

The influence of prior predictiveness on human incidental learning

Thomas Beesley

A thesis submitted for the degree of Doctor of Philosophy.



UMI Number: U584334

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U584334

Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against
unauthorized copying under Title 17, United States Code.



ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Summary

A common distinction made by theorists examining the mental processes contributing to human learning is that between the automatic formation of associations and controlled reasoning about beliefs. On the other hand, it is widely believed that animal cognition is fundamentally associative in nature. Over the last 25 years, the wealth of data from studies on animal learning has begun to shape our understanding of associative learning in humans. Yet it seems that whether automatic or controlled processes govern human learning is likely to be determined by the context in which learning takes place.

Recent research has highlighted cue-predictiveness as an important component modulating the rate of human learning. Findings consistent with those seen in animals have suggested that an interpretation in terms of associative mechanisms is justified. However, the use of explicit learning paradigms – in which participants are encouraged to engage in hypothesis-testing – makes these data open to alternative explanations.

In this thesis changes in cue-predictiveness were examined under incidental learning conditions: experimental tasks were used in which there was no instruction to learn, which should minimise the contribution to learning of controlled reasoning processes. In Chapters 2 and 3, a series of experiments provides evidence for a change in cue-associability under these conditions, primarily in a sequence learning task. Chapter 4 describes the application of several models of animal conditioning to the data generated in these experiments, and highlights a need for associative models to incorporate changes in cue-associability. The results of these simulations then provide a basis for modifications to a more complex model of sequence learning, the Simple Recurrent Network. Given the parallel between changes in

associability and the allocation of attentional resources, Chapter 5 examines the possibility of changes in attention during sequence learning by measuring eye gaze.

Acknowledgements

Firstly, I would like to thank my supervisor Mike Le Pelley for his tireless help and support throughout this project. It has been a privilege to work with Mike over the last 4 years and I would like to thank him for making this period an immensely enjoyable experience. There are so many things I could mention, but perhaps I'm most grateful that he encouraged me to explore the research ideas I put forward and allowed me to make my own mistakes. He must have known it was the only way I'd learn and that I'd get there in the end.

I would like to thank the School of Psychology for funding this research, with special thanks to Rob Honey for his part in sorting this out. I'd like to thank Mark Suret and Guglielmo Calvini for their helpful advice during weekly M4 meetings, as well as their consistently awful performances on the golf course and, therefore, offering some competitive and enjoyable rounds. Thanks also to my wonderful office mates: to Rose for her kind words of support when my paper was rejected; to Laura for her critical evaluation of the previous night's TV; to Adam for eco-babble, games of squash and keeping my music collection well stocked; and to Suzie for making sure I was "put in perspective".

Several people within the department have offered helpful suggestions and comments on this research, notably John Pearce, Dominic Dwyer, Rob Honey, Mark Haselgrove, Lewis Bott, Christina Howard and Ulrike Hahn. Many researchers outside of the department responded to my theoretical ramblings by email or in person, notably David Shanks, Andy Wills, Ian McLaren and Fergal Jones. Special thanks go to Axel Cleeremans and colleagues at ULB in Brussels, for their theoretical suggestions and hospitality during my stay.

Thanks to all my friends and family who have offered kind words of support along the way. I'd like to thank my parents for repeatedly telling me that I can do anything I put my mind to: it turns out they may have been right all along.

Finally, I would like to thank George. Despite her limited technical knowledge she would listen patiently to my incessant moaning about ropey data, always offering much needed consolation and encouragement. Her love and support kept me going.

19.11.2008

Table of Contents

Summary	i
Acknowledgements	iii
Table of Contents.....	v
Chapter 1 – Introduction	1
Chapter overview.....	1
Associative learning and models of animal conditioning.....	1
<i>The Rescorla-Wagner model (1972)</i>	2
<i>Early evidence from animal learning for changes in associability</i>	5
<i>The Mackintosh model (1975)</i>	8
<i>The Pearce-Hall model (1980)</i>	11
Evidence for changes in associability in humans	14
Non-associative accounts of human contingency learning.....	21
Contingency learning, task demands and mental processes	25
Dissociable learning and memory systems.....	26
Implicit learning	27
<i>Examining awareness in implicit learning</i>	28
<i>The role of attention in implicit learning</i>	36
Chapter 2 – Examining changes in cue-associability in incidental learning tasks.	41
Rationale.....	41
Experiment 1.....	45
<i>Method</i>	48
<i>Results</i>	55
<i>Discussion</i>	59
Experiment 2.....	61
<i>Method</i>	63
<i>Results</i>	67
<i>Discussion</i>	73

Experiment 3.....	75
<i>Method</i>	76
<i>Results</i>	79
<i>Discussion</i>	84
Chapter 3 – Reducing between-subject confounds: evidence for changes in cue-associability in within-subject designs.....	87
Rationale.....	87
Experiment 4.....	87
<i>Method</i>	89
<i>Results</i>	93
<i>Discussion</i>	99
Experiment 5.....	101
<i>Method</i>	103
<i>Results</i>	106
<i>Discussion</i>	115
Experiment 6.....	117
<i>Method</i>	120
<i>Results</i>	124
<i>Discussion</i>	133
Experiment 7.....	135
<i>Method</i>	136
<i>Results</i>	136
<i>Discussion</i>	143
Summary and discussion of the empirical evidence presented in Chapters 2 and 3	144
Chapter 4 – Computational simulations of associability effects in sequence learning.....	149
Computational modelling of psychological processes	149
Simulations with the Rescorla-Wagner (1972) model	149
Simulations with the Mackintosh (1975) model.....	158
<i>Relative predictiveness (with respect to context) as a determinant of cue-associability</i>	158
<i>Relative predictiveness (across trials) as a determinant of cue-associability</i>	164

<i>Absolute predictiveness as a determinant of cue-associability</i>	166
Interim summary	170
Evidence for the acquisition of remote contingencies in sequence learning	172
The Simple Recurrent Network	175
<i>Simulations with the standard SRN</i>	184
<i>Simulations with the Alpha SRN</i>	192
Summary of the simulation results	201
Chapter 5 – Examining the attentional component to sequence learning by measuring eye gaze	204
The relationship between associability and attention	204
Eye gaze as a measure of perceptual learning	207
Rationale for the current experiment	211
Experiment 8	212
<i>Method</i>	217
<i>Results</i>	219
<i>Discussion</i>	241
Chapter 6 – Conclusions	246
Examining conscious awareness	248
Predictive history and the partial reinforcement extinction effect	250
Mechanisms for changes in cue-associability	252
Associability and attention	254
Future directions and final remarks	256
Appendix	258
References	259

Chapter 1 – Introduction

Chapter overview

In this thesis I will present new empirical data and theoretical ideas which aim to further our understanding of human learning. Specifically the research assesses evidence for changes in the ‘associabilities’ of stimuli within an incidental learning paradigm. The theoretical background to these processes stems from work in animal conditioning, work which has led to the formulation of important theoretical models of how animals learn to associate stimuli which co-occur in their environment. It is therefore imperative that these associative learning models be introduced first of all, before moving on to discuss more recent work conducted with human participants. I shall then introduce an important theme of the thesis: to what extent are associability processes in human learning governed by automatic learning mechanisms? I shall propose that one method of investigating this issue is to examine whether associability effects occur under conditions of incidental learning, before introducing the field of implicit learning and discussing related methods and findings.

Associative learning and models of animal conditioning

When two stimuli are presented together – more specifically, when a neutral stimulus (or conditioned stimulus; CS; e.g. a light) signals the presence of a stimulus of some importance (an unconditioned stimulus; US; e.g. an electric shock) – a common consequence of this pairing is that the CS will come to elicit a conditioned response (CR; e.g. the animal will freeze) even when the US is not presented. Following the pioneering work of Pavlov (1927), this phenomenon has come to be known as Pavlovian (or Classical) Conditioning. Building

on the early empirical and theoretical work by Thorndike (1911), Pavlov described this process of conditioning as the formation of an associative connection between the representations of the CS and US, such that excitation of the CS representation automatically activates the representation of the US, which in turn automatically elicits the unconditioned response (UR). The introduction of controlled techniques for collecting conditioning data, in turn led to a more formal description of the manner in which such procedures lead to learned behaviour (e.g. Estes, 1950; Bush & Mosteller, 1951). These early examples provided researchers with a common language for describing conditioning effects, as well as allowing for the generation of transparent, testable predictions. This technique of mathematical description was hugely influential and provided the basic framework for the following associative learning models.

The Rescorla-Wagner model (1972)

A general premise of many early associative learning models (e.g. Bush & Mosteller, 1951) is that if a CS and a US are paired together consistently, and within close temporal proximity, the strength of the association between the representations of these stimuli will increase: learning of the CS-US association will occur. However, this fundamental assumption was called into question when Kamin (1969) produced findings with rats to suggest that the contingency (the frequency of co-occurrence of two events) between the CS and the US was not the only influential factor in associative learning. In Kamin's experiment, a CS was initially pretrained consistently with a US (e.g. a light with a shock). In a subsequent training phase, this CS was presented in compound with a novel CS, again reinforced by the same US (e.g. a light and a tone paired with shock). Despite the fact that

the contingency between this new CS and the US was consistent (tone was consistently paired with shock), very little was learnt about the new CS, as compared with a control condition, for which neither CS was pretrained. In our example, it is as if successful learning of the light-shock relationship blocks any learning about the new tone-shock relationship. This effect of “blocking” has become one of the most widely examined conditioning phenomena of the last 40 years. Kamin suggested that the lack of conditioning to the blocked stimulus occurs as a result of an unsurprising US: the US is already well predicted by the pretrained CS.

‘Cue-competition’ effects such as blocking, were a driving force behind the associative learning model proposed by Rescorla and Wagner (1972) – commonly referred to as the Rescorla-Wagner (hereafter R-W) model. Changes in the associative strength of each CS on trial N are made on the basis of the error derived from the prediction of all cues present on trial N:

$$\Delta V_A = \alpha_A \cdot \beta \cdot (\lambda - \Sigma V) \quad 1.1$$

where α_A is a CS-specific learning rate, β is a US-specific learning rate, λ is the amount of conditioning that the US can support, and ΣV is the sum of the associative strength of all CSs present on that trial. The model inherits the basic functions of previous models based on the linear operator equation (e.g. Bush & Mosteller, 1951). For example, it produces negatively accelerating acquisition and extinction curves which are commonly seen in conditioning data. However, it is the use of a ‘summed error term’ (i.e. $\lambda - \Sigma V$), and therefore its treatment of cues trained in compound, which makes the model a particularly powerful learning theory. We can see how an account of blocking falls naturally out of the cue-competition generated by the summed error term employed in the R-W model. Over the course of pretraining with the CS-US contingency (e.g. the pairing of a light with shock),

the associative strength of the CS_L (the light) will increase. Thus, as V_L increases, it will tend to λ , and so the error in the prediction will reduce (i.e. $\lambda - \Sigma V = 0$). On initial compound trials (light and tone), the summed associative strength of both CSs will also be close to λ (driven entirely by the associative strength of the light). As such the error in prediction will be close to zero, and the associative strength of the tone will remain relatively unchanged. In a similar manner, the model predicts that less associative strength will accrue to stimuli trained in compound (e.g. AB+) as opposed to stimuli trained apart (e.g. A+, B+): the overshadowing phenomenon (e.g. Pavlov, 1927; Kamin, 1969). The basic principle of the R-W model, therefore, is that there is a fixed amount of learning that the US supports (that is, λ) and the amount of conditioning is distributed equally between present CSs. If the US is no longer surprising given the present CSs (i.e. $\Sigma V = \lambda$), no conditioning with that US will occur.

In addition to accounting for well established phenomena, the R-W model made several interesting and somewhat counter-intuitive predictions which, more often than not, have turned out to occur in animal learning. For example, the R-W model predicts that following asymptotic conditioning with two stimuli (A+, B+), pairing these stimuli in compound (AB+) will produce stronger responding than when each is presented separately (e.g. Whitlow & Wagner, 1972; Keyhoe, Horne, Horne, & Macrae, 1994; although see e.g. Aydin & Pearce, 1994; Rescorla & Coldwell, 1995). The R-W model provides a simple explanation of this 'summation' effect. Since the associative strength of each CS following individual training will approach λ , the summed associative strength for the compound will approach 2λ . As a result, when these two perfect predictors are presented in compound and paired with the US, the summation of associative strength will ensure the error term will be

negative (i.e. $\lambda < 2\lambda$), and hence there will be a decrement in the associative strength of each CS (e.g. Kremer, 1978).

Miller, Barnet and Grahame (1995) provide a comprehensive review of the successes and failures of the R-W model, concluding that although it by no means captures all aspects of animal behaviour, since the model makes clear testable ordinal predictions it will continue to serve as a “standard” formulation of associative learning. Whilst it is not the purpose of this summary to discuss all the failures of the R-W model, a group of these failures highlighted a shortcoming of the R-W model and formed a basis for future associative models. This shortcoming was that the associabilities (i.e. the learning rates) of the CS and US (Equation 1: α and β , respectively) remained constant over the course of conditioning. In fact, empirical data from a number of conditioning procedures had suggested that these values were able to change. That is to say, the readiness with which a stimulus will engage in associative learning seems to be a function of the prior associative history of that stimulus.

Early evidence from animal learning for changes in associability

One of the earliest findings to suggest a change in cue-processing was the CS preexposure effect (Lubow and Moore, 1959; Lubow, 1989), more commonly referred to as “latent inhibition”. Latent inhibition refers to the demonstration of attenuated conditioning of a CS with a US, following a period of nonreinforced preexposure with that CS. This effect lies beyond the R-W model. According to this model, preexposure will have no influence on conditioning with the CS: no change in associative strength occurs during preexposure, and consequently, conditioning of that CS should proceed from a state akin to no preexposure.

Despite it being a relatively simple procedure, there is still considerable debate as to the component processes producing latent inhibition (Hall, 1991). For example, it seems that latent inhibition results, in part at least, from a failure in the retrieval of the CS-US association (Revusky, 1971). However, changes in the associability of the preexposed CS are also likely to contribute to the latent inhibition effect (e.g. Pearce, Kaye & Hall, 1982).

A complementary strand of research has examined the phenomenon of “learned irrelevance” (Baker & Mackintosh, 1977): retarded conditioning with a CS following a preexposure phase in which that CS is paired in an entirely uncorrelated manner with the US. Much debate has surrounded the mechanisms responsible for learned irrelevance (for a review, see Bonardi & Ong, 2003), in particular the extent to which learned irrelevance can be attributed to the summation of the effect of latent inhibition and the ‘US-preexposure effect’ (i.e. reduced conditioning as a result of habituation to the US). However, certain experimental procedures have suggested that learned irrelevance equates to more than this summation (e.g. Bennett, Wills, Oakeshott, & Mackintosh, 2000), and therefore that learned irrelevance should be attributed as a change in the ‘reinforcer-specific’ associability of a CS (e.g. Baker, 1976).

Further evidence for changes in α comes from studies on discrimination learning. For instance, Reid (1953; see also Mackintosh, 1963; 1969) trained rats on a simple discrimination in which a black stimulus signalled reinforcement, whilst a white stimulus did not (A+, B-). After a relatively short period of training on this discrimination, rats were able to accurately discriminate between these two stimuli. Reid then transferred one group of rats to the reverse discrimination, in which white now signalled reinforcement, whilst black did not (A-, B+). Another group of rats – group ‘overtrained’ – received additional training on the original discrimination, before being transferred to the reverse

discrimination. Surprisingly, Reid observed faster conditioning with the reversed discrimination in group overtrained.

Again, the ‘overtraining reversal effect’ (ORE) is not predicted by the R-W model. Instead the R-W model would predict that overtraining on the initial discrimination would only strengthen these associations, leading to greater proactive interference, and hence impairing the acquisition of the reverse discrimination. Instead, the ORE suggests that discrimination learning leads the α values of relevant (e.g. colour) and irrelevant cues (e.g. experimental context, stimulus shape and position, etc.) to differentiate over the course of training on the first discrimination. The period of overtraining on the initial discrimination simply ensures further differentiation in these associabilities, which in turn leads to rapid extinction of the original pretraining associations for relevant cues, as well as rapid conditioning for these cues with the reversed pattern of reinforcement.

There is also some empirical evidence to suggest that animals learn to attend to relevant dimensions during discrimination learning (e.g. Mackintosh & Little, 1969; George & Pearce, 1999; Oswald et al., 2001). For example, following training with a discrimination in which the colour of the stimuli accurately signals the pattern of reinforcement whilst the shape of the stimuli does not, conditioning on a new discrimination with novel stimuli, but for which the relationships between the dimensions and reinforcement are consistent (i.e. colour is predictive; shape is non-predictive), proceeds more readily than for one in which these relationships are reversed (i.e. shape is predictive; colour is non-predictive). In other words, a new discrimination resulting from an “intradimensional shift” (IDS) from a previously learned discrimination is solved more readily than one which involves an “extradimensional shift” (EDS). Sutherland and Mackintosh (1971) suggested that these

data provide evidence that attentional mechanisms operate at the dimensional level, acting to “tune out” irrelevant dimensions, and “tune in” relevant ones.

The Mackintosh model (1975)

From the brief summary presented here, it would seem that the associability of a CS (α) can vary as a consequence of its predictive history, and it was therefore necessary for theorists to approach how these changes may be operationalised within the framework of the R-W model. Mackintosh (1975) proposed that all stimuli on a given trial compete for finite attentional resources and changes in α reflect the relative predictive validity of a particular stimulus (CS_A) relative to all other stimuli present on that trial (CS_X). If CS_A is a better predictor of the US than CS_X – formally:

$$|\lambda - V_A| < |\lambda - V_X| \quad 1.2$$

then α_A will increase. Whilst if the relative predictive value of CS_A is less than or equal to that of all other stimuli – formally:

$$|\lambda - V_A| \geq |\lambda - V_X| \quad 1.3$$

then α_A will decrease. Mackintosh (1975) did not specify the exact means by which α will change, but nevertheless the preceding mechanism was sufficient to establish how the model would cope with the limitations we have observed in the R-W model.

As a starting example, consider how the model accounts for blocking by fundamentally different means than the R-W model. Initial pretraining (A+) will cause the associative strength of the CS to rise. When paired in compound with a novel CS (AB+), since A is a relatively good predictor of the US occurring relative to B, according to

equation 1.2 the associability of A will remain high, whilst according to equation 1.3, the associability of B will decrease. Since further compound training will ensure rapid decline in the associability of B, and since the associability of the stimulus controls the rate at which learning proceeds, little associative strength will accrue to this stimulus on subsequent trials.

Since the model explains cue competition effects (e.g. blocking, overshadowing) as resulting from changes in the associability of the CS as opposed to the limited associative strength supported by the US (as in the R-W model), Mackintosh (1975) stated the model could operate with a 'separate', rather than a summed error term:

$$\Delta V_A = \alpha_A \cdot \beta \cdot (\lambda - V_A) \quad 1.4$$

So we can see that although the model closely resembles the R-W model and its predecessors in the use of a linear error correction mechanism, the model was the first to rely so heavily on changes in the processing power devoted to the CS.

