results from ANOVA for planned comparisons on average RT and proportion of errors colour difference scores, comparing the three Groups to one another across the two blocks of training in Experiment 7.

	RT difference score: Training			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual vs Sequence (control)	43.0	< .001	57.3	.489
Colour vs Sequence (control)	11.9	.001	57.3	.209
Dual vs Colour	9.70	.003	57.3	.177
	Error difference score: Training			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual vs Sequence (control)	8.91	.005	.001	.165
Colour vs Sequence (control)	3.83	.056	.001	.078
Dual vs Colour	1.06	.309	.001	.023

Test phase. At test the participants showed a main effect of Group, RT difference score, $F(2,45) = 8.07$, $p = .001$, $MSE = 1253$, $\eta_p^2 = .264$; proportion of error difference score, $F(2,45) = 1.98$, $p = .149$, $MSE = .011$, $\eta_p^2 = .081$, see Figure 6.4. The Dual group showed evidence at test of significantly faster responding to consistent colours over inconsistent colours than the Sequence group (see Table 6.4). This gives us strong evidence of learning about colour contingencies in the Dual group, now in the absence of confounding sequential learning or effects, as at test there was no correlation between the two stimuli. The Colour group showed no difference from the Sequence group in speed, but responded with a near-significant trend towards greater error difference scores compared to Sequence groups, and numerically more than Dual groups (see Table 6.4). However the Dual group was also significantly better than the Colour group, RT difference score (see Table 6.4). Therefore we have some evidence of colour learning in the Colour group at test, and strong evidence for colour learning in the Dual group at test.

There was an effect of Block in errors, RT difference score, $F(1,45) = .000$, $p = .993$, $MSE = .089$, $\eta_p^2 = .000$; proportion of error difference score, $F(1,45) = 6.26$, $p = .016$, $MSE = .051$, $\eta_p^2 = .122$, with a higher error difference in Block 1 ($M = .040$, $SE = .011$) than 2 ($M = .008$, $SE = .008$). This demonstrates the rapid decline in learning because of extinction at test. This did not interact with Group, RT difference score, $F(2,45) = .872$, $p = .425$, $MSE = 900$, $\eta_p^2 = .037$; proportion of error difference score, $F(2,45) = 2.29$, p

= .113, *MSE* = .019, η_p^2 = .092 and as with Experiment 5 this suggests that colour learning was quick to suffer from extinction.

Table 6.4. Results from ANOVA for planned comparisons on average RT and proportion of errors colour difference scores, comparing the three Groups to one another across the two blocks of test in Experiment 7.

	RT difference score: Test			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual vs Sequence (control)	10.8	.002	313	.194
Colour vs Sequence (control)	.123	.728	313	.003
Dual vs Colour	13.3	.001	313	.228
	Error difference score: Test			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual vs Sequence (control)	1.60	.212	.003	.034
Colour vs Sequence (control)	3.86	.056	.003	.079
Dual vs Colour	.488	.489	.003	.011

6.2.3 Discussion

The results clearly show that both Dual and Sequence groups learned sequential contingencies well across training, and that this learning was well established and robust across extinction at test. The overall amount of sequence learning was largely unaffected, it seems, by the presence of colours. Whilst some numerical differences existed between Dual and Sequence groups in learning of subsequences, there was no interaction and therefore no concrete evidence of differential subsequence learning. The only caveat to these conclusions stems from the strong trend towards better performance on consistent sequences during test in the RT difference measure for the Dual group relative to the Sequence group. But this in no way suggests that learning of sequences was poorer in that group.

There was convincing evidence of colour learning, with the Colour group showing evidence of learning across training and, to a weaker extent, in proportion of errors at test. It is clear that this colour learning is far weaker than the sequence learning, as across the two blocks of test we see the effect of extinction for colours only. With the evidence of learning from the Colour group enabling a comparison: the Dual group performed significantly better at test on the RT difference score than the Colour group. It seems that the presence of the sequential contingencies might have affected how well

the Dual group learned about colours and performed on them during test. Across training it is not possible to compare the groups, as the measurements for consistent colours are confounded with consistent sequences and therefore the two are impossible to disentangle. At test when the colour and sequence information are separated, we can see that the Dual group are above zero on both of their difference score measures, and reliably higher than the Colour group in RT difference scores which might suggest that the Dual group learn more about colours than the Colour group itself. However, as the score for the colour group was numerically higher in the error differences we must be cautious in claiming that this was necessarily the case.

Whilst McLaren et al. (2013) reported that two sets of contingencies in such a task produced an overshadowing effect on colour learning, Experiment 7 provides evidence of either no or the opposite effect. The Dual group performed numerically better when examining colour learning than the Colour group. Indeed the studies reported here may support the opposite conclusion, that colour learning may be potentiated in the Dual group. The source of these contentious results could arise from numerous factors, as the differences between McLaren et al. (2013) and the experiments reported here include: different colour contingencies; longer training; white response stimuli; and a longer test phase. However, the previous work reported in Chapter 4 suggests that the relationship between colours and sequences could produce this effect. As we observed that a contingency between the stimuli (sequences and additional concurrent stimuli) in Chapter 4 produced increased sequence learning, it is possible that a relationship between sequences and cues could also produce an increase in learning, or protection from overshadowing. As participants in McLaren et al.'s (2013) study had no relationship between the two sets of contingencies, colour learning was simply overshadowed by the strong sequence learning.

The current study did not, however, provide by any means a conclusive result that suggests the Dual group learn more than the Colour group. The Colour group performed better in terms of their accuracy to consistent versus inconsistent colours than the Dual group at test, and thus the difference could be explained by an attentional or motivational difference between the groups, as they may have focused more on speed of responding rather than accuracy to trained colours. Further to this, the groups had a different number of colour contingencies to learn: whilst the proportion of consistent to inconsistent trials was matched across groups, the Dual group had twice as many

colours to learn about than the Colour group. This could have encouraged participants to notice the contingencies in the Dual condition and employ different strategies or systems whilst responding to the task, causing perhaps an increase in colour learning as observed in the Intentional condition of Experiment 5. Therefore the groups were matched and efforts made to examine the participant's reportable knowledge about the contingencies in play in a further experiment.

6.3 Experiment 8: Cue-response and sequence learning when correlated and uncorrelated

A further experiment was run with the aim of better matching the Dual and Colour groups. As Experiment 7 involved four colours in the Dual group and two in the Colour group, I wanted to ensure that a comparison could be made between these conditions to assess whether the Dual group exhibited learning comparable to or greater than the Colour group. Experiment 8 therefore aimed to match the number of Predictive colours in the Colour and Dual groups, which was reduced to two, alongside two Non-Predictive control colours. Thus the groups were matched exactly in both the amount and proportion of Predictive consistent, inconsistent and Control trials. Further to this, a prediction task was added to the end of the experiment to better assay explicit knowledge. I was interested to know whether participants in the Dual group on direct tests were different to the Colour group regarding colour cue-response learning.

Experiment 7 also provided evidence that the Dual group exhibited at least the same amount of learning as the Colour group, the opposite result to the McLaren et al. (2013) study. To investigate this further an additional Dual group was run, one to match the Dual group in Experiment 7 (Dual Correlated) and one to match McLaren et al. (2013; Dual Uncorrelated). The first of these was a two-Predictive-colour version of the Dual group described in Experiment 7 – Dual Correlated, alongside a Dual Uncorrelated group that matched the colour contingency construction method involved in the McLaren et al. (2013) experiment. The other differences that existed between the two Dual groups in Experiment 7 and McLaren et al. (2013): training length; test length; colour-response contingency; and colour of response stimuli, were all equated. By matching these parameters it was possible to investigate whether manipulation of Dual group correlation (the relationship between colours and sequences) itself produced cue-competition in one case, and an absence of cue-competition in the other.

6.3.1 Method

6.3.1.1 Participants

48 participants (aged between 18 and 25 [M = 19.1]; 46 female and 2 male) were recruited from first year psychology undergraduate students and were awarded one credit in return for participation. Participants provided informed consent prior to taking part in one session lasting roughly one hour. Participants were allocated into one of three between subject groups: Dual Correlated; Dual Uncorrelated; or Colour.

6.3.1.2 Materials and Stimuli

The materials and stimuli used were the same as for Experiment 7 (see 6.2.1.2).

6.3.1.3 Design

The experiment was a two-choice SRT task comprising of one session of twenty blocks. These blocks each contained 120 trials, with the first eighteen blocks acting as training and the final two blocks acting as test. Depending on the group that participants were assigned to, across training participants received either blocks containing sequential contingencies (Dual Correlated, Dual Uncorrelated) or no sequential contingencies (Colour). All participants received blocks during training that contained colour contingencies, and neither colour nor sequential contingencies for the two blocks of test.

Sequence construction. Sequences of rights and lefts were constructed as for Experiment 7 (see 6.2.1.3).

Colour construction. All groups experienced all four colours equally across and within training and test blocks, as previously described. For all groups, Colour 1 and 2 were both Predictive, meaning that across training blocks they had a 67% contingency with a response stimulus location. 67% of the time that Colour 1 was present in the cue and response stimuli it was on an X trial. Colour 2 had the complementary contingency, and when it was present in the cue and response stimuli, 67% of the time this would be a Y trial. Colour 3 and 4 were Non-Predictive stimuli and occurred with equal likelihood on an X or Y trial. On test trials all colours occurred with equal likelihood on either X or Y trials and were allocated randomly across these trials. There was a programming error in this experiment that resulted in a non-random appearance of colours at test. Whilst the colours occurred equally across right and left trials, the programming error made it

more likely for the same colour to repeat on consequent trials. This meant that the test phase was not entirely random and participants may have noticed a difference in the order that colours occurred. This does not violate the use of the controlled test phase, where the colours are all equally likely to occur before a right or left required response, this may have introduced some confound to the task.

Colour construction: Dual Correlated group and Colour group. Across training blocks, the Dual Correlated group's trials were allocated colours according to the subsequence pairs that made up the right and left responses, similar to what was described in Experiment 7 for the Dual group (see 6.2.1.3). To some extent, the Dual Correlated group was the same as the Dual group reported in Experiment 7, as after a certain subsequence pair (e.g. XX) one Colour would always occur (e.g. Colour 1). However, only two total subsequences were predictive in this way, either: XXX & XYY; XXX & YXY; YYX & XYY; or YYX & YXY. This was counterbalanced across participants in the Dual Correlated group. These pairs were chosen as each possible combination of two subsequences that follow the exclusive-or rule that result in opposite (both X and Y) responses. Non-Predictive Colours 3 and 4 were distributed across the other half of trials with a 50% contingency with X or Y response locations. The Colour group was allocated colours across training in the same way as described for Experiment 7 (see 6.2.1.3). One consequence of this change in training schedule is, of course, that the number of predictive colours contributing to an assessment of colour learning has been halved, quite possibly reducing power, but this was a necessary compromise to allow proper experimental control.

Colour construction: Dual Uncorrelated group. Training blocks for participants in the Dual Uncorrelated group were assigned colours according to a different system, which involved first placing Colour 1 on every third trial that was an X and Colour 2 on every third trial that was a Y. The remaining 80 first and second trials in each block were then randomly assigned ten instances of Colour 1 if a Y trial and ten instances of Colour 2 if an X trial. The remaining 60 trials (of which an equal amount were X and Y trials) were randomly assigned Colour 3 or Colour 4, each balanced across the X and Ys equally. Colour 1 and 2 therefore had no relationship with a specific subsequence (although Colour 1 Consistent trials were 100% contingent with 'exclusive' cases and Colour 2 Consistent trials were 100% contingent with 'or' cases) and overall had a 67%

contingency with the required response. Colours 3 and 4 again had no contingency with either a subsequence nor a particular response.

6.3.1.4 Procedure

The procedure followed was the same as described for Experiment 7 (see 6.2.1.4) with the addition of a prediction task following the structured interview, as outlined in Chapter 5. In this prediction task participants began each block with a "press any key to begin" command and on each trial experienced an RSI of 250 ms. The cue stimulus would appear and participants were instructed to respond to this cue stimulus with the key press response compatible with the response stimulus that they thought would have filled during the experiment. A response to a response stimulus was not possible, as the white circle outlines did not fill at all during the prediction task. These trials had no time out, and participants were told to take as long or short as they liked to make a response and they could base this response on either: a random guess; some intuition; or any knowledge they had about the task. Participants were finally debriefed and thanked for their participation.

6.3.2 Results

RTs and proportion of errors were recorded and sequence learning difference scores calculated as described in Experiment 7 (see 6.2.2). Each participant had an average RT and proportion of errors calculated for the three types of colour trials (Predictive Consistent; Predictive Inconsistent; and Non-Predictive) and no difference score was calculated across training or test for colours.

6.3.2.1 Sequence learning

An analysis of variance was conducted on both RT and error difference scores across training and test with the factors Block (training: 18; test: 2)) by Subsequence (4) by Group (3).

Training phase. The variable of interest, Group, demonstrated a main effect across both measures, see Figure 6.5: RT difference score, $F(2,45) = 10.2$, $p < .001$, $MSE = 8484$, $\eta_p^2 = .311$; proportion of error difference score, $F(2,45) = 4.93$, $p = .012$, $MSE = .033$, $\eta_p^2 = .180$.

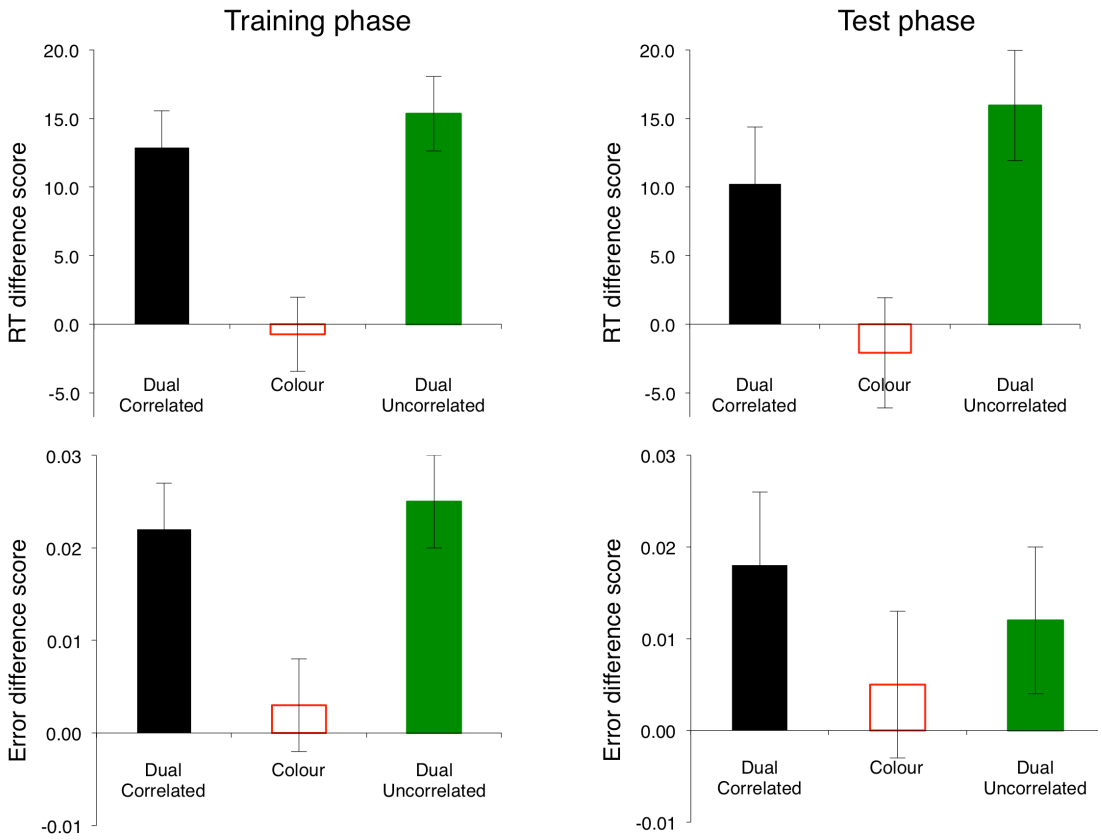

Figure 6.4. RT (top panel) and error (bottom panel) sequence learning difference scores for Dual Correlated (black filled bars); Colour (red open bars) and Dual Uncorrelated (green filled bars) Groups across training (left panel) and at test (right panel). Error bars show standard errors.

A series of planned comparisons unpacks this effect, shown in Table 6.5, demonstrating first a significant advantage for Dual Correlated over Colour groups in both RT difference score and proportion of error difference score. This demonstrates the learning of sequential contingencies by the Dual Correlated group compared to the control Colour group who received no contingencies. The Dual Uncorrelated group also showed significant sequence learning compared to the Colour group in both RT difference score and proportion of error difference score. There was no difference between the Dual Correlated and Dual Uncorrelated groups across training. Thus, both Dual groups trained on sequential exclusive-or contingencies showed strong evidence of learning these across training and did not themselves differ from one another.

Table 6.5. Results from ANOVA for planned comparisons on average RT and proportion of errors sequence difference scores, comparing the three Groups to one another across the eighteen blocks of training in Experiment 8.

	RT difference score: Training			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual Correlated vs Colour (control)	12.5	.001	118	.218
Dual Uncorrelated vs Colour (control)	17.5	< .001	118	.281
Dual Correlated vs Dual Uncorrelated	.419	.521	118	.009
	Error difference score: Training			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual Correlated vs Colour (control)	6.01	.018	.001	.670
Dual Uncorrelated vs Colour (control)	8.55	.005	.001	.816
Dual Correlated vs Dual Uncorrelated	.223	.639	.001	.075

Across training participants showed no main effect of Block: RT difference score, $F(17,765) = 1.55$, $p = .156$, $MSE = 9383$, $\eta_p^2 = .033$; error difference score, $F(17,765) = 1.17$, $p = .291$, $MSE = .016$, $\eta_p^2 = .025$, the errors participants made differed across Blocks depending on the Group (see Figure 6.6), proportion of error difference score, $F(34,765) = 1.57$, $p = .031$, $MSE = .016$, $\eta_p^2 = .065$. Whilst the Colour group remain relatively flat around zero, both Dual Correlated and Dual Uncorrelated show a trend towards higher proportion of error difference scores and therefore better performance across training, providing further evidence of learning.

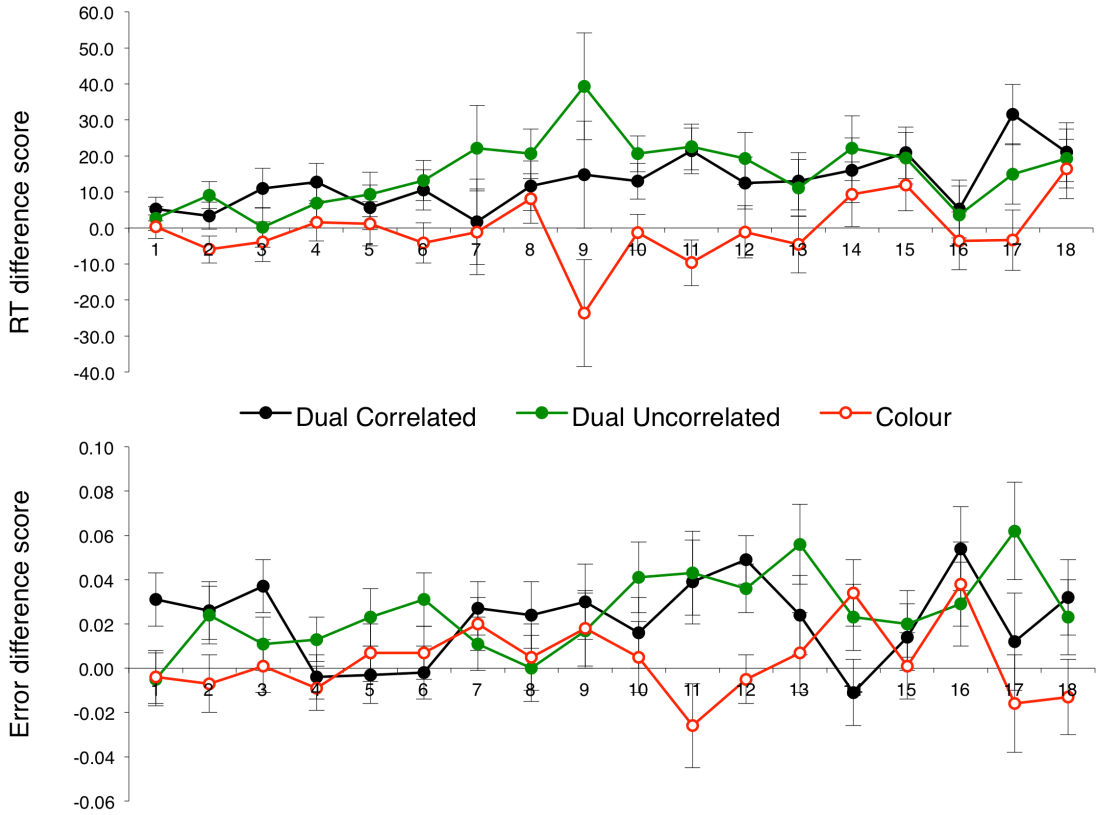

Figure 6.5. RT (top panel) and proportion of error (bottom panel) sequence learning difference scores across training for the three different groups. Filled circles indicate groups that were trained with sequential contingencies (both Dual Correlated [black] and Dual Uncorrelated [green]), with unfilled circles indicating those groups who had no sequential contingencies (Colour [red]) and therefore were not expected to demonstrate learning about sequences. Error bars show standard error.