By allowing for changes in α to reflect the predictive history of the CS, the Mackintosh model is also able to provide a simple account of many of the findings described above. For instance, the model is able to account for the finding of latent inhibition and the related finding of learned irrelevance. Mackintosh (1975) suggested that the preexposure phases for both latent inhibition (non-reinforced CS preexposure) and learned irrelevance (uncorrelated CS-US pairings) can be thought of as involving the presentation of a compound stimulus of the CS and the experimental context. Since in each case the CS is no better a predictor of the US than the context itself (or its absence, in the case of latent inhibition), according to equation 1.3 α will decline over the course of the pretraining phase, resulting in a retardation in subsequent conditioning.

The Mackintosh model is also able to provide a simple account of the ORE. Training on an initial discrimination will result in an increase in the α value for relevant features, since it is these features that produce the least prediction error. Conversely, irrelevant features will generate a greater amount of prediction error, and therefore the α values for these features will decline. Hence, according to equations 1.2 and 1.3, the α values of relevant and irrelevant stimuli will diverge over the course of training on a discrimination. Provided these α values have not reached asymptotic levels, overtraining simply allows the associabilities of the relevant stimuli to differentiate beyond what is possible using the shorter pretraining schedule. Thus, when the response assignments are reversed, since additional processing power has been devoted to the relevant features of the stimuli (at the expense of the irrelevant stimuli) in the overtrained condition, learning of the new discrimination occurs at a faster rate (e.g. Mackintosh, 1963; 1969).

The model also provides an explanation of facilitated discrimination learning following an IDS, than following an EDS (see above), as a result of a generalisation in associability. Consider an initial discrimination in which a blue circle and a blue square are reinforced (BC+; BS+), whilst a red circle and a red square are non-reinforced (RC-; RS-). Sutherland & Mackintosh (1971) suggested that the solution to the discrimination involves attention being directed towards the colour of stimuli, and away from the shape of the stimuli. Thus, on a subsequent discrimination, in which a green triangle and a green diamond are reinforced (GT+; GD+), whilst a yellow triangle and a yellow diamond are non-reinforced (YT-; YD-), attention to the dimension of colour aids discrimination. However, Mackintosh (1975) suggested that one could also account for these results by assuming that the associabilities of the colours in the original discrimination (the relevant stimuli for discrimination learning) will generalise more readily to physically similar stimuli

(the new colours green and yellow), and less well to physically dissimilar stimuli (the new shapes triangle and diamond). The faster conditioning observed after an IDS than an EDS occurs, therefore, because stimuli belonging to the same dimension are inherently more similar than those belonging to different dimensions.

The Pearce-Hall model (1980)

We have seen in the theory and application of the Mackintosh model how important changes in CS processing are in associative learning. Since the model makes such specific predictions about the rules governing changes in associability (i.e. that α increases if a CS is a better predictor of the US than all other presented CSs; α decreases if the CS is a poorer predictor), it is somewhat surprising then, that data arose shortly afterwards to suggest that there might be exceptions to this pattern. Hall and Pearce (1979, Experiment 3; see also Pearce & Hall, 1979; Kaye & Pearce, 1984) pretrained two groups of rats with either a tone or a light paired with a weak electric shock. In a second stage, all animals received conditioning of the tone with a stronger electric shock. Extinction trials (i.e. CS-) revealed that less conditioning had occurred in those animals given tone-weak shock pretraining, than those given light-weak shock pretraining. This result is particularly surprising, since the Mackintosh (1975) model predicts that if a CS consistently predicts reinforcement (i.e. the tone in the case of the group pretrained with tone-weak shock pairings), the associability of that stimulus should rise over the course of pretraining. However, the results of Hall and Pearce (1979) suggest that the associability of a CS can decline whilst that stimulus is still a good predictor of reinforcement. The interpretation of this effect as a reduction in the associability of the CS was supported by the finding that nonreinforced CS exposure

between pretraining and training was sufficient to abolish the retardation in learning. That is, learning that the tone was no longer a good predictor of the weak shock (it now did not predict this US), restored the associability of the CS (Hall & Pearce, 1982).

As a result of these findings, and in stark contrast to Mackintosh (1975), Pearce and Hall (1980) argued that the associability of a stimulus will decrease if the consequences of that stimulus are well known. Conversely, the associability of a CS will be high if it is not a good predictor of the US. Given this simple inverse relationship between associability and associative strength, Pearce and Hall (1980; see also Pearce, Kaye & Hall, 1982) proposed that changes in associability are made in the following way:

$$\alpha_A^N = |\lambda^{N-1} - \sum V^{N-1}| \quad 1.5$$

The superscripts show that the value of α on trial N is determined by the error in prediction on the previous trial, N – 1. If the outcome on the trial is well predicted by CS_A, the error term will be small (or zero) and therefore the associability of CS_A will be low¹. Changes in the associative strength of a CS are driven largely by the salience and associability of the CS:

$$\Delta V_A = S_A \cdot \alpha_A \cdot \lambda \quad 1.6$$

¹ Pearce, Kaye and Hall (1982) adapted the equation governing changes in α (1.5), so that it was calculated not on the basis of the immediately preceding trial alone, but instead changed gradually across trials. They also proposed that more recent trials would have a greater impact on changes in associability than older trials.

where S is the salience of CS_A . We can see that the prediction error of the model has no direct bearing on changes in associative strength², only via its impact on associability in equation 1.5.

The model is able to account for blocking by the employment of a summed error term in equation 1.5. Since pretraining with CS_A will ensure ΣV is close to λ , once CS_B is presented in compound with CS_A the associability of CS_B will also be close to zero. As a result little associative strength will accrue to CS_B . We can also see how the model produces overshadowing. During compound training, ΣV will quickly approach λ , and so the associability of each stimulus will fall rapidly. Whilst, in the case of single cue conditioning, since ΣV will approach λ more slowly the associability of the CS will also remain high for longer, allowing for greater associative strength to be accrued.

Since the model states that the associabilities of predictive stimuli will decline (see equation 1.5), it provides no clear explanation of the ORE, or the facilitated discrimination after an IDS, over that of an EDS. For example, in the case of the ORE, the model predicts that the associabilities of relevant features for the discrimination will decline over the course of training on the initial discrimination. Provided these associability values have not reached a minimal level, any overtraining that occurs will lead these associabilities to decline

² In fact this is a simplified version of that presented by Pearce and Hall (1980), who detail mechanisms by which inhibitory learning proceeds, and how this learning also interacts with associability. In fact, the effectiveness of the US plays more of a role than that which has been attributed here, particularly with respect to the inhibitory learning mechanism. However, the basic assumption of the model – that associability decreases as the CS comes to perfectly predict the US – is the same in the complete version, and is the issue of primary concern here.

further, which in turn will lead to attenuated conditioning on the reversed discrimination, compared to a condition in which this overtraining does not occur.

Evidence for changes in associability in humans

The beauty of the associative models described above lies in the formal notation in which they are expressed, and we have seen how these models make specific predictions and provide precise accounts of a range of phenomena in animal conditioning. Several years after this revolution in associative learning theory, based on ideas put forward by Alloy and Abramson (1979), Dickinson, Shanks and Evenden (1984) provided some of the first evidence to suggest that the same associative mechanisms might also underlie human causal learning. In a typical human contingency learning (HCL) experiment, participants experience a number of trials in which they are presented with a cue (or cues) and are asked to make judgments on whether they believe a certain outcome will occur on that trial. For example, one widely used procedure sees participants play the role of an allergist, deciding whether a fictitious patient will suffer an allergic reaction (the outcome) after eating certain foods (the cues). Participants are presented with a series of discrete trials in which they experience the contingencies between certain cues and outcomes. Participants are then typically given test procedures in which they must make overt judgements regarding the causal status of certain cues. In a series of contingency learning studies³, Dickinson, Shanks and Evenden (1984) observed that humans produced blocking effects akin to those seen in

³ These studies used a slightly different scenario to the food-allergy paradigm described here, but for all intents and purposes the tasks can be assumed to operate in the same way.

animal conditioning. As a result of this important finding, researchers within the field of HCL began to interpret the behaviour observed in these tasks as reflective of the associative mechanisms attributed to animal conditioning. It is perhaps unsurprising, therefore, that shortly after this amalgamation of research areas, results from HCL research started to have significant impact on associative theory (e.g. Shanks, 1985).

Although the application of associative learning theory to HCL has been extensive (for a review see Dickinson, 2001), the purpose of this section is to review specifically those studies that have investigated the role associability plays in human causal learning. Studies by Lochman and Wills (2003) and Le Pelley and McLaren (2003) provided some of the first demonstrations that the associability of a cue can be modulated in an HCL task (see also Kruschke, 1996). Le Pelley and McLaren's study employed the allergy prediction scenario described above; their design is shown in Table 1.1. Throughout the experiment, participants were presented with compounds of two foods (e.g. eggs and onions), based on which they made a prediction as to which of two reactions the patient would suffer (e.g. dizziness or sweating). In Stage 1, for each compound, one cue was a perfect predictor of the outcome (cues A-D, which consistently predict the occurrence of a single outcome), whilst the other cue was non-predictive (cues V-W, which are paired with each outcome on an equal number of occasions). In Stage 2, cues were paired in novel compounds, each containing one good predictor and one poor predictor from Stage 1.

Stage 1	Stage 2	Test
AV – 1	AX – 3	AC?
BV – 2	BY – 4	BD?
AW – 1	CV – 3	VX?
BW – 2	DW – 4	WY?
CX – 2		
DX – 1		
CY – 2		
DY – 1		

Table 1.1. The design of Le Pelley and McLaren (2003). Letters A-Y are cues (foods).

Numbers 1-4 are outcomes (allergic reactions).

Let us consider what the various associative learning theories would predict about the acquisition of the Stage 2 cue-outcome contingencies. The Rescorla-Wagner (1972) model, in which associability is a fixed parameter, would predict that since all cues are paired with novel outcomes the associative strengths for all cues will be acquired at the same rate. The Pearce-Hall (1980) model makes a similar prediction: since a summed error term is used to determine associability on each trial (equation 1.5), the associabilities of each element of the compound (the ‘good’ and ‘poor’ predictor cues) will be equal at the end of Stage 1. As such learning about these different types of cue will proceed at the same rate. The Mackintosh (1975) model on the other hand, predicts that during Stage 1 the associabilities of cues A-D will increase, since they are better predictors of the outcome on each trial than cues V-Y, for which associability will decrease. As a result, acquisition of the

Stage 2 associations will proceed at a faster rate for good predictor cues than for poor predictor cues.

A test phase was used to assess the extent to which participants had learnt about the various cues during Stage 2 training. Participants were asked to provide causal ratings for a set of novel compounds, using a scale from 0 (the foods are very unlikely to cause the reaction) to 10 (the foods are very likely to cause the reaction). Compounds AC and BD consisted of cues that predicted outcomes 3 and 4 respectively in Stage 2, and all of the cues could be considered good predictors of their Stage 1 outcomes. Compounds VX and WY consisted of cues that predicted outcomes 3 and 4 respectively in Stage 2, and all cues were poor predictors of their Stage 1 outcomes. The extent to which the associabilities of these different sets of cues modulated the rate at which cue-outcome associations were formed during Stage 2, would be shown in the causal ratings given on test. Indeed, Le Pelley and McLaren found that participants gave higher causal ratings for the good predictor compounds AC and BD (to outcomes 3 and 4 respectively) than for the poor predictor compounds VX and WY. The authors concluded that the results were consistent with the predictions of the Mackintosh (1975) model: increased associability of good predictor cues leads attentional processes to focus on these cues for the purposes of future learning.

Bonardi, Graham, Hall and Mitchell (2005) examined the role of attentional processes in a human discrimination learning task. The design of the experiment followed that conducted by Delamater (1998), and is presented in Table 1.2. In a first stage participants saw stimuli from two distinct sets (i.e. snowflakes – sn1 & sn2, and different shades of red – co1 & co2) paired with nonsense syllables (e.g. *wug* and *zif*). In group ‘Same’, both stimuli within each set were followed by the same nonsense syllable. In group ‘Different’, the two stimuli within each set were followed by a different nonsense syllable.

The authors reasoned that the training in group Same should result in behaviour being controlled by the common features shared between the stimuli within each set (since this is all that is required to solve the discrimination). On the other hand, the more difficult discrimination in group Different requires that participants attend to the unique features of the stimuli within each set. In other words, during Stage 1, the features which might be used to discriminate between the members of each group will become less salient (will lose associability) in group Same, but since these features are key to predicting the correct outcome for group Different, these features should remain salient in this group.

	Stage 1	Stage 2
Group Same	Sn1 – “wug”	Sn1 – A
	Sn2 – “wug”	Sn2 – B
	Co1 – “zif”	Co1 – B
	Co2 – “zif”	Co2 – A
Group Different	Sn1 – “wug”	Sn1 – A
	Sn2 – “zif”	Sn2 – B
	Co1 – “wug”	Co1 – B
	Co2 – “zif”	Co2 – A

Table 1.2. The design used by Bonardi et al. (2005). Sn1 and Sn2 were different snowflake stimuli. Co1 and Co2 were rectangular stimuli of light and dark shades of red, respectively. In Stage 1 these stimuli were paired with nonsense syllables “wug” and “zif”. In Stage 2 participants had to categorise each stimulus as belonging to one of two categories, A and B.

In Stage 2 participants were asked to decide which of two categories each stimulus belonged to, with feedback provided. In order to solve the discrimination, participants must be able to differentiate both between the two sets of stimuli, and between the stimuli within each set. In line with their hypothesis, Bonardi et al (2005) found that participants in group Different solved this discrimination faster than those in group Same.

Bonardi et al suggested that the results can be explained by changes in the associabilities of the unique and common features of the stimuli. In group Different, the associabilities of the unique elements of the stimuli would have reached a high level by the end of Stage 1, since these elements were the best predictors of the outcomes during this stage. For group Same, for whom these unique elements were no better at predicting the outcomes than the common features of the stimuli, the associabilities decreased during Stage 1. When discrimination comes to rely on these unique features in the future (Stage 2), since the associabilities of these features are higher in group Different than in group Same, better discrimination performance is observed. The results of Bonardi et al (2005) suggest that the associabilities of certain features within a stimulus will change as a result of these features being good predictors of outcomes during training, thus providing evidence that associability can alter cue processing at an elemental level (see also Suret and McLaren 2003; 2005; McLaren & Suret, 2000).

The results of Le Pelley and McLaren (2003) have direct relevance to findings of blocking in human learning (e.g. Dickinson, Shanks & Evenden, 1984) as they suggest a role for attentional processes in cue competition. Moreover, recent work by Kruschke, Kappenman and Hetrick (2005; see also Wills, Lavric, Croft & Hodgson, 2007) has demonstrated that during compound training in the blocking procedure (AB+, following A+), participants' eye gaze is drawn predominantly towards the pretrained cue and away from the blocked cue, providing indirect evidence for associability processes and moreover an attentional basis to these effects (see Chapter 4 for a detailed discussion of the use of eye gaze as a measure of attention). If blocking in HCL results from a reduction in associability for the blocked cue, then we would expect future learning to be retarded for blocked cues relative to controls.

A recent study by Le Pelley, Beesley and Suret (2007) tested this directly by initially training participants on a standard two stage within-subjects blocking design: participants were presented with cues in either blocking procedures (A+, AB+), or overshadowing procedures (I+, JK+). Participants demonstrated a standard blocking effect on test – lower ratings to cue B than to cue K. In a third stage, participants received compounds of Stage 2 cues paired with novel outcomes, each containing one blocked cue and one overshadowing control cue (e.g. BK+). Le Pelley et al found that participants learnt less about the causal effect of previously blocked cues with respect to these new outcomes than they learnt about previous control cues. These results therefore suggest that blocking leads to a reduction in the associability of the blocked cue (for similar findings in animals, see Mackintosh & Turner, 1971)

Non-associative accounts of human contingency learning

At the time of the seminal work of Dickinson, Shanks and Evenden (1984), research on human learning had rarely been considered with respect to the associative framework emerging from animal conditioning. Up until that time theorists in human learning tended to espouse more cognitive theories of human learning that described these processes as the product of ‘higher-level’ reasoning (e.g. Brewer, 1974). In the case of human contingency judgements, the ΔP statistic seemed to offer an accurate statistical model of the relationship between the co-occurrence of a cue and an outcome:

$$\Delta P = P(E | C) - P(E | \sim C) \quad 1.7$$

Equation 1.7 states that the contingency (ΔP) between a cause (C; the cue) and an event (E; the outcome), is determined by the probability of observing the event given the cause, $P(E|C)$, minus the probability of the event occurring in the absence of the cause, $P(E|\sim C)$ (see Allan, 1980).

The ΔP model, along with other variations (such as ΔD), was able to provide an impressive fit to much of the empirical data on human contingency judgements (see Allan, 1993). As such, these models provided the contemporary challenge to associative accounts of human learning. Initially, however, it wasn't clear how statistical models would be able to account for the cue-competition effects demonstrated by Dickinson et al. (1984). Equation 1.7 states that the contingency between each cue and outcome is calculated independently of other cue-outcome contingencies. As such, learning about the blocked cue (cue B in AB+) should proceed unimpaired. This shortcoming of standard statistical models, such as ΔP , led Cheng and colleagues to propose the probabilistic contrast model (e.g. Cheng & Novick, 1990; 1992; see also Cheng, 1997). This model focused on situations in which multiple independent cues could act as separate or interactive causes of an outcome: as the name suggests, a judgement of causality for each cue is made within the context of (or by means of contrast with) the 'focal set' of potential causes. Thus, the model is able to provide an explanation of cue-competition effects in human contingency judgements by assuming that the causal status of the blocked cue is assessed by contrasting the probability of the outcome on pretraining (A+) and compound (AB+) trials. Since the probability of the outcome occurring when B is present (AB+) is the same as when B is absent (A+), the probabilistic contrast model states that B will be attributed low causal status.

Statistical models, such as the probabilistic contrast model can be described as normative, in the sense that these models provide equations which accurately fit the pattern

of observed data, but make no specific commitments to the underlying psychological processes at work (Shanks, 1995; De Houwer, Beckers & Vandorpe, 2005). A simple extension of these models, and one which begins to move towards a description of the psychological processes, has been to describe the operation of probabilistic contrast as a process of higher-order cognitive inference. Several researchers have suggested that participants will use the cue-outcome contingency information in a given task to make 'rational' causal judgements. For instance, De Houwer, Beckers and Glautier (2002) gave a contingency learning task in which participants used weapons (cues) to damage tanks with various levels of impact (outcomes). A standard blocking design was used, but in one condition, participants were told that the maximum impact of the weapons was 10, whilst in another condition they were told that the maximum impact was 20. After seeing weapon A cause impact 10, followed by weapons A and B causing impact 10, participants were asked to rate how much damage they thought weapon B would produce. Ratings for B were lower in condition 'max 20' than in condition 'max 10'. De Houwer et al. argued that in condition 'max 20' participants inferred that B was non-causal since there was scope to observe a greater outcome intensity on trials on which both A and B were present, whilst in condition 'max 10' a ceiling effect limits the ability to make such an inference – it is possible that B has a causal influence but the ceiling in outcome magnitude means that this influence cannot be observed, and hence participants might be less willing to state that B is non-causal.

Similar findings have been reported by Lovibond et al. (2003) who gave participants a series of pretraining trials before a standard blocking design was presented. In one condition this pretraining demonstrated to participants that compounding separate causal cues resulted in additive effects (I+; J+; IJ++), whilst in another condition the effect was shown to be non-additive (I+; J+; IJ+). Stronger blocking effects were observed in the

additive pretraining condition. Lovibond et al. suggested that additive pretraining leads participants to believe the blocked cue is non-causal, since if it was causal, additive effects would have occurred when it was paired with the pretrained cue (i.e. AB++, not the observed AB+).

Further evidence to support a higher-order account of cue-competition in contingency learning was put forward by De Houwer (2002). In a standard blocking experiment participants were given A+ trials that preceded AT+, but were told that information regarding T was missing during pretraining, such that it may, or may not have been present during the A+ trials. Only those participants who thought T was absent during A+ showed a blocking effect. In a second study, after the AT+ trials, participants were told that T was either present or not during A+ trials. Only those participants who were told it was absent showed a blocking effect.

To summarise, it has been suggested that the results presented in this section are somewhat problematic for associative theories of HCL (e.g. Mitchell & Lovibond, 2002; Lovibond, 2003; Mitchell, De Houwer, & Lovibond, in press). For example, the results of Lovibond et al. (2003) suggest that additivity and non-additivity training procedures lead participants to form a rule about the size of the effect produced by stimuli presented in compound. Participants are then able to apply this rule to assess the causal efficacy of novel cues. Mitchell and Lovibond (2002) argue that "...there is nothing in the associative theories to explain the failure of the general learning mechanism when non-additivity instructions are given prior to training", instead proposing that these, and similar findings in HCL tasks (for a comprehensive review see De Houwer, Beckers & Vandorpe, 2005), are likely to be controlled by higher-order reasoning processes. The sorts of mental operations that are assumed to be involved in such processes (e.g. effortful cue comparison; logical

deduction; the evaluation of beliefs) are in-keeping with the probabilistic contrast model (e.g. Cheng, 1997). As such, the general framework of this approach – that these processes involve conscious, effortful decision making – is at odds with the idea that common mechanisms of contingency learning underpin animal and human behaviour.

Contingency learning, task demands and mental processes

In the preceding sections I have described two contrasting theories of the type of psychological processes driving human contingency learning. On the one hand, associative learning theorists suggest that since much of the data from human learning bears striking resemblance to behaviour seen in animal conditioning experiments, a parsimonious interpretation is to assume that automatic, error driven learning mechanisms underlie learning in humans and non-human animals. In contrast, higher-order reasoning accounts would suggest that humans approach contingency learning as an effortful, controlled process of logical deduction. In fact, there is a wealth of support for both associative and inferential accounts of human learning, and it has been suggested that both associative and higher-order reasoning processes might contribute to HCL (e.g. Shanks & Darby, 1998; Le Pelley, Oakeshott & McLaren, 2005; Shanks, 2007).

The extent to which automatic or controlled processes will influence behaviour is likely to be largely dependent on the task demands (Shanks, 2007). Since the large majority of studies on human contingency learning have used explicit learning paradigms (e.g. the food-allergy paradigm) – in which task instructions are likely to promote the use of conscious, higher-order reasoning strategies in order to ‘solve’ the task at hand – it is perhaps not surprising that the application of associative learning theory to human

contingency learning using these paradigms has received such strong challenges. Perhaps better evidence for associative learning processes in humans might come from demonstrations using tasks which are likely to limit the operation of higher-order reasoning processes. In the final part of this introductory chapter I shall present a brief summary of the literature on implicit learning. I suggest that the incidental learning procedures used in this research area may provide a useful tool for examining the contribution of associative mechanisms to human contingency learning.

Dissociable learning and memory systems

A longstanding distinction in the psychology of human learning and memory is that between the learning and knowledge of facts, and that of procedural skills (e.g. Cohen & Squire, 1980). The primary reason for this separation of these two sub-systems of human learning and memory comes from examination of memory performance in amnesic patients, in which selective deficits to a single sub-system have been observed. One patient in particular, H.M., was shown to be able to learn a task which required motor-learning (i.e. learning a procedural skill), whilst on subsequent days of testing the patient himself showed no explicit knowledge or recognition of having come into contact with the procedural task before (Milner, 1962).

Findings such as these in amnesic patients, led researchers to suggest that some forms of learning – particularly the learning of actions through procedural experience – may be subserved by a different mechanism to that underlying semantic knowledge. Various different descriptions of these two systems have been given over the years, but modern approaches to the issue (for a review see Squire, 2004) have settled on the terms

‘declarative’ (e.g. facts and events) and ‘non-declarative’ (e.g. procedural skills or habits, priming, conditioning and stimulus-response learning). Given the evidence provided by amnesic patients such as H.M., some authors have also described these two systems as ‘explicit’ and ‘implicit’ memory stores (e.g. Graf & Schacter, 1985). This differentiation may also be applied to the acquisition of these memories: is it necessary to be consciously aware in order to learn?

Implicit learning

Over 40 years ago, Arthur Reber (1967; 1969; 1989) coined the term “implicit learning” to describe behaviour he had observed in an ‘artificial grammar learning’ (hereafter AGL) task. In a typical version of this task, participants are initially asked to memorise strings of letters (e.g. MXMR; MRTTRX etc). Following the memorisation of this initial set of strings, participants are informed that these strings have been created from a grammatical structure (but are not told what these grammatical rules are). Participants are then given a test phase in which they have to discriminate between novel letter strings, some of which adhere to the grammatical structure and some of which are randomly created. Reber (1969) found that participants in this condition are able to discriminate between novel grammatical and non-grammatical strings in a subsequent test phase, whilst control subjects who are either asked to memorise non-grammatical strings at the start of the experiment, or who have no memorisation phase, typically perform at chance on this test. Reber (1967; 1969) also noted that those participants shown grammatical strings during the first phase tended to have poor explicit knowledge of the underlying rules governing grammatical strings, which led him to conclude that the incidental training procedure leads to a tacit knowledge base

which is automatically retrieved at test. Since the grammatical strings presented during the test phase are novel, Reber (1989) suggested that participants acquire an abstract representation of the underlying structure – a representation that may not be identical to the structure used to create the strings, but which nevertheless is sufficient to produce above-chance performance.

Examining awareness in implicit learning

Since Reber's (1967) original claim that memorisation in AGL tasks leads to implicit learning of grammatical structure, many other researchers have also observed poor explicit knowledge in the AGL task. For instance Reber and Lewis (1977) asked participants to solve anagram puzzles based on an artificial grammar structure, whilst simultaneously reporting the rules they were using to do so. Over the course of 4 days participants came to solve these puzzles with increasing speed, using more complex rules each session, yet Reber and Lewis (1977) claimed that the explicit knowledge participants reported always emerged after evidence for this knowledge was shown by implicit performance measures (time to solve the anagram).

Reber and Allan (1978) went some way towards quantifying the amount of conscious knowledge participants obtain in AGL tasks. In a standard AGL task, participants were asked to accompany each classification decision on test with a report of any explicit strategy or rule used when making that decision. The authors noted that the proportion of trials on which participants reported explicit knowledge of a classification rule (59%) was significantly less than the proportion of trials on which their classification was accurate (81%). However, Dienes (1993) notes that the authors failed to examine whether the rules

provided by participants could have been used to classify the strings for which no rule was reported. Furthermore, Dienes noted that the statement of a rule may not directly map onto a grammatical judgement if other aspects of the string – which have not been reported on – signalled the string as non-grammatical.

Dulany, Carlson and Dewey (1984) circumvented this problem by asking participants to mark those parts of the test strings that they thought ensured the test string was either grammatical or ungrammatical. From these data it was possible to calculate the validity of each rule: the probability of classifying a string correctly given it contains that rule. If participants' responses were based solely on the rules they report, then percent correct scores (on grammaticality decisions) should correlate perfectly with mean rule validity scores. Dulany et al found the correlation coefficient to be .83, which suggests that participants' explicit knowledge of grammatical rules is largely consistent with the knowledge driving grammaticality decisions.

Although claims that AGL performance is based on implicit mechanisms have continued to emerge (e.g. Matthews, Buss, Stanley & Blanchard-Fields, 1989; Stanley, Matthews, Buss, Kotler-Cope, 1989; Dienes, Broadbent, & Berry, 1991), in an influential review of the field Shanks and St. John (1994) argued that these demonstrations of 'implicit learning' fail to meet two important criteria. The first is the *Information Criterion*: any test of awareness must probe for the same knowledge driving performance on the implicit task. For example, we may initially assume that participants' inability to produce grammatical rules following an AGL task suggests that they have implicitly acquired these rules. Yet they may know (explicitly) that all grammatical strings begin with a certain letter, and this knowledge alone may be driving performance on the classification task. The second important consideration is the *Sensitivity Criterion*: a test of awareness must be sensitive to

all the relevant conscious knowledge. On the basis of this criterion it might be argued that simply asking for free verbal report at the end of an experiment is an insensitive measure of explicit knowledge, which will fail to elicit the complex knowledge-base participants may have acquired. In order to satisfy this latter criterion it is necessary to demonstrate that our test of awareness is exhaustive, which, it has been argued, may be logically impossible (Reingold & Merikle, 1988).

Although these criteria seem somewhat pessimistic with regards to a successful demonstration of implicit learning (Merikle, 1994), Shanks and St. John (1994) believe it is possible for these conditions to be met. For instance, in order to meet the information criterion it is simply necessary to have perfect knowledge (as an experimenter) of how the task is being performed: the criterion would be easily met for simple tasks comprising basic stimuli, but more difficult in say, AGL tasks, for which performance can be driven by multiple features of the stimuli simultaneously. As for the sensitivity criterion, Shanks and St. John suggest this too can be met, so long as the procedural aspects of the chosen test of awareness are akin to those on which implicit performance has been demonstrated. We can see immediately that free recall / verbal report procedures for testing awareness do not meet this criterion. For example in AGL tasks, reports are given in a different context to that in which participants make grammaticality decisions.

In many ways the AGL task is ill-suited as a procedure for investigating implicit learning. Incidental memorisation of letter strings may seem *prima facie* a task which limits the development of explicit learning strategies, but this sort of memorisation is likely to lead to conscious encoding of exemplars, or parts of strings, which may be sufficient to produce test performance on novel items (e.g. Perruchet & Pacteau, 1990). Traditionally, implicit processes are thought to operate in non-declarative procedural tasks, and so tasks that

require motor skill learning would seem better suited for the investigation of implicit learning. The serial reaction time (SRT) task (Nissen and Bullemer, 1987) has been suggested to offer a procedure in which implicit learning processes might play a more prominent role. In this task participants are required to respond to a target stimulus (e.g. a grey circle) which can appear in one of a fixed number of positions (usually 4 or 6) on a computer screen. These positions are commonly arranged in a horizontal line on the screen and each position is assigned a corresponding response key. Participants are usually informed that the task is designed to examine the effect of practice on their reaction times. In fact, the movement of the target stimulus is actually determined by a set sequence, such that participants should be able to learn where the target will appear next given the previous locations. In order to differentiate a more general effect of practice (i.e. participants should get faster on the task as they get used to both responding to the stimulus positions and the timing of the task) from learning of the underlying sequence, control trials are included in which the target moves in a non-sequenced (often random) way. If participants have learnt the specific sequence then reaction times should be longer on these control trials, than on sequenced trials.

Nissen and Bullemer (1987) showed that participants can learn these sequences, as evidenced through their reaction time data, yet when asked to give a verbal report of any knowledge they have of the underlying sequence, this explicit memory seems to be absent. In a similar experiment Willingham, Nissen and Bullemer (1989) performed a post-hoc analysis showing that even when participants were split into 'aware' and 'unaware' groups, the unaware group still exhibited a reaction time (RT) benefit, leading Willingham et al. (1989) to conclude that this group had learned the sequence implicitly.

Although the lack of knowledge expressed through verbal report tests is suggestive of implicit learning, these tests would fail to meet the criteria laid out by Shanks and St. John (1994). Evidence of intact SRT performance in amnesic patients would, however, seem to offer evidence for implicit learning. Nissen and Bullemer (1987; Nissen, Willingham & Hartman, 1989) used the SRT task to examine learning and memory performance in patients with Korsakoff's syndrome: a neurological condition characterised by damage to the diencephalic brain regions, often as a result of chronic alcoholism. These patients are profoundly amnesic, and therefore demonstrate extremely poor explicit learning and memory function, when measured on standard neuropsychological tests (see Brooks & Baddeley, 1976). However, Nissen and Bullemer demonstrated that these patients show unimpaired performance on the SRT task, whilst verbal reports given after the task showed no evidence of awareness for the underlying sequence. Although results with amnesic patients are undeniably useful in understanding the neurological basis of performance in the SRT task (for a review see Curran, 1995), they should not necessarily be taken as direct evidence for the exclusive operation of implicit processes. Shanks and St. John (1994) argued that poor explicit knowledge of amnesic patients on post-training awareness tests could simply be due to poor memory retrieval during the awareness test, rather than SRT performance being unconscious. That is, participants may be aware of task-relevant information while performing the SRT task and may use this information to aid performance, but may then subsequently forget that information prior to the awareness test.

What factors might contribute to a suitably sensitive test of awareness in the SRT task; what features would allow a test to satisfy the information and sensitivity criteria? One test which Shanks and St. John (1994) commented favourably on is the 'Generation task', first used by Nissen and Bullemer (1987). In the generation task, participants are asked to

make an explicit prediction as to where the target stimulus will appear on the next trial. This task uses the same procedural features (i.e. stimulus locations, response keys) as the 'implicit' learning task, with the only essential difference being the instruction to predict the location in which the target will appear, rather than simply to respond to it. Willingham et al. (1989) used this task, providing feedback on each trial – the stimulus appeared in the correct location regardless of whether a prediction was correct or not. Since this feedback allowed participants to learn the sequence during the generation task, any savings in learning (as compared to a control group) would be evidence of conscious awareness of the sequence. Willingham et al. (1989) found that those participants classified as unaware on the verbal report measure did not demonstrate any savings in learning on the generation task. However, when participants were not split dichotomously as aware or unaware, generation performance did correlate positively with performance on the 'implicit' learning task: those participants with higher learning scores (RT on sequenced trials minus RT on unsequenced control trials) tended to generate more sequenced transitions.

Use of the generation task as a means of examining awareness in the SRT task has become a common tool in sequence learning. However, evidence from the majority of studies points towards an association, rather than a dissociation, between implicit performance measures and explicit awareness (e.g. Willingham et al. 1989; Cleeremans & McClelland, 1991; Perruchet & Amorim, 1992; Shanks & Johnstone, 1999; although see Jiménez, Mendez & Cleeremans, 1996). Although we might therefore conclude that learning cannot proceed without awareness, one potential problem with the generation task is that the task is not 'process pure' (Destrebecqz & Cleeremans, 2001). That is, given the generation task is very similar to the one used in the learning task, performance (i.e. the sequence of responses generated by the subject) is likely to be driven not only by

participants' explicit knowledge of the sequence, but also any 'implicit' knowledge they have gained. Therefore, if we make the assumption that participants' knowledge in the task comprises both implicit and explicit forms, then the standard generation task is unable to offer a pure test of only explicit knowledge: participants could be entirely unaware of the sequence, yet their implicit knowledge could automatically initiate a movement towards a certain location, therefore leading to above chance levels on the 'explicit' measure.

In order to address this issue, Destrebecqz and Cleeremans (2001; 2003) modified the generation task by adopting the logic of the 'process dissociation procedure' (PDP) from implicit memory research (e.g. Jacoby, 1991). Following a standard SRT learning stage, participants performed the generation task under one of two instructional conditions: *inclusion*, in which participants had to try and generate the sequence they were trained with (i.e. the normal generation instructions), and *exclusion*, in which participants had to try not to generate the sequence they had been trained with (i.e. they had to try and produce novel or non-sequenced movements of the target). The logic behind the PDP is that if participants have conscious knowledge of the sequence then this aspect of their knowledge should be entirely controllable (Jacoby, 1991), and so participants should be able to control, and hence prevent, the expression of this knowledge under exclusion instructions. Participants' implicit knowledge however, is not controllable, and thus will drive generation of the sequence under either instruction set. A finding of above chance generation performance on the exclusion instruction condition would be evidence of there being some implicit knowledge driving performance. Indeed, under certain conditions⁴, Destrebecqz and

⁴ In this study, participants were placed in either a "no-RSI group" (the response to stimulus interval was 0 milliseconds) or an "RSI group" (250 milliseconds). Generation performance of participants in the no-RSI

Cleeremans (2001; 2003) found equal (and above chance) performance in the two instruction sets.

Wilkinson and Shanks (2004) sought to replicate the findings of Destrebecqz and Cleeremans (2001; 2003), whilst controlling for a potential confound of motivation on generation performance, namely that in these prior studies the exclusion instruction condition was always tested after the inclusion condition. Furthermore, to increase motivation on the generation task, participants were incentivised with a monetary reward for accurate performance (see also Fu, Fu, & Dienes, 2008). In Experiments 1 and 2 they used both deterministic (i.e. a fixed repeating sequence) and probabilistic sequences (i.e. a sequence that includes random noise) and found in each case that performance on the inclusion test was greater than that on the exclusion test. In Experiment 3 they used a one-trial generation test: participants saw a small run of 5 target presentations and had to generate the next position in the sequence. Again, performance was greater on inclusion than exclusion conditions. These results are clearly inconsistent with those of Destrebecqz and Cleeremans (2001; 2003). Wilkinson and Shanks (2004) suggested that the earlier findings of no difference between inclusion and exclusion scores might have been due to low power. In fact, the results of Destrebecqz and Cleeremans' replication (2003) showed a marginally significant difference ($p = .06$) between inclusion and exclusion scores.

The results of Wilkinson and Shanks (2004) are consistent with the notion that knowledge is acquired explicitly in the SRT task, and performance on the 'implicit' and

group suggested they had implicitly learnt the sequence, whilst performance of the RSI group suggested explicit learning: they could suppress generation of the sequenced pattern.

'explicit' measures results from a unitary knowledge base. One caveat to this, however, is that although inclusion rates were consistently higher than exclusion rates, in all 6 conditions examined by Wilkinson and Shanks (2004; see also Shanks, Rowland & Ranger, 2005), exclusion rates were never significantly lower than baseline generation performance (that expected by chance). If knowledge on the task was entirely explicit we would expect participants to be able to fully control this knowledge and generate no sequenced transitions at all – their exclusion score would be zero. Therefore, Wilkinson and Shanks (2004) conceded that although under exclusion conditions explicit knowledge can be suppressed, it is possible that the presence of implicit knowledge leads to a generation score above zero.