Test phase. Group also produced a main effect at test in the RT difference score (see Figure 6.5): RT difference score, $F(2,45) = 5.39$, $p = .008$, $MSE = 2101$, $\eta_p^2 = .193$; but not in the proportion of error difference score, $F(2,45) = .612$, $p = .547$, $MSE = .009$, $\eta_p^2 = .026$. A series of planned comparisons were conducted (see Table 6.6) demonstrating a significant advantage for Dual Correlated over the Colour group in RTs at test, which demonstrates the learning of sequential contingencies by the Dual Correlated group compared to the control Colour group. The Dual Uncorrelated group also showed significant sequence learning compared to the Colour group in the RTs but not in proportion of errors (see Table 6.6). There was no difference between the Dual Correlated and Dual Uncorrelated groups at test, thus both Dual groups who had sequential contingencies showed some evidence of learning at test, and were not different from one another in their sequence learning.

Table 6.6. Results from ANOVA for planned comparisons on average RT and proportion of errors sequence difference scores, comparing the three Groups to one another across the two blocks of test in Experiment 8.

	RT difference score: Test			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual Correlated vs Colour (control)	5.69	.021	263	.112
Dual Uncorrelated vs Colour (control)	9.90	.003	263	.180
Dual Correlated vs Dual Uncorrelated	.578	.451	263	.013
	Error difference score: Test			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual Correlated vs Colour (control)	1.22	.275	.001	.026
Dual Uncorrelated vs Colour (control)	.341	.562	.001	.008
Dual Correlated vs Dual Uncorrelated	.272	.604	.001	.006

6.3.2.2 Colour Learning

An analysis of variance was conducted on both RT and error difference scores across training and test with the factors Block (Training: 18; Test: 2) x Colour Type (3) x Group (3). Note: there was no problem with weighted averages and therefore the data was not collapsed across Epochs as it could be analysed across the full range of blocks. The difference score used in Experiment 7 was not used here as each group had all three Colour Types and so, in order to better understand the learning processes that may have occurred, these raw average RT and proportion of error scores were analysed as in Chapter 5.

Training phase: Group comparison. When comparing the groups to one another, across training there was no main effect of Group, average RT, $F(2,45) = 2.18$, $p = .125$, $MSE = 218732$, $\eta_p^2 = .088$; proportion of error, $F(2,45) = .115$, $p = .892$, $MSE = .227$, $\eta_p^2 = .005$. There was also no interaction between Group and Colour Type in average RT, $F(4,90) = 2.37$, $p = .091$, $MSE = 3802$, $\eta_p^2 = .095$; nor proportion of error, $F(4,90) = 1.66$, $p = .189$, $MSE = .014$, $\eta_p^2 = .069$. This suggests that the groups are not different from one another across training, and do not differ in the extent to which they are learning the trained colour contingencies. This was further supported by a main effect of Colour Type in both measures: average RT, $F(2,90) = 9.09$, $p = .002$, $MSE = 3802$, $\eta_p^2 = .168$; and proportion of error, $F(2,90) = 5.46$, $p = .014$, $MSE = .014$, $\eta_p^2 = .108$, suggesting the overall pattern of learning is apparent across the three groups (see Figure 6.8).

There was a main effect of Block in the proportion of errors across training, suggesting that participants were making more errors as the task went on, $F(17,765) = .15,0$, $p < .001$, $MSE = 12461$, $\eta_p^2 = .250$; average RT, $F(17,765) = .911$, $p = .465$, $MSE = .017$, $\eta_p^2 = .561$. However, Block did not interact with Group; Colour Type; or both Group and Colour Type. Therefore, whilst participants may have begun to make more errors, participants did not differ in terms of their learning or across Groups across the course of training.

Separate analyses of variance were then conducted for each group, with the factors Block (Training: 18; Test: 2) x Colour Type (3) to assess learning within the group, as no control group was available to compare against as baseline because all groups were trained with colour cue-response contingencies. Therefore the difference between Predictive Consistent, Predictive Inconsistent and Control colours is the assay of colour learning for each group.

Training phase: Dual Correlated group. The Dual Correlated group demonstrated a trend towards evidence of colour learning across training in RTs with a strong trend but no main effect of Colour Type (see Figure 6.8): average RT, $F(2,30) = 4.01$, $p = .058$, $MSE = 10658$, $\eta_p^2 = .211$; proportion of errors, $F(2,30) = 3.27$, $p = .080$, $MSE = .033$, $\eta_p^2 = .179$. The pattern of results, however, was consistent with those that one would expect if learning of colours had occurred. Planned comparisons showed that participants in this group had significantly slower RTs, $F(1,15) = 9.52$, $p = .008$, $MSE = 4832$, $\eta_p^2 = .388$ and more errors, $F(1,15) = 4.89$, $p = .043$, $MSE = .031$, $\eta_p^2 = .246$, to Predictive Inconsistent colours compared to Non-Predictive. The differences between Predictive Consistent colours and both Inconsistent and Non-Predictive colours were in the expected direction, with faster and more accurate responding to Predictive Consistent colours, but were not significant.

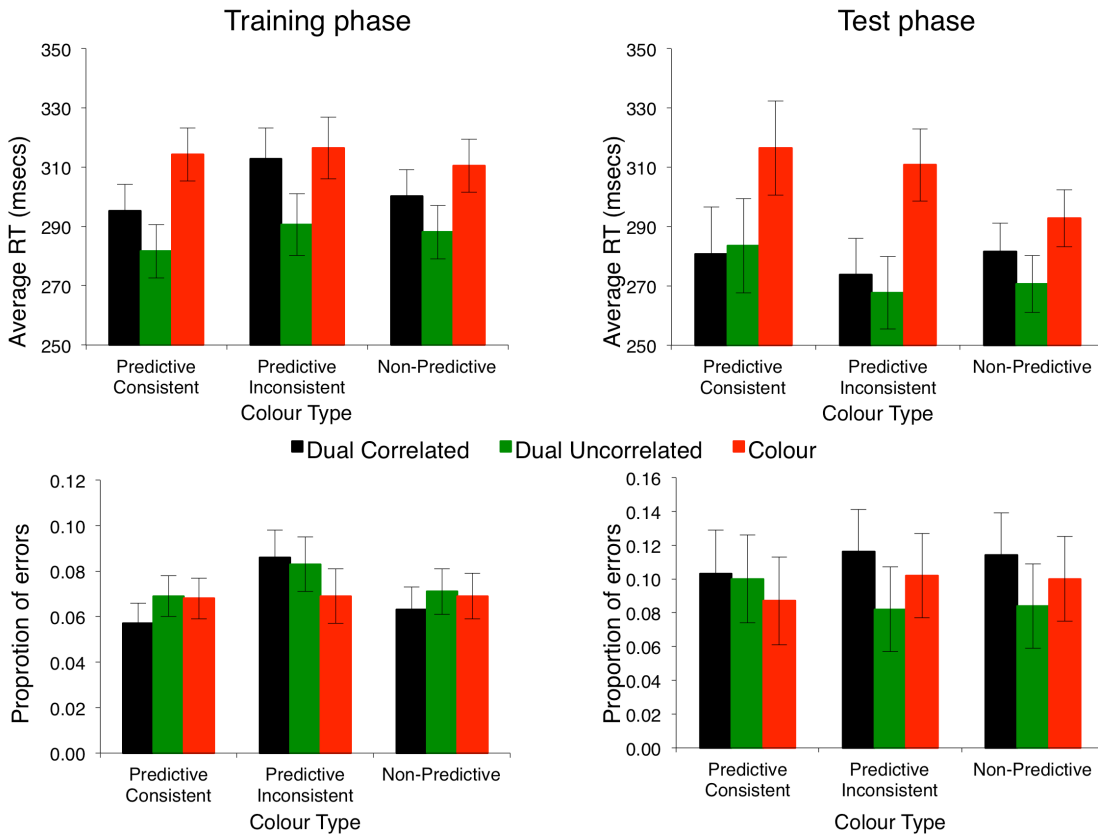

Figure 6.6. Average RT (top panel) and proportion of errors (bottom panel) for the different Colour Type trials: Predictive Consistent, Predictive Inconsistent and Non-Predictive across training (left panel) and test (right panel) for the three groups: Dual Correlated (black bars); Dual Uncorrelated (green bars); and Colour (red bars). Error bars show standard error.

Training phase: Dual Uncorrelated group. The Dual Uncorrelated group also demonstrated learning of colour contingencies across training in RTs and errors (see Figure 6.8) as seen in the main effect of Colour Type, average RT, $F(2,30) = 8.81$, $p = .005$, $MSE = 1124$, $\eta_p^2 = .370$; proportion of error, $F(2,30) = 4.40$, $p = .026$, $MSE = .004$, $\eta_p^2 = .227$. Planned contrasts showed that Predictive Consistent colours were responded to reliably faster, $F(1,15) = 9,75$, $p = .007$, $MSE = 2399$, $\eta_p^2 = .370$; and more accurately than Predictive Inconsistent colours, $F(1,15) = 5.25$, $p = .037$, $MSE = .004$, $\eta_p^2 = .259$. RTs also showed a large effect when comparing Predictive Consistent and Control colours, $F(1,15) = 28.7$, $p < .001$, $MSE = 419$, $\eta_p^2 = .657$; but not proportion of errors, $F(1,15) = .266$, $p = .613$, $MSE = .006$, $\eta_p^2 = .017$. There was also evidence of learning from the Predictive Inconsistent versus Control colour comparison in participants' errors, average RT, $F(1,15) = 1.33$, $p = .267$, $MSE = 1414$, $\eta_p^2 = .082$; proportion of error, $F(1,15) = 8.07$, $p = .012$, $MSE = .005$, $\eta_p^2 = .350$.

Training phase: Colour group. The Colour group demonstrated learning of colour contingencies (see Figure 6.8), although the main effect of Colour Type was confined to the RTs: $F(2,30) = 6.00$, $p = .006$, $MSE = 435$, $\eta_p^2 = .286$; proportion of error, $F(2,30) = .025$, $p = .932$, $MSE = .006$, $\eta_p^2 = .002$. This effect, similar to the Dual Correlated group, was driven by the large effect of slower responses to Predictive Inconsistent versus Control colours, $F(1,15) = 16.2$, $p = .001$, $MSE = 629$, $\eta_p^2 = .520$. Whilst the Predictive Consistent colours were numerically faster and more accurate than the Predictive Inconsistent colours these were not significant differences. Similarly, the difference between Predictive Consistent and Control colours was not reliable.

Test phase: Dual Correlated group. The Dual Correlated group showed no main effect of Colour Type at test (see Figure 6.8), average RT, $F(2,30) = 1.398$, $p = .263$, $MSE = 453$, $\eta_p^2 = .085$; proportion of error, $F(2,30) = .578$, $p = .567$, $MSE = .003$, $\eta_p^2 = .037$. However, the numerical pattern in the errors and the difference between Predictive Consistent and Control colours in the RTs followed the pattern shown during training, although the Predictive Inconsistent trials were surprisingly fast. This suggests that if learning has occurred it may have either: extinguished very rapidly at test; been simply a confound of sequence learning; or that the two blocks do not contain enough trials to capture the sensitive colour learning effect.

Test phase: Dual Uncorrelated group. The only Group to demonstrate a main effect of Colour Type when analysed separately to assess learning at test was Dual Uncorrelated in RTs (see Figure 6.8), average RT, $F(2,30) = 5.036$, $p = .026$, $MSE = 649$, $\eta_p^2 = .251$; proportion of error, $F(2,30) = .832$, $p = .435$, $MSE = .004$, $\eta_p^2 = .053$. This, however, showed the opposite pattern to that obtained during training, providing evidence against the Dual Uncorrelated group learning about the colour contingencies. Both RTs and errors for this group followed the same pattern, with significantly slower RTs, $F(1,15) = .568$, $p = .031$, $MSE = 1416$, $\eta_p^2 = .275$ (and numerically more errors) to Predictive Consistent colours compared to Inconsistent, and to Control colours: average RT, $F(1,15) = 5.33$, $p = .036$, $MSE = 984$, $\eta_p^2 = .262$. Predictive Inconsistent trials were not significantly different from Control colours. This provides evidence that the Dual Uncorrelated group showed no evidence of the expected colour learning at test, and the reason for this reverse pattern is not clear.

Test phase: Colour group. The Colour group also showed no main effect of Colour Type at test (see Figure 6.8), average RT, $F(2,30) = 2.33$, $p = .114$, $MSE = 3748$, $\eta_p^2 = .135$; proportion of error, $F(2,30) = .717$, $p = .497$, $MSE = .003$, $\eta_p^2 = .160$. However, similar to the Dual Correlated group, the pattern of results in the errors showed the same numerical pattern as during training. The RT data at test also demonstrated a significant difference between Predictive Inconsistent and Control colours, $F(1,15) = 6.97$, $p = .019$, $MSE = 1483$, $\eta_p^2 = .317$, with participants making reliably more errors on Predictive Inconsistent trials compared to Control colour trials.

Test phase: Group comparison. The groups did not differ at test in RTs or errors, average RT, $F(2,45) = 2.17$, $p = .126$, $MSE = 41474$, $\eta_p^2 = .088$; proportion of error, $F(2,45) = .228$, $p = .797$, $MSE = .163$, $\eta_p^2 = .010$. When collapsed the groups showed a main effect of Colour Type in RTs at test, average RT, $F(2,90) = 3.85$, $p = .037$, $MSE = 1296$, $\eta_p^2 = .079$; proportion of error, $F(2,90) = .092$, $p = .912$, $MSE = .003$, $\eta_p^2 = .002$. This effect could be unpacked as significantly slower responses to Predictive Consistent trials at test compared to Predictive Inconsistent, average RT, $F(1,45) = 4.54$, $p = .039$, $MSE = 1899$, $\eta_p^2 = .092$; and slower responding to Predictive Consistent trials than to Control colour trials, average RT, $F(1,45) = 4.38$, $p = .042$, $MSE = 3093$, $\eta_p^2 = .089$. There was no difference between Predictive Inconsistent and Control colour trials, average RT, $F(1,45) = .603$, $p = .442$, $MSE = 919$, $\eta_p^2 = .013$. As the numerical pattern is not the same at test, I can conclude little from these results but a possible speed accuracy trade-off, where participants were slower and more accurate to Predictive Consistent colours, and faster but more likely to make a mistake to Predictive Inconsistent colours. Consequently, there is no evidence for colour learning at test in this experiment.

6.3.2.3 Prediction task

Participants performed just below chance on the prediction task, with an average amount of correct choices on the prediction task to Predictive colours of 49.4% ($M = 7.89$, $SE = .527$). Using a prior based on a sample of participants who were instructed to attend to colours and were made aware of relationships between colours and responses from Experiment 6, albeit at a slightly higher contingency (80% rather than 66%), a Bayes factor was calculated. As the prior was taken from a sample who were not only likely to learn more by nature of having information made explicit to them, but also had higher contingencies in play between colours and responses a uniform distribution was

chosen, from 0 (chance performance on the task) to 6.44 (the average number of trials correct above chance on the prediction task in the intentional group used as a prior). The Bayes factor produced from the average number of correct responses greater than 0 is 0.09, providing evidence for the null and suggesting that there was no evidence of explicit knowledge of colour-response contingencies across this sample, demonstrating that the Colour groups are showing evidence of learning about colours without awareness across training. The Dual groups may not be aware of colour contingencies but without evidence of a colour learning effect at test any learning about colours across training may have been purely sequence based.

6.3.3 Discussion

The pattern of results demonstrated in Experiment 8 provides evidence of sequence learning in Dual Correlated and Dual Uncorrelated groups, both of which were trained with sequential contingencies. This learning was apparent across training and during test, with no differences apparent between these two groups. The differences arise when considering the colour learning results, which across training provided evidence of learning by all three groups, having all been trained with colour contingencies. At test however, no groups showed good evidence of colour learning. The Dual Correlated and Colour groups demonstrated a similar pattern of results, following the pattern observed in training across errors and in part in RTs. The Dual Uncorrelated group however, showed an absence of any colour learning at test, with both RT and error data showing the opposite numerical pattern to what one would expect if showing evidence of learning, supported by significant differences between Colour Types in the opposite direction. This could suggest that in the Dual Uncorrelated group any weak learning about colour that could have occurred was overshadowed by the presence of the sequential contingencies, but I have to acknowledge that the evidence for this conclusion is rather weak.

This suggests that the Dual Correlated group had somehow been protected from overshadowing by the contingencies in force between the colours and subsequences themselves. This could be explained by a variety of learning theories in which either the representation of colour and stimuli were bound together whereby the association between colour and sequence was built up, and the presentation of one activates a representation of the other.

Does this effect only apply to the colours and not the sequences? This might suggest that the weaker or less salient stimulus (in this case colour) has little to add to the robust sequence learning that we observe in both Dual groups, which may already be at ceiling. If this were the case, Experiments 3 and 4 would not have produced evidence of the potentiation of sequence learning. It could be that the effect of colour on sequence learning is minimal, especially and not detected by the group comparison here, as there are more colours than in Experiment 4. The lack of sequence learning potentiation may have been affected by the presence of Non-Predictive colours, which may reduce the potentiating effect of the 100% sequence pair correlated Predictive colours in the Dual Correlated group. Further to this, the potentiation of sequence learning seen in Experiments 3 and 4 in the Previous group involved a contingency with the previous element in the sequence, not the previous two elements in the sequence. Perhaps participants are able to learn these simple sequence-stimuli relationships only about the previous trial, or within a certain temporal window. Additionally, it could be possible that the presence of the Previous stimuli in Experiments 3 and 4 improved learning in some other, non-associative way.

This experiment did not provide definitive evidence of colour learning in any group, however, as there were no significant effects at test save one difference in the Colour RTs to suggest that *any* group can learn about colours during extinction. However, as the Colour group demonstrated evidence of learning across training, which is not confounded or produced by sequential contingencies, the test phase may not be sensitive enough to capture the effect. Extinction of colour contingency learning may be occurring very rapidly. To assay colour learning fully against a control, a Sequence group with no colour contingencies was run in the next experiment. In addition to this, the problems with the allocations of colours during test in this experiment due to a programming error may have influenced the results. Experiment 9 addresses this issue, provides a control group for colour learning, and seeks to confirm the tentative conclusions drawn from Experiment 8.

6.4 Experiment 9

The final experiment in this thesis examined once more the difference in sequence and colour learning across two Dual groups, replicating Experiment 8. Experiment 9 also incorporated two groups that received only one contingency with the other stimuli

simply random – a Colour and Sequence group. The issues with the test sequencing in Experiment 8 were corrected and a comparison to a control group for colour learning (the Sequence group) as a comparison at test was essential. When measuring sequence learning the Colour group forms a control across training as well as at test, as no sequences were trained in this group. When measuring colour learning the Sequence group forms the appropriate control across training and test, as no positive colour-cue response contingencies were trained. Therefore, the sequence group offers a control for the unavoidable confound of measuring the colour learning across training on trials that also measure sequence learning in the Dual groups, as I can compare the Sequence group to the Dual and Colour groups to assess colour learning; and compare the Colour group to the Dual and Sequence groups to assess sequence learning.

6.4.1 Method

6.4.1.1 Participants

64 participants (aged between 18 and 49 [M = 21.5]; 53 female and 11 male) were recruited from first year psychology undergraduate students and were awarded one credit (N = 11) or £5 (N = 53) in return for participation. Participants provided informed consent prior to taking part in one session lasting roughly one hour. Participants were allocated into one of three between subject groups: Dual Correlated; Dual Uncorrelated; Sequence or Colour.

6.4.1.2 Materials and Stimuli

The materials and stimuli were the same as for Experiment 7 (see 6.2.1.2).

6.4.1.3 Design

The experiment followed the design of Experiment 8 (see 6.3.1.3), with the addition of a Sequence group whose sequence construction and colour construction was the same as Experiment 7 (see 6.2.1.3).

6.4.1.4 Procedure

The procedure was a direct replication of Experiment 8 (see 6.3.1.4).

6.4.2 Results

RTs and errors were measured and sequence learning difference scores were calculated as described in Experiment 7 (see 6.2.2). Colour learning was assayed using raw RTs and proportion of errors as described in Experiment 8 (see 6.3.2).

6.4.2.1 Sequence learning

An analysis of variance was conducted on both RT and error difference scores across training and test with the factors Block (Training: 18; Test: 2) x Subsequence (4) x Group (4).

Training phase. There was a main effect of the variable of interest, Group, in both RT and proportion of error difference scores across training (see Figure 6.9), RT difference score, $F(3,60) = 13.8$, $p < .001$, $MSE = 2432$, $\eta_p^2 = .409$; error difference score, $F(3,60) = 20.7$, $p < .001$, $MSE = .027$, $\eta_p^2 = .256$.

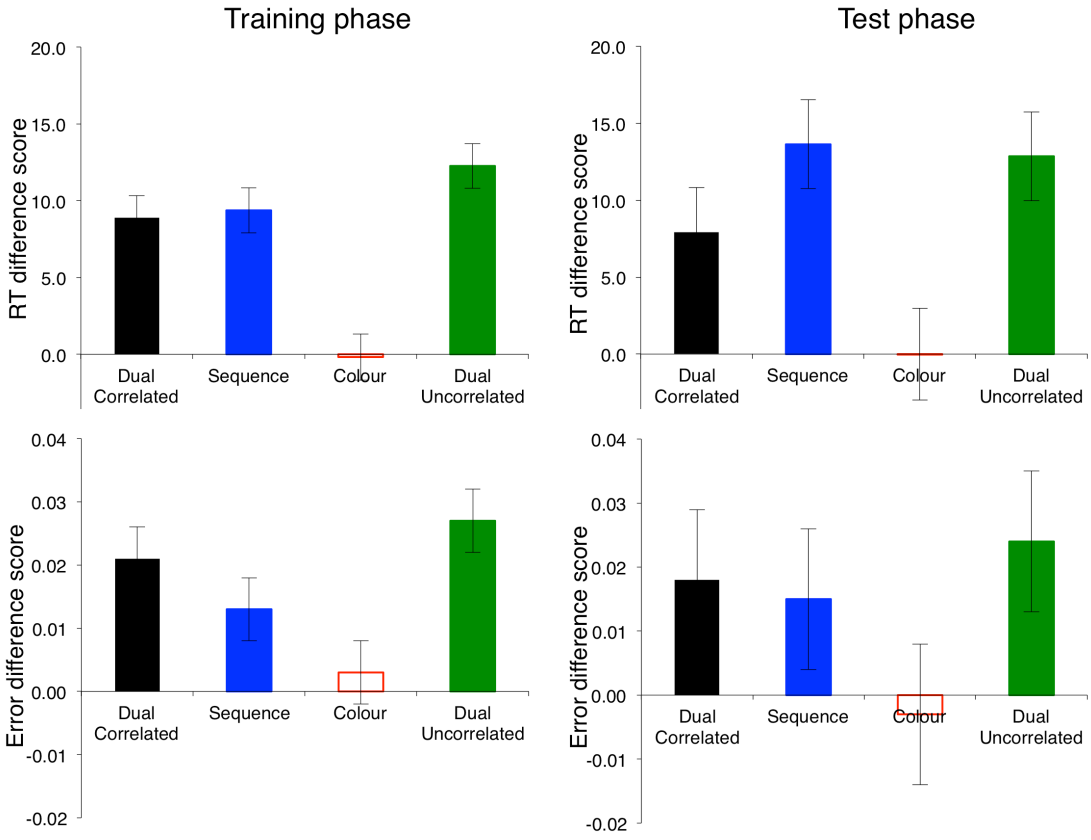

Figure 6.7. RT (top panel) and proportion of error (bottom panel) sequence learning difference scores across training (left panel) and test (right panel) for the four groups in the study: Dual Correlated (black filled bars); Sequence (blue filled bars); Colour (red open bars); and Dual Uncorrelated (green filled bars). Filled bars indicate those groups who were trained with exclusive-or sequences, open bars the Colour group who were not trained with exclusive-or sequences. Error bars show standard error.

Using a series of planned comparisons shown in Table 6.7 we can break this down and show that the Dual Correlated group showed significant learning of sequential contingencies, with higher RT and proportion of error difference scores compared to the Colour group who were not trained with any sequential contingencies. The Dual Uncorrelated group also showed significant learning across both difference score measures compared to the control Colour group (see Table 6.7). The Sequence group showed numerically higher scores in errors and a significantly higher RT difference score compared to controls (see Table 6.7). Therefore each group that was trained with sequential contingences showed evidence of learning about these contingencies. There was no difference across training between either of the Dual groups; nor was the Dual Correlated group different to the Sequence group, see Table 6.7. The Dual Uncorrelated group was no different to the Sequence group in RT difference score, but was however, in the proportion of errors, see Table 6.7. Similar to the account of the Random group given in Chapter 3, it is possible that the random noise generated by the non-predictive colours in the Sequence group may have an adverse effect on some adaptive learning rate, thus reducing sequence learning.

Table 6.7. Results from ANOVA for planned comparisons on average RT and proportion of errors sequence difference scores, comparing the four Groups to one another across the eighteen blocks of training in Experiment 9.

	RT difference score: Training			
	F (df = 1,45)	p	MSE	η_p^2
Dual Correlated vs Colour (control)	19.6	< .001	33.7	.246
Dual Uncorrelated vs Colour (control)	36.9	< .001	33.7	.381
Sequence vs Colour (control)	21.8	< .001	33.7	.266
Dual Correlated vs Sequence	.057	.812	33.7	.001
Dual Uncorrelated vs Sequence	1.99	.164	33.7	.032
Dual Correlated vs Dual Uncorrelated	2.72	.105	33.7	.105
	Error difference score: Training			
	F (df = 1,45)	p	MSE	η_p^2
Dual Correlated vs Colour (control)	6.74	.012	.001	.101
Dual Uncorrelated vs Colour (control)	12.8	.001	.001	.175
Sequence vs Colour (control)	2.17	.146	.001	.035
Dual Correlated vs Sequence	1.26	.266	.001	.021
Dual Uncorrelated vs Sequence	4.40	.040	.001	.068
Dual Correlated vs Dual Uncorrelated	.952	.333	.001	.016

There was a significant main effect of Block across training in the RTs, RT difference score, $F(17,1020) = 2.30$, $p = .019$, $MSE = 3731$, $\eta_p^2 = .037$; proportion of error difference score, $F(17,1020) = 1.020$, $p = .432$, $MSE = .014$, $\eta_p^2 = .017$. This did not show evidence of an interaction with Group, which suggests that participants overall improved their speed responding to consistent subsequence versus inconsistent subsequences across the experiment.

Test phase. There was a main effect of Group across test in the RT difference scores (see Figure 6.9), RT difference score, $F(3,60) = 5.65$, $p = .002$, $MSE = 1294$, $\eta_p^2 = .220$; proportion of error difference score, $F(3,60) = 1.21$, $p = .314$, $MSE = .014$, $\eta_p^2 = .057$. The difference between both the Dual and Sequence groups (who have sequential contingences present throughout training) compared to Control colours provides numerical support for the claim that all groups have learned at test, see Table 6.8. When comparing the Dual Correlated, Dual Uncorrelated and Sequence groups at test there were no significant differences in either difference score measure, see Table 6.8.

Table 6.8. Results from ANOVA for planned comparisons on average RT and proportion of errors sequence difference scores, comparing the four Groups to one another across the two blocks of test in Experiment 9.

	RT difference score: Test			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual Correlated vs Colour (control)	5.68	.020	162	.086
Dual Uncorrelated vs Colour (control)	12.1	.001	162	.168
Sequence vs Colour (control)	13.3	.001	162	.182
Dual Correlated vs Sequence	1.61	.210	162	026
Dual Uncorrelated vs Sequence	.030	.864	162	.001
Dual Correlated vs Dual Uncorrelated	1.20	.278	162	.020
	Error difference score: Test			
	F (df = 1,45)	*p*	*MSE*	η_p^2
Dual Correlated vs Colour (control)	1.98	.165	.002	.032
Dual Uncorrelated vs Colour (control)	3.25	.076	.002	.051
Sequence vs Colour (control)	1.51	.224	.002	.025
Dual Correlated vs Sequence	.031	.860	.002	.001
Dual Uncorrelated vs Sequence	.330	.568	.002	.005
Dual Correlated vs Dual Uncorrelated	.158	.692	.002	.003

6.4.2.2 Colour learning

An analysis of variance was conducted on RT and error difference scores across training and test with the factors Block (Training: 18; Test: 2) x Group (4). These scores were calculated from the difference between average RT and proportion of errors to Predictive Inconsistent trials minus performance on Predictive Consistent trials.

Training phase: Difference scores. There was a main effect of Group across training RT difference scores, $F(3,60) = 8.23$, $p < .001$, $MSE = 3743$, $\eta_p^2 = .291$; and error difference scores, $F(3,60) = 4.66$, $p = .005$, $MSE = .032$, $\eta_p^2 = .189$, see Figure 6.8.

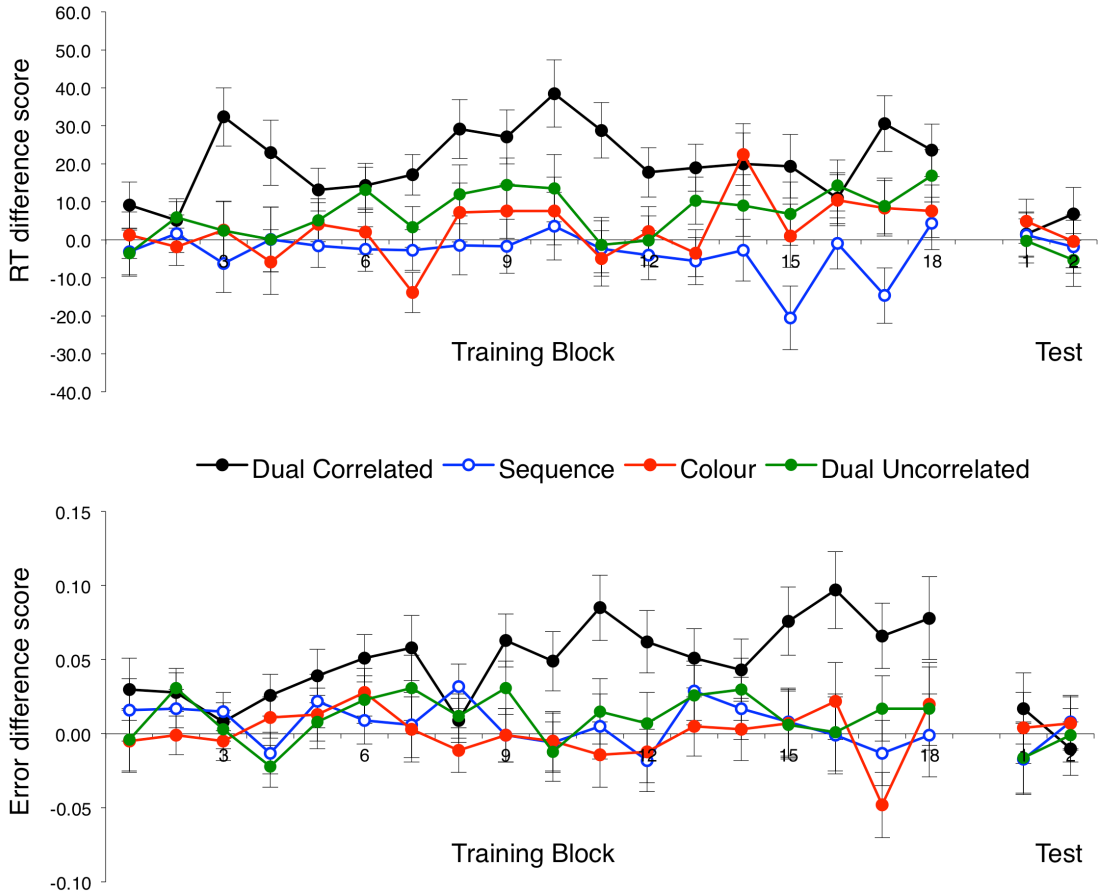

Figure 6.8. RT (top panel) and error (bottom panel) colour learning difference scores across training and Blocks for the four Groups: Dual Correlated (black filled circles); Dual Uncorrelated (green filled circles); Colour (red filled circles) and Sequence (blue open circles). Error bars show standard error.

There was a significant effect of Group when comparing Dual Correlated to the Sequence group, RT difference score, $F(1,30) = 12.6$, $p = .001$, $MSE = 6815$, $\eta_p^2 = .296$; proportion of error difference score, $F(1,30) = 4.76$, $p = .037$, $MSE = .059$, $\eta_p^2 = .137$,

providing evidence for Dual Correlated colour learning. The Dual Uncorrelated group show a large effect compared to the Sequence group only in speed of response, RT difference score, $F(1,30) = 21.1$, $p < .001$, $MSE = 773$, $\eta_p^2 = .413$; proportion of error difference score, $F(1,30) = .775$, $p = .386$, $MSE = .005$, $\eta_p^2 = .025$. Of course, in both Dual groups colour learning is to some extent confounded with sequence learning. Importantly, the Colour group show an effect in RTs compared to the Sequence group, $F(1,30) = 6.54$, $p = .016$, $MSE = 892$, $\eta_p^2 = .179$; proportion of error difference score, $F(1,30) = .889$, $p = .353$, $MSE = .006$, $\eta_p^2 = .029$. Therefore the Colour group score significantly higher RT difference scores than the Sequence (control for colour learning) group. As sequences are not confounded with colours in this group this provides evidence of colour cue-response learning.

Separate analyses of variance were conducted for raw average RT and proportion of errors for each group, with the factors Block (Training: 18; Test: 2) x Colour Type (3) to assess learning within the group. This enables us to further examine the pattern of responding compared to the Non-Predictive colours also.

Training phase: Dual Correlated group. The Dual Correlated group demonstrated a main effect of Colour Type (see Figure 6.9), average RT, $F(2,30) = 9.78$, $p = .005$, $MSE = 5921$, $\eta_p^2 = .395$; proportion of error, $F(2,30) = 6.58$, $p = .019$, $MSE = .054$, $\eta_p^2 = .305$. This was supported by differences between all of the Colour Types across both RTs and errors, firstly with significantly faster and more accurate responding to Predictive Consistent trials compared to Predictive Inconsistent trials, average RT, $F(1,15) = 10.1$, $p = .006$, $MSE = 12635$, $\eta_p^2 = .402$; proportion of error, $F(1,15) = 6.72$, $p = .020$, $MSE = .112$, $\eta_p^2 = .316$. Participants were also significantly faster and more accurate in training on Predictive Consistent trials compared to Control colour trials, average RT, $F(1,15) = 4.89$, $p = .043$, $MSE = 3941$, $\eta_p^2 = .246$; proportion of error, $F(1,15) = 5.20$, $p = .038$, $MSE = .026$, $\eta_p^2 = .257$. Finally, Predictive Inconsistent trials resulted in slower and less accurate responses compared to Control colour trials, average RT, $F(1,15) = 14.3$, $p = .002$, $MSE = 49943$, $\eta_p^2 = .489$; proportion of error, $F(1,15) = 7.19$, $p = .017$, $MSE = .034$, $\eta_p^2 = .324$. This provides strong evidence for learning of colour contingencies across training in the Dual Correlated group.

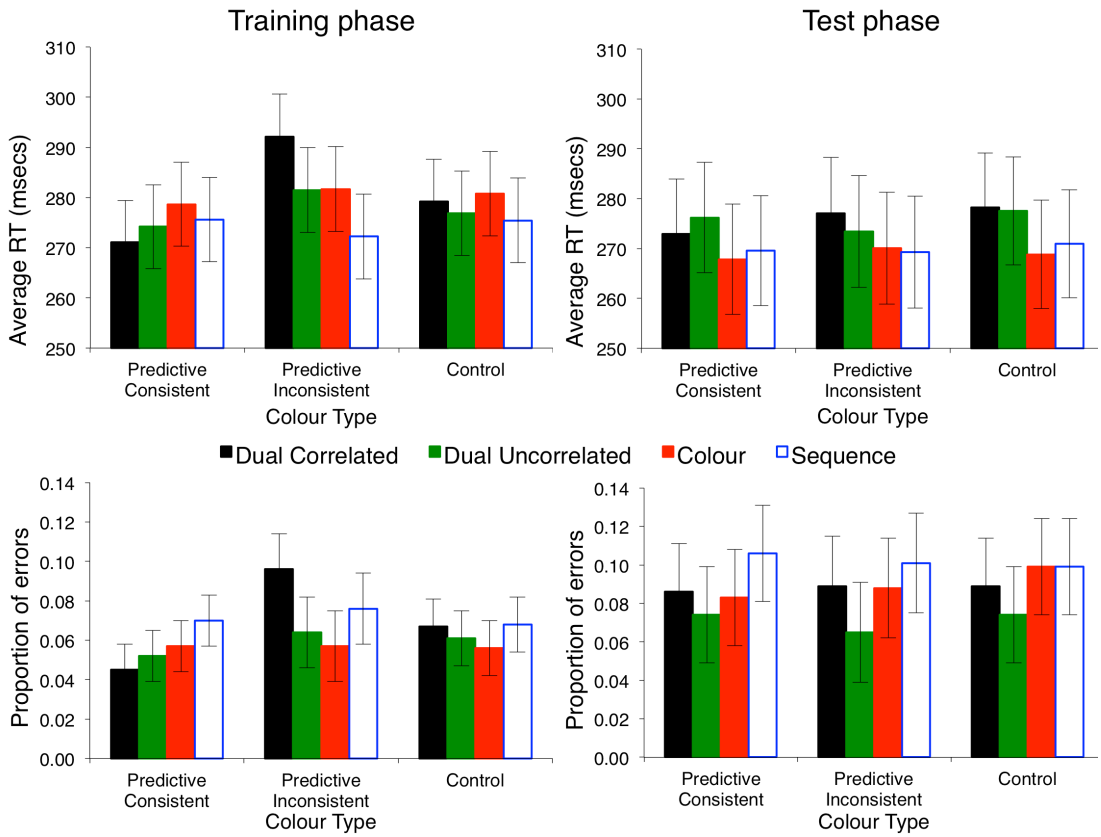

Figure 6.9. Average RT (top panel) and proportion of errors (bottom panel) across training (left panel) and test (right panel) for the four Groups: Dual Correlated (black filled bars); Dual Uncorrelated (green filled bars); Colour (red filled bars) and Sequence (blue open bars) on the three different Colour Types: Predictive Consistent; Predictive Inconsistent; and Control. Error bars show standard error.

Training phase: Dual Uncorrelated group. The Dual Uncorrelated group also demonstrated a large main effect of Colour Type (see Figure 6.9), average RT, $F(2,30) = 17.9$, $p < .001$, $MSE = 218$, $\eta_p^2 = .545$; proportion of error, $F(2,30) = 8.08$, $p = .002$, $MSE = .002$, $\eta_p^2 = .350$. Predictive Consistent trials were responded to more quickly and accurately than both: Predictive Inconsistent trials, average RT, $F(1,15) = 27.8$, $p < .001$, $MSE = 550$, $\eta_p^2 = .650$; proportion of error, $F(1,15) = 10.9$, $p = .002$, $MSE = .004$, $\eta_p^2 = .420$; and Control trials, average RT, $F(1,15) = 5.58$, $p = .032$, $MSE = 371$, $\eta_p^2 = .270$; proportion of error, $F(1,15) = 15.1$, $p = .001$, $MSE = .002$, $\eta_p^2 = .510$. Predictive Inconsistent trials were responded to significantly slower and numerically with more errors than Control colour trials, average RT, $F(1,15) = 15.8$, $p = .001$, $MSE = 5815$, $\eta_p^2 = .512$; proportion of error, $F(1,15) = .864$, $p = .367$, $MSE = .003$, $\eta_p^2 = .054$.

Training phase: Colour group. The Colour group did not demonstrate a main effect of Colour Type, average RT, $F(2,30) = 1.64$, $p = .217$, $MSE = 510$, $\eta_p^2 = .098$; proportion of error, $F(2,30) = .088$, $p = .884$, $MSE = .002$, $\eta_p^2 = .006$, however the RTs showed the pattern one would expect if colour learning had occurred (see Figure 6.8). Colour Type did interact with Block in RTs (see Figure 6.10), average RT, $F(34,510) = 1.82$, $p = .029$, $MSE = 556$, $\eta_p^2 = .108$; proportion of error, $F(34,510) = 1.33$, $p = .193$, $MSE = .004$, $\eta_p^2 = .082$, with Control trials showing little change over the experiment, whereas slowing occurs across the experiment to both Predictive Colours, regardless of whether Consistent or Inconsistent with the trained contingency.

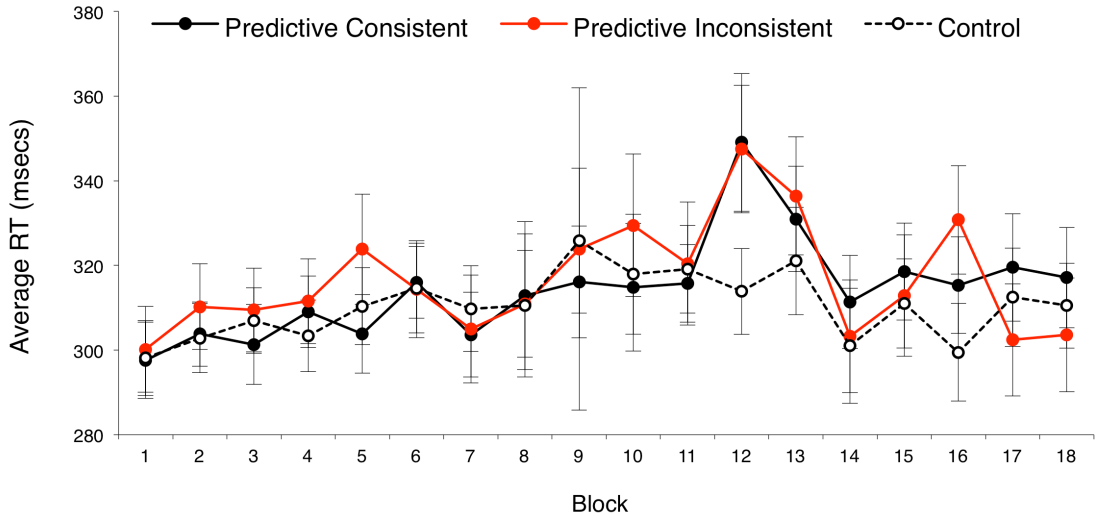

Figure 6.10. Average RTs of participants in the Colour group only across training blocks for the three Colour Types. Predictive Consistent trials (black filled circles); Predictive Inconsistent (red filled circles); and Non-Predictive (black open circles) are plotted. Error bars show standard error.

Training phase: Sequence group. Finally, the Sequence group showed no evidence of learning, with no main effect of Colour Type, average RT, $F(2,30) = 2.97$, $p = .070$, $MSE = 373$, $\eta_p^2 = .166$; proportion of error, $F(2,30) = 2.20$, $p = .129$, $MSE = .003$, $\eta_p^2 = .128$. There was a difference between Predictive Inconsistent and Control colour trials, but in the opposite direction in RTs to the differences observed in the other groups, average RT, $F(1,15) = 5.02$, $p = .041$, $MSE = 8843$, $\eta_p^2 = .251$; proportion of error, $F(1,15) = 5.02$, $p = .041$, $MSE = .004$, $\eta_p^2 = .251$, therefore Predictive Inconsistent trials were responded to significantly faster yet less accurately than Control trials. This provides evidence that the sequence group learned nothing about colour-response contingencies; indeed they had no colour contingencies from which to learn.