The role of attention in implicit learning

Several studies, mostly using the SRT task, have examined the extent to which implicit learning is dependent on attentional resources. A demonstration of learning in a situation where participants are asked to simultaneously engage in an attention-demanding secondary task would be consistent with a view of implicit learning as an automatic learning process. For example, Nissen and Bullemer (1987) modified a standard SRT task to include the concurrent presentation of low- and high-pitched tones. Participants were required to count the number of high-pitched tones and report the total at the end of the block. Nissen and Bullemer showed there was no difference in performance on the SRT task between participants trained on sequenced transitions and participants trained on random transitions, when participants were required to perform a secondary tone-counting task. This result suggests that the use of a dual-task places a limit on the attentional resources available for sequence learning.

Cohen, Ivry and Keele (1990) unpacked the results of Nissen et al. in more detail, showing that simple ‘unique’ sequences, in which each sequence element predicts a unique element on the next trial (e.g., 12341234) could be learnt alongside a concurrent, secondary task, implying that the learning of these transitions can proceed without attention. Cohen et al. also examined learning of ‘ambiguous’ sequences under dual task conditions.

Ambiguous transitions, more commonly referred to as second-order conditional (SOC) transitions, are those in which the location of the target on trial N is contingent on the previous two elements from the preceding sequence (e.g. 13 predicts 4, whilst 23 predicts 1). Cohen et al. found that sequences containing SOC transitions could not be learned under dual-task conditions, which suggests that this learning requires attentional resources. They suggested that knowledge about SOC transitions must be “coded hierarchically” (p. 28) and the storage of several sequence elements increases the demand on short-term memory, which is likely to be disrupted by the secondary task (although see Reed and Johnson, 1994).

The extent to which the studies described above have investigated specifically the attentional demands that tone-counting place on sequence learning has been challenged by several researchers. For instance Stadler (1995), suggested that, rather than placing limits on the learning mechanism per se, the presence of a secondary tone-counting task leads to a disruption of the ‘organization’ of the sequence in memory. Stadler showed similar decrements in learning between a condition in which participants had to count tones, and a condition in which the response-stimulus-interval (RSI) between trials was frequently disrupted. Stadler concluded that attention is required to organize runs of sequenced trials, and this is disrupted by pauses during the processing of secondary-task stimuli.

Similarly, Frensch, Lin and Buchner (1998), argued that the use of a dual task in the SRT might affect the expression of what has been learned. When tested under single-task conditions, the proportion of training given under dual-task conditions had no effect on the amount of learning expressed. This result suggests that learning proceeds unimpaired in the dual-task condition and is therefore consistent with a non-attention demanding implicit learning mechanism.

These studies highlight the difficulties in selecting an appropriate dual-task design for examining the contribution of attention to sequence learning. Several researchers have attempted to circumvent these problems by using a continually changing shape as the target stimulus whilst asking participants to keep track of the number of times a particular shape appears in each block. Jiménez and Mendez (1999; see also Jiménez & Mendez, 2001) trained participants for 31,000 trials using a probabilistic sequence, in which the location of the target on each trial was determined by an artificial grammar (like that used in AGL tasks). Equivalent learning was observed in both single- and dual-task conditions, even when the difficulty of the dual-task was increased. However, recent data from Shanks, Rowland and Ranger (2005) are inconsistent with these findings. Using probabilistic sequences – but over a much shorter training period than that used by Jiménez and Mendez – Shanks et al. found greater learning in single- than dual-task subjects, even when the learning of these groups was tested under single- and dual-task conditions. Shanks et al. suggested that the discrepancy in the results of these two studies may lie in the fact that extended training (as given by Jiménez and Mendez) might lead to the dual-task performance becoming automatic, in turn allowing more attentional resources to be dedicated to the primary task.

A complimentary strand of research within the field of implicit learning has investigated whether two sets of task contingencies can be learnt simultaneously. An observation of equivalent performance in the learning of single and multiple sequences would certainly be in keeping with the notion of an automatic implicit learning mechanism. An early study by Mayr (Experiment 2, 1996) examined this issue by using an SRT task in which the colour of the target stimulus changed on each trial. Although the stimulus could appear in one of four positions on the screen, participants were asked to respond only to the colour of the stimulus on each trial, using different buttons corresponding to each possible colour. For one group of participants the target was presented in a fixed sequence of stimulus colours as well as in a fixed sequence of target locations. In two additional 'single sequence' conditions, either the target colour or the target location was randomly determined on each trial. In all conditions, participants were trained with an attention demanding tone-counting secondary task. Mayr was able to assess learning of each sequence in the dual-sequence condition by introducing periods in which one sequence was replaced with random transitions, whilst the other sequence was held constant. The results suggested that the learning of each sequence (motor and perceptual) in the dual-sequence condition was at least as strong as learning in each of the single-sequence conditions. The results therefore suggest that concurrent learning of a sequence of perceptual locations (i.e. where the target would appear) and a sequence of motor responses (i.e. which colour would appear) is possible, even under dual-task conditions. I shall return to this interesting distinction between different forms of sequence learning in Chapter 5, where I will review additional literature supporting the distinction made by Mayr.

Supporting evidence for unimpaired dual sequence learning was provided in studies by Jiménez and Mendez (1999) and Cock, Berry and Buchner (2002). For instance, in the

Jiménez and Mendez study described above, the location of the target on the next trial could be predicted by both the prior sequence of target shapes and the prior sequence of target locations. Jiménez and Mendez demonstrated unimpaired learning of both contingencies within the task, but only when the attention to the shapes was integral to the task demands (i.e. counting as part of the dual-task). Rowland and Shanks (2006) noted that the tasks used by Jiménez and Mendez (1999) and Cock et al. (2002) had used sequences which correlated, such that it was possible to show ‘learning’ of the secondary sequence from learning a configuration of both sequences, and as such these tasks had not investigated learning of two *purely* independent sequences. In a study designed to address this, Rowland and Shanks (2006) gave participants a dual-SRT task in which two separate four-item stimulus arrays were presented, one above the other, each with a separate target stimulus. Participants were required to concentrate on, and respond to, one of the stimulus arrays. Participants were unaware that the movement of either target was governed by a sequence. In test phases designed to assess the learning of each sequence individually, learning was evident for both the attended and unattended sequence. In a second experiment, when the attentional demand of the primary task was increased, there was no evidence for learning of the secondary sequence. The learning of a secondary sequence seems, therefore, to require the spill-over of attentional resources from the primary sequence to the secondary sequence, supporting the idea that learning in the SRT is impaired when attentional resources are limited.

Chapter 2 – Examining changes in cue-associability in incidental learning tasks.

Rationale

The literature on animal conditioning and human learning demonstrates that the predictive history of a stimulus has a profound influence on the rate of associative learning that will proceed for that stimulus in the future. An associative analysis seems particularly suitable here, since formal mechanisms for these changes in associability have been proposed (e.g. Mackintosh, 1975; Pearce & Hall, 1980), and can be easily applied to existing data from human contingency learning (e.g. Le Pelley & McLaren, 2003). We have seen in Chapter 1 that a simple statistical account (e.g. ΔP) is unable to show cue-competition effects such as blocking, but that these effects are within the scope of more elaborate models, such as the probabilistic contrast model (e.g. Cheng & Novick, 1990; Cheng & Holyoak, 1995).

However, because these models are governed by the statistical relationships between causes and effects, they are unable to account for demonstrations of non-normative causal judgments resulting from differences in the predictive history of causes. For instance, probabilistic models are unable to account for the demonstration of greater learning about good predictor cues over poor predictor cues in HCL (e.g. Le Pelley & McLaren, 2003; Le Pelley, Beesley, & Suret, 2007), since these models do not permit the predictive history of a cue to modulate future learning involving novel outcomes. It is somewhat more difficult, however, to dismiss an inferential or higher-order reasoning account of these data. On the basis of such higher-order reasoning accounts, performance in human contingency learning tasks is presumed to be based on the operation of rational, controlled processes: participants

make deductive inferences regarding the causal status of the cues within the task, on the basis of explicit knowledge of the relationships between cues and outcomes (Mitchell, Lovibond, Minard, & Lavis, 2006).

On first inspection, it seems difficult for such processes to account for effects of prior predictiveness on novel learning. For example, in Le Pelley and McLaren's (2003) demonstration it seems irrational for participants to rate poor predictors as having low causal efficacy in Stage 2, given all outcomes are now novel in this stage, and so effectively all cues should start this stage with equal causal strength with respect to these new outcomes. However, one can imagine how participants might, through generalization, use the causal effectiveness of cues from Stage 1 to draw inferences about how effective those cues will be in Stage 2. That is, whilst an associability account (e.g. Mackintosh, 1975) suggests that little attention is paid to the poor predictor cues in Stage 2, an inferential account might suggest that these cues will be equally attended as the good predictor cues (memory for the cue-outcome pairing is complete), but that participants will actively reject the poor predictor cues as causally ineffective as a result of their causal status within the experimental context as a whole. For example, a participant's inference may be along the lines of "...since cues X and Y were less effective than cues A and B in Stage 1, cues X and Y are also likely to be ineffective with respect to the new Stage 2 outcomes."

Alternatively, one can imagine how attentional processes could operate in combination with higher-order reasoning processes, such that those cues which are causally effective within the experimental context are allocated more attentional resources. For instance, although inferential processes lead to the accurate encoding of all cue-outcome relationships in memory, attentional changes might lead certain cues to influence the decision process more than others. Alternatively, attention might impact the encoding

process. For example, one can imagine that during a blocking procedure (A+ followed by AB+), inferential processes might act on the first compound trial to establish B as causally ineffective, and therefore on subsequent compound trials attention will be allocated primarily towards the pretrained cue (A). Although this latter account in particular might be seen as simply a re-description of the associative position (see Le Pelley, Beesley & Suret, 2007), the mechanisms underlying attentional modulation in inferential reasoning are assumed to be fundamentally dissimilar to those impacting on associative learning:

“...whereas associative attentional models of learning emphasize the role of automatic (in the sense of effortless, unconscious, unintentional, or uncontrollable) processes, an inferential account implies that the impact of attentional processes on contingency judgments will be mediated by controlled processes (i.e., conscious propositional knowledge).” De Houwer, Beckers and Vandorpe (2005, p.245)

According to De Houwer et al. (2005) therefore, associative and inferential accounts make contrasting predictions about whether associability effects would occur in situations in which participants' learning is not mediated by controlled processes using conscious propositional knowledge – that is, situations involving implicit learning. Although there remains some debate over whether the tasks typically claimed to demonstrate implicit learning really demonstrate learning without awareness (see Chapter 1) there remains good reason for using such tasks to study the contribution of higher-order reasoning processes to associability effects in human learning. It seems reasonable to assume that in a standard HCL paradigm (e.g. the food-allergy task) participants are likely to see the task as a problem

with an end solution. That is to say, the nature of the corrective feedback in the task suggests to participants that there is an end solution that can be solved through effortful deductive reasoning. Assuming participants approach the task in this way, one effective strategy in the task would be to assume that the task will contain valid and invalid sources of information (i.e. good and poor predictor cues).

In stark contrast, the incidental learning instructions in an implicit learning paradigm are unlikely to evoke conscious strategies. For instance, in the SRT task, participants are asked to respond to the location of the current target. Note that participants are not asked to make a specific prediction about the location of the target on each trial. Furthermore, in the SRT task participants commonly respond to the location of each stimulus within about 600 milliseconds. When coupled with a response-stimulus interval of say 250 milliseconds, successive stimuli will appear within the space of 1 second of each other. This rapid presentation mode, coupled with the use of incidental learning instructions, will reduce the likelihood that higher-order reasoning processes will be employed in the task, if not abolish them completely. As such, a demonstration of an associability effect in the SRT task would suggest that such processes can occur automatically in human learning and are not reliant on controlled reasoning processes.

As discussed in Chapter 1, the extent to which the knowledge can be described as implicit in the SRT task has sparked a considerable amount of research and debate within the field (for a review see Shanks, 2005). However, there exists no conclusive evidence that learning can proceed without awareness (e.g. Wilkinson & Shanks, 2004), or that methods for assessing conscious knowledge are sensitive enough to reveal a dissociation between implicit and explicit knowledge (see Shanks & St. John, 1994). It is not my intention that the current experiments will contribute directly to this debate, and strong claims as to

whether participants can be classified as unaware in the experiments shall be avoided. Nevertheless, it is generally accepted that learning in the SRT task is incidental (see e.g. Cleeremans, Destrebecqz & Boyer, 1998). That is, there is no directed instruction to participants to learn in the SRT task – participants are asked to respond to, rather than predict, the outcome on each trial, such that the task can be performed with perfect accuracy in the absence of any learning. Consequently there is no explicit requirement for participants to intentionally exploit knowledge of the situation. It is therefore the use of an incidental learning procedure that is of primary interest here, and the extent to which effects observed in more “standard” HCL tasks can also be observed under these conditions.

Experiment 1

As a starting point for examining associability processes in incidental learning tasks, I adapted a condition from one of Reber’s early AGL experiments (Reber, 1969). In Reber’s experiment, participants were asked to memorise three strings of letters at a time. The strings were then removed and participants were asked to try and recall the strings as best they could; Stage 1 consisted of six sets of three strings, which had to be learnt to criterion. Participants were then presented with the same task again in Stage 2, but the content of the strings changed in certain ways. In the condition of interest, the underlying rule structure of the strings changed during Stage 2; while the Stage 2 strings were generated from a grammatical structure, that structure was different to the one used to generate the strings that participants had studied in Stage 1. Nevertheless, participants in this condition still showed ‘savings in learning’ during Stage 2: they made fewer errors in Stage 2 than in Stage 1. However, it is unclear whether this effect was simply due to a general practice effect on the

task, as opposed to transfer of knowledge between the two structures, since Reber (1969) did not include a control condition as a means of comparison (i.e. a condition in which participants experienced unstructured strings during Stage 1; see Redington & Chater, 2002). Reber (1993) suggests that during an AGL task participants might acquire an abstract representation of the grammatical structure used to create strings, which "...captures the patterns of covariation between the various stimulus types that are displayed." (p. 121). However, Reber (1969) made no strong claims about participants' ability to show a transfer of knowledge from one grammatical structure to another.

To my knowledge, this 'syntactic transfer' condition has received no further examination in the AGL literature. However, this condition seems particularly relevant to the aims of this project, since one explanation of this result could be that learning grammatical strings (as opposed to non-grammatical strings) results in a change in the associabilities of the stimuli in the task. In order to explain such an account, consider an AGL task in which participants are presented with the following three grammatical strings: XTRM; XTMV; MRMV. We can see that certain letters in certain positions are good predictors of other letters in the string – e.g. X in position 1 is a good predictor of T in position 2. In general, since grammatical strings are constrained by the underlying structure, the components of these strings will naturally be good predictors of the letters occurring in other positions. In the case of ungrammatical strings, which do not have such constraints, each letter will be a poor predictor of the letters appearing in other locations. According to Mackintosh (1975), if one element (e.g. X in position 1) is a good predictor of another element (e.g. T in position 2), relative to all other elements (e.g. M in position 3 is a poor predictor of T in position 2, relative to X in position 1), the associability of the good predictor element will increase. In the case of ungrammatical strings, the associabilities of

all elements will be low. If the associabilities of the stimulus elements are high at the end of Stage 1, according to Mackintosh (1975), acquisition of a novel structure in Stage 2 should be rapid. On the other hand, if the associabilities are low at the end of Stage 1, acquisition will be slow. For example, if in the new grammatical structure letter X in position 1 is consistently paired with letter R in position 2, then acquisition of the association between these two elements should be rapid when the associabilities of the elements are high (i.e. for participants receiving Stage 1 training with grammatical strings), but much slower when the associabilities are low (i.e. for participants receiving Stage 1 training with ungrammatical strings). Thus, in line with previous work in HCL (e.g. Le Pelley & McLaren, 2003; Bonardi et al. 2005), we would expect acquisition to be rapid when the elements of the stimuli have been predictive (for participants trained with grammatical strings), and slower when they have been non-predictive (in participants trained with ungrammatical strings).

Experiment 1 sought to replicate Reber's (1969) 'grammatical transfer' effect. In an experimental condition, participants were asked to memorise a set of grammatical strings during an initial phase. Instead of using a savings in learning test phase, the more common discrimination test phase was used: participants were asked to decide which strings were grammatical from a new set of strings, half of which were grammatical and half of which were non-grammatical. Grammatical strings in the test phase were created from a different grammar to those presented during the training (memorisation) phase. Any 'grammatical transfer' observed in the experimental condition can be contrasted with performance in a control condition, for whom the training strings were non-grammatical (i.e. randomly created). This control group was used to assess whether discrimination performance in the experimental group resulted from training on structured strings.

A large body of evidence has suggested that during the training phase of an AGL task, participants ‘chunk’ strings into small fragments to aid memorisation (e.g. Gomez & Schvanevelt, 1994; Knowlton & Squire, 1994; Perruchet & Pacteau, 1990; Redington & Chater, 1996; Servan-Schreiber & Anderson, 1990). Grammaticality judgements at test are then assumed to be based on the fragment content of the strings; those test strings containing a higher proportion of fragments encoded during training will have a higher probability of being endorsed as grammatical (Perruchet, Vinter, Pacteau, & Gallego; 2002). Furthermore, participants might learn a host of ‘microrules’ relating to what makes strings grammatical (Dulany et al., 1984). These might, for instance, consist of the simple distributional statistics of grammatical strings: information about string length, legal starting and ending letters, or letter frequency (see Shanks, Johnstone, & Staggs, 1997).

It was imperative, therefore, that in the design of Experiment 1, all these possible factors governing AGL performance were carefully considered. The aim was to design stimuli that, as much as possible, controlled for the influence of these well known factors governing AGL performance. By carefully controlling the stimuli it should be possible to rule out certain interpretations of any observed discrimination performance by statistical knowledge, or fragment memorisation.

Method

Participants, apparatus and display


Fifty-six Cardiff University undergraduate students participated in the experiment for payment of £4. Twenty-eight participants were assigned to the experimental condition, and twenty-eight to the control condition. Testing was conducted in a quiet room, using PCs

with 17" CRT monitors; participants sat approximately 80cm from the monitor. The experiment was run using software written in Visual Basic, which recorded participants' responses during the test phase. Responses were made with the keys 'z' and 'm'.

Design

The two grammatical structures (Grammar A and Grammar B, hereafter GA and GB) used in the experiment are shown in Figure 2.1. Grammatical strings are created by making a pass through the structure from left to right, following the direction of the arrows between nodes. An example grammatical string from GA is "MRVRV". The two grammatical structures were carefully designed to be distinct in several important ways. Firstly, the grammars have unique starting and ending letters. Secondly, the grammars create sets of strings that are unique in their 'bigram' content (the pairs of letters, e.g. MM, XT). In other words, each bigram could feature in strings created by only one of the two grammatical structures. By extension, the structures therefore also contained no overlapping chunks of any greater length. Using the 5 letters M, R, X, T and V, it is possible to create twenty-five bigrams. Twelve of these were assigned to GA and thirteen to GB. For each set of five bigrams starting with the same letter, two were used in one grammatical structure and three for the other. Repetition bigrams (e.g. "MM") were also as equally distributed as possible, with two assigned to GA and three to GB.

Grammar A



Grammar B




Figure 2.1. The two grammatical structures used in Experiment 1. Grammatical strings are created by making a pass from node 1 through to node 8, in the direction of the arrows between nodes.