Test phase. At test, there was no main effect of group across the difference score measures in either RT, $F(3,60) = .508$, $p = .678$, $MSE = 576$, $\eta_p^2 = .025$; or errors, $F(3,60) = .180$, $p = .909$, $MSE = .008$, $\eta_p^2 = .009$. No groups show any significant effects across any variable or comparison in raw average RTs or proportion of errors, with the main effect of interest, Colour Type, showing no evidence of learning in the Dual Correlated group: average RT, $F(2,30) = .871$, $p = .429$, $\eta_p^2 = .055$; proportion of error, $F(2,30) = .025$, $p = .968$, $\eta_p^2 = .002$; Dual Uncorrelated group: average RT, $F(2,30) = .966$, $p = .388$, $\eta_p^2 = .060$; proportion of error, $F(2,30) = .332$, $p = .630$, $\eta_p^2 = .022$; Colour group: average RT, $F(2,30) = .231$, $p = .795$, $\eta_p^2 = .015$; proportion of error, $F(2,30) = 2.170$, $p = .132$, $\eta_p^2 = .126$; nor Sequence group: average RT, $F(2,30) = .056$, $p = .940$, $\eta_p^2 = .004$; proportion of error, $F(2,30) = .163$, $p = .778$, $\eta_p^2 = .011$. Whilst not significant, both Dual Correlated and Colour groups show numerically faster and more accurate responding to Predictive Consistent trials over Predictive Inconsistent trials. The opposite pattern is demonstrated in both Dual Uncorrelated and Sequence groups, providing no support for the suggestion that either Dual Uncorrelated or Sequence groups may have learned colours at test.

When compared to one another there is no main effect of Group at test, average RT, $F(3,60) = .122$, $p = .947$, $\eta_p^2 = .006$; proportion of error, $F(3,60) = .841$, $p = .014$, $\eta_p^2 = .100$, suggesting no difference in the response times or errors made by the groups. There is also no evidence of a Colour Type effect, average RT, $F(2,120) = .631$, $p = .534$, $\eta_p^2 = .010$; proportion of error, $F(2,120) = .230$, $p = .756$, $\eta_p^2 = .004$, nor an interaction between Colour Type and Group, average RT, $F(6,120) = .363$, $p = .901$, $\eta_p^2 = .018$; proportion of error, $F(6,120) = .369$, $p = .870$, $\eta_p^2 = .011$.

6.4.2.3 Prediction task and post-experimental interview

Participants performed at chance on the prediction task, with an average amount of correct choices on the prediction task to Predictive colours at 51.8% ($M = 8.27$, $SE = .438$). Using a prior based on a sample of participants from Experiment 6, as for Experiment 8 a Bayes factor was calculated. As the prior was taken from a sample who were not only likely to learn more by nature of having information made explicit to them, but also had higher contingencies in play between colours and responses, a uniform distribution was chosen: from 0 (chance performance on the task) to 6.44 (the average number of trials correct above chance on the prediction task in the intentional group used as a prior). The Bayes factor produced is 0.15, giving evidence for the null

and suggesting that there is no evidence of learning across this sample. Examining only the groups who were trained with colour (M = 8.33, SE = .492) participants scored 0.33 correct on average above chance. Using only these groups in the same analysis as previously described, the Bayes factor was 0.18, still providing evidence for the null and demonstrating that the Dual Correlated, Dual Uncorrelated and Colour groups showed evidence of learning without awareness.

When guessing which colours had a relationship with a response in the structured interview participants also responded with chance accuracy, with 34 out of 64 participants choosing the correct colour that predicted a left response and 31 out of 64 choosing the correct colour that predicted a right response. These participants, however, did not all identify both the colour and the correct response location – with 25 of the 34 participants correct on the left colour suggesting that it predicted a right response and only 9 reporting the correct contingency. The right colour was predicted with the same degree of accuracy, with 19 participants suggesting that it predicted a left and 12 accurately reporting that it predicted a right response. None of the participants could accurately describe the sequential rule, but some reported being aware of strings of responses on one side, or runs of trials where responses would alternate from one side to the other.

6.4.3 Discussion

This study provided evidence of learning about sequences, with Dual Correlated, Uncorrelated and Sequence groups showing strong evidence for sequence learning across training and at test compared to the control Colour group. This is not surprising, given the body of evidence provided so far that suggests that participants demonstrated strong learning of sequences. Sequence learning was unaffected by manipulations of the colour stimuli and their presence or relationships with the sequence.

The colour learning observed in this study was again weaker than the learning of sequences, and whilst Dual Correlated, Dual Uncorrelated and Colour groups all showed evidence of learning across training, none of these effects remained at test in either RT or error measures. As the Dual groups experience colour and sequence learning that is confounded across training, we therefore have no way to compare these groups as colour learning is obviously too weak to show up reliably at test. This may

also suggest as previous studies have, that participants did not or could not demonstrate learning of colour contingencies. However, learning was evident across training of colour contingences in the Colour group, who had no sequential information across training to confound, potentiate or give any advantage to any particular response. Therefore, it seems likely that colour learning itself suffers rapidly from extinction.

6.5. Simulation 11: RASRN simulation

To further investigate the processes at play, Experiment 9 was simulated using the RASRN. Whilst previous chapters suggest that the model cannot capture the detailed pattern of learning, it remains one of few models able to both include a representation of the stimulus conditions as well as accounting for human incidental sequence learning. As Chapter 4 was inconclusive regarding the interaction of the processes involved in human sequence learning with concurrent stimuli; this chapter seeks to integrate the results of the human experiments within the simulation context. This simulation therefore was not assumed to be an excellent model of the human data, given the performance of the RASRN in Chapter 4. However, due to the correlation between sequences and concurrent stimuli in Simulation 9 (Chapter 4), it was impossible to analyse learning about the stimuli separately. Whilst Experiments 8 and 9 may be inconclusive, the mechanisms underlying sequence and cue-response learning do seem to interact.

6.5.1. Simulation details

The RASRN was run for 64 networks following the procedure outlined for Experiment 9 and the parameters for Simulation 10 (Chapter 5). Again the input units represented the two previous required responses (on activation = 0.75); two current response stimuli locations (on activation = 0.1); four possible cue colours (on activation = 0.4); and the context units (1.3 times the activation of the hidden units on the previous trial). Each network was run for the same number of trials as each participant, with 16 networks in each of the four conditions.

6.5.2. Results

6.5.2.1. Sequence learning

Sequence learning was assessed in the same way as for participants using the inconsistent minus consistent sequence learning difference score for the trained exclusive-or sequences in the Dual Correlated, Sequence and Dual Uncorrelated networks. Colour networks acted as a control for sequence learning. The MSE difference scores from training and test were analysed by an ANOVA on Block (18 training; 2 test), Subsequence (4) and Group. There was a main effect of Group across training, $F(3,60) = 28.0$, $p < .001$, $MSE = .002$, $\eta_p^2 = .583$; but not test, $F(3,60) = 2.20$, $p = .097$, $MSE = .001$, $\eta_p^2 = .099$, which is shown in Figure 6.11.

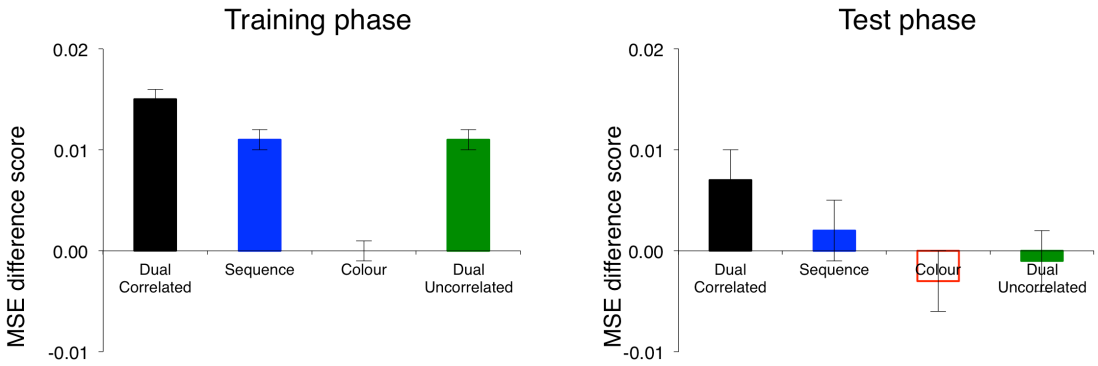

Figure 6.11. MSE sequence learning difference scores across training (left panel) and test (right panel) for the four groups of networks in Simulation 11: Dual Correlated (black filled bars); Sequence (blue filled bars); Colour (red open bars); and Dual Uncorrelated (green filled bars). Filled bars indicate those groups who were trained with exclusive-or sequences, open bars the Colour group who were not trained with exclusive-or sequences. Error bars show standard error.

As we can see from the Figure, networks learned significantly more in all three groups that were trained with sequences than in the Colour group as shown across training by Bonferroni corrected comparisons: Dual Correlated versus Colour, $p < .001$; Sequence versus Colour, $p < .001$; Dual Uncorrelated versus Colour, $p < .001$. No other comparison was significant, nor were any groups significantly different at test. This suggests that sequence learning progressed regardless of the presence of colour-response contingencies, although the Dual Correlated networks approach a significantly higher difference score compared to the control Colour networks, $p = .123$. This, to some extent, follows the performance of human participants as sequence learning was not significantly different depending on the colours. On the other hand, the model does

not provide evidence of sequence learning at test, which is in clear contrast to the performance of human participants.

6.5.2.2. Colour cue-response learning

The MSE scores for networks were analysed in the same was as for participants in Experiment 9, with difference scores taken between Predictive Inconsistent and Predictive Consistent trials. These were analysed in an ANOVA across Block (training: 18, test: 2). A main effect of Group occurs across both training, $F(3,60) = 89.5$, $p < .001$, $MSE = .002$, $\eta_p^2 = .817$, and test, $F(3,60) = 377.3$, $p < .001$, $MSE = .001$, $\eta_p^2 = .950$, shown in Figure 6.12. Networks in Colour, $p < .001$, and Dual Uncorrelated groups, $p < .001$ performed better than the control Sequence group across training, however there was no colour cue-response learning evident in the Dual Correlated group, $p > .9$. This might suggest that there was no colour learning occurring, but what it actually reflects is, like in the human participants, training data is confounded by sequence learning. Whilst in the case of humans the sequence learning is far greater than the colour cue-response learning, in the case of the RASRN sequence learning is far smaller than for humans (see the scales on Figures 6.11 and 6.12). Therefore the expression of colour learning across training is being restricted by the correlation with sequences, which are significantly learned in the Dual Correlated group at test, but these MSE difference scores are not large.

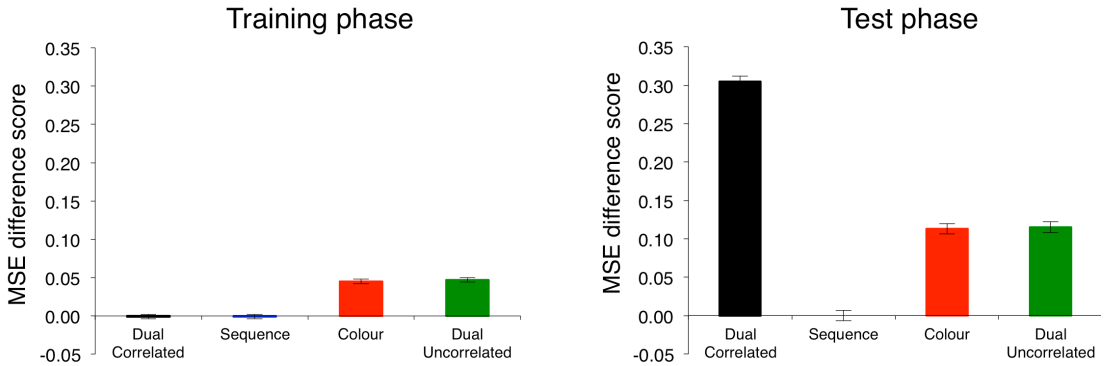

Figure 6.12. MSE colour learning difference scores across training (left panel) and test (right panel) for the four groups of networks in Simulation 11: Dual Correlated (black filled bars); Sequence (blue open bars); Colour (red filled bars); and Dual Uncorrelated (green filled bars). Filled bars indicate those groups who were trained with exclusive-or sequences, open bars the Colour group who were not trained with exclusive-or sequences. Error bars show standard error.

At test, however, when colours and sequences no longer correlate we can see that the Dual Correlated group have learned significantly more than not only the control Sequence networks, $p < .001$; but also both the Colour, $p < .001$ and Dual Uncorrelated networks, $p < .001$. Both Colour, $p < .001$ and Dual Uncorrelated, $p < .001$ groups show evidence of learning about colour cue-response contingencies at test; therefore this provides evidence not of overshadowing but potentiation of colour learning in the Dual Correlated group.

6.5.3. Discussion

The results of the RASRN simulations clearly show the incorrect pattern of learning effects, as firstly Colour learning difference scores are far higher than those for sequence learning. This suggests that the model is learning the colour cue-response contingencies far better (or sequences far worse) than human participants. This seems to result in the opposite pattern of results observed in humans, namely the overshadowing of sequence learning in the Dual Uncorrelated networks; whereas colour learning was overshadowed in the experiment run by McLaren et al. (2013) which are supported somewhat by the numerical scores in Experiments 8 and 9. Furthermore, the Dual Correlated networks produce evidence of colour potentiation, evidence of which was not supported by the experimental human results from this chapter, as the lack of overshadowing was the extent of the facilitatory effect of sequences on colour learning.

This could suggest that the stimulus units receive an activation value that is far too high, which may go some way to explain the terrible performance of the Current group when simulated in Chapter 4. The stimulus-response learning was such that it interfered with sequence learning, causing decreasing difference scores for all groups – but more rapidly for the Current group. By reducing the activation of these stimulus units the performance of the Current group may be improved, but this would not produce an increase in the sequence learning produced by the Previous networks. That the model does produce an overshadowing effect in the Dual Uncorrelated group and a potentiation effect in the Dual Correlated group does, however, go some way to support the idea that stimuli that are related may not suffer from cue-competition and may indeed come to facilitate learning about the other. In summary, however, it is clear that the RASRN in its current form is not a good model of human performance on these SRT tasks.

6.6 General Discussion

All of the studies presented in Chapter 6 have in common a robust demonstration of sequence learning of the exclusive-or rule. Each of the groups in each experiment, regardless of whether sequences were trained in the presence of colour contingencies or random colours, demonstrate learning of sequences. This indicates that human learning systems are well equipped to learn these sequential contingencies. Experiment 7 provided evidence of Colour cue-response learning in both the Colour and Dual groups at test, allowing us to observe that this colour learning was at least no worse in the Dual group, providing no evidence for cue-competition. Taken with the work of McLaren et al. (2013) this suggests that an absence of cue-competition may occur if the competing stimuli are themselves related.

Crucially, the between-subject Experiments 8 and 9 did not find evidence of colour cue-response learning at test, which meant that a comparison was not possible between the groups. Whilst we find evidence of colour learning in Chapter 5, it may be possible that the learning was not strong enough to survive extinction, especially as the contingencies in this task were reduced from 80% in Chapter 5 to 67% in Experiments 7 to 9. This was in order to match the Dual and Colour groups, as well as to match the contingency between sequences and colour cues with response stimuli. Experiment 7 provided evidence that Colour learning did occur, although this was only in the errors at test, and therefore the experiment was not altered to account for this. Using the Sequence group as a control for performance effects across training we saw evidence for Colour learning in Experiment 9 across training, however it is at test where this is of crucial importance. It is a requirement of the task design to compare the Dual groups to the Colour group in the absence of sequential contingencies (at test), but the colour cue-response learning obviously suffers from rapid extinction. This may suggest that cue-response learning is an unsuitable form of learning to attempt to investigate cue-competition.

The lack of colour learning evident in Experiment 7, as well as the overshadowing observed in McLaren et al. (2013) and Dual Uncorrelated groups in Experiments 10 and 11 suggest that colour cue-response contingencies are more difficult for people to learn. When matched for frequency and relationship with an outcome, humans may find sequential contingencies easier to learn than simple cue-outcome contingencies, a result not demonstrated outside of these studies. This suggests that the system learning these contingencies must, in some way, give preference to the learning of sequential

dependencies. This is discussed further in the context of the other results provided by this thesis in Chapter 7, as it may suggest any number of possibilities. The role of time, spatial locations, trial order, responding and complexity are all features of sequences. Any one of these features may be preferred or given priority to by automatic learning, or the many facets of sequences may sum together to increase learning.

Experiment 7 provided evidence contrary to McLaren et al. (2013), who found that colour-response contingency learning would be overshadowed by sequence learning. Evidence for the opposite effect on colour-response contingency learning was found, with the potentiation of colour learning in the Dual group at test compared to the Colour group. Experiments 8 and 9 were unable to confirm that the key difference between McLaren et al. (2013) was the between cue relationships, but the numerical differences do not disagree with the suggestion that Experiment 7 protected colour learning due to the relationship between sequence and colours. This can be explained associatively, as in the Dual group of Experiment 7 sequence learning itself provides a between-cue association with colour learning that means that as the sequence and colours became associated, so increasingly do the colours and responses.

Beesley and Shanks (2012) initially dismissed the absence of cue-competition in their visual search tasks as evidence for a single process, propositional account of human learning. Indeed they saw learning and perhaps some evidence for an increase in learning about the blocked cue. The cue competition studies in this chapter suggest that this is not clear cut evidence for a propositional account of learning, as it suggests that cue-competition is less likely to occur when stimuli are related. Beesley and Shanks (2012) use a number of distractor patterns and cue locations, which are different for each subject, so it seems unlikely that these patterns are all somehow related to one another. However, it is possible that because these are the same modality, or share the common feature of the target stimulus location, they are represented configurally as a whole or as a shared set of elements whose between-cue associations are strong. Both elemental and configural accounts could explain these results and the results found in Experiment 7.

These studies therefore suggest that, similar to the proposals of Urcelay and Miller (2009), there is no need to propose that potentiation is a result of a different learning process when the procedures and stimuli themselves may interact to facilitate learning

or result in cue-competition. Urcelay and Miller (2009) suggest that this is a result of the configural processing of stimuli that overlap in time. They found that trace conditioning (where there is a temporal gap between CS and US) results in potentiation whereas delay conditioning (where CS and and US overlap) results in overshadowing. These studies do not find evidence for this mechanism per se, as the differences between Dual groups are not in the temporal overlap of sequential and colour contingencies. This suggests that a between-cue association serves to facilitate learning, which may be based on either time or frequency or between stimulus contingency.

Whilst not by any means an exhaustive investigation into cue-competition effects, this chapter demonstrated that an absence of cue-competition may not indicate the absence of associative processes. A lack of overshadowing and some evidence for the potentiation of cue-response learning was found in at least one experiment. However, the other tasks suffered from a lack of colour learning in general and solid conclusions can not be made. Taken with the results of previous work (McLaren et al., 2013), this chapter provides an indication that cue competition may arise from some difference in the ability to learn sequences and simple cue-response contingencies, as one is clearly learned more easily than the other. While I do not find reliable evidence for cue competition in this chapter, this is due to the difficulty that participants have in learning these simple cue-response contingencies under incidental conditions. This reflects the importance of understanding sequence learning in order to understand automatic learning processes as it further reveals a preference in human incidental learning that is contrary to the intuitive assumptions of a variety of authors (Cleeremans, 1997; Jiménez & Méndez, 1999). This is hard to reconcile propositionally as it suggests that learning has an advantage for more complex, sequential stochastic contingencies. We may be able to further understand human associative learning by investigating the absence of cue-competition when the two contingencies themselves relate.

Chapter 7. General discussion and conclusions

The experiments in this thesis have provided evidence of both sequence and cue-response learning under incidental conditions that go some way to provide evidence for automatic, implicit learning processes in humans. An associative model supplied predictions regarding how learning about these contingencies would interact, although the model proposed in the thesis falls short of accounting for all of the observed phenomena. In this chapter I will collect and summarise the key findings across my experimental and computational work and discuss them in terms of what this can tell us about sequence learning under incidental conditions. I will finally discuss: implications for studies of incidental sequence learning; models of sequence learning; an associative account of incidental learning; the role of automatic processes within human learning; and future research directions.

7.1. Summary of findings

In Chapter 2 I found that humans can learn sequences under incidental conditions that involve a probabilistic sequential rule based on the element in the sequence two trials before ($t - 2$) the current stimulus (t). The central finding of this chapter was that whilst learning was observed both for the Different rule sequential contingency (the location of t was more likely to be in the opposite location to $t - 2$) and the Same rule contingency (the location of t was more likely to be in the same location as $t - 2$), there was a difference in the amount of learning of these two rules, with participants learning the Different rule more than the Same rule. Why this happened was not clear, and the following chapter aimed to investigate this.

Chapter 3 provided evidence that a model (RASRN) adapted to include a better representation of the task, including the on-screen response stimuli that would occur on each trial, was able to simulate the Same vs. Different group difference as well as the sequential effects observed in Experiment 1. Therefore, I noted that the inclusion of a stimulus representation in the model resulted in better simulation of human performance on the task, which suggests that stimulus-response learning can affect sequence learning.

Taking on board the conceptual addition of current stimulus representation and using this to derive a prediction of the RASRN formed the basis for Chapter 4, in which the role of stimulus-response associations in a sequence learning task was investigated. Both the RASRN and human participants demonstrated that sequence learning differed when cues were introduced to the task that related to the sequence of response-stimuli. Humans clearly demonstrated greater learning of sequences in the Previous group: where response locations were accompanied by concurrent cues that corresponded with the previous stimulus element. The Previous group showed more learning than the Current group, which showed some evidence of better learning than the Random group but was similar to it. The model predicted an overshadowing effect of cue-response learning over sequence learning (that the Current group would learn less than both Random and Previous, which were predicted to be the same), which was not found in participants, and therefore the model was falsified. An associative account would still predict that concurrent cues had some effect on sequence learning, however the precise interaction between cue-response and sequence learning mechanisms in humans remain unclear; and it was not possible to investigate human cue-response learning and assess its impact on sequence learning.

In order to assess whether simple cue-response learning could occur, Chapter 5 investigated whether participants were able to learn cue-response contingencies under incidental conditions. Participants demonstrated faster and more accurate responding to colour cues that were trained to partially predict one response-stimulus location than to both control colour cues and trials that were inconsistent with the trained colour cue contingencies. Participants who completed the task intentionally provided a prior probability for responding on two direct measures of explicit knowledge: identifying the predictive colour cues and a forced-choice prediction task that replicated the context and sensitivity of the task. Participants in the incidental condition were at chance on these tasks and evidence for the null provided support for the absence of any explicit cue-response contingency knowledge.

Given that I found evidence for colour cue-response learning in Chapter 5, I introduced these contingencies with concurrent sequential contingencies in Chapter 6, which provided evidence that participants learned exclusive-or sequential contingencies (that t could sometimes be predicted depending on whether $t - 1$ and $t - 2$ were the same or

different) under incidental conditions regardless of the presence of the colour cues. This provided evidence that sequence learning was not affected by the presence of cue-response learning. In the Dual groups trained with both sequence and colour cue-response contingencies I found evidence that the two cues do not compete when they were related to one another. When attempting to compare this Dual Correlated group to an Uncorrelated group there was, unfortunately, little evidence that colour cue-response contingencies were learned at all in these experiments. As training performance for sequence and colour cue-response learning was confounded in the groups who were trained on both contingencies, a comparison was required at test. However, as the colour cue-response learning experienced rapid extinction at test, evidence for differential colour cue-response learning was minimal and it is therefore difficult to draw any firm conclusions regarding the interaction of these learning processes. However, taken alongside previous work these results suggest that we can find evidence of cue competition in incidental human learning (McLaren et al., 2013) and that this overshadowing effect can be avoided (and perhaps even reversed) if the cues themselves are related.

7.2. Evidence for implicit learning in humans

7.2.1. Qualitative differences between explicit and implicit learning

Chapter 2 compared human sequence learning under incidental and intentional conditions with the aim of finding evidence for implicit sequence learning. A variety of authors find a qualitative difference between sequence learning under the two instructions (Dominey, Lelekov, Ventre-Dominey, & Jeannerod, 1998; Guo et al., 2011; Jimenez, Vaquero, & Lupianez, 2006; Jones & McLaren, 2009; Kuhn & Dienes, 2006) and this can be taken as evidence for functionally different processes: one under our control and one not. Chapter 2 found no evidence of such a difference as there was no difference found on indirect RT and error measures of sequence learning between the incidental and intentional groups. Whilst one could take the view that the similarity between the groups reflects the intentionality of the incidental group, I argued that whilst the intentional group seemed to try to look for patterns, they struggled to find them. They also reported that intentional effort was unhelpful in completing the task quickly and accurately and were therefore not motivated to search and use sequential knowledge and subsequently reverted to automatic responding. Some participants

reporting that they attempted to count certain patterns, and indeed the comparison of the instructed conditions by means of a state-trace analysis are consistent with the possibility of multiple learning processes. However, I suggest there were no qualitative differences between conditions because participants were unable to work out and apply the rule.

However, is rule learning the only way that participants could have learned, and do we expect participants to learn more when they have explicit knowledge? Jones & McLaren (2009) suggested that participants were able to learn that the subsequence XXX was more likely; but no other subsequence trained – suggesting that not the rule itself but a specific instance was learned (Dienes & Fahey, 1995; Logan, 1988). Under intentional conditions in Experiment 2 in this thesis there was evidence for learning of each subsequence – which suggests that they were not learning one specific exemplar. Participants may instead have been able to abstract some rule, or some general property about the rule from their experience of it, for example that there were 'quite a lot of repeating chunks and runs of alternations' rather than hypothesis testing specific trial orders or particular contingencies and responded with more chunks or runs. Indeed, a specific rule-based representation is not a requirement for behaviour that appears to follow such a rule (Redington & Chater, 2002). It seems unlikely though that this knowledge would result in less learning, which was the case for the intentional participants in Experiment 2.

It is suggested that certain stimuli under intentional conditions are learned because of their salience (Jones & McLaren, 2009; Lee & Livesey, 2013) which suggests that participants may not have experienced any salient subsequence elements in this design, as the subsequences (e.g. in the Same group RRR, LLL, RLR and LRL) were all of similar salience and therefore no one particular element stood out. This is consistent with the work of a variety of authors (Frensch et al., 2003; Rünger & Frensch, 2008) who suggest that explicit learning occurs when participants are able to identify stimuli that are unexpected and produce anticipatory responses. Therefore, participants may be unable to isolate specific instances because of the similarity of the subsequences in this task compared to in other tasks (Jones & McLaren 2009; Lee & Livesey, 2013) where specific instances may have stood out to participants.

An increase in learning might be expected to occur when participants experience explicit contingency knowledge (Curran & Keele, 1993). Participants in Experiment 6 (the Intentional cue-response learning study) were clearly aware of colour cue-response contingencies and could report these confidently and with 100% accuracy. These participants in the Intentional condition demonstrated far greater learning across the training phase of the task compared to the Incidental condition, as participants responded far quicker and more accurately to trained locations of predictive colours compared to the inconsistent and control colours. Therefore, on this task we have evidence of explicit knowledge, which resulted in evidence of stronger learning. This supports the proposal that there was a lack of explicit knowledge across the sequence learning task in Experiment 1, regardless of the instructed conditions. If we take either significantly greater or qualitatively different learning as evidence of *explicit* learning itself, we clearly have very little in Experiment 2. Indeed, this may suggest that the sequence learning task is unlikely to be learned explicitly, either through instance or some rule based approach. This provides indirect support for Experiment 1 as a demonstration of implicit learning.

A further implication of this result is that defining explicit and implicit learning processes simply in terms of their volitional properties is not as simple as when dealing with explicit and implicit knowledge (e.g. Jacoby, 1991). Participants are not responding under incidental conditions with no control over learning, and with complete control of learning under intentional conditions. Indeed, in either case participants are able to freely think, ignore or invent tasks for themselves, as well as choose to look for or ignore patterns that they may notice throughout the experiment. Therefore, whilst the manipulation has had some success in producing qualitatively different results (Dominey et al., 1998; Jimenez et al., 2006; Jones & McLaren, 2009; Lee & Livesey, 2013) it seems that even when providing participants with the explicit rule that the sequence follows that they do not use this information on all trials (Lee & Livesey 2013), and therefore do not apply a controlled, explicit response to each stimulus in an SRT task. This seems to be a consequence of the SRT task design, which by its nature is used to discourage reasoned and controlled responding as the task demands require quick responses.

As studies of incidental sequence learning hope to capture implicit processes, direct tests of knowledge are taken after training on these sequences so as to avoid participants noticing or using these sequences (Cleeremans et al., 1998; Shanks & St John, 1994). However, higher order probabilistic sequences, as well as complex conditional sequences are hard to verbalise or notice as participants tend to learn and rely on the simple surface features of a sequence (Dominey et al., 1998). Therefore it may be possible to investigate the interaction between the development of sequence learning and explicit knowledge using the guessing criterion (Dienes & Berry, 1997), where participants are required to report the degree to which they believe they had any knowledge about the task. This would not require that participants were informed of contingencies, as participants could be instructed at the start of the task that there may or may not be contingencies in the experiment. Measures of both participant's knowledge of contingencies as well as their confidence in this knowledge could be taken on each block to track participant's performance across a task. Participants may or may not notice; attempt to notice; or use these contingencies, but rather than assuming that they have full volitional control of their learning as a result of a between-subject manipulation based on a single task instruction, it would be possible to investigate more sensitively when and how these strategies are used and affect sequence learning throughout a task.

7.2.2. Evidence for implicit cue-response learning

Participants in Experiment 5 and to some extent 6 provide good evidence for implicit learning, as participants demonstrate clear learning of colour cue-response contingencies across training in the absence of being able to correctly identify those colours (let alone the response location that they predict) more than chance would allow. When completing the task intentionally, with the nature of the contingencies provided as a hint to ensure explicit knowledge would develop, all participants were able to correctly identify both colours with 100% response-location accuracy.

Taking the zero-correlation criterion (Chan, 1992; Dienes & Berry, 1997) regarding confidence in knowledge and directly measured task performance, participants in the intentional group demonstrated this knowledge with full confidence, whereas only four incidental participants mentioned that they may have noticed something during the experiment. Classifying these participants as confident and participants who were

surprised that there were contingencies as not confident produced no correlation between confidence judgements and colour identification accuracy, suggesting further evidence for an absence of awareness. This measure of implicitness does not entirely follow the suggestions of Dienes and Berry (1997), however, as confidence ratings regarding participants' colour guesses were not taken, with the questions instead referring to the extent that participants felt they had noticed relationships within the experiment. Whilst this is a conceptually similar question, it does not refer to confidence in identifying the correct colour cue-responses.

Performance on a prediction task provided further evidence for the absence of explicit awareness, using the intentional group performance as a prior to produce evidence for the null – that incidental subjects seemed to have no explicit contingency knowledge. However, as the prediction task was calculated after test blocks had been given, extinction may have destroyed contingency knowledge that was previously explicitly available to incidental subjects (Cleermans et al., 1998; Shanks & St John, 1994). Evidence that this is not the case comes from intentional subjects, who also suffered from extinction, but were able to produce perfect performance on the identification task. However, participants may have simply learned more (and enough) in the explicit condition for this to survive the two block test phase. Intentional condition participants also had experience of the sheet with the colours they thought were predictive (which was removed for the interview and prediction task); as well as any memory for this sheet and their answers, which might have increased the explicit knowledge that occurred in both groups. As mentioned previously, memory is a consideration in assessing implicit and explicit knowledge post-training; and this issue may have perhaps been modified by post-test training blocks (e.g. Jones & McLaren, 2009), or confidence judgments within the prediction task (e.g. Destrebecqz & Cleeremans, 2003). Whilst I considered using confidence judgments for each response in the prediction task, I wanted the prediction task to replicate the SRT training context as closely as possible to avoid producing a new task in which participants experienced differing demands and were unable to express their contingency knowledge from the training setting.

7.2.3 Evidence for incidental cue-competition in humans

Following the suggestions of Beesley and Shanks (2012) Chapter 6 of this thesis aimed to demonstrate cue-competition effects as observed across the animal associative

learning literature in humans under, crucially, incidental conditions. There was no evidence that any of the groups in Chapter 6 were aware of colour or sequential contingencies with responses, which provides a strong foundation for investigating the interaction of the cues, and providing evidence for associative processes in humans. In the first experiment of Chapter 6, we find similar results to Beesley and Shanks (2012), that on an apparently implicit task participants show a trend towards the *potentiation* of learning of the cue that should have been competing for associative strength. In Experiment 7 we see that participants demonstrate learning about both sequences and colour cure-response contingencies. Beesley and Shanks (2012) suggest that this result suggests an absence of cue-competition and therefore that learning under incidental conditions is not associatively driven.

I suggest that this should not be the conclusion drawn. The reason, I suggest, for this is that the sequences and colours were themselves perfectly related, as each colour was 100% likely to follow the four possible second order transitions in the task (RR, RL, LL and LR). Therefore, while the sequence and colours were both themselves predicting a certain response that could be separately analysed at test, they were also themselves associated. It is possible that these between cue associations were learned and protected the model from cue competition. Taken with the evidence provided by McLaren et al. (2013) that this task can produce an overshadowing effect on colours by sequences, this provides support for the presence of associative processes on such tasks and may provide an explanation for Beesely and Shanks' (2012) results. Whilst a number of patterns were used as cues in these visual search tasks, if the between-element associations between items were strongly learned then participants may not have experienced cue-competition. This is a possibility, firstly if the stimuli are encoded configurally (Urcelay & Miller, 2009) or because the patterns were not trained to predict a particular response, but the location of the target stimulus. Therefore the contingency between A and B distractor patterns was higher than the contingency between either stimuli and a required response and the between cue associations would therefore be stronger.

Conclusions that the relationship between stimuli is the mechanism through which cue-competition effects are reduced or even reversed can, however, only be made tentatively. This thesis was unable to provide reliable evidence of any colour cue-response learning

across both tasks (Experiments 8 and 9) that attempted to investigate the role of between cue associations in predictive cue competition effects. Whilst the numerical pattern of colour learning at test followed these predictions, there was both no evidence of colour learning in the colour group at test and no evidence for a significant difference between the Dual groups at test. This thesis therefore offers the possibility that an associative account of human learning processes cannot be dismissed based on the absence of a cue-competition effect, as the frequency with which these cues co-occur as well as their temporal overlap (Urcelay & Miller, 2009) may encourage between stimulus learning that protects from overshadowing. These mechanisms suggest that sequence and cue-response learning may indeed share the same, incidental, associative learning system. Further work, however, is required to demonstrate and further characterise the nature of how between stimulus associations form and themselves (differentially) contribute to sequence and cue-response learning.

7.3. Review of the RASRN

Whilst Chapter 4 ultimately falsifies the RASRN as it stands, the model was instrumental in understanding the influence of stimulus-response associations on sequence learning that formed the predictions and results outlined in this thesis. Clearly I cannot suggest that it is a satisfactory account of learning under incidental conditions, nor of human sequence learning under incidental conditions and further work is needed to investigate the precise involvement of stimulus representation within recurrent models. Whilst other authors have attempted to model additional stimuli and non-sequential learning using some version of the SRN (e.g. Cleeremans, 1993; 1997; Destrebecqz & Cleeremans, 2003), no models have considered that the learning that occurs about stimulus-response contingences on each trial may interact with sequential learning across trials; therefore the RASRN is unique in this aspect. This provides predictions regarding cue competition effects that are supported by the results of Chapter 2, and to some extent Chapters 4 and 6; although the precise mechanisms and relationships between sequences and cues within the model is not clear. Whilst the model as it is reported is therefore incapable of replicating the human learning observed in Chapters 4 and 6, the model principles predicted at least some of the results of a cue-sequence learning interaction.

Models of sequence learning do not require stimulus-response associations, as contingencies are formed across time. Therefore, humans do not need to make these associations, as they are instructed to respond to the stimuli explicitly (no associative link required) and therefore performance on SRT tasks of sequence learning have progressed happily without these adaptations. I suggest that not only does including these stimulus-response links better represent task conditions (Destrebecqz & Cleeremans, 2003), but that this has an important influence on how sequences are learned (this thesis, Chapters 2 and 3). This approach could still be criticised, as from an associative perspective, that humans or animals encode and form associations between every element of the environment is not considered to be adaptive for the purposes of learning, nor realistic (Pearce & Bouton, 2001).

The RASRN attempts to represent the within-trial time course by a very rough approximation using a simple activation difference at input level approximating the differential influence of stimuli depending on their temporal relationship to the trial to be predicted. More sophisticated ways of doing this are available and discussed in section 7.5.3, but in attempting to do so the RASRN has captured some element of the increased learning of sequences compared to stimulus-response learning on each trial. However, this is clearly not fully represented in the model as the results of the final simulation suggest that the RASRN learns cue-response relationships better than sequences, as sequence learning is overshadowed in the Dual Uncorrelated simulations. This suggests that the activation of current stimulus units should be further reduced in comparison to the previous required response activation to increase sequence learning relative to colour cue-response learning. This may also suggest that there is more to sequence learning compared to simple cue-response learning that is currently not captured by the model (for example, spatial location). Indeed, it could be that sequential effects or stimulus-response learning are restricting sequence learning in the model in a way that does not occur in humans.

The model is able to represent individual differences on the task to some extent, as the learning rates and number of hidden units, as well as their starting connection weights can all be altered to provide some degree of variability, however it is clear that individual error between networks in these stimulations does not reflect human performance differences, and regarding motivational and attentional influences that

differ across participants and the task itself, the RASRN falls short. The model is capable of producing these effects however, with random noise to the hidden layer used in the AugSRN to produce distraction on a similar sequence learning task (Cleeremans & McClelland, 1991) for example; however this level of detail was not the aim of the thesis. So whilst the model was unable to precisely mimic human responses to the task, the specificity of these predictions were not taken to be as important as the predictive value of simple, relational changes between stimulus relationships within the task.

Further to this, however, the role of responding and feedback is an important component of the SRT task that was not modeled by the RASRN. The model simply produces an MSE approximating a human RT to the next trial, it assumes that no incorrect responses are made, and no error feedback given. In all tasks an incorrect response was followed by a beep, which participants often reported as being highly salient and frustrating. Whilst this may have had an effect on explicit processes of attention or motivation, error feedback is an instrumental reinforcer that may produce significant effects on learning. Whilst trials following an error are excluded from the analysis, their impact on learning is likely to be significant and to my knowledge error feedback is not represented in any version of an SRN. That the RASRN promotes the accurate representation of stimulus conditions but does not represent error feedback is a strong criticism of the model.

Therefore, in conclusion the RASRN accounted for the different sequential learning observed in human data in a simple, constant RSI two-choice spatial SRT task by introducing stimulus-response associations that, uniquely, could compete with sequence learning and led to the simulation of the Different rule sequence learning effect. The SRN and AugSRN were unable to account for these results and therefore the RASRN was instrumental in understanding how humans may have learned these sequences. The model predicted that additional stimuli within the SRT task would interact with sequence learning, which was found to be the case in humans across Chapters 4 and 6, although functionally representing the precise mechanisms by which this occurs remains a challenge for any model of sequence learning.

7.4. Implications for human sequence learning

7.4.1. Trial order and sequence learning

Evidence has been provided that participants can learn probabilistic relationships based on a trial experienced before the previous sequence element ($t - 2$, Chapters 2 and 4) as well as the previous two elements together ($t - 1$ and $t - 2$, Chapter 6) under incidental conditions. That participants are sensitive to, and can learn a variety of probabilistic sequential structures under incidental conditions is by no means a novel contribution to the literature as experiments training participants on first (D'Angelo, Jiménez, Milliken, & Lupiáñez, 2013; Jiménez, Lupiáñez, & Vaquero, 2009; Shanks, Wilkinson, & Channon, 2003) and higher order (Cleeremans & McClelland, 1991; Jiménez, Méndez, & Cleeremans, 1996; Jones & McLaren, 2009; Lee & Livesey, 2013) probabilistic sequences are numerous. However, the number of studies that compare trained performance with control groups matched for sequential effects (Anastapolou & Harvey, 1999) are limited (Jones & McLaren, 2009). Therefore, to some extent, these studies contribute to a small body of sequential learning research that adequately controls and considers the effect of the previous stimuli in the sequence.

The results of this thesis further support the claim that sequences are learned differently depending on their structure (Jones & McLaren, 2009). Chapter 2 provides evidence that under incidental conditions participants learned two sequential rules differently despite that they essentially involved the same probabilistic rule: that $t - 2$ predicts t on two thirds of trials. A rule-based account would find it hard to reconcile this result, unless one assumed that learning that $t - 2$ equals t is a harder rule to learn than $t - 2$ does not equal t, which at first reading is not intuitively plausible. One might suggest that some explicit heuristic is more amenable to the Different rule, for example a gambler's fallacy (Jarvik, 1951; Kahneman & Tversky, 1982). Whilst this effect is usually considered to be confined to the preceding trial, participants may expect trial alternations and therefore be more sensitive to rules embedded in a sequence that involve alternations. It is possible therefore that some explicit expectancy enabled better learning that $t - 2$ does not equal t and made it easier to acquire than the Same rule, but this is based on a heuristic that is restricted to first-order effects; the influence of stimulus order prior to $t - 1$ has been suggested to have a benefit only for repeats (Soetens, Boer, & Hueting, 1985).

The results of the RASRN simulation in Chapter 3 suggest that this difference is caused by the influence of stimulus-response learning, which at $t - 2$ is the same stimulus-response mapping as at t in the Same group, therefore an association between the two trials is blocked to some extent by their stimulus-response associations. This suggests that the stimulus-response associations can differentially interact with sequence learning and sequential effects to produce different patterns of both learning and performance. Similarly, Jones & McLaren (2009) observed the absence of learning about the subsequence XXX under incidental conditions and suggested that blocking of the trained contingency (that XX predicts an X) occurred as a result of transient learning that X predicts X reducing the error term for the final X, where learning about the trained sequential contingency occurs. An instance or exemplar-based account might suggest that this subsequence was simply harder to learn (e.g. Shanks & St John, 1994), but Experiment 1 demonstrated that XXX was learned by participants in the Same rule group under incidental conditions. This suggests that the effect of trial order on sequence learning is itself dependent on the statistical regularities of the sequence to be learned. This provides support for of a complex, highly interconnected learning process; rather than a simple exemplar based system that can store and retrieve information based on the number of occurrences.

7.4.2. Representing stimulus conditions in sequence learning

A central aim of this thesis was to investigate the relationship between stimuli presented to participants on each trial and sequence learning. Experiment 1 and the simulations of Chapter 3 suggested that stimulus-response associations played an important role in incidental sequence learning as the overall Different versus Same learning effect as well as the sequential effects observed in humans under incidental conditions were reproduced when the AugSRN included representations of the current trial stimulus (RASRN). Models do not *require* a representation of both t -1 and t when predicting t, which is not entirely surprising, as sequence learning is based on contingencies between t and previous trials, and not based on the influence of concurrently presented stimuli. However, even the low activations (0.1) of current stimuli have a qualitative impact on the way a model learns and responds to sequences, and therefore suggests that an associative model of sequence learning is highly sensitive to the stimulus conditions that occur in between using the previous trial to predict the location of the current trial.