A large set of strings was created for each grammar, from which strings were selected at random until the two sets of grammatical strings were matched with respect to the following two characteristics. Exactly half of the grammatical strings in each set contained a double (e.g. TT), and no more than 1 double was allowed in each string. String length was limited to between 4 and 7 letters. Non-grammatical strings were created by randomly selecting letters for strings of length 4 to 7. No more than one double could appear in a non-grammatical string. From a large set of these non-grammatical strings, those strings containing bigrams from both GA and GB were selected for use in the experiment. Non-grammatical strings of lengths 5 and 7 contained an equal proportion of legal GA and GB bigrams. For strings of length 4 and 6 this was not possible (strings of these lengths comprise 3 and 5 bigrams, respectively), and so for strings of this length half contained one more GA bigram, and half one more GB bigram. Table 2.1 shows the descriptive statistics for each set of grammatical and non-grammatical strings. The aim was to match as closely as possible the statistics of all four sets of letter strings. All strings are presented in Appendix A.

Set	Length	Doubles	% Unique Letters
GA	5.85 (1.04)	.50 (0.51)	0.72 (0.11)
GB	5.85 (0.88)	.50 (0.51)	0.66 (0.11)
NG ₁	5.84 (1.02)	.45 (0.51)	0.65 (0.10)
NG ₂	5.80 (1.11)	.50 (0.51)	0.67 (0.14)

Table 2.1. Descriptive statistics for the sets of strings used in Experiment 1. GA and GB are grammatical sets of strings produced by the grammars shown in Figure 2.1. NG₁ and NG₂ are non-grammatical sets of strings. Statistics shown are mean length, proportion of strings containing doubles (e.g. TT), and the mean percentage of unique letters for strings in each of the four sets of 20 strings. Standard deviations are given in parentheses.

Procedure

Participants in the experimental condition were trained on one grammatical set of strings, before receiving a test phase containing strings from the ‘untrained’ grammar (that which was not presented during the training phase) and non-grammatical strings (NG₂). Half of these participants were given GA strings in the training phase, and GB strings in the test phase, whilst for the other half of the participants this arrangement was reversed. All of the control participants received training with strings from NG₁. Half of the participants in the control condition received a test phase comprising strings from GA and NG₂, whilst the other half received strings from GB and NG₂. In addition, assignment of response keys in the test phase was counterbalanced across participants; half the participants in each sub-condition made ‘non-grammatical’ responses with their left hand and ‘grammatical’ responses with their right hand, and vice versa for the remaining participants. All participants were given the following instructions at the start of the experiment:

“This is a memory experiment. You will see a list of strings. Strings are defined as a series of letters. Your task is to try to remember as much as you can from the list of strings. There are 20 strings in total. You will have 5 seconds on each string and you may use any method you like to try to memorize the string. After 5 seconds the next string in the list will appear until all 20 strings are present. Try to concentrate on the most recent string. Please wait for the experimenter to tell you to start.”

In the memorization phase itself, strings were presented one after the other, at a rate of one string every five seconds, until all 20 were displayed simultaneously in one column

on the screen. Following this training phase, participants were given instructions regarding the nature of structured strings and the details of the test phase. Experimental participants were informed that they had been memorizing structured strings. Control participants were informed that the strings they had been memorizing were created randomly. All participants were then told that they would be asked to try and pick grammatical strings from randomly created strings in a test phase. Both conditions were told that structured strings were created from a complex set of rules governing the order of letters, and that these rules state which letters can follow and precede other letters. Participants in the experimental condition were told that the grammatical strings in the test phase were created from a different grammatical structure to that which created the strings presented during the training phase. Both conditions were told that this task was very hard, that they were not expected to make 100% correct classifications, and that they should not be concerned by acting on impulse. All participants were told that half of the test strings were grammatical and half were non-grammatical.

Before commencing the test phase, participants were required to complete a key-training phase in which they practised responding to the words ‘grammatical’ and ‘non-grammatical’ with the appropriate keys. Feedback was given after each response; on each trial the phrase “correct response” or “incorrect response” was displayed. Ten ‘grammatical’ and ten ‘non-grammatical’ trials were presented in a random order. Following the completion of the key-training phase, participants received a brief summary of the instructions for the test phase.

The test phase contained 80 trials, comprising two presentations of each of the 40 strings. Each string was repeated only after all test stimuli had been presented once. Each half of the test phase involved a randomized ordering of the test strings, and this randomized

order was different for each participant. Some constraints were placed on the randomization procedure. A test item in the last three trials of the first presentation half could not be presented in the first three trials of the second half. In addition, each sub-block of 10 test items contained exactly five grammatical strings and five non-grammatical strings.

During the test phase participants were required to respond to each test string with the appropriate key for grammatical and non-grammatical responses. After participants had responded, the response “G” or “NG” appeared next to the test item (participants were not informed as to whether this response was correct or incorrect) and the next stimulus appeared after a response-stimulus interval of 1 second. Participants were allowed 3 seconds in which to make a response after which a response of “T/O” (timeout) was presented next to the string and was recorded as such. Items appeared successively in 2 columns on the screen, with 20 test strings appearing in each column, and the first column being filled before the second one was started. Once the first 40 test strings had been presented the screen was cleared and the second presentation commenced. Once participants had responded to a test string they could not change their response. Following the test phase participants were debriefed on the purpose of the study.

Results

Three experimental participants and three control participants produced accuracy below 75% on the simple key-training phase prior to the test. These six participants were removed from subsequent analyses as it is unclear whether they had learned the response configuration sufficiently. The overall pattern of data is the same with these participants included.

The left hand panel of Figure 2.2 shows the grammatical endorsement rates for grammatical and non-grammatical strings during the test phase (i.e. hits vs. false-alarms), for both the experimental and control conditions. An analysis of variance (ANOVA) with a within-subject factor of string type (grammatical vs. non-grammatical) and a between-subject factor of condition (experimental vs. control), revealed a main effect of string type, $F(1, 48) = 4.16, p < .05$, indicating that grammatical strings were endorsed as grammatical more often than non-grammatical strings. There was no main effect of condition, $F < 1$, but importantly there was a significant interaction between string type and condition, $F(1, 48) = 4.89, p < .05$, indicating that participants in the experimental condition showed better discrimination of the different types of test strings, than did control participants. t-tests revealed that participants trained on grammatical strings had higher endorsement rates for grammatical than non-grammatical strings, $t(24) = 3.09, p < 0.01$, and hence showed sensitivity to the different types of test string. This sensitivity was not observed in control participants trained on non-grammatical strings, $t < 1$.

Unfortunately, a mistake was made in the presentation of the test strings in the experimental program: the string VTMVVTR (a legal GB string) was presented incorrectly as VTVVTR. Inspection of Figure 2.1 reveals that the string VTVVTR is not a legal string in GB. Moreover, the string actually contains a legal bigram in GA, namely “TV”. That is, for participants in the experimental condition, the bigram TV appeared during both the training and test phases. It is possible, therefore, that participants in both experimental sub-conditions viewed the bigram TV as a signal of grammaticality in the test phase. In other words participants could have acquired specific item-knowledge from the training phase that could directly facilitate the test phase discrimination.

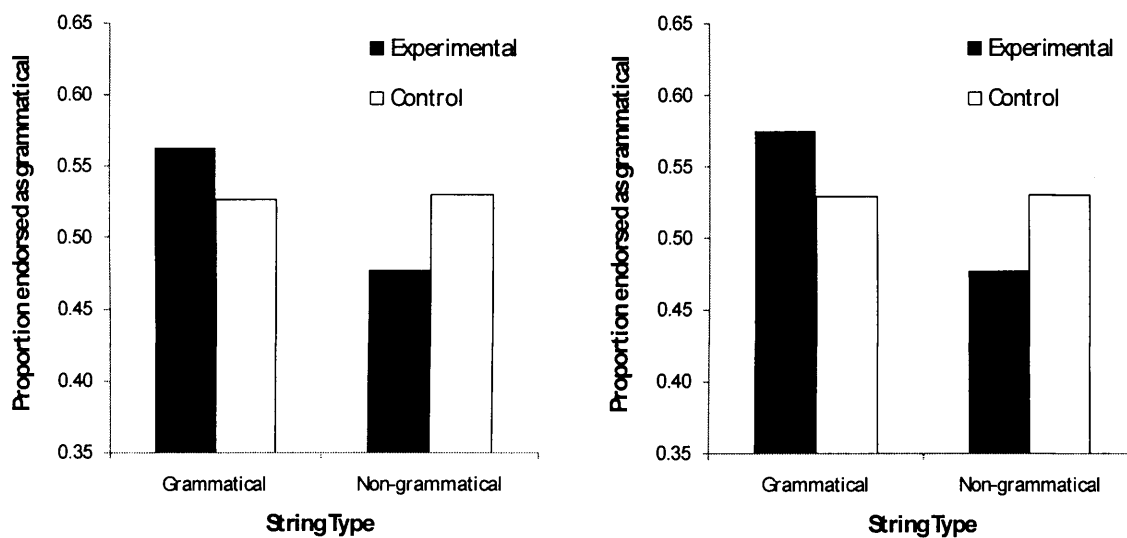


Figure 2.2. *Left panel:* Data from the test phase of Experiment 1. Data show the proportion of grammatical and non-grammatical strings endorsed as grammatical for the experimental and control conditions. Participants in the experimental condition were trained with grammatical strings; participants in the control condition were trained with non-grammatical strings. *Right panel:* the same analysis conducted on the data after the endorsements for a subset of the grammatical test strings were excluded (see text).

Whilst it seems unlikely that such a small feature of one string could have a significant impact on test performance, one must also bear in mind that discrimination performance in AGL tasks is not particularly strong (see Figure 2.2). One must take the conservative view that a single feature – such as a memorised bigram from the training phase – may be sufficient to produce above chance performance on the test discrimination.

As a means of examining whether this erroneous string led to the discrimination performance in the experimental condition, a new analysis of the test phase data was conducted. For participants trained on strings from GB – for whom the erroneous string was present in the training phase – all grammatical test strings containing the bigram TV were removed from the analysis. For participants trained on strings from GA – for whom the erroneous string was present in the test phase – this test string alone (VTVVTR) was removed from the analysis. The same test strings were removed from the test data for participants in the control condition. This new analysis therefore removed any potential for this item-specific knowledge to facilitate the correct endorsement of a grammatical string in the test phase.

The right-hand panel of Figure 2.2 shows the data from this new analysis: the overall pattern of results is unchanged. An ANOVA with factors of string type and condition produced a main effect of string type, $F(1, 48) = 5.65, p < .05$, no main effect of condition, $F < 1$, and importantly a significant interaction between string type and condition, $F(1, 48) = 5.75, p < .05$. These results suggest that the presence of this erroneous string was not responsible for the difference in discrimination performance between experimental and control participants. As previously, t-tests revealed that participants trained on grammatical strings had higher endorsement rates for grammatical than non-grammatical strings, $t(24) = 3.56, p < 0.01$, and hence showed sensitivity to the different types of test strings. This

sensitivity was not observed in control participants trained on non-grammatical strings, $t < 1$.

Discussion

In Experiment 1, participants were trained incidentally on either grammatical or non-grammatical strings during an initial memorisation phase. In a subsequent test phase, all participants were asked to discriminate between grammatical and non-grammatical strings. Only those participants who were trained on grammatical strings were able to successfully discriminate between these different types of test string, despite the fact that the grammatical test strings used at test were created from an entirely different grammatical structure to those that had been presented in the training phase.

The use of two highly controlled grammatical structures makes it unlikely that this ‘grammatical transfer’ effect is due to a simple fragment learning account of AGL performance (e.g. Perruchet & Pacteau, 1990; Redington & Chater, 1996). The grammatical structures were entirely unique in terms of their bigram (and therefore trigram) membership, such that for all participants in the experimental condition, half of the bigram set (the set of 25 possible bigrams from using 5 letters) was used in the grammatical training strings and half of the bigram set was used in grammatical test strings. In fact, if participants were to encode fragmentary components from the training strings and then apply a strategy using this knowledge to make discriminations during the test phase (as suggested by Perruchet and Pacteau, 1990; Perruchet et al., 2002), they would be more likely to classify a non-grammatical test string as grammatical than those from the novel grammatical structure, due to the higher overlap of shared features between the non-grammatical test strings and the

training set. Participants in this task are clearly not relying on such knowledge to make their discriminations during the test phase.

One interpretation of the results of Experiment 1 is that the associabilities of stimuli changed during the training period, as a result of the underlying structure of the letter strings. For those participants trained on grammatical strings, the associabilities of the stimulus elements should be relatively high at the end of the training phase, whilst for those participants trained on non-grammatical strings, the associabilities should be somewhat lower. When these two conditions are then given a discrimination test phase involving strings created from a novel grammatical structure and non-grammatical strings, this difference in associability should lead to faster acquisition of the new grammatical structure at test, relative to the control group. Thus, if one assumes that it is possible to learn about a new grammatical structure during the test phase in the absence of any prior exposure to that grammatical structure (see Redington & Chater, 1996; Reber & Perruchet, 2003), one would expect this learning to be facilitated in a condition in which the associabilities of the stimuli are high, compared to one in which they are low.

It should be stressed, however, that the attribution of the results of Experiment 1 to a change in the stimulus associabilities should be taken as a tentative conclusion. There are, of course, other factors that could also have led to the observed difference in test performance between experimental and control participants. For instance, the overall variability of the grammatical strings is less than that of the non-grammatical strings. As a result, one can imagine that after memorising a set of grammatical strings, participants in the experimental condition might have a better sense of the extent to which grammatical strings vary. When presented with the test strings, participants in the experimental condition could have then used this knowledge about 'set-variance' to pick out strings that are more similar to one

another. Participants in the control condition, who would not be able to gain this appreciation of set-variance from the training phase, would have been less able to group test strings in this way.

The current data do not allow for the separation of these two theoretical accounts of the observed performance in Experiment 1. It is also worth noting another factor that could have contributed to the different performance in the experimental and control conditions, namely the impact of the task instructions on participant motivation. For participants in the control condition, one can imagine that the instructions given before the test phase might result in a lack of belief in their ability to perform the task: the strings they have memorised during the training phase are (at least from the participant's perspective) useless for making test decisions. Conversely, for participants in the experimental condition, the memorisation phase will seem far more relevant, and therefore the level of motivation should be far higher in this group. One can imagine how this difference in the level of motivation in the two groups of participants could impact on the rate of learning in the task, and thus produce the observed difference in classification performance in the two groups.

Experiment 2

For the purpose of examining associability processes in incidental learning conditions, the SRT task seems more suitable than the AGL task for several reasons. Firstly, it is widely accepted that the SRT task is less likely than the AGL task to promote the use of explicit processes (e.g. Shanks & St. John, 1994). As discussed above, the different task instructions given to experimental and control conditions in an AGL task, are likely to lead to different levels of motivation prior to the test phase. Secondly, it is far easier to control the

processing of task information – and consequently the sorts of contingencies participants might learn – in the SRT task, than in the AGL task. There are many ways in which participants might process grammatical letter strings in an AGL task (e.g. left to right, right to left, in chunks, or as a whole, etc.), whilst stimulus presentation in the SRT task is under greater control by the experimenter, and hence is far more uniform and therefore less susceptible to deviations in processing strategy. Finally, whilst AGL performance is usually assessed on the basis of a single test phase, learning in the SRT task can be assessed continuously as training progresses, which allows for a more detailed examination of the time-course of learning (see Method).

Experiment 2 was designed as an SRT analogue of the AGL task used in Experiment 1. Two between-subjects conditions were used: in the ‘sequenced pretrained’ condition, participants were presented with a repeating 12 element second-order conditional (SOC – see Chapter 1) sequence for two blocks. For this group, the target location on each trial can be predicted on the basis of the previous elements in the sequence. In the ‘unsequenced pretrained’ condition, each target location was pseudo-randomly selected on each trial for the first two blocks of the experiment. The Mackintosh (1975) model would predict that the associability of the cues in group sequenced pretrained will be higher than the associability of the cues in group unsequenced pretrained at the end of the pretraining phase. In the subsequent training phase, both groups experienced a novel repeating 12 item SOC sequence. As was the case for the AGL design used in Experiment 1, one can imagine that pretraining with sequenced transitions will lead the associabilities of the stimuli to be at a high level by the end of this phase, since each stimulus is paired in consistent relationships with other stimuli. Conversely, in group unsequenced pretrained, one might expect the associabilities of the stimuli to decrease during pretraining, since for this group each

stimulus is a poor predictor of the location of the target on subsequent trials. To the extent that the associabilities of the stimuli change as a result of these different pretraining procedures, one would expect faster learning of the training sequence in group sequenced pretrained than group unsequenced pretrained.

Method

Participants, apparatus and display

Thirty-four Cardiff University undergraduate students took part for payment of £4 or for course credit. Twenty participants were assigned to the ‘sequenced pretrained’ condition, and 14 were assigned to the ‘unsequenced pretrained’ condition. Testing was conducted in a quiet room divided into two booths to allow two participants to be tested at the same time, using PCs with 17” TFT monitors; participants sat approximately 80cm from the monitor. The experiment was run using software written in Visual Basic. Reaction times (hereafter RTs) were recorded with Windows API functions `QueryPerformanceCounter` and `QueryPerformanceFrequency` for millisecond resolution. Responses were made with a standard keyboard using the keys X, C, V, B, N, M. Error signals were presented over headphones. The stimulus array consisted of six grey response circles (3mm in diameter), evenly spaced 25mm apart in a horizontal line across the middle of the screen. The target stimulus was a larger grey circle, 12mm in diameter. On each trial the target stimulus would appear 20mm below one of the response circles.

Design

Two SOC sequences were created for use in the pretraining and training phases. These will be referred to as SOCA (which was the sequence 124625135634) and SOCB (152653143642). For each participant the elements 1-6 in the two sequences were randomly assigned to stimulus locations on the screen. The two sequences were designed such that both the first- and second-order contingencies within the two sequences were entirely different from one another. For example, in SOCA element 1 has a first-order relationship with 2 and 3 (i.e. it is followed immediately by each of these positions on different occasions), whereas in SOCB element 1 has a first-order relationship with 5 and 4. Likewise, in SOCA element 1 has a second-order relationship with 4 and 5 (i.e. these elements occur two positions after element 1), whereas in SOCB element 1 has a second-order relationship with 2 and 3. Item frequency was consistent between both sequences, with each stimulus position occurring twice in a full sequence run. By using sequences that contain entirely different first- and second-order contingencies, it is possible to rule out the direct influence of pretraining associations on performance during the training phase. Therefore, by manipulating the content of the sequences in a similar manner to how the letter strings were designed in Experiment 1, one can be sure that the associations formed during sequence pretraining (e.g. with SOCA) will themselves not be of direct benefit for performance on the training sequence (e.g. SOCB). Thus, any facilitation in sequence learning during the training phase in group sequenced pretrained cannot be attributed directly to the transfer of associations.

The sequences were presented in a repeating fashion, such that the last element of the sequence was followed by the first element (so SOCA could also be presented as 563412462513 for example). In order to provide an 'online' means of assessing sequence

learning at each stage of the experiment, a short set of unsequenced (random) trials was interspersed between each run of sequenced trials. Sequence learning can be assessed by comparing accuracy and reaction time (RT) on sequenced and unsequenced trials; if participants have learnt the sequence they will be both faster and more accurate on sequenced, than unsequenced trials. Each block (of 174 trials) started with 4 unsequenced trials, followed by ten, 13-item sequence runs (each run started and ended in the same position), with each run interspersed with a set of 4 unsequenced trials. Although there are 12 starting positions for the SOC sequences, to keep the blocks to a reasonable length a subset of 10 starting positions was randomly selected for the 10 sequence runs in each block. The order of these 10 starting positions within each block was randomised. The 4 unsequenced trials at the start of each block were constrained such that no first-order conditional transitions from the current sequence were created. Neither repetitions (e.g., 1-2-2) nor alternations (e.g., 1-2-1) were permitted, as these have been shown to give speeded responses (Soetens, Boer & Heuting, 1985).

The unsequenced pretrained condition received pseudo-random sequences during the pretraining stage. For these blocks, each trial was generated randomly with the constraint that repetitions and alternations could not occur.

Earlier pilot work with these SOC sequences had shown that a ratio of 13 sequenced to 4 unsequenced trials (as described above) was sufficient to observe sequence learning after 2 blocks. As such, participants received 2 blocks of pretraining and a subsequent 3 blocks of training.

Procedure

The assignment of the two sequences to the two stages was counterbalanced, such that half the sequenced pretrained participants experienced 2 blocks of SOCA, followed by 3 blocks of SOCB, whilst the other half experienced 2 blocks of SOCB, followed by 3 blocks of SOCA. All participants in the unsequenced pretrained condition were given 2 blocks of unsequenced trials, with half the participants then experiencing 3 blocks of SOCA, and half experiencing 3 blocks of SOCB. Participants were given the following verbal instructions:

“The aim of this study is to examine the effect of practice on motor control. In this task you are required to follow a grey circle as it moves between six positions on the screen. The six positions are situated across the middle of the screen and are represented by six smaller grey circles. Each of the six positions corresponds to a key on the keyboard, these keys are X, C, V, B, N and M, along the bottom row of the keyboard. X is used to respond to the far left position, M the far right position, and the others for the positions in-between. Each time the larger circle changes position you are required to press the appropriate key to identify its new location. Once you have pressed a key, the circle representing that position will turn red to indicate where you have responded. Should you respond incorrectly you will hear a beep in the headphones. After you have responded the larger grey circle will disappear and reappear in a new location. In summary, your task is to follow the circle as it moves between the six positions. Each time the stimulus moves you should respond as fast, but as accurately as you can. In particular, we want you to avoid making errors in this task. The experiment is split into 5 blocks, each of which lasts for approximately 3 minutes. At the end of each block you will be given a rest of 15 seconds before the next block starts. During this break a countdown will appear to show you when the next block of trials will start.”

Results

The first four trials of each block were not analysed; nor were trials with RTs under 100 milliseconds (ms) or over 1000 ms. RTs to inaccurate trials were also discounted. Figure 2.3 shows RT data for the both the sequenced pretrained (top-left panel) and unsequenced pretrained (top-right panel) conditions. An analysis of variance (ANOVA) was conducted on the pretraining data from group sequenced pretrained, with within-subject factors of trial type and block. This revealed a main effect of trial type, $F(1, 19) = 7.84, p < .05$, with significantly faster responding on sequenced trials than unsequenced trials. It is important to note that sequence learning will affect RTs on both sequenced and unsequenced trials. Whilst on sequenced trials, anticipation of the next target location will result in decreases in RT, this same anticipation will have the effect of increasing RT when that sequenced location is replaced by an unsequenced location. The ANOVA also revealed a main effect of block, $F(1, 9) = 5.27, p < .05$, with mean RT decreasing across pretraining. The interaction between these factors was non-significant, $F < 1$. In line with the pilot data, sequence learning was observed after two blocks of pretraining. An ANOVA conducted on the data from the unsequenced pretrained condition revealed that RTs decreased across the pretraining phase, $F(1, 13) = 5.58, p < .05$.

In order to assess any overall RT differences between the two conditions, RTs in the sequenced pretrained condition (an average of the RTs on sequenced and unsequenced

trials⁵) were compared to RTs in the unsequenced pretrained condition. An ANOVA with condition as a between subject-factor and block as a within-subject factor revealed no effect of condition, $F < 1$, a significant effect of block, $F(1, 32) = 9.57$, $p < .01$, but no interaction, $F < 1$. This suggests that the general rate of responding in the two conditions was similar during the pretraining phase.

Figure 2.4 shows accuracy data for the sequenced pretrained (top-left panel) and unsequenced pretrained (top-right panel) conditions. An ANOVA on the data from the sequenced pretrained condition, revealed a marginally significant effect of trial-type, $F(1, 19) = 4.02$, $p = .059$ (with higher accuracy on sequenced than unsequenced trials), but no significant effect of block, $F < 1$, and no significant interaction, $F(1, 19) = 2.57$, $p = .13$. As with the RT data, an ANOVA was conducted on the accuracy data from the sequenced pretrained (an average of the accuracy on sequenced and unsequenced trials) and unsequenced pretrained conditions. This revealed a main effect of condition, $F(1, 32) = 6.71$, $p < .05$, but no effect of block, $F(1, 32) = 1.65$, $p = .21$, nor an interaction effect, $F < 1$. The data suggest that participants in the unsequenced pretrained condition might have been performing more accurately than participants in the sequenced pretrained condition during the pretraining phase.

⁵ Although a crude measure, the average of these two trial-types is arguably the most accurate estimate of the overall speed of responding available.




Figure 2.3. RT data for pretraining (blocks 1-2) and training (blocks 3-5) phases in Experiment 2. ‘Seq. 1’ and ‘Seq. 2’ refer to sequenced transitions, whilst ‘Unseq. 1’ and ‘Unseq. 2’ refer to unsequenced transitions. *Top-Left panel*: Data for participants pre-trained with sequenced transitions. *Top-Right panel*: Data for participants pre-trained with unsequenced (pseudo-random) transitions. *Bottom panel*: The training phase data expressed as learning scores – RT on unsequenced trials minus RT on sequenced trials.




Figure 2.4. Accuracy data for pretraining (blocks 1-2) and training (blocks 3-5) phases in Experiment 2. ‘Seq. 1’ and ‘Seq. 2’ refer to sequenced transitions, whilst ‘Unseq. 1’ and ‘Unseq. 2’ refer to unsequenced transitions. *Top-Left panel*: Data for participants pre-trained with sequenced transitions. *Top-Right panel*: Data for participants pre-trained with unsequenced (pseudo-random) transitions. *Bottom panel*: The training phase data expressed as learning scores – accuracy on sequenced trials minus accuracy on unsequenced trials.

During the training phase (blocks 3-5) all participants experienced both sequenced and unsequenced transitions. Figure 2.3 shows that the pattern of RT data for the sequenced pretrained condition (top-left panel) is somewhat different to that for the unsequenced pretrained condition (top-right panel). Participants in the unsequenced pretrained condition show a steady level of learning throughout training, with the difference between sequenced and unsequenced trials never exceeding the level shown in Block 4. In comparison, although participants in the sequenced pretrained condition show no evidence for learning initially in Block 3, learning proceeds rapidly and exceeds learning in the unsequenced pretrained condition by the end of the training phase (Block 5).

These data were subjected to an ANOVA with trial-type and block as within subject factors, and condition as a between-subject factor. This revealed a significant effect of trial type, $F(1,64) = 39.31$, $p < .001$, indicating that participants learnt the new sequence presented during the training phase. There was no main effect of block, however the interaction between block and trial type was significant, $F(2,64) = 4.71$, $p < .05$, demonstrating that, on average, the magnitude of the learning effect increased across the 3 training blocks. There was no main effect of condition, nor an interaction between condition and trial type, both $F_s < 1$, and no interaction between condition and block, $F(2,64) = 1.67$, $p = .20$. However, the three-way interaction was significant, $F(2,64) = 5.45$, $p < .01$, indicating that participants in the two conditions differed in the rate at which they learnt about the training sequence. Inspection of Figure 2.3 shows that the effect is driven largely by the different pattern of RTs on unsequenced trials in the two groups. As mentioned above, sequence learning will lead not only to faster RTs on sequenced trials, but also to slower RTs on unsequenced trials, in which an anticipated location does not occur.

The training phase effect is illustrated clearly in the bottom panel of Figure 2.3, which shows the data expressed as learning scores (RTs on unsequenced trials minus RTs on sequenced trials) for the two conditions across the 3 blocks of the training phase. In order to establish the nature of this interaction, pre-planned independent samples t-tests were carried out on the learning scores for each block during the training phase. The difference between the two conditions in Block 3 was significant, $t(32) = 2.35$, $p < .05$, indicating that for participants who were pretrained with sequenced transitions learning of the second sequence was initially attenuated, compared to participants pretrained with unsequenced transitions. There was no significant difference on Block 4, $t < 1$, but there was a significant difference in Block 5, $t(32) = 2.31$, $p < .05$, indicating that participants pretrained with sequenced transitions demonstrated greater learning of the second sequence by the end of the training phase, compared to participants pretrained with unsequenced transitions.

Figure 2.4 shows the accuracy data from the training phase for both the sequenced and unsequenced pretrained conditions. There was evidence for learning in the two conditions by the end of the training phase (Block 5). These data were also subjected to ANOVA, however all main effects were non-significant: trial type, $F(1,64) = 1.05$, $p = .31$; block, $F < 1$; and condition, $F(1,32) = 2.56$, $p = .12$. The only interaction which reached significance was that of block and trial type, $F(2,64) = 4.00$, $p < .05$, indicating that the difference in error rate between unsequenced and sequenced trials increased as training progressed. Learning scores (accuracy on sequenced trials minus accuracy on unsequenced trials) are plotted in the bottom panel of Figure 2.4. There was no significant difference between the two conditions in the amount of learning shown in the accuracy data for any of the training blocks, all $t_s(32) < 1.11$, $p_s > .28$.

Discussion

The results of Experiment 2 demonstrate that participants pretrained with sequenced elements learnt more about a novel sequence by the end of the training phase than did participants pretrained with pseudo-random transitions. There was a striking difference in the pattern of data in the two groups. Although sequenced pretrained participants showed impaired learning during the early stages of training (Block 3), by the end of training (Block 5) these participants demonstrated greater sequence learning than participants in the unsequenced pretrained condition. The data are in keeping with the results of Experiment 1 and Reber's (1969) 'syntactic transfer' condition. The data also provide preliminary evidence in support of a role for associability in sequence learning: the different pretraining schedules may have resulted in the associabilities of the stimulus elements being different in the two groups by the end of the pretraining phase, which in turn would have modulated the rate at which learning proceeded for the novel sequence during the training phase.

One notable similarity between the two sequences was that they share similar third-order contingencies (e.g. in both sequences, if stimulus 1 occurs on trial N, stimulus 6 often occurs on trial N+3), which therefore provides a means by which knowledge of the first sequence could have been directly applied to performance of the second sequence.

However, it seems unlikely that the transfer performance observed in the sequenced pretrained condition was influenced by this similarity between the two structures, since research has shown that the learning of such third-order contingencies requires a period of training far beyond that provided by the pretraining phase of the current experiment (e.g. Cleeremans & McClelland, 1991; Remillard & Clarke, 2001).

The between subjects manipulation in Experiment 2 is somewhat akin to the phenomenon of learned irrelevance (Baker & Mackintosh, 1977), in which retarded conditioning is observed for a CS that has received uncorrelated pairings with a US in an initial preexposure phase. One way of analysing the transitions of a sequence learning procedure is to take each successive pairing of elements as a solitary CS-US pairing (see Experiment 3 for a detailed discussion of this method). Therefore, the pretraining phase in the random pretraining condition can be seen as somewhat akin to uncorrelated preexposure of all of the possible stimulus pairings.

As discussed in Chapter 1, learned irrelevance was a key finding that led to models incorporating changes in cue-processing. For example, Mackintosh (1975) provided an explanation of learned irrelevance by assuming that the experimental context plays the role of an ever-present competing stimulus: on some trials it is presented alone and on some trials it forms a compound stimulus with the preexposed stimulus. Since the preexposed stimulus is no better a predictor of the likelihood of reinforcement than the context, the associability of the preexposed stimulus will decline over the course of preexposure. A similar argument can be put forward as an explanation of the current results. In group unsequenced pretrained, each stimulus position was trained initially as no better a predictor of the location of the target on the next trial, than was the experimental context. Thus, in accordance with Mackintosh's explanation of learned irrelevance, we would expect an attenuation in the subsequent rate of learning in this group.

It is worth noting that although the rate of learning in the sequenced pretrained condition eventually exceeded that seen in the unsequenced pretrained condition, learning was, at least initially, significantly worse in the former condition. One possible explanation

of this initial attenuation in group sequence pretrained is that for this condition, the change in the task contingencies at the start of the training phase would have been particularly disruptive. As a result, it is possible that for this group the pretraining sequence needed to be extinguished from memory, before the learning of the training sequence could commence. In some sense then, these findings are still compatible with an account based on changes in cue-associability: following an initial period in which the pretraining associations were extinguished, the high associabilities of the stimuli in group sequenced pretrained eventually led to a greater rate of learning.

Experiment 3

The purpose of Experiment 3 was to provide a replication of the effects shown in Experiment 2. In Experiment 2, sequence training involved sequences made up of second-order conditional transitions. As such, each stimulus position was not itself a good predictor of the location of the target on the next trial, rather the combination of the previous two positions predicted the next target location. Although one can imagine how training with second-order sequences would lead to the associabilities of the stimulus positions being high (relative to the unsequenced control condition in which each position is non-predictive of the location of the target on all of the subsequent trials), a more direct method of manipulating associability would be to use a first-order sequence in which each stimulus position is a good predictor of the target location on the immediately subsequent trial. In this sense, one is able to draw a comparison between first-order sequences and cue-outcome training in human contingency learning (HCL) tasks. For example, the sequence 1-3-5 can be thought of as containing two cue-outcome pairings: cue 1 is followed by outcome 3, and

cue 3 is followed by outcome 5. Each stimulus in a sequence therefore acts as the outcome for one transition and the cue for the next transition.

The first-order sequences used in Experiment 3 were generated in a ‘probabilistic’ fashion: each position was followed by one other position with a high probability, and also by two other positions, each with a low probability (see Method). The use of probabilistic sequences achieves two things. Firstly, sequence learning will proceed more rapidly for first-order contingencies than for second-order contingencies (e.g. Cleeremans & McClelland, 1991; Remillard & Clarke, 2001). One can imagine that an associability effect on sequence learning might be somewhat short-lived; any influence associability has on new learning will cease once learning reaches an asymptotic level. Therefore, a procedure in which sequence learning proceeds too quickly might fail to detect an associability effect. The use of a probabilistic method of sequence generation will weaken the first-order contingencies in the task and therefore reduce the rate of learning in the task. Secondly, as was the case in Experiment 2, the difference in performance on sequenced (high-probability) and unsequenced (low-probability) trials will enable a continuous measure of sequence learning.

Method

Participants, apparatus and display

Twenty Cardiff University undergraduate students took part in the experiment either for course credit or for payment of £4. Ten participants were assigned to the ‘sequenced pretrained’ condition and ten to the ‘random pretrained’ condition. The apparatus was

identical to that used in Experiment 2, except only four stimulus positions were used which mapped to the keys C, V, B and N.

Sequence generation

Experiment 3 used a 4-choice SRT task. Sequence training involved exposure to first-order probabilistic sequences; the contingencies between the four stimulus positions are shown in Table 2.2. The location in which the target appeared on a given trial (trial N+1) was determined by the position in which it had appeared on the previous trial (trial N). During the pretraining phase, each stimulus position is followed by one other position with a probability of .75, and the other two positions with a probability of .125 each (immediate repetitions never occurred). Consider the sequence 124. The first transition is 1-2, which occurs with a probability of .75. The extent to which participants have learnt that position 1 is followed by position 2 will be shown in the RT and accuracy to position 2. This response will contribute to the average for 'sequenced' trials. The next transition is 2-4, and since position 4 has a low probability of occurring after position 2, this response will contribute to the 'unsequenced' trial type. As in Experiment 2, the amount of learning about this probabilistic sequence can be assessed by comparing performance on sequenced (high probability) and unsequenced (low probability) trials. Participants in the random pretrained condition were exposed to pseudo-random transitions - repetitions were not allowed, and runs of alternations were restricted to a maximum length of four items (i.e. 1212 could

occur, but 12121 could not). Due to a programming error, this limit on alternations was not enforced for blocks of sequenced trials⁶.

Following two blocks of pretraining all participants received a further four blocks of sequenced training. The training sequence was adapted from the pretraining sequence, such that of the two low probability outcomes following a cue in the pretraining phase, one of these outcomes became a high probability outcome for that cue in the training phase. The other low probability outcome from the pretraining phase remained a low probability outcome for that cue during the training phase, whilst the original high probability outcome from the pretraining phase did not occur during the training phase. As in the pretraining phase, learning in the training phase can be assessed by comparing performance on high and low probability outcomes. Each sequenced block started with six random trials, created in the same way as for the pseudo-random blocks. The instructions and procedure were identical to those used in Experiment 2, except that the training phase was extended to four blocks and the task was described as a 4-choice task using the keys C, V, B and N.

⁶ Whilst this error is unfortunate, for participants in the sequenced pretrained condition the average number of alternation runs greater than length 5 (2.6 per block) or length 6 (1.6 per block), is fairly small. Therefore this potential difference in the length of permitted alternations during pretraining is unlikely to have had a significant influence on training performance.

		Trial N (cue) – Pretraining				Trial N (cue) – Training			
		1	2	3	4	1	2	3	4
Trial N+1 (outcome)	1		.125	.125	.75		.75	.25	
	2	.75		.125	.125			.75	.25
	3	.125	.75		.125	.25			.75
	4	.125	.125	.75		.75	.25		

Table 2.2. Conditional probabilities for the two probabilistic sequences (pretraining and training) used in Experiment 3. For clarity, those transitions which have a zero probability have been left blank.

Results

The first six trials of each block were not analysed. Trials following an error were not analysed: given the use of first order-conditional sequences we would expect only consecutive correct responses to truly reflect a participant's learned response for each transition. RTs under 100ms and over 1000ms were not analysed, nor were RTs on inaccurate trials. Figure 2.5 shows the RT data for sequenced pretrained (top-left panel) and unsequenced pretrained (top-right panel) conditions. An ANOVA on the RT data for the sequenced pretrained condition, with within-subject factors of block and trial-type (sequenced vs. unsequenced), revealed a significant effect of trial-type, $F(1, 9) = 16.31, p <$

.01, but no main effect of block, $F < 1$. The interaction between trial-type and block was significant, $F(1, 9) = 13.43$, $p < .01$, indicating that sequence learning was greater in Block 2 than in Block 1. There was no difference between RTs in blocks 1 and 2 for the random pretrained condition, $t(9) = 1.05$, $p = .32$. As in Experiment 2, in order to evaluate any differences in RT between the two conditions, RTs from the sequenced pretrained (the average of RTs on sequenced and unsequenced trials) and unsequenced pretrained conditions were analysed by means of a mixed ANOVA. There was no within-subject effect of block, $F(1, 18) = 1.06$, $p = .32$, nor a between-subject effect of condition, $F(1, 18) = 1.81$, $p = .20$, and no interaction between condition and block, $F < 1$. Thus, we can tentatively conclude that there was no significant difference in overall RT between the conditions during the pretraining phase, although the non-significant trend suggests that there was faster responding in the unsequenced pretrained condition.




Figure 2.5. RT data for pretraining (blocks 1-2) and training (blocks 3-5) phases in Experiment 3. ‘Seq. 1’ and ‘Seq. 2’ refer to sequenced transitions, whilst ‘Unseq. 1’ and ‘Unseq. 2’ refer to unsequenced transitions. *Top-Left panel*: Data for participants pretrained with sequenced transitions. *Top-Right panel*: Data for participants pretrained with unsequenced (pseudo-random) transitions. *Bottom panel*: The training phase data expressed as learning scores – RT on unsequenced trials minus RT on sequenced trials.

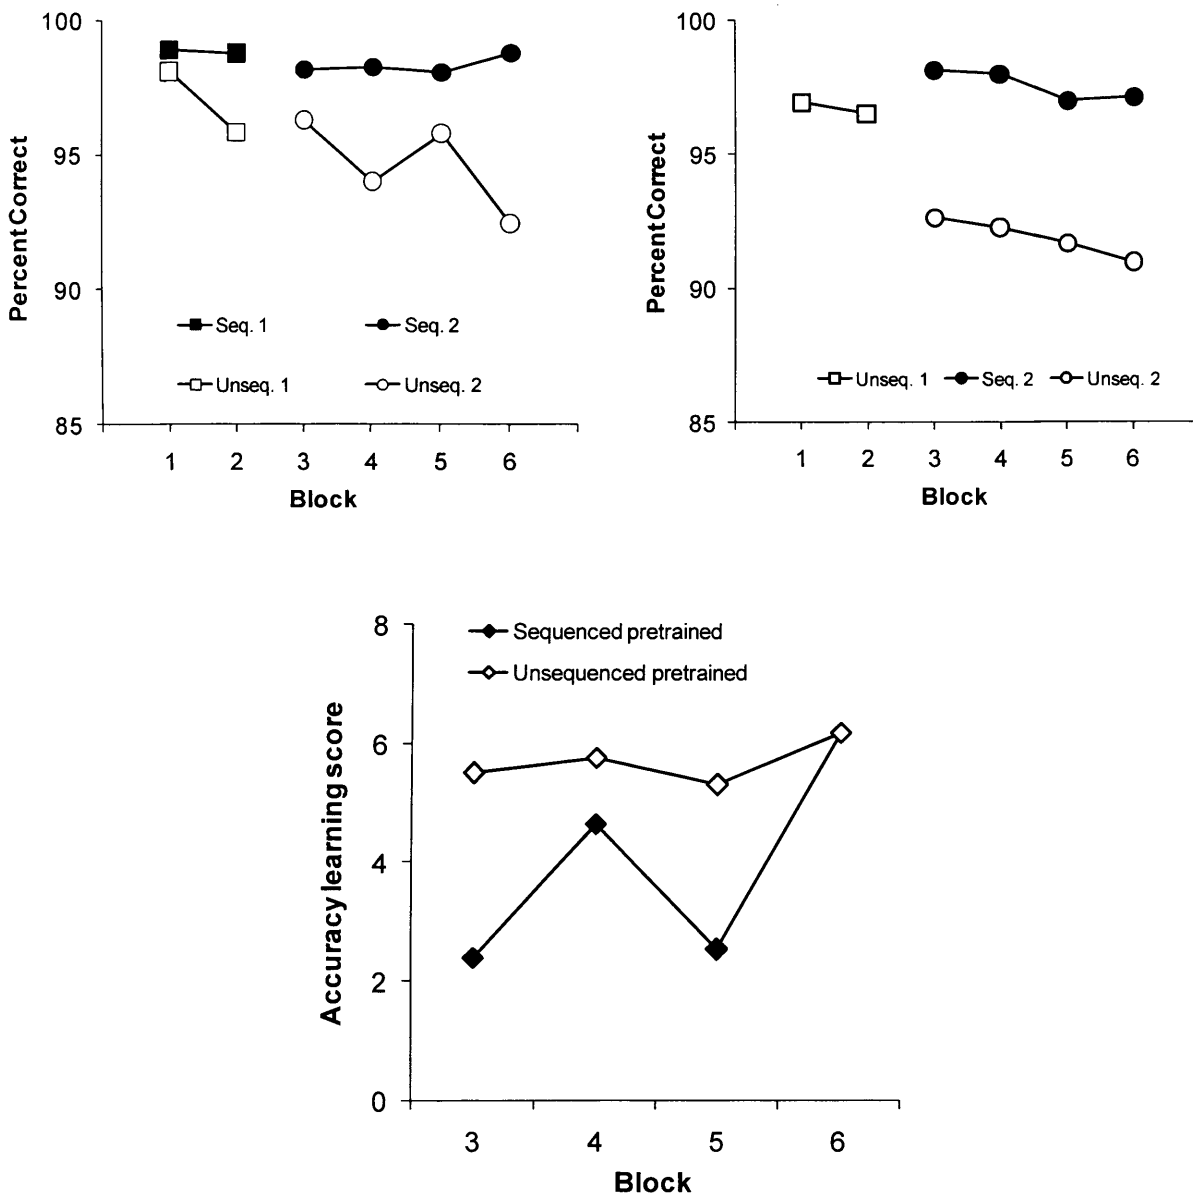


Figure 2.6. Accuracy data for pretraining (blocks 1-2) and training (blocks 3-5) phases in Experiment 3. ‘Seq. 1’ and ‘Seq. 2’ refer to sequenced transitions, whilst ‘Unseq. 1’ and ‘Unseq. 2’ refer to unsequenced transitions. *Top-Left panel*: Data for participants pretrained with sequenced transitions. *Top-Right panel*: Data for participants pretrained with unsequenced (pseudo-random) transitions. *Bottom panel*: The training phase data expressed as learning scores – accuracy on sequenced trials minus accuracy on unsequenced trials.

The accuracy data are shown in Figure 2.6. The pretraining data from sequenced pretrained participants were subjected to an ANOVA with factors of trial-type and block. The main effect of trial-type approached significance, $F(1, 9) = 4.41, p = .065$, which suggests participants were more accurate on sequenced than unsequenced trials. There was no main effect of block, $F(1, 9) = 2.12, p = .18$, and the interaction between block and trial-type was also not significant, $F(1, 9) = 1.70, p = .23$. As in the RT data, a between-subjects comparison was made between accuracy in the sequenced pretrained and unsequenced pretrained conditions. There was no within-subject effect of block, $F(1, 18) = 2.62, p = .12$, no between-subject effect of condition, $F < 1$, and no interaction between condition and block, $F < 1$.

It is clear from the RT data in Figure 2.5 that rapid sequence learning occurred in both conditions during the training phase (blocks 3 to 6). These data were subjected to a mixed ANOVA, with a between subject factor of condition, and within-subjects factors of trial-type and block. This revealed a main effect of trial-type $F(1, 54) = 75.98, p < .001$, which indicates strong sequence learning during the training phase. There was no main effect of block, $F(3, 54) = 1.06, p = .38$. While the main effect of condition did not reach a conventional level of significance, $F(1, 18) = 3.05, p = .098$, there was a trend towards faster responding in the unsequenced pretrained condition. There was no interaction between trial-type and condition, $F < 1$, but the interaction between trial-type and block was significant, $F(3, 54) = 3.68, p < .05$, which indicates that sequence learning emerged across the training phase. There was no interaction between block and condition, $F(3, 54) = 1.88, p = 1.44$, nor a three-way interaction, $F < 1$. The bottom panel of Figure 2.5 shows the RT data expressed as learning scores (RTs on unsequenced trials minus RTs on sequenced trials) for both the sequenced pretrained and unsequenced pretrained conditions. Pre-planned independent

samples t-tests revealed that there was no difference in the learning scores on any block, all $t_s < 1$.

The accuracy data from the training phase are shown in Figure 2.6. These data were subjected to a mixed ANOVA with trial-type and block as within-subjects factors and condition as a between-subjects factor. This revealed a significant effect of trial-type, $F(1, 54) = 23.41, p < .001$, which suggests participants were more accurate on sequenced than unsequenced trials. There was no main effect of block, $F(3, 54) = 1.23, p = .31$, nor of condition, $F(1, 18) = 1.57, p = .23$, and none of the interaction effects were significant, all $F_s < 1$. The bottom panel of Figure 2.6 shows the accuracy data expressed as learning scores (accuracy on sequenced trials minus accuracy on unsequenced trials). There was no difference between the learning scores in the two conditions in any block, all $t_s(18) < 1.22, p_s > .26$.

Discussion

The results of Experiment 3 are inconsistent with those of Experiment 2. Although sequence learning progressed rapidly during the training phase, the rate of learning was equivalent in both conditions. There were two important differences between the current procedure and that employed in Experiment 2, which may have resulted in the different pattern of data observed. Firstly, the six-choice task used in Experiment 2 was reduced to a four-choice task in Experiment 3. The use of fewer target locations will lead to less response competition and therefore faster, more accurate responses. This, in turn, may lead to floor effects in RTs and ceiling effects in accuracy. Secondly, second-order conditional sequences were used in Experiment 2, whilst first-order sequences were used in the current

experiment. Learning of first-order contingencies will occur at a faster rate than second-order contingencies, as can be seen by comparing the data in Figures 2.3 and 2.5. As a result of this rapid learning, we might expect any advantage for group sequenced pretrained to be relatively short-lived in Experiment 3, as learning for this group would rapidly reach ceiling, allowing the unsequenced pretrained group to catch up. Given there is no evidence for an advantage for group sequenced pretrained even in the initial blocks of the training phase, this account would have to assume that learning is already approaching ceiling by the end of the very first training block. The finding of a significant interaction between block and trial-type during the training phase, however, would seem to rule this interpretation out.

It is worth noting that in both Experiment 2 and 3 some unexpected between-subject differences occurred. For example, the RT data from Experiment 3 (shown in Figure 2.5) suggested a general difference (albeit not quite significant) between the two conditions during the training phase. Although it is difficult to interpret the direct impact of these differences on the measure of learning during the training phase, the fact that participants in one condition seem to be making generally faster response means any direct comparison between the two conditions should be made with caution. Similarly, a between-subject difference was also observed in the accuracy data from the pretraining phase of Experiment 2. As such, the data from Experiments 2 and 3 illustrate the magnitude of the between-subject variability in both behavioural measures. Possibly this difference was driven by the different pretraining procedures used, or alternatively it may simply reflect natural variation between the (relatively) small samples used. Furthermore, it is difficult to discount the possible influence of motivational differences in the two groups. Although pretraining with random material may reduce the associabilities of the cues within the task, it is also possible that the different pretraining procedures might result in motivational differences in the two

conditions, in turn resulting in an attenuation in the rate of learning in this condition. For example the less engaging pretraining phase in the unsequenced pretrained condition might lead to a reduced level of motivation in this group, as compared to participants in the sequenced pretrained condition. Alternatively, in group sequenced pretrained – for whom the learned contingencies from the pretraining phase are a hindrance to performance in the training phase – the transfer between stages may well result in a short period of frustration. Indeed, evidence for an initial attenuation in the learning of the training phase contingencies was evident in the RT data for Experiment 2 (see Figure 2.3)

Overall, the data from Experiments 1 and 2 provided some preliminary support for changes in stimulus associability in the AGL and SRT procedures. However, the data presented in Experiment 3 in particular, suggest that a between-subjects approach to examining associability effects in sequence learning will require large sample sizes in order to minimise the noise caused by the large amount of between-subject variability, rendering this approach inefficient.

Chapter 3 – Reducing between-subject confounds: evidence for changes in cue-associability in within-subject designs.

Rationale

As was highlighted at the end of Chapter 2, any evidence for changes in cue-associability provided by a between-subjects design is also open to a variety of alternative interpretations, particularly in relation to the motivational state of the participants. For instance, training with random material in the AGL or SRT tasks may well lead to a state akin to ‘learned helplessness’. Learned helplessness describes an emotional and motivational deficit that arises as a result of exposure to uncontrollable events, often leading to an inability to learn contingent events in the future (Hiroto & Seligman, 1975). Whilst learned helplessness in human cognition is more commonly discussed with respect to tasks in which hypothesis testing is encouraged, the sorts of cognitive processes disrupted by training with random, non-contingent events are likely to be quite extensive (see Sedek & Kofta, 1990). Furthermore, it is apparent from the data presented in Chapter 2 (especially that in Experiment 3), that there is considerable between-subject variability in the dependent measures. Both of these factors, can however, be eliminated by examining changes in cue-associability on a within-subject basis.

Experiment 4

The objective of Experiment 4 was to manipulate the associability of the cues within the SRT task, such that some cues were pretrained as good predictors of outcomes, whilst others

were simultaneously pretrained as poor predictors of outcomes. This within-subject manipulation will allow a direct comparison of responses to outcomes following these different types of cue, therefore allowing for an examination of the effect of these different cue pretraining procedures. One potential problem with this manipulation is that, as a result of the differential pretraining of good and poor predictor cues, any measure of new learning about these cues will be affected by different amounts of response interference. Consider the case in which participants learn that cue 1 consistently predicts outcome 3 during pretraining, whilst cue 2 predicts outcomes 3 and 4 with equal probability. During this pretraining, participants will presumably build up a strong association between cue 1 and the next appropriate response, 3, whilst the associations between cue 2 and outcome 3, and cue 2 and outcome 4, will be somewhat weaker. In the training phase, in which each cue is paired with a novel outcome, each of the pretraining associations will impair responding to these new outcomes: the tendency to make the (now-inappropriate) pretraining response will lead to a slower RT to the current training response. Consequently, if this tendency to make the pretraining response is greater for cues that were previously good predictors than those that were previously poor predictors, there will be a greater slowing of responses for good predictors than poor predictors. Hence this influence of response interference could feasibly mask any advantage for good predictors that derives from differences in associability⁷.

⁷ The issue of interference discussed here is of course applicable to any sequence learning experiment in which the contingencies are changed following a transfer from one learning phase to another. As such, this issue is also relevant to the designs of Experiments 2 and 3. However, for reasons discussed in the introduction to Experiment 4, the issue of response interference does not alter the conclusions initially drawn from these experiments.

In order to try to minimize the effect of response interference on the measure of sequence learning, in Experiment 4 pretraining was conducted with dedicated pretraining outcomes and responses. Following pretraining, these dedicated pretraining outcomes were removed from the task, in order to reduce participants' tendency to make the corresponding responses.

Method

Participants, apparatus and display

Fifteen participants were paid £4 for participation. The apparatus was identical to that used in Experiment 2 (a six-choice SRT) with the addition of a dedicated pretraining outcome stimulus and response key used during the pretraining phase. A 5-pointed star and an equilateral triangle were used as these pretraining outcome stimuli. Both were grey, and were both approximately 3cm wide and 3cm high. These stimuli were presented in the middle of the screen, 2.5cm above the row of standard SRT positions.

Sequence generation

As in Experiment 3, the position of the target stimulus on any given trial was determined by its location on the previous trial. The contingencies for the pretraining and training phases are shown in Table 3.1. During the pretraining phase, following each response to the target in one of the six standard SRT positions, one of the dedicated pretraining outcomes appeared. Three cues were pretrained as good predictors of the star appearing on the next trial: following the target appearing in one of these locations the star occurred on the subsequent trial with a probability of .9. The remaining cues were equally

predictive of the star and triangle appearing on the subsequent trial. Therefore, participants should learn to anticipate the star appearing after a good predictor cue, but not after a poor predictor cue, and this should be borne out in the data collected to responses to the dedicated pretraining outcomes during this phase. Following either a star or triangle occurring, a cue stimulus always appeared and there was an equal probability of each cue stimulus appearing.

		Trial N (cue) – Pretraining								Trial N (cue) – Training						
		1	2	3	4	5	6	S	T	1	2	3	4	5	6	
Trial N+1 (outcome)	1							.166	.166		.5					.5
	2							.166	.166			.5	.5			
	3							.166	.166	.5					.5	
	4							.166	.166			.5				.5
	5							.166	.166	.5			.5			
	6							.166	.166		.5				.5	
	S	.9	.9	.9	.5	.5	.5									
	T	.1	.1	.1	.5	.5	.5									

Table 3.1. Sequence contingencies between cues during pretraining and training phases in Experiment 4. Numbers 1-6 refer to cues in the stimulus array, although the mapping of numbers in this table to stimulus positions was randomized for each participant. S and T refer to a star shape and a triangle shape, respectively, appearing in the dedicated pretraining outcome position (see text). These pretraining outcome stimuli did not appear in the training phase. Blank cells indicate zero probability.

During the training phase each cue was paired consistently with two outcomes (see the right-hand panel of Table 3.1). This aspect of the design reflects the fact that when low probability contingencies were used in Experiment 2 an associability effect was observed, whilst when high probability contingencies were used in Experiment 3 no corresponding effect was observed. As such, training contingencies were used which were low probability outcomes: each cue predicted two outcomes with an equal probability of .5.

Procedure

The procedure was amended from that described in Experiment 2, such that during the pretraining phase, on every other trial a star or triangle shape appeared above the stimulus array, centred horizontally on the screen. Additional instructions were added to those used in Experiment 2 to explain the presentation of the neutral stimuli:

“For the first half of the experiment, in addition to following the grey circle, on every other trial either a star or a triangle will appear on the screen. These stimuli will appear in the same position, in the centre of the screen. When a star appears we want you to press the spacebar with your thumb. When a triangle appears we want you to not press the spacebar. If you incorrectly press the spacebar when the triangle is on the screen you will get a beep in the headphones. Similarly, if you do not press the spacebar within one second of the star appearing on the screen you will hear a beep.”

Participants were only able to respond with the spacebar when the neutral stimulus was presented (i.e. the keys X-M were disabled), and similarly were only able to use the

keys X-M when the standard cue appeared (i.e. the spacebar was disabled). Pretraining consisted of 4 blocks of 150 trials. Since the inclusion of the neutral stimulus lengthened the time taken to complete a pretraining block, in order to keep training blocks to a similar length, 200 trials were used in the 4 blocks of training. Participants were given a rest break of 20 seconds at the end of each block. At the end of the 4th block (the end of the pretraining phase), the following message appeared during the rest break: “Note: From now on the star/triangle will no longer appear. Just respond to the 6 positions using keys x to m.”

Results

The first 5 trials of each block were not analysed. Trials were excluded on the same basis as in Experiment 3. The responses of interest during the pretraining phase were reaction times to the star, and error data for both star and triangle trials (no RT data are available for the triangle stimulus since participants were specifically instructed not to respond to this stimulus). RT data for the star is shown in the top panel of Figure 3.1, which shows that on average RTs are (numerically) shorter to the star following poor predictor cues than good predictor cues, in each block. However, an ANOVA with cue (good versus poor predictor cue) and block as within-subject factors yielded no significant main effects or interactions, all $F_s < 1$.