Destrebecqz and Cleeremans (2003) criticise the SRN for two reasons, firstly that it does not represent time by any other means than a trial-by-trial time-step; and secondly because it only examines a prediction of one trial based on the previous, which is inconsistent with task demands. I will discuss timing and sequence learning in section 7.5, but first I will discuss briefly the models presented in the literature that do encode some representation of trial *t* when predicting trial *t*, starting with the model produced by Destrebecqz and Cleeremans (2003). This involved three components, a simple SRN, a set of response units and an auto-associator; which represented learning about the previous trial; the influence of responding; and simple stimulus-response learning, respectively. Competition between sequence learning and stimulus-response learning could occur going into the response units, which were activated based on the accumulated (and competing) strength of the SRN or auto-associative prediction. Learning about the stimulus-response and previous-current trial contingencies was, however, conducted separately, and Destrebecqz and Cleeremans (2003) suggest that they are inevitably separate.

Similarly, Cleeremans (1993) included a representation of the trial to be predicted in the Dual SRN (DSRN) when he attempted to account for the sequence learning results of Curran and Keele (1993), who themselves trained participants on a simple six-item repeating sequence and found that explicit knowledge and intentional learning interacted with implicit sequence learning. The SRN was able to simulate these explicit results when a buffer network that contained a memory of the sequence trained to predict *t*, produced an output activation of *t* which was used alongside $t - 1$ as input into a hidden layer that predicted *t*. Whilst this model includes a *t* representation at input to predict *t*, this did not represent the actual stimuli, but was itself a prediction of a model that had a memory for the sequence, representing learned explicit knowledge and not simply the stimulus presented to participants that enabled them to make a response.

Both of these models include a *t* representation, in that one is supplied in order to predict *t* (Destrebecqz & Cleeremans, 2003); as well as one that produces a prediction of *t* that is used to better predict *t* (Cleeremans, 1993). Whilst the RASRN is not the first model to attend to and represent these stimulus conditions within a model of sequence learning, it is the first model that allows contingencies between these stimulus conditions and representations of the sequence to interact and compete at a learning

level. The RASRN also deals with the representation of time and stimulus conditions in a simple way; using lower activation values for the current stimuli that does not require a further learning system. By restricting a recurrent loop to sequential information only as is done in Cleeremans work, this presupposes that memory only occurs for certain stimuli and not others, which seems a large assumption to make given that the simplest account of associative learning would suggest that implicit learning is an entirely automatic process whereby associations are formed indiscriminately between regularities in the environment (Shanks, 2010).

7.4.3. Temporal effects on sequence learning

As mentioned previously, the SRN provides a model of learning that can occur in series, but beyond trial-by-trial order it is unable to represent time (Destrebecqz & Cleeremans, 2003). Activations are calculated once per trial, and cannot represent between- or within-trial temporal effects. The RASRN goes some way to address these shortcomings by altering the activations of locally represented units to represent their temporal influence on learning. These activations were based on an assumption that a prediction will receive greater influence from representations that can accrue strength over time, with stimuli presented just before an event consequently producing less activation. This reflects the observations of Destrebecqz and Cleeremans (2001; 2003) that sequence learning increases with a function of RSI length and is supported by the work of McClelland (1979), who proposed a functional implementation of this with incremental propagation of activation in his cascade algorithm. This was successfully applied within the adapted SRN of Destrebeceqz and Cleeremans (2003) to represent RSI influences on sequence learning. However, this is in contrast to a variety of studies that suggest that RSI increases have a negative impact on sequence learning (Frensch & Miner, 1994; Stadler, 1995; Willingham, Greenberg, & Thomas, 1997). Frensch and Miner (1994) propose that learning of sequences is based on decaying memory activations of previous trial representations, and therefore participants are less likely to learn with greater time between stimuli.

Whilst previous studies have attempted to examine how the length of RSI influences learning (e.g. Shin, 2009; Willingham, Green, & Thomas, 1997), converging on a general dual-process account of increasing explicit preparation and decreasing implicit learning (e.g. Frensch & Miner, 1994), inconsistencies appear, perhaps as a result of

sequence complexity and intervening stimuli (Destrebecqz & Cleermeans, 2001; 2003). Further to this, there are a number of studies that suggest that temporal and spatial information is encoded separately when learning about sequences (Miyawaki, 2006; Rünger, 2012). The RASRN represents time as a constant influence at the input level; and Destrebecqz & Cleeremans (2003) can alter the influence of activation according to time by incrementally increasing overall activation. Both models improve on the representation of time as a trial-by-trial series of discrete events in the SRN, however both of these models also represent time as some influence on the amount of input.

These models therefore do not represent the possibility that time itself may be encoded within the model and associations between stimuli and time might occur. Indeed, the work of Shin (2008) suggests that participants can learn sequences with a constant, patterned or random RSI; however learning is stronger for the constant group, consistent with an account where time is associated with sequential elements and stronger learning produced when RSI is constant and does not vary and therefore interfere with learning. Rather than simply using time as an index of how much activation, learning or performance effects may occur on a subsequent trial, it can therefore be integrated within a model. Whilst time is thus considered important in the sequence learning literature, models of sequence learning are yet to appropriately reproduce these real-time effects and this requires further work. Rather than concentrating solely on the RSI, however, future sequence learning research should consider also the interaction between the time course of particular stimuli within a trial.

A further explanation for the lack of colour potentiation of sequence learning in the Dual Correlated groups in Chapter 6 is the encoding of time. Whilst participants experienced somewhat variable intervals between previous response-stimuli and the following colour (response latencies plus RSI); their responses themselves were always followed after 500msec by these tones or colours in Experiments 3 and 4, which were the concurrent stimuli experiments in Chapter 4. Chapter 6 involved 250msec RSIs between responses and colours, therefore less learning could have occurred in this shorter time (e.g. Dominey, 1998; McClelland, 1979), which is perhaps why no sequential learning differences were observed.

7.4.4. The role of responding in the SRT task

The separate encoding of time and stimuli within a model brings me to a discussion of the separate encoding of responses. Whilst Chapter 5 provides evidence that participants can learn simple stimulus-response contingencies, the sequence learning literature converges on the idea that responses are key in learning contingencies (Goschke, 1998). Perceptual and motor sequences have been found by a number of studies to produce different amounts of sequence learning (Bischoff-Grethe, Goedert, Willingham, & Grafton, 2004; Willingham, 1999; Willingham, Wells, Farrell, & Stemwedel, 2000), which suggests that making responses in sequence is what drives the robust and automatic sequence learning demonstrated across Chapters 2, 4 and 6 in this thesis. This is, perhaps, why sequence learning has become the "best behavioral paradigm through which to study implicit learning" (Destrebecqz & Cleeremans, 2003, p. 181); as it is extremely reliable, replicable and whilst it may differ depending on certain stimulus conditions (Chapter 4 in this thesis; Nissen & Bullemer, 1987; Stadler, 1995) or task parameters such as RSI, salience, instructions, there are few studies that report a lack of sequence learning. Indeed, this thesis provides little evidence that sequence learning was damaged by increased or random RSI (Chapter 6); additional random stimuli (Chapters 4 and 6); or additional contingencies to learn about (Chapter 6). It is clearly a very robust form of learning.

7.5. Implications for how additional cue stimuli interact with sequence learning

Research has considered the influence of additional *tasks* on sequence learning (e.g. Nissen & Bullemer, 1997; Stadler, 1995), but only a small number of researchers have examined how additional perceptual stimuli interact with sequence learning (e.g. Cleeremans, 1997; Clegg, 2005; Deroost & Soetens, 2006). Few authors have found evidence of learning about other contingencies present in the data when sequences are in play (Cleeremans, 1997) and those that have provide no evidence of an interaction between learning of these two contingencies (Robertson & Pascaul-Leone, 2001). However, these studies also provided no evidence that the additional learning also occurs incidentally or implicitly, hence separate systems may underlie these data and we might expect them not to interact. Chapter 5 provides good evidence that participants can learn stimulus-response contingencies without awareness or intention, and Chapter

6 some evidence that these are differentially affected by the presence of and relationship to sequences themselves.

Chapter 4 suggests that incidental cue-response learning can interact with the same processes involved in sequence learning. These stimuli were all task irrelevant, insomuch as participants were not required to process these stimuli in order to complete the task, nor were they required to classify, recall, count or otherwise interact with the stimuli, over and above being instructed to attend to them perceptually, which suggests that implicit learning is not restricted to active features of the task set (Abrahamse et al., 2012). However, some stimuli in my experiments did provide additional information on each trial about either: the previous stimulus element (Previous condition: Chapter 4) or the current trial itself (Dual groups: Chapter 6), although participants were not informed of this, nor were required to attend to this to compete the task. Only the Previous condition in both experiments in Chapter 4 showed evidence that additional concurrent stimuli had an effect on sequence learning. This suggests that sequence learning was largely unaffected by additional stimuli, but that it could be significantly enhanced.

As mentioned in the previous sections, this suggests that a model of associative sequence learning should encode for the stimuli presented to participants and that these stimuli should not be processed by a separate system (e.g. Cleermeans, 1993; 1997; Destrebecqz & Cleeremans, 2003). The RASRN was ultimately unsuccessful within the confines of the parameters that produced the predictions regarding the role of these cues (Chapters 3 and 4), which would have provided excellent *a priori* support for the model (Boucher & Dienes, 2003). However, I will continue to argue that stimuli and cues must be represented within a model of sequence learning; and that whilst separating cue-response learning from recurrence in a model may possess *prima facie* simplicity, this is neither more parsimonious nor based on legitimate assumptions regarding recurrence.

7.5.1. How do additional cue stimuli affect sequence learning?

Given that the RASRN predicted that stimulus-response learning would have a qualitative effect on sequence learning, the results of Chapter 4 go some way to support and expand on this prediction. When considering the influence of response-stimuli and their associations with the required response, Chapter 4 demonstrated that in a sequence learning task, the response-cue associations between the previous stimulus element and

colour or tone presented concurrently with the stimulus on the current trial produced increased sequence learning of Same rule sequences. This suggests that participants were able to learn a relationship between the concurrent cue stimulus (*t*) and the response-stimulus location at time ($t - 1$), which resulted in higher activation of the $t - 1$ element, and therefore the $t - 2$ element in the sequence.

This associative explanation of the result is not the only account, however, and it is possible that instead participants were able to use the colours to somehow rehearse or reinforce the sequence experienced, giving participants in this condition double the memory strength, or giving an associative model that is able to extract statistical regularities from sequences two shots at extracting the statistical structure of the rule itself. Further to this, it possible that these additional concurrent cues made the sequential structure more salient, which the cues matching the current trials and random cues could not; as participants were able to both find and exploit patterns in the response stimuli locations and colours or tones. It seems unlikely that this is case as participants were unable to report a contingency between response stimuli or additional cues, and were surprised when these were explained.

That concurrent cues can interact with sequence learning by no means suggests that they can eradicate sequence learning, which seems to progress robustly in the presence and absence of additional stimuli that are or are not related to the sequence. All of the experimental work in this thesis containing sequential contingencies (Chapters 2, 4, and 6) show that humans are able to learn these well, regardless of an absence of explicit knowledge or a volitional intention to learn. This is supported by the results of Chapter 6, which suggest that sequence learning is stronger than cue-response learning, which naturally leads to questions about the origin of this discrepancy.

7.5.2. How does sequence learning interact with cue-response learning?

Previous research on cue-response learning and sequence learning by Jiménez, Méndez and Lorda (1993) and Jiménez and Méndez (1999) has been unsuccessful in producing a cue-response effect on sequence learning. The authors emphasised that an incidental blocking effect was expected (Jiménez, Méndez, & Lorda, 1993) as unlike Cleeremans's (1997) work, they did not make the contingencies between cues and responses explicit to participants. Expecting that implicit learning of these simple

contingencies would occur, Jiménez and colleagues were surprised that participants did not show learning about these more simple relations whilst still showing strong sequence learning of a complex probabilistic sequence. These studies are, from the perspective adopted in this thesis, flawed as neither provides evidence that participants could learn cue-response contingencies without awareness. Without a demonstration that cue-response contingencies can be learned, it is not possible to make any conclusions about the interaction it my or may not have alongside a sequential rule.

This problem of demonstrating cue-response learning independently aside, Jiménez and Méndez (1999) discussed cue-competition and expected it to occur; however they fail to consider the possibility that the lack of an effect on sequence learning occurred *because* of cue-competition. Chapter 6 and the work of McLaren et al. (2013) provides evidence that this may not be the case, and when contingencies were matched between a simple delay conditioned cue-response relationship on the same trial and sequential contingencies with the same probabilities across trials this sequence learning overshadowed cue-response learning. It seems that the cues in the work of Jiménez and Méndez (1999) were not related to the sequence itself, and therefore this thesis offers the explanation that (if we assume that these relations could show a learning effect in the absence of sequential contingencies) cue-response learning in these studies was overshadowed by the presence of sequential contingencies.

The work of McLaren et al. (2013) suggesting that overshadowing of colour cue-response learning can occur by sequence learning provides an indication that sequential contingencies are somehow prepotent under incidental conditions. Whilst colour cue-response learning was not found in Chapter 6, this was a result of a general absence of colour learning in all groups. Whilst we cannot conclude therefore that colours were overshadowed, it is clear that humans over the same length of training and within the same task are able to learn sequential contingencies that occur with the same frequency, and that in some sense carry the same information. It seems that humans demonstrate some advantage for sequence learning over simpler cue-response contingencies. This could be due to the increased activation of the stimuli within an associative system across time, or indeed the role of the response in reinforcing an additional motor component to the perceptual stimulus-based sequence learning. This could also provide evidence that incidental learning is more sensitive to complex, statistical regularities

than simple ones; which follows from the suggestion by Reber (1989) that an implicit system is designed to learn the information that an explicit system would find more difficult to acquire. This account is supported to some extent by the RASRN, which simulated the overshadowing of the weaker learning (sequences) by the stronger learning (colour cue-response) in the Uncorrelated group. Although this effect is the wrong way round (sequences overshadowed by colours), the model also learned far more about colours in all of the networks trained on colours. This suggests that the relative weighting of these input activations is wrong, or that an additional component of sequences as mentioned here (e.g. time, motor-responses, spatial locations) is absent from the model.

The preferential learning of sequences could be explained propositionally even though these are more complex, if one suggests perhaps that sequences comprise of multiple elements that give participants more time between them to entertain an explicit hypotheses about these contingencies. The discrepancy between an explicit trial order expectation based on a gambler's fallacy heuristic, for example, could cause participants to notice some difference from what they expect; which may lead to more learning than about cue-response contingencies for which they may have no pre-existing expectation. Either account would suggest that the results in this thesis provide good evidence that, at the very least, the two learning processes are not independent, as I have provided demonstrations where each affects the other.

7.5.3. Further questions about the interaction between sequences and cue-response learning

A question about the role of additional cue-response learning in the SRT task worth investigating is: what if the cue perfectly predicted the location of the next stimulus? If two colours, for example, gave a perfect prediction of the next trial location would participants use this information instead of sequence learning? A Mackintosh (1975) approach to associability suggests that increased predictiveness will increase attention to these cues and therefore the associative strength of this learning. Cleeremans (1997) found evidence that both sequence and cue learning occurred in such an experiment involving additional concurrent stimuli (that predicted the next trial). There was no evidence for any interaction between the two types of learning and Cleeremans (1997) represented the additional cues used in his experiment in an adapted version of the SRN

on trial *t* alongside the current response-stimuli information at *t* to predict *t* + 1. He avoided the issue of interference between cue-response and sequence learning by giving them each their own separate set of hidden units and indeed this produced no interference in the simulations produced by this adapted SRN as was also the case in humans. However, this account is flawed as whilst it produced the isolated cue-response and stimulus learning observed in the experiment, participants in the experiment were instructed in the cue-response contingencies and had explicit knowledge of these. A dual-process explanation of this data would suggest that these learning processes could therefore occur independently, and this could be why separating their internal representations (hidden units) produced results consistent with the data.

7.6. Implications for an associative account of sequence learning under incidental conditions

Central to this thesis is the question as to whether incidental sequence learning can be explained by an associative account (Cleeremans, 1993; Jones & McLaren, 2009). The evidence provided by my experimental and computational work may not provide a definitive answer, but it adds to the body of existing evidence that supports the presence of automatic learning processes in humans that follow the predictions of an associative account. I suggest that in the larger context of human learning, this supports the presence of dual processes as suggested by McLaren and colleagues (McLaren et al., 2014; McLaren, Green and Mackintosh, 1994). That humans are, of course, able to use explicit knowledge and propositions about events to learn relationships between them; but that a functionally separate system based on the automatic formation of associations exists that can learn complex contingencies between events across time. Further questions remain, however, in situating sequence learning within the context of a dual process account of human learning.

7.6.1. The interaction of explicit and implicit processes

The interaction between explicit and implicit processes is a further level of complexity not considered in this thesis, which simply assumes that associative processes may underlie automatic human learning; but that these can be overruled by our explicit intentions. In reality, the case may not be so simple as this (Sun, Slusarz, & Terry,

2005) as evidenced in this thesis by the differential influence of intentional instructions on sequential (Experiment 2) and cue-response contingency (Experiment 6) learning. Participants demonstrated an increase in learning consistent with the idea that explicit learning can produce superior knowledge and performance opposed to incidental learning; however Experiment 2 provides very little evidence for a learning increase.

This may suggest that volitional learning has a qualitatively different effect on sequence learning to simple cue-response learning; supported by the effect of intentional learning in Jones & McLaren (2009), which while not matched for training length, did not show greater learning but qualitatively different learning (see also, Dominey et al., 1998; Jimenez et al., 2006). How this impacts upon a dual-process account is a critical question for sequence learning research in humans, which might attempt to isolate the two, possible processes, but as previously outlined may do better to consider them both together and then attempt to disentangle them in some way. Whilst a qualitative difference in sequence learning might suggest different learning processes were activated, it also suggests that in Jones and McLaren's (2009) experiment that learning of subsequences that had been acquired across training incidentally was reduced. Indeed this aligns with the reduction in overall learning observed in Experiment 2, which as discussed earlier suggests that participants were using up resources by searching for sequences similar to effects observed in dual task sequence learning situations (Nissen & Bullemer, 1987). This perhaps provides an indication that the essential difference between sequence and cue-response contingency learning in the explicit sense is the difficulty that participants find in both identifying (Jones & McLaren, 2009; this thesis, Experiment 1) and applying (Lee & Livesey, 2013) complex probabilistic rules.

Therefore whilst explicit learning may increase as a function of the simplicity of what is to be learned; assuming that this is the case for incidental learning may not be wise. The incidental learning in this thesis seems, if anything, to follow the opposite pattern and increase with (or at least be relatively unaffected by) increased contingency complexity. The effect of explicit learning on implicit learning does not suggest that the automatic system can actually be turned on and off under our control, but that processing resources attributed to encoding stimulus relationships, order and time may be interrupted by explicit attention to other external or internal representations (Stadler, 1995). The influence of implicit, automatic processes on explicit processes requires a

further level of analysis that this thesis does not address, as whether participants can acquire explicit knowledge as a result of strengthening associative representations (Cleeremans, 2006) or some other theory regarding the construction of explicit knowledge (e.g. Mitchell et al., 2009; Rünger & Frensch, 2008) was not investigated; only in Experiment 6 did any participants provide evidence of explicit knowledge. Models such as CLARION (Sun, Slusarz, & Terry, 2005) and ACT-R (Anderson, 1993; applied to sequence learning, Lebiere, Wallach, & Taatgen, 1998) are not discussed here, but have provided convincing accounts of both implicit and explicit sequence learning through hybrid connectionist and symbolic or procedural systems.

7.7. Further research

7.7.1. Behavioural predictions

One outcome of the experimental and modeling work in this thesis is the claim that stimulus-response associations produced the greater Different rule learning above Same rule learning in Experiment 1. A simple experiment to test whether this was the case would involve training participants with these two sequential rules, but with a variety of different stimuli across both response locations. Participants could be presented with any number of different stimuli (different shapes, colours etc.) on the left or right hand side of the screen, in any number of locations. Therefore there would be no specific place on the screen, colour or shape that could build up an association with either response key. Whilst participants would be able to follow response instructions regarding a left or right response to any stimuli in that area of the screen, specific stimulus-response associations would not be able to interfere with the sequence or right and left responses, which would, if my theories are correct, alter sequence learning by increasing Same rule learning in the absence of strong stimulus-response associations.

Further behavioral predictions of this thesis involve the role of time, which may decrease the amount of sequence learning observed compared to the influence of stimulus-response mappings if indeed sequence learning is reduced. Therefore, in the experiments in Chapter 2, I expect that a bigger difference between Different and Same group learning would be observed with shorter RSIs despite less sequence learning in this condition. However, this is highly controversial as some authors (e.g. Fresnch & Miner, 1994) suggest that shorter RSIs lead to greater learning. Therefore a short (or no)

RSI condition could be compared to a longer RSI condition. I predict that less learning will occur in the longer RSI condition, according to the predictions of the RASRN. Further to this, shorter RSIs are implicated in increasing the influence of short term priming of the previous response, therefore there may be a bigger impact of sequential effects in the short RSI condition. Whilst there was no interaction between subsequences and learning in the experiments in this thesis, greater or less time between trials may alter this and produce less Same rule learning in the short RSI condition and less Different rule learning in the long RSI condition as a result of the influence of sequential effects.

The thesis also predicts that under the same task instructions experienced in Experiment 1, that Jones and McLaren's (2009) result may disappear, as the sequential effects and learning under incidental conditions does not match those found in any of the sequence learning experiments in this thesis. The best explanation of these differences, especially between the control group who should not show any difference whatsoever, is in the feedback given at the end of each block. Jones & McLaren (2009) provided participants with monetary bonuses, which may have encouraged them to perform faster and more accurately, reducing the influence (in terms of learning) of the current on-screen stimulus.

7.7.2. Model development

The challenge remains to develop a model of sequence learning that can account for the role of stimulus-response relationships within a task, and considerations for this have been mentioned throughout the discussion. The representation of associability, time, error feedback and responding are all important when attempting to represent the task conditions with even greater specificity. Indeed, if a model was able to learn simple and sequential associations as observed in humans under incidental conditions it would provide huge power in motivating and enabling further research. The challenge for researchers will be centred on the trade-off between increasing explanatory power and increasing the number of processes (e.g. associability) and free-parameters in a model, which is not to say that this is a criticism in itself (Boucher & Dienes, 2003). With a simple, parsimonious model (e.g. the SRN, Elman, 1990) that has extraordinary emergent properties (Beesley, Jones and Shanks, 2012) and can roughly simulate a huge number of tasks we are able to generate a large body of evidence that roughly supports

some association formation in perceptual-motor sequence learning as well as a number of other human learning tasks such, for example artificial grammar learning (Dienes, 1993); speech perception (Gaskell & Marslen-Wilson, 1997); and learning musical sequences (e.g. Altmann, Dienes, & Goode, 1995). However increasing specificity based on accurately simulating human performance encourages specific predictions that can be falsified, leading to a better understanding of detailed learning processes that may occur.

7.7.3. Other suggestions

The first question this thesis asked that remains unanswered and merits further experimental attention concerns the overshadowing of cue-response learning by sequence learning in Chapter 6. Without definitive evidence for colour cue-response learning at test in the group trained only with colours, the conclusions about the interaction between learning about the two sets of relations are by no means definitive. As the current results suggest that learning about these contingencies can interact – with some evidence of overshadowing in work not done as part of this thesis (McLaren et al., 2013) and possibly potentiation of colour cue-response learning in the first experiment of Chapter 6 (as well as in Beesley & Shanks, 2012, although not with sequences); this suggests that a system that learns stimulus-response contingencies is not separate from one that learns sequential contingencies (e.g. Cleeremans, 1997; Destrebecqz and Cleeremans, 2003) if the two cues become associated when correlated.

This has obvious implications for our understanding of automatic learning, and as discussed previously further work must demonstrate the presence of learning of contingencies without the presence of the other to show that they are in fact learned under incidental conditions. Rather than simply training one dual stimulus condition and suggesting that participants can or cannot learn about the contingencies (e.g. Cleeremans, 1997), we can examine whether this was a result of cue competition effects or whether participants are simply unable to learn contingencies. Whilst the contingencies in Experiments 7, 8 and 9 were matched across colour cues and sequences (all 67%), increasing these colour-cue response contingencies to the 80% that produced evidence for implicit learning in Chapter 5 would not devalue any conclusions about the nature of the interaction between the two contingencies and simply make it

more likely to observe colour learning without the presence of sequential contingencies and therefore provide a suitable point of comparison at test.

The role of additional stimuli is further predicted to increase sequence learning when it follows the previous stimulus element in the series. This may be a specific effect on the Same rule sequential contingencies, which is worthy of further investigation. Whilst the RASRN predicted that this learning would reduce the Same rule sequences, further experiments are required to ascertain what sequences this may have an effect on. Nevertheless, this may more generally be a result of some potentiating influence of the representation of $t - 1$ as part of the sequence on the next trial. It also may provide participants with two opportunities to create a stronger memory for the sequence, which may in turn lead to strong representation of sequential elements, thus making these easier to explicitly recall (Perruchet & Vinter, 2002). This would perhaps, in turn, lead to better recall or recognition of these sequence elements which may be captured in some direct test of sequence learning (Rünger, 2012; Shanks, 2005), and therefore this task might be repeated with a battery of such tests (e.g. Destrebecqz & Cleermans, 2003; Dienes & Berry, 1997; Wilkinson & Shanks, 2004). If we see increased explicit knowledge in the Previous group compared to the other groups, or a correlation between subjective and objective measures of explicit knowledge, this may account for the increased learning observed across training.

Furthermore, Chapter 6 provided no evidence that colours which followed the previous two response-stimulus locations increased sequence learning of the exclusive-or rule, but there was a strong trend in that direction in Experiment 7. Therefore, this suggests that perfectly correlated colours with the previous *two* sequential elements did not potentiate sequence learning, but perhaps this effect was reduced by the increased complexity of this relationship, or the increased number of colours to learn about. This provides a number of predictions: either that potentiation of sequence learning is restricted to a greater $t - 1$ representation at t; that participants are unable to strongly associate both two previous trials with a current cue; or that the exclusive-or sequence learning is unaffected by additional cues.

It is of interest then to assess the extent to which participants learn about these trial-by-trial response-stimulus-colour-cue relationships; which could be done in a test phase

where additional cues were random. This would lead to trials that were both consistent and inconsistent with the trained Previous-contingency, on which a significant difference in responding could be expected. Further to this, consistent and inconsistent trials may also provide a further level with which to analyse the learned relationship between cues and response-locations, as it may be possible that a high tone-left response contingency is generalised across to current trials, further providing evidence for the learning of these contingencies as the positive influence on sequence learning; rather than some memory or explicit knowledge of the sequence.

Similarly, it would be of interest to investigate to what extent participants learned or indeed were aware of the $t - 1$ cue contingency between tones or colours on the current trial and the previous response-stimulus location. Whilst participants were asked whether they were aware of any contingency and none reported one we assume that it was an incidental training of this relationship that came to facilitate greater learning. It could just as easily been some explicit knowledge. It is therefore an important design feature of further investigations to try and isolate the influence of one stimulus from the other.

In the task as it stands, this is difficult as the response location sequence is responded to whereas the concurrent tones or colours are not. However, it may be possible in the case of the tone experiment to investigate response location only responding, and tone only responding. In this tone only condition participants would be instructed to make either response when they heard a tone, similar to a prediction task instruction. If participants have learned about the relationship between Previous tones and the sequences of required responses and were using this to improve performance, I predict that they will perform better than the Random group on such a task. However, it might be that the Previous group were learning about the contingency between their *response* and the next tone, which suggests that in the Previous group they may become worse than the Random group in a tone only condition when their responses and next trial tone do not match.

7.8. Concluding comments

It was the intention of this thesis to improve our understanding of how humans learn sequences under incidental conditions by testing predictions of an associative account. I have presented both experimental and computational contributions that provide support for the role of associations in sequence learning tasks, with the influence of stimulus-response associations implicated in the learning of different sequences. Whilst it remains a challenge to accurately model these results, they are consistent with associative predictions based on competition between predictive contingencies within the environment. The work reported in this thesis provides strong evidence that humans can learn complex probabilistic rules as well as simple stimulus-response contingencies automatically and outside of the influence of any explicit knowledge or controlled, intentional learning processes. Therefore, in the wider context of cognition, this thesis offers a better understanding of sequential associative learning processes within the context of a dual process account of human learning.

References

Abrahamse, E. L., Lubbe, R. H. J., Verwey, W. B., Szumska, I., & Jaskowski, P. (2012). Redundant sensory information does not enhance sequence learning in the serial task. *Advances in cognitive psychology*, *8*(2), 109-120.

Aitken, M. R. F. (1996). *Peak shift in pigeon and human categorization.* Unpublished doctoral dissertation, University of Cambridge, UK.

Anastasopoulou, T., & Harvey, N. (1999). Assessing sequential knowledge through performance measures: The influence of short-term sequential effects. *The Quarterly Journal of Experimental Psychology: Section A*, *52*(2), 423-448.

Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Lawrence Erlbaum.

Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977). Distinctive features, categorical perception, and probability learning: some applications of a neural model. *Psychological Review*, *84*(5), 413-451.

Ashby, F. G. (2014). Is state-trace analysis an appropriate tool for assessing the number of cognitive systems? *Psychonomic Bulletin & Review, 21*, 935-946.

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

Bamber, D. (1979). State-trace analysis: A method of testing simple theories of causation. *Journal of Mathematical Psychology*, *19*(2), 137-181.

Baum, L. E. (1972). An equality and associated maximization technique in statistical estimation for probabilistic functions of Markov processes. *Inequalities*, *3*, 1-8.

Beauchamp, A. J., Gluck, J. P., Fouty, H. E., & Lewis, M. H. (1991). Associative processes in differentially reared rhesus monkeys (Macaca mulatta): blocking. *Developmental psychobiology*, *24*(3), 175-189.

Beckers, T., De Houwer, J., Pineno, O., & Miller, R. R. (2005). Outcome additivity and outcome maximality influence cue competition in human causal learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(2), 238-249.

Beckers, T., Miller, R. R., De Houwer, J., & Urushihara, K. (2006). Reasoning rats: forward blocking in Pavlovian animal conditioning is sensitive to constraints of causal inference. *Journal of Experimental Psychology: General*, *135*(1), 92-102.

Beesley, T., Jones, F. W., & Shanks, D. R. (2012). Out of control: An associative account of congruency effects in sequence learning. *Consciousness and cognition*, *21*(1), 413-421.

Beesley, T., & Le Pelley, M. E. (2010). The effect of predictive history on the learning of sub-sequence contingencies. *The Quarterly Journal of Experimental Psychology*, *63*(1), 108-135.

Beesley, T., & Shanks, D. R. (2012). Investigating cue competition in contextual cuing of visual search. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*(3), 709-725.

Berry, D. C., & Broadbent, D. E. (1984). On the relationship between task performance and associated verbalizable knowledge. *The Quarterly Journal of Experimental Psychology*, *36*(2), 209-231.

Bertelson, P. (1961). Sequential redundancy and speed in a serial two-choice responding task. *Quarterly Journal of Experimental Psychology, 13,* 90-102.

Bertelson, P. (1963). S-R relationships and reaction times to new versus repeated signals in a serial task. *Journal of Experimental Psychology, 65,* 478-484.

Bischoff-Grethe, A., Goedert, K. M., Willingham, D. T., & Grafton, S. T. (2004). Neural substrates of response-based sequence learning using fMRI. *Journal of cognitive neuroscience*, *16*(1), 127-138.

Bliss-Moreau, E., & Barrett, L. F. (2009). What's reason got to do with it? Affect as the foundation of learning. *Behavioral and Brain Sciences*, *32*(2), 201-202.

Boakes, R. A. (2009). Learning without thinking. *Behavioral and Brain Sciences*, *32*(2), 202-203.

Boucher, L., & Dienes, Z. (2003). Two ways of learning associations. *Cognitive Science, 27*, 807-842.

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

Broadbent, D. E., FitzGerald, P., & Broadbent, M. H. (1986). Implicit and explicit knowledge in the control of complex systems. *British Journal of Psychology*, *77*(1), 33-50.

Carew, T. J., Walters, E. T., & Kandel, E. R. (1981). Classical conditioning in a simple withdrawal reflex in *Aplysia californica*. *The Journal of Neuroscience*, *1*(12), 1426-1437.

Carpenter, G. A., Grossberg, S., & Reynolds, J. H. (1991). ARTMAP: Supervised real-time learning and classification of nonstationary data by a self-organizing neural network. *Neural networks*, *4*(5), 565-588.

Chaiken, S., & Trope, Y. (Eds.), (1999). *Dual-process theories in social psychology*. Guilford Press.

Chambers, L., Mobini, S., & Yeomans, M. R. (2007). Caffeine deprivation state modulates expression of acquired liking for caffeine-paired flavours. *The Quarterly Journal of Experimental Psychology*, *60*(10), 1356-1366.

Chan, C. (1992). *Implicit cognitive processes: Theoretical issues and applications in computer systems design.* Unpublished doctoral dissertation, University of Oxford, UK.

Cheesman, J., & Merikle, P. M. (1984). Priming with and without awareness. *Perception & Psychophysics*, *36*(4), 387-395.

Cheng, P. W., & Novick, L. R. (1992). Covariation in natural causal induction. *Psychological Review*, *99*(2), 365-382.

Clark, R. E., & Squire, L. R. (1998). Classical conditioning and brain systems: the role of awareness. *Science*, *280*, 77-81.

Cleeremans, A. (1993). *Mechanisms of implicit learning.* Cambridge, MA: MIT Press.

Cleeremans, A. (1997). Sequence learning in a dual-stimulus setting. *Psychological Research*, *60*, 72-86.

Cleeremans, A. (2006). Conscious and unconscious cognition: A graded, dynamic, perspective. In Q. Jing, M. R. Rosenzweig, G. d'Ydewalle, H. Zhang, H. C. Chen & K. Zhang (Eds.), *Progress in Psychological Science around the world. Vol I. Neural, Cognitive and Developmental Issues* (pp. 401-418). Hove: Psychology Press.

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

Cleeremans, A., & Dienes, Z. (2008). Computational models of implicit learning. *Cambridge handbook of computational psychology*, 396-421.

Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, *120*(3), 235-253.

Clegg, B. A. (2005). Stimulus-specific sequence representation in serial reaction time tasks. *The Quarterly Journal of Experimental Psychology Section A*, *58*(6), 1087-1101.

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

Colgan, D. M. (1970). Effects of instructions on the skin resistance response. *Journal of Experimental Psychology*, *86*(1), 108-112.

Cook, S. W., & Harris, R. E. (1937). The verbal conditioning of the galvanic skin reflex. *Journal of Experimental Psychology*, *21*(2), 202-210.

Curran, T. (1995). On the neural mechanisms of sequence learning. *Psyche*, *2*(2). Retrieved from http://www.theassc.org/files/assc/2349.pdf.

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

Curran, T., & Hintzman, D. L. (1995). Violations of the independence assumption in process dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(3), 531-547.

Curran, T., & Keele, S. W. (1993). Attentional and nonattentional forms of sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*(1), 189-202.

D'Angelo, M. C., Jiménez, L., Milliken, B., & Lupiáñez, J. (2013). On the specificity of sequential congruency effects in implicit learning of motor and perceptual sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(1), 69-84.

De Houwer, J. (2009). The propositional approach to associative learning as an alternative for association formation models. *Learning & Behavior*, *37*(1), 1-20.

DeCoster, J., & O'Mally, J. (2011). Specific sequence effects in the serial reaction time task. *Journal of motor behavior*, *43*(3), 263-273.

Dehaene, S., Changeux, J. P., & Nadal, J. P. (1987). Neural networks that learn temporal sequences by selection. *Proceedings of the National Academy of Sciences*, *84*(9), 2727-2731.

Dennis, N. A., Howard Jr, J. H., & Howard, D. V. (2006). Implicit sequence learning without motor sequencing in young and old adults. *Experimental Brain Research*, *175*(1), 153-164.

Deroost, N., & Soetens, E. (2006). Perceptual or motor learning in SRT tasks with complex sequence structures. *Psychological research*, *70*(2), 88-102.

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

Destrebecqz, A., & Cleeremans, A. (2003). Temporal effects in sequence learning. *Advances in Consciousness Research*, *48*, 181-214.

Destrebecqz, A., Peigneux, P., Laureys, S., Degueldre, C., Del Fiore, G., Aerts, J., Luxen, A., Van Der Linden, M., Cleeremans, A., & Maquet, P. (2005). The neural correlates of implicit and explicit sequence learning: Interacting networks revealed by the process dissociation procedure. *Learning & Memory*, *12*(5), 480-490.

Dickinson, A. (1988). Intentionality in animal conditioning. In L. Weiskrantz (Ed.), *Thought without language* (pp. 305-325). Oxford: Clarendon Press.

Dickinson, A., Shanks, D., & Evenden, J. (1984). Judgement of act-outcome contingency: The role of selective attribution. *The Quarterly Journal of Experimental Psychology*, *36*(1), 29-50.

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

Dienes, Z. (2004). Assumptions of subjective measures of unconscious mental states: Higher order thoughts and bias. *Journal of Consciousness Studies, 11* (9), 25-45.

Dienes, Z. (2011). Bayesian versus Orthodox statistics: Which side are you on? *Perspectives on Psychological Sciences, 6*(3)*,* 274-290.

Dienes, Z. (2012). Conscious versus unconscious learning of structure. In P. Rebuschat & J. Williams (Eds.), *Statistical Learning and Language Acquisition*. (pp. 337-364). Mouton de Gruyter Publishers.

Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology, 5*, 781.

Dienes, Z. (in press). How Bayesian statistics are needed to determine whether mental states are unconscious. In M. Overgaard (Ed.), *Behavioural methods in consciousness research.* Oxford: Oxford University Press.

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

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

Dienes, Z., Broadbent, D., & Berry, D. C. (1991). Implicit and explicit knowledge bases in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*(5), 875-882.

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

Dodson, C. S., & Johnson, M. K. (1996). Some problems with the process-dissociation approach to memory. *Journal of Experimental Psychology: General*, *125*(2), 181-194.

Dominey, P. F., Lelekov, T., Ventre-Dominey, J., & Jeannerod, M. (1998). Dissociable processes for learning the surface structure and abstract structure of sensorimotor sequences. *Journal of cognitive neuroscience*, *10*(6), 734-751.

Dunn, J. C. (2003). The elusive dissociation. Cortex, 39, 177-179.

Dunn, J. C. (2008). The dimensionality of the remember-know task: A state-trace analysis. *Psychological Review, 115,* 426-446.

Dunn, J. C., & Kirsner, K. (1988). Discovering functionally independent mental processes: The principle of reversed association. *Psychological Review, 95,* 91-101.

Durlach, P. J., & Rescorla, R. A. (1980). Potentiation rather than overshadowing in flavor-aversion learning: an analysis in terms of within-compound associations. *Journal of Experimental Psychology: Animal Behavior Processes*, *6*(2), 175-187.

Dwyer, D. M., Le Pelley, M. E., George, D. N., Haselgrove, M., & Honey, R. C. (2009). Straw-men and selective citation are needed to argue that associative-link formation makes no contribution to human learning. *Behavioral and Brain Sciences*, *32*(2), 206-207.

Ebbinghaus, H. (1964). *Memory: a contribution to experimental psychology.* New York: Dover.

Ellis, R., & Humphreys, G. W. (1999). *Connectionist psychology: A text with readings*. London: Psychology Press.

Elman, J. L. (1990). Finding structure in time. *Cognitive science*, *14*(2), 179-211.

Evans, J. S. B. (2003). In two minds: dual-process accounts of reasoning. Trends in cognitive sciences, 7(10), 454-459.

Feigenbaum, E. A., & Simon, H. A. (1962). A theory of the serial position effect. *British Journal of Psychology, 53,* 307-320.

Feigenbaum, E. A., & Simon, H. A. (1984). EPAM-like models of recognition and learning. *Cognitive Science, 8,* 305-336.

Forrest, C. L. D. (2012). An associative approach to task-switching. Unpublished doctoral dissertation, University of Exeter, UK.

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

Frensch, P. A., Haider, H., Rüinger, D., Neugebauer, U., Voigt, S., & Werg, J. (2003). Verbal report of incidentally experienced environmental regularity: The route from implicit learning to verbal expression of what has been learned. *Advances in Consciousness Research*, *48*, 335-366.

Fu, Q., Bin, G., Dienes, Z., Fu, X., & Gao, X. (2013). Learning without consciously knowing: Evidence from event-related potentials in sequence learning. *Consciousness and cognition*, *22*(1), 22-34.

Fu, Q., Fu, X., & Dienes, Z. (2008). Implicit sequence learning and conscious awareness. *Consciousness and Cognition*, *17*(1), 185-202.

Gaskell, M. G., & Marslen-Wilson, W. D. (1997). Integrating form and meaning: A distributed model of speech perception. *Language and cognitive Processes*, *12*(5-6), 613-656.

Gilbert, D. T. (1999). What the mind's not. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 3–11). New York: Guilford Press.

Gordon, P. C., & Meyer, D. E. (1987). Hierarchic representation of spoken syllable order. In A. Allport, D. MacKay, W. Prinz, & E. Scheerer (Eds.), *Language perception and production* (pp. 445-462). Orlando: Academic Press.

Goschke, T. (1998). Implicit learning of perceptual and motor sequences: Evidence for independent learning systems. M.A. Stadler & P.A. Frensch (Eds.), *Handbook of implicit learning* (pp. 401-444). Thousand Oaks, CA: Sage Publications.

Graf, P., & Komatsu, S. I. (1994). Process dissociation procedure: Handle with caution! *European Journal of Cognitive Psychology*, *6*(2), 113-129.

Guo, X., Zheng, L., Zhu, L., Yang, Z., Chen, C., Zhang, L., Ma, W., & Dienes, Z. (2011). Acquisition of conscious and unconscious knowledge of semantic prosody. *Consciousness & Cognition, 20*, 417-425.

Gureckis, T. M., & Love, B. C. (2010). Direct associations or internal transformations? Exploring the mechanisms underlying sequential learning behavior. *Cognitive science*, *34*(1), 10-50.

Hartley, T., & Houghton, G. (1996). A linguistically constrained model of short-term memory for nonwords. *Journal of Memory and Language*, *35*(1), 1-31.

Haselgrove, M. (2010). Reasoning rats or associative animals? A common-element analysis of the effects of additive and subadditive pretraining on blocking. *Journal of Experimental Psychology: Animal Behavior Processes*, *36*(2), 296-306.

Hawkes, T. (2003). *Structuralism and semiotics*. Berkley, CA: University of California Press.

Heathcote, A., Bora, B., & Freeman, E. (2010). Modeling choice-similarity effects in episodic recognition. *Journal of Memory and Language*, *62*, 183-203.

Heyes, C. M., & Foster, C. L. (2002). Motor learning by observation: Evidence from a serial reaction time task. *The Quarterly Journal of Experimental Psychology: Section A*, *55*(2), 593-607.

Hikosaka, O., Nakahara, H., Rand, M. K., Sakai, K., Lu, X., Nakamura, K., Miyachi, S. & Doya, K. (1999). Parallel neural networks for learning sequential procedures. *Trends in neurosciences*, *22*(10), 464-471.

Hinton, G. E., & Plaut, D. C. (1987). Using fast weights to deblur old memories. In *Proceedings of the ninth annual conference of the Cognitive Science Society* (pp. 177-186). Erlbaum.