The accuracy data for the star and triangle trials following both good and poor predictor cues are shown in the bottom panel of Figure 3.1. Overall accuracy was greater to the star stimulus, presumably because this stimulus occurred with a greater overall probability than did the triangle: the probability of the star following a good predictor was .9, and following a poor predictor was .5, giving an average probability of the star appearing

of .7. These data were subjected to ANOVA with outcome (star versus triangle), cue (good versus poor) and block as within-subject factors. This revealed a significant effect of outcome, $F(1, 14) = 15.29, p < .01$, which indicates fewer errors were made on star outcomes than on triangle outcomes. The main effect of cue failed to reach significance, $F(1, 14) = 3.55, p = .080$, and there was also no main effect of block, $F(3, 42) = 2.08, p = .12$. None of the two-way interaction effects were significant: cue by outcome, $F(1, 14) = 2.24, p = .16$; cue by block, $F(3, 42) = 1.63, p = .20$, outcome by block, $F(3, 42) = 2.12, p = .11$. The three-way interaction was marginally significant, $F(3, 42) = 2.68, p = .059$. Inspection of Figure 3.1 suggests that this latter interaction effect is likely to be due to the factors of cue and block interacting in different ways for star and triangle outcomes: accuracy following poor predictors is fairly constant across pretraining for both star and triangle outcomes, whereas for good predictors, accuracy to the star increases across pretraining, whilst accuracy to the triangle decreases dramatically during the first two blocks, before increasing across the last two blocks of the pretraining phase.

As pretraining progresses, we would expect participants to become more sensitive to the contingencies between the cues and the neutral stimuli. This seems to be the case for responses to the triangle stimulus (see the bottom panel of Figure 3.1), for which responses following a good predictor are less accurate than those following a poor predictor cue across the last three blocks of pretraining. In order to assess learning across the latter part of the pretraining phase more directly, a three-way ANOVA was conducted on the final three blocks of the pretraining phase only. This revealed a marginally significant effect of cue, $F(1, 14) = 4.45, p = .054$, a main effect of outcome, $F(1, 14) = 13.43, p < .01$, but no main effect of block $F(2, 28) = 2.35, p = .11$. There was no interaction between cue and block, $F(2, 28) = 1.14, p = .33$, nor between outcome and block, $F(2, 28) = 1.41, p = .26$. The

interaction between cue and outcome did, however, reach significance, $F(1, 14) = 4.94$, $p < .05$, which is likely to reflect the different pattern of errors to the triangle stimulus following good and poor predictor cues. The three-way interaction was not significant, $F(2, 28) = 1.01$, $p = .38$.






Figure 3.1. Data from the pretraining phase of Experiment 4. *Top panel:* Reaction time (in milliseconds) to the star outcome following cues that were good predictors of the star occurring (GP) and poor predictors of either outcome (PP). *Bottom panel:* Accuracy data to both star and triangle outcomes following GP and PP cues.

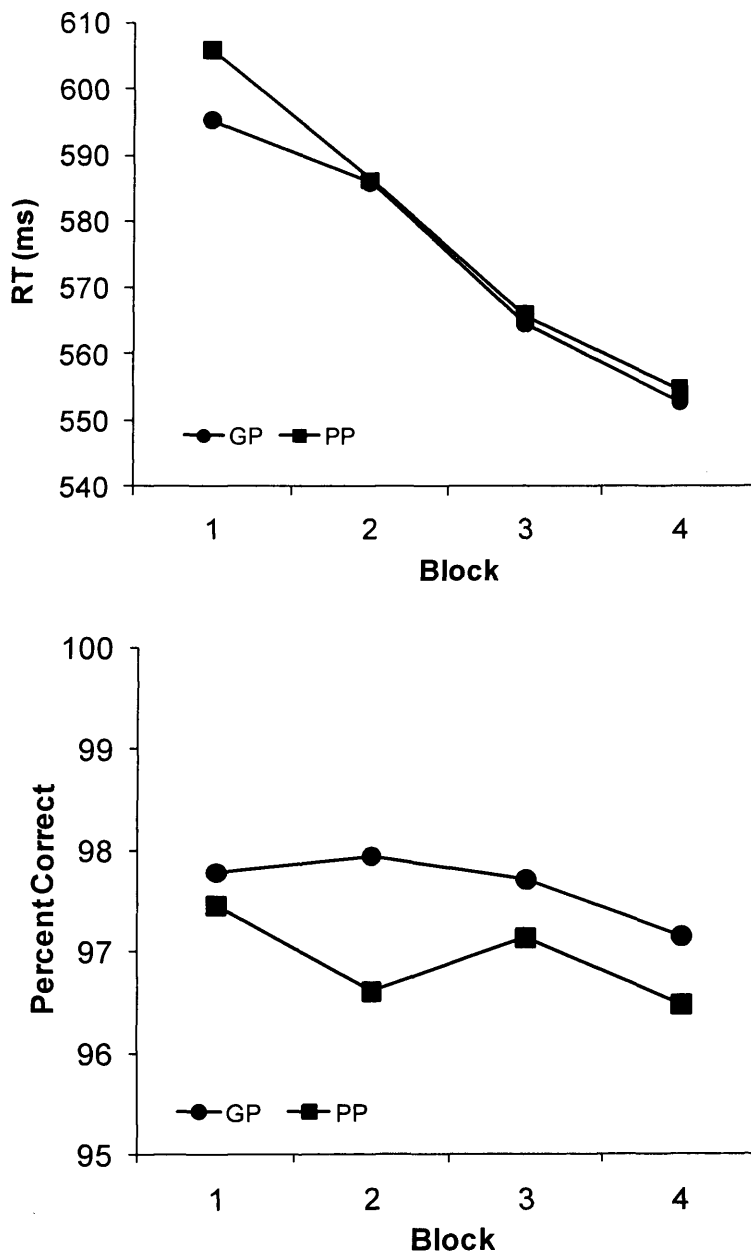


Figure 3.2. Data from the training phase of Experiment 4. *Top panel:* Reaction time to outcomes following cues that were pretrained as good predictors (GP) and poor predictors (PP). *Bottom panel:* Accuracy to outcomes following GP and PP cues.

The RT data for the training phase are shown in the top panel of Figure 3.2. The extent to which participants have learnt about a given transition will be shown in the reaction time and accuracy to the outcome stimulus of that transition. For example, the extent to which participants have learnt the sequenced transition 1-3 can be assessed by measuring the RT and accuracy to position 3 when it follows stimulus 1. Similarly the RT and accuracy to position 3 when it follows stimulus 5 will provide a measure of how much participants have learned about the sequenced transition 5-3. Although all cues are equally predictive of outcomes during the training phase (see Table 3.1), I will continue to use the terms ‘good’ and ‘poor’ predictor to refer to how these cues were pretrained during the first phase. It is clear that RTs to new outcomes following good predictors were very similar to RTs to new outcomes following poor predictors during the training phase. These data were subjected to an ANOVA with prior predictiveness (whether the cue was a good or poor predictor of the neutral outcomes in the pretraining phase) and block as within-subject factors. There was a main effect of block $F(3, 42) = 19.80, p < .001$, but no main effect of prior predictiveness, $F < 1$, nor was there an interaction between block and prior predictiveness, $F(3, 42) = 1.09, p = .36$. In order to evaluate whether there was a short lived effect of prior predictiveness at the start of the training phase a pre-planned related-samples t-test was conducted on the data from the first block only. This indicated that responding was significantly faster to outcomes following cues pretrained as good predictors than cues pretrained as poor predictors during the first block of training, $t(14) = 2.42, p < .05$.

The accuracy data during the training phase are shown in the bottom panel of Figure 3.2. These data were subjected to an ANOVA with factors of prior predictiveness and block. While accuracy for good predictors was numerically greater than that for poor predictors over all blocks, the main effect of prior predictiveness failed to reach significance, $F(1, 14)$

= 2.48, $p = .14$. There was also no main effect of block, $F(3, 42) = 1.17$, $p = .33$, and no interaction between cue and block, $F < 1$. Paired t-tests revealed there was no effect of prior predictiveness in any block, all $t_s(14) < 1.30$, all $p_s > .21$.

Discussion

The data from the pretraining phase showed that only accuracy was useful in determining the extent to which participants had learnt that good predictor cues selectively predicted the occurrence of the star stimulus. The RT data showed no significant difference between responses to the star following good and poor predictor cues. Since the star outcome was overall more likely to occur than the triangle outcome, irrespective of which cue preceded it, it is likely there was a general level of expectancy for this stimulus to occur. It is also likely that the use of easily discriminable stimuli (a star and a triangle) for the dedicated pretraining outcomes would have led to rapid identification of the stimulus on each trial. Furthermore, responding to the star required a spacebar response with the thumb, which is a common response for anyone with a basic typing ability and therefore may be a well practised response in many participants. Indeed, the data suggest that this response may have reached an optimum level early on in training: RTs to the star during the first block were not significantly different from any others, all $t_s(14) < 1.30$, $p_s > .21$. There was also no difference in the accuracy of responding to the star following good and poor predictor cues, which is perhaps unsurprising, since these errors reflect a failure to respond to the star within 1000ms and hence are generally very infrequent. The final measure of learning during the pretraining phase was errors made to the triangle stimulus: an inappropriate response with the spacebar. With this measure, at least in the last 3 blocks of pretraining,

more incorrect responses were made following cues that were good predictors of the star than after poor predictors – a difference we would expect if participants had learnt the consistent contingency between the good predictor cues and the star stimulus.

The RT data from the training phase demonstrate a short lived effect of prior predictiveness: faster responses were initiated to outcomes following good predictor cues than to outcomes following poor predictor cues during the first block of training. This RT benefit is unlikely to be due to a speed-accuracy trade-off, since there was no statistical difference in accuracy between outcomes following good and poor predictor cues (with accuracy, if anything, greater for good predictors than for poor predictors). The data therefore suggest that cue-outcome learning proceeded faster for associations pertaining to cues that were pretrained as good predictors, in comparison to those pretrained as poor predictors.

To conclude, the pretraining data provide weak evidence that the procedure used had been successful in manipulating cue-predictiveness. Although it might be argued that the associability effect shown in the training phase (RT data) could only have resulted from the differential pretraining procedures used, it would be more satisfactory to replicate the training effect with stronger evidence of differential learning from the pretraining phase. Therefore before concluding that cue-associability modulated sequence learning in Experiment 3, the results should be replicated in unison with more robust pretraining results.

Experiment 5

In Experiment 4 only one dedicated pretraining response was used in the pretraining phase: participants responded to the star and inhibited this response when the triangle was presented. I suggested that, amongst other factors, the high probability of the star outcome occurring led to a ceiling effect in accuracy and a floor effect in RTs for this stimulus. The pretraining procedure in Experiment 5 avoids this asymmetry by using two dedicated pretraining outcomes. One cue was trained as a good predictor of one of these outcomes, whilst another cue was pretrained as a good predictor of the other outcome. Two more cues were pretrained as poor predictors of these dedicated pretraining outcomes: they were each as likely to be followed by one outcome as another. By using two dedicated pretraining outcomes, that both require responses (rather than one requiring a response and one requiring the inhibition of that same response), it was hoped that the effectiveness of all the dependent measures would increase for the pretraining phase.

Although it would have been possible to amend the procedure used in Experiment 4 to include the pretraining procedure described above, it was decided that this design could also be implemented in a standard SRT task. Thus, during the pretraining phase a 6-choice SRT task was used, with the two outermost stimuli acting as the dedicated pretraining stimuli. The four centre cues were divided into two good predictor cues and two poor predictor cues (see Method). Whilst in Experiment 4 the two dedicated pretraining outcomes (the star and the triangle) were particularly salient stimuli, the standard SRT procedure used in the current experiment has the advantage of ensuring that the dedicated pretraining outcomes are no more salient than the other task cues. As a result, it seems less likely that participants will actively search for the contingencies between the task cues and the

dedicated outcome stimuli in the current procedure, than in the procedure used in Experiment 4.

Although I argued in Experiment 4 that the use of dedicated pretraining outcomes should reduce the amount of response interference affecting training performance, it is possible that despite the explicit instruction that these stimuli would no longer appear during the training phase, associations between cues and the neutral outcomes interfered with responding during the training phase. It is unlikely that the association between the good predictor cues and the spacebar response would have had anything but an inhibitory effect on responding during the training phase. However, it is possible that the poor predictor cues had a strong inhibitory association with the spacebar response, and that this association could have generalised to inhibit responses in general during the training phase. This account would therefore predict longer RTs to outcomes following cues pretrained as poor predictors than to those pretrained as good predictor cues – the result found in Experiment 4 (see Figure 3.2).

Although it is difficult to imagine a situation in which a pretraining association will not have at least a short-lived impact on training performance in the SRT task, it is possible to accommodate this response interference, such that the measure of learning during the training phase is unconfounded from response interference. One method of eliminating the impact of response interference is to compare responses for predicted sequenced outcomes to those of unsequenced outcomes. Since the response interference generated by any previously acquired pretraining contingency will affect responses to new sequenced and unsequenced outcomes to the same degree (during the training phase), the difference between these two responses provides a measure of learning that is unconfounded from

response interference (this issue is discussed in more detail in the context of the relevant data in the Results section).

As a result of eliminating the impact of response interference, it is also feasible to extend the pretraining phase of the current design beyond that used in Experiment 4 (which did not eliminate response interference entirely). It was hoped that extending the pretraining phase would allow more time for the associabilities of the task cues to differentiate before participants are transferred to the training phase.

Method

Participants, Apparatus and Stimuli

Twenty Cardiff University undergraduates, who had not participated in an SRT task before, participated for course credit or payment. The stimuli and apparatus were identical to those used in Experiment 2. In the pretraining phase a standard 6-choice SRT task was used, whilst in the training phase a 4-choice SRT task was used (see below).

Sequence generation

The left-hand side of Table 3.2 shows the conditional probabilities of the sequence used during the pretraining phase. As in Experiment 3, the location in which the target appeared on a given trial (trial N+1) was determined by the position in which it had appeared on the previous trial (trial N). Position 2 was a good predictor of the target appearing in position 1 on the next trial, whilst position 3 was a good predictor of the target appearing in position 6. When the target appeared in position 4 or 5, it would appear in

positions 1 and 6 on the next trial with equal probability of .5. Thus, positions 4 and 5 were relatively poor predictors of the location of the target on the next trial. After an outcome trial (positions 1 & 6) the location of the next target was selected at random from the four cue positions (2-5), with the caveat that the sequence could feature the same pair of locations twice in a row (e.g. 121213), but not three times (e.g. 121212). For the training phase, the two stimulus locations used as outcomes during the pretraining phase (locations 1 & 6) were removed from the stimulus array, creating a four-choice task. The right-hand side of Table 3.2 shows the conditional probabilities for the training sequence. Each of the remaining four stimulus locations acted as both a cue and an outcome during this phase. It is clear from Table 3.2 that the possible training transitions were entirely different from those used in the pretraining phase. During the training phase, all cues predicted one location with a probability of .8, and the two other locations with a probability of .1 each (repetitions were not permitted). All cues were equally valid predictors of outcomes during training. Nevertheless, as in Experiment 4, I will continue to use the terms good and poor to refer to the way in which these sets of cues had been pretrained.

		Trial N (cue) – Pretraining						Trial N (cue) – Training						
		1	2	3	4	5	6	1	2	3	4	5	6	
Trial N+1 (outcome)	1		.9	.1	.5	.5								
	2	.25					.25			.1	.1	.8		
	3	.25					.25		.8		.1	.1		
	4	.25					.25		.1	.8		.1		
	5	.25					.25		.1	.1	.8			
	6		.1	.9	.5	.5								

Table 3.2. Transition contingencies for the sequences used during the pretraining and training phases in Experiment 5.

Numbers 1-6 refer to cues in the stimulus array, although the mapping of numbers in this table to stimulus positions was randomized for each participant. Stimulus positions 1 and 6 did not appear in the training phase. Blank cells indicate zero probability.

Procedure

The procedure was similar to that of Experiment 2. Since the measure of sequence learning (the difference in RT on low and high probability outcome trials) was unconfounded from any response interference generated as a result of pretraining (for a more detailed discussion see the Results section, below), the pretraining phase was extended to 10 blocks of 150 trials each. Following the 10th block the following message was displayed during the rest break: “PLEASE NOTE! From now on the circle will only appear in the middle 4 positions, using keys C, V, B and N.” This was the first time participants were informed that the task would change in this way. The two outermost positions (1 & 6) were removed from the stimulus array for the remainder of the experiment. The training phase comprised 2 blocks of 150 trials.

Results

The first six trials of each block were not analysed. Trials were excluded on the same basis as Experiment 4. During the pretraining phase, three types of trial were of interest: high probability outcomes following good predictors (GPH; occurring with a probability of .9); low probability outcomes following good predictors (GPL; occurring with a probability of .1); and medium probability outcomes following poor predictors (PPM; occurring with a probability of .5). As in the training phase of Experiment 4, the analysed trial was that occurring on trial N+1 (the outcome). For example, consider the sequence 214631 and the conditional probabilities for these transitions in Table 3.2. The first transition, 2-1, is of type GPH since outcome 1 is predicted by cue 2 with a probability of .9. The extent to which participants have learnt that cue 2 predicts outcome 1, will be shown in the RT and accuracy

to outcome 1, and it is therefore this trial (outcome 1) which contributes to the average for trial type GPH. The next transition, 1-4, is from outcome location 1 to cue 4, is not a trial of interest, and is not analysed. The next transition, 4-6, is of trial type PPM, since outcome 6 is predicted by cue 4 with a probability of .5. The response to outcome 6 will therefore contribute to the average for trial type PPM. Transition 6-3 is from outcome location 6 to cue 3, and is not analysed. Finally, transition 3-1 is of type GPL, as outcome 1 is predicted by cue 3 with a probability of .1. The response to outcome 1 will contribute to the average for trial type GPL.

Due to the low probability of an inconsistent outcome occurring after a good predictor cue during the pretraining phase, a participant would occasionally produce no data for the GPL trial-type in a given block. There were 7 instances of missing RT data, and 7 instances of missing accuracy data. In order to conduct a full ANOVA, missing data were estimated using an average from the two blocks immediately before and after the block with missing data. When this was not possible (in Blocks 1 & 10) the nearest adjacent data were used.

The top panel of Figure 3.3 shows RTs across the ten blocks of pretraining for the three trial types. By the end of training, participants were fastest on outcomes that could be predicted with a high probability (GPH), slowest on those that could be predicted with lowest probability (GPL), and are of intermediate speed on those with a medium probability (PPM). These data were subjected to repeated measures ANOVA, with factors of outcome probability (GPH, GPL and PPM) and block. There was a significant effect of outcome probability, $F(2,38) = 13.49, p < .001$, and block, $F(9, 171) = 6.15, p < .001$. The interaction between outcome probability and block was also significant, $F(18, 342) = 1.83, p < .