Hird, M. J. (2000). Gender's nature Intersexuality, transsexualism and the 'sex'/'gender'binary. *Feminist theory*, *1*(3), 347-364.

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), 1029-1039.

Hsiao, A. T., & Reber, A. S. (2001). The dual-task SRT procedure: Fine-tuning the timing. *Psychonomic Bulletin & Review*, *8*(2), 336-342.

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of memory and language*, *30*(5), 513-541.

Jarvik, M. E. (1951). Probability learning and a negative recency effect in the serial anticipation of alternative symbols. *Journal of Experimental Psychology, 41,* 291-297.

Jiménez, L., Lupiáñez, J., & Vaquero, J. M. (2009). Sequential congruency effects in implicit sequence learning. *Consciousness and cognition*, *18*(3), 690-700.

Jiménez, L., & Méndez, C. (1999). Which attention is needed for implicit sequence learning? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*(1), 236-259.

Jiménez, L, Méndez, C, y Cleeremans. (1996). Comparing direct and indirect measures of sequence learning. *Journal of Experimental Psychology: Learning, Memory and Cognition, 22,* 948-969.

Jiménez, L., Vaquero, J. M., & Lupiáñez, J. (2006). Qualitative differences between implicit and explicit sequence learning. *Journal of experimental psychology: Learning, Memory, and Cognition*, *32*(3), 475-490.

Jiménez, L. & Vázquez, G. A. (2005). Sequence learning under dual-task conditions: Alternatives to a resource-based account. *Psychological Research, 69*, 352-368.

Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness.* Harvard University Press.

Jones, F. W., Le Pelley, M. E., & McLaren, I. P. L. (2002). The APECS-SRN: Towards a model of SRT sequence learning. In D.B. Fogel (Chair), *Proceedings of the 2002 international joint conference on neural networks* (pp. 692-696). Piscataway, NJ: IEEE.

Jones, F. W., & McLaren, I. P. L. (2009). Human sequence learning under incidental and intentional conditions. *Journal of Experimental Psychology: Animal Behavior Processes*, *35*(4), 538-553.

Jones, F. W., Wills, A. J., & McLaren, I. P. L. (1998). Perceptual categorization: Connectionist modelling and decision rules. *Quarterly Journal of Experimental Psychology, 51*B, 33–58.

Jordan, M.I. (1986). Serial order: A parallel distributed processing approach (Tech. Rep. No. 8604). San Diego: University of California, Institute for Cognitive Science.

Juola, P. & Plunkett, K. (1998). Why double dissociations don't mean much. In M. A. Gernsbacher & S. J. Derry (Eds.), *Proceedings of the 20th Annual Conference of the Cognitive Science Society* (pp. 561-566). Mahwah, NJ: Lawrence Erlbaum Associates.

Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus and Giroux.

Kahneman, D., & Tversky, A. (1982). On the study of statistical intuitions. *Cognition*, *11*(2), 123-141.

Kamin, L. J. (1969). Predictability, surprise, attention and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behaviour* (pp. 279-296). New York: Appleton-Century-Crofts.

Karazinov, D. M., & Boakes, R. A. (2007). Second-order conditioning in human predictive judgements when there is little time to think. *The Quarterly Journal of Experimental Psychology*, *60*(3), 448-460.

Keele, S. W., Ivry, R., Mayr, U., Hazeltine, E., & Heuer, H. (2003). The cognitive and neural architecture of sequence representation. *Psychological review*, *110*(2), 316-339.

Kuhn, G., & Dienes, Z. (2006). Differences in the types of musical regularities learnt in incidental and intentional learning conditions. *Quarterly Journal of Experimental Psychology, 59,* 1725-1744.

Lashley, K. S. (1951). The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior: The Hixon symposium* (pp. 112-136). New York: Wiley.

Le Pelley, M. E., Beesley, T., & Suret, M. B. (2007). Blocking of human causal learning involves learned changes in stimulus processing. *The Quarterly Journal of Experimental Psychology*, *60*(11), 1468-1476.

Le Pelley, M. E., Oakeshott, S. M., & McLaren, I. P. (2005). Blocking and unblocking in human causal learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *31*(1), 56-70.

Lebiere, C., Wallach, D. P., & Taatgen, N. A. (1998). Implicit and explicit learning in ACT-R. In F. Ritter & R. Young (Eds.), *Cognitive modeling* (Vol. II, pp. 183–193). Nottingham, UK: Notthingam University Press.

Lee, J. C., & Livesey, E. J. (2013). Effects of explicit abstract knowledge and simple associations in sequence learning. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 2843-2848). Austin, TX: Cognitive Science Society.

Levey, A. B., & Martin, I. (1975). Classical conditioning of human 'evaluative' responses. *Behaviour Research and Therapy*, *13*(4), 221-226.

Lewicki, P. (1986). Processing information about covariations that cannot be articulated. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*(1), 135-146.

Lewicki, P., Hill, T., & Czyzewska, M. (1992). Nonconscious acquisition of information. *American psychologist*, *47*(6), 796-801.

Lewicki, P., Czyzewska, M., & Hoffman, H. (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*(4), 523-530.

Loftus, G. R. (1978). On interpretation of interactions. *Memory & Cognition*, *6*, 312-319.

Loftus, G. R., Oberg, M. A., & Dillon, A. M. (2004). Linear theory, dimensional theory, and the face-inversion effect. Psychological Review, 111, 835-863.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological review*, *95*(4), 492-527.

Lovibond, P. F. (2003). Causal beliefs and conditioned responses: retrospective revaluation induced by experience and by instruction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(1), 97-106.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*(4), 276-298.

Marcus, D. J., Karatekin, C., & Markiewicz, S. (2006). Oculomotor evidence of sequence learning on the serial reaction time task. *Memory & Cognition*, *34*(2), 420-432.

Marinkovic, K., Schell, A. M., & Dawson, M. E. (1989). Awareness of the CS-UCS contingency and classical conditioning of skin conductance responses with olfactory CSs. *Biological psychology*, *29*(1), 39-60.

Maskara, A., & Noetzel, A. (1993). Sequence recognition with recurrent neural networks. *Connection Science*, *5*(2), 139-152.

Mathews, R. C., Buss, R. R., Stanley, W. B., Blanchard-Fields, F., Cho, J. R., & Druhan, B. (1989). Role of implicit and explicit processes in learning from examples: A synergistic effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(6), 1083-1100.

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*(2), 350-364.

Mayr, U., Kliegl, R., & Krampe, R. T. (1996). Sequential and coordinative processing dynamics in figural transformations across the life span. *Cognition*, *59*(1), 61-90.

McAndrew, A., Jones, F. W., McLaren, R. P., & McLaren, I. P. L. (2012). Dissociating expectancy of shock and changes in skin conductance: An investigation of the Perruchet effect using an electrodermal paradigm. *Journal of Experimental Psychology: Animal Behavior Processes*, *38*(2), 203-208.

McClelland, J. L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, *86*(4), 287-330.

McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, *114*(2), 159-188.

McGlinchey-Berroth, R., Carrillo, M. C., Gabrieli, J. D., Brawn, C. M., & Disterhoft, J. F. (1997). Impaired trace eyeblink conditioning in bilateral, medial-temporal lobe amnesia. *Behavioral neuroscience*, *111*(5), 873-882.

McLaren, I. P. L., Forrest, C. L. D., McLaren, R. P., Jones, F. W., Aitken, M. R. F., & Mackintosh, N. J. (2014). Associations and propositions: The case for a dual-process account of learning in humans. *Neurobiology of learning and memory*, *108*, 185-195.

McLaren, I. P. L., Green, R. E. A., & Mackintosh, N. J. (1994). Animal learning and the implicit/explicit distinction. In N. C. Ellis (Ed.), *Implicit and explicit learning of languages* (pp. 313-332). London: Academic Press.

McLaren, R. P., Jones, F. W., Yeates, F., & McLaren, I. P. L. (2013). Cue competition in human associative learning [Abstract]. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 4044). Austin, TX: Cognitive Science Society.

McLeod, P., Plunkett, K., & Rolls, E. T. (1998). *Introduction to connectionist modelling of cognitive processes*. Oxford University Press.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological review*, *85*(3), 207-238.

Merikle, P. M., & Reingold, E. M. (1991). Comparing direct (explicit) and indirect (implicit) measures to study unconscious memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*(2), 224-233.

Merikle, P. M., & Reingold, E. M. (1992). Measuring unconscious perceptual processes. In R. F. Bornstein & T. S. Pittman (Eds.), *Perception without awareness: Cognitive, clinical, and social perspectives* (pp. 55-80). New York: Guilford.

Minsky, M. L., & Papert, S. A. (1969). *Perceptrons,* Cambridge, MA: MIT Press.

Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Sciences*, *32*(2), 183-198.

Miyawaki, K. (2006). The influence of the response-stimulus interval on implicit and explicit learning of stimulus sequence. *Psychological Research, 70*(4), 262–272.

Mowrer, O. (1947). On the dual nature of learning—a re-interpretation of" conditioning" and "problem-solving.". *Harvard educational review*.

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

Neumann, D. L., & Waters, A. M. (2006). The use of an unpleasant sound as an unconditional stimulus in a human aversive Pavlovian conditioning procedure. *Biological Psychology*, *73*(2), 175-185.

Newell, B. R. (2012). Levels of explanation in category learning. *Australian Journal of Psychology, 64*, 46-51.

Newell, B. R., & Dunn, J. C. (2008). Dimensions in data: Testing psychological models using state-trace analysis. Trends in Cognitive Science, *12*, 285-290.

Newell, B. R., Dunn, J. C., Kalish, M. (2010). The dimensionality of perceptual category learning: A state-trace analysis. *Memory & Cognition, 38,* 563-581.

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

Norman, E., Price, M. C., & Duff, S. C. (2006). Fringe consciousness in sequence learning: The influence of individual differences. *Consciousness and Cognition*, *15*(4), 723-760.

Nosofsky, R. M., Clark, S. E., & Shin, H. J. (1989). Rules and exemplars in categorization, identification, and recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(2), 282-304.

Pavlov, I. P. (1927). *Conditioned reflexes* (F. V. Anrep, Trans.) London: Oxford University Press.

Pearce, J. M., & Bouton, M. E. (2001). Theories of associative learning in animals. *Annual review of psychology*, *52*(1), 111-139.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological review*, *87*(6), 532-552.

Perruchet, P. (1985). A pitfall for the expectancy theory of human eyelid conditioning. *The Pavlovian journal of biological science: official journal of the Pavlovian*, *20*(4), 163-170.

Perruchet, P., & Amorim, M. A. (1992). Conscious knowledge and changes in performance in sequence learning: evidence against dissociation. *Journal of experimental psychology: learning, memory, and cognition*, *18*(4), 785-800.

Perruchet, P., Cleeremans, A., & Destrebecqz, A. (2006). Dissociating the effects of automatic activation and explicit expectancy on reaction times in a simple associative learning task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(5), 955-965.

Perruchet, P., & Pacton, S. (2006). Implicit learning and statistical learning: One phenomenon, two approaches. *Trends in cognitive sciences*, *10*(5), 233-238.

Perruchet, P., & Vinter, A. (2002). The self-organizing consciousness: a framework for implicit learning. In R. M. French & A. Cleeremans (Eds.), *Implicit learning and consciousness* (pp. 41–67). Hove, UK: Psychology Press.

Posner, M.I. & Synder, C.R.R. (1975). Attention and Cognitive Control. In Robert L. Solso (Ed.), *Information Processing and Cognition: The Loyola Symposium*. Lawrence Erlbaum.

Povel, D. J., & Collard, R. (1982). Structural factors in patterned finger tapping. *Acta Psychologica, 52,* 107-123.

Prince, M., Hawkins, G., Love, J., & Heathcote, A. (2012). An R package for state-trace analysis. *Behavioural Research, 44,* 644-655.

Prince, M., & Heathcote, A. (2009). State-trace analysis of the face-inversion effect. In N. A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 685-690). Austin, TX: Cognitive Science Society.

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

Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of experimental psychology: General*, *118*(3), 219-235.

Reber, A. S., & Millward, R. B. (1968). Event observation in probability learning. *Journal of Experimental Psychology*, *77*(2), 317-327.

Reber, P. J., & Squire, L. R. (1994). Parallel brain systems for learning with and without awareness. *Learning & Memory*, *1*(4), 217-229.

Reber, P. J., & Squire, L. R. (1998). Encapsulation of implicit and explicit memory in sequence learning. *Journal of Cognitive Neuroscience*, *10*(2), 248-263.

Redington, M. and Chater, N. (2002). Knowledge representation and transfer in artificial grammar learning. In R. French & A. Cleeremans (Eds.), *Implicit learning and consciousness* (pp. 121–143). Hove, Psychology Press.

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

Restle, F. (1970). Theory of serial pattern learning: structural trees. *Psychological Review*, *77*(6), 481-495.

Restle, F., & Burnside, B. L. (1972). Tracking of serial patterns. *Journal of Experimental Psychology, 118,* 219-235.

Robertson, E. M., & Pascual-Leone, A. (2001). Aspects of sensory guidance in sequence learning. *Experimental Brain Research*, *137*, 336-345.

Robertson, E. M., Tormos, J. M., Maeda, F., & Pascual-Leone, A. (2001). The role of the dorsolateral prefrontal cortex during sequence learning is specific for spatial information. *Cerebral Cortex*, *11*, 628-635.

Rohrmeier, M., Dienes, Z., Guo, X., & Fu, Q. (2014). Implicit learning and recursion. In F. Lowenthal & L. Lefebvre (Eds.), *Language and Recursion* (pp. 67-85). New York: Springer.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland

(Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (Vol. 1, pp. 318-362). Cambridge, MA: Bradford Books.

Rünger, D. (2012). How sequence learning creates explicit knowledge: the role of response–stimulus interval. *Psychological research*, *76*(5), 579-590.

Rünger, D., & Frensch, P. A. (2008). How incidental sequence learning creates reportable knowledge: the role of unexpected events. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(5), 1011-1026.

Rusconi, E., Kwan, B., Giordano, B. L., Umilta, C., & Butterworth, B. (2006). Spatial representation of pitch height: the SMARC effect. *Cognition*, *99*(2), 113-129.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*(5294), 1926-1928.

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

Seth, A. K., Dienes, Z., Cleeremans, A., Overgaard, M., & Pessoa, L. (2008). Measuring consciousness: relating behavioural and neurophysiological approaches. *Trends in cognitive sciences*, *12*(8), 314-321.

Shanks, D. R. (1985). Forward and backward blocking in human contingency judgement. *The Quarterly Journal of Experimental Psychology*, *37*(1), 1-21.

Shanks, D. R. (1995). Is human learning rational?. *The Quarterly Journal of Experimental Psychology*, *48*(2), 257-279.

Shanks, D. R. (2005). Implicit learning. In K. Lamberts & R. Goldstone (Eds.), *Handbook of cognition* (pp. 202-220). London: Sage.

Shanks, D. R. (2007). Associationism and cognition: Human contingency learning at 25. *The Quarterly Journal of Experimental Psychology*, *60*(3), 291-309.

Shanks, D. R. (2009). The associative nature of human associative learning. *Behavioral and Brain Sciences*, *32*(2), 225-226.

Shanks, D. R. (2010). Learning: From association to cognition. *Annual review of psychology*, *61*, 273-301.

Shanks, D. R., & Dickinson, A. (1988). Associative accounts of causality judgment. *Psychology of learning and motivation*, *21*, 229-261.

Shanks, D. R., & Dickinson, A. (1990). Contingency awareness in evaluative conditioning: A comment on Baeyens, Eelen, and Van den Bergh. *Cognition and Emotion*, *4*(1), 19-30.

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

Shanks, D. R., & Perruchet, P. (2002). Dissociation between priming and recognition in the expression of sequential knowledge. *Psychonomic Bulletin & Review*, *9*(2), 362-367.

Shanks, D. R., & St John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral and brain sciences*, *17*(3), 367-395.

Shanks, D. R., Wilkinson, L., & Channon, S. (2003). Relationship between priming and recognition in deterministic and probabilistic sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(2), 248-261.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological review*, *84*(2), 127-190.

Shin, J. C. (2008). The procedural learning of action order is independent of temporal learning. *Psychological research*, *72*(4), 376-386.

Simon, H.A., & Lea, G. (1974). Problem solving and rule induction: A unified view. In L.W. Gregg (Ed.), *Knowledge and cognition* (pp. 105-127). Hillsdale, NJ: Erlbaum.

Soetens, E., Boer, L. C., & Hueting, J. E. (1985). Expectancy or automatic facilitation? Separating sequential effects in two-choice reaction time. *Journal of Experimental Psychology: Human Perception and Performance*, *11*(5), 598-616.

Soetens, E., Melis, A., & Notebaert, W. (2004). Sequence learning and sequential effects. *Psychological research*, *69*(1-2), 124-137.

Spiegel, R., & McLaren, I. P. L. (2006). Associative sequence learning in humans. *Journal of Experimental Psychology: Animal Behavior Processes*, *32*(2), 150-163.

Squire, L. R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychological review*, *99*(2), 195-231.

Stadler, M. A. (1992). Statistical structure and implicit serial learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(2), 318-327.

Stadler, M. A. (1995). Role of attention in implicit learning. *Journal of experimental psychology: Learning, memory, and Cognition*, *15,* 1061-1069.

Stadler, M., & Neely, C. B. (1997). Effects of sequence length and structure on implicit serial learning. *Psychological Research*, *60*, 14-23.

Staresina, B. P., Fell, J., Dunn, J. C., Axmacher, N., & Henson, R. N. (2013). Using state-trace analysis to dissociate the functions of the human hippocampus and perirhinal cortex in recognition memory. *Proceedings of the National Academy of Sciences*, *110*(8), 3119-3124.

Stevenson, R. J., & Boakes, R. A. (2004). Sweet and sour smells: learned synesthesia between the senses of taste and smell. In G. A. Calvert, C. Spence, & B. E. Stein (Eds.), *The handbook of multisensory processing* (pp. 69-83). Cambridge, MA: MIT Press.

Sun, R., Slusarz, P., & Terry, C. (2005). The interaction of the explicit and the implicit in skill learning: a dual-process approach. *Psychological review*, *112*(1), 159-192.

Ungerleider, L. G. (1982). Two cortical visual systems. In Ingle, D. I., Goodale, M. A., & Manfield, R. J. W. (Eds.), *Analysis of visual behavior* (pp. 549-586). Cambridge, Mass.: MIT Press.

Urcelay, G. P., & Miller, R. R. (2009). Potentiation and overshadowing in Pavlovian fear conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *35*(3), 340-356.

Vadillo, M. A., Beesley, T., & Shanks, D. R. (2014, January). *The role of configural representations in contextual cuing*. Paper presented at the meeting of the Experimental Psychology Society, London, UK.

Vadillo, M. A., & Matute, H. (2010). Augmentation in contingency learning under time pressure. *British Journal of Psychology, 101,* 579-589.

Vadillo, M. A., Orgaz, C., & Matute, H. (2008). Overshadowed cues have reduced ability to retroactively interfere with other cues. *Learning and Motivation*, *39*(4), 313-322.

Vervliet, B., Vansteenwegen, D., Baeyens, F., Hermans, D., & Eelen, P. (2005). Return of fear in a human differential conditioning paradigm caused by a stimulus change after extinction. *Behaviour research and therapy*, *43*(3), 357-371.

Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: asymmetries in cue competition. *Journal of Experimental Psychology: General*, *121*(2), 222-236.

Wallace, D. G., & Fountain, S. B. (2002). What is learned in sequential learning? An associative model of reward magnitude serial-pattern learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *28*(1), 43-63.

Wan, L. L., Dienes, Z., & Fu, X. L. (2008) Intentional control based on familiarity in artificial grammar learning.. *Consciousness and Cognition, 17,* 1209-1218.

Watson, J. B. (1913). Psychology as the behaviorist views it. *Psychological review*, *20*(2), 158-177.

Weidemann, G., Tangen, J. M., Lovibond, P. F., & Mitchell, C. J. (2009). Is Perruchet's dissociation between eyeblink conditioned responding and outcome expectancy evidence for two learning systems? *Journal of Experimental Psychology: Animal Behavior Processes*, *35*(2), 169-176.

Wiens, S., & Öhman, A. (2002). Unawareness if more than a change event. Comment on Lovibond and Shanks (2002). *Journal of Experimental Psychology: Animal Behavior Processes, 28*(1), 27-31.

Wilkinson, L., & Shanks, D. R. (2004). Intentional control and implicit sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*(2), 354-369.

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

Willingham, D. B., Greenberg, A. R., & Thomas, R. C. (1997). Response-to-stimulus interval does not affect implicit motor sequence learning, but does affect performance. *Memory & Cognition*, *25*(4), 534-542.

Willingham, D. B., Nissen, M. J., & Bullemer, P. (1989). On the development of procedural knowledge. *Journal of experimental psychology: learning, memory, and cognition*, *15*(6), 1047-1060.

Willingham, D. B., Salidis, J., & Gabrieli, J. D. (2002). Direct comparison of neural systems mediating conscious and unconscious skill learning. *Journal of Neurophysiology*, *88*(3), 1451-1460.

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

Wills, A. J., & McLaren, I. P. L. (1997). Generalization in human category learning: A connectionist account of differences in gradient after discriminative and non-discriminative training. *Quarterly Journal of Experimental Psychology, 50*A, 607–630.

Witnauer, J. E., Urcelay, G. P., & Miller, R. R. (2009). A one-system theory that is not propositional. *Behavioral and Brain Sciences*, *32*(2), 228-229.

Ziessler, M., & Nattkemper, D. (2001). Learning of event sequences is based on response-effect learning: Further evidence from a serial reaction task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*(3), 595-613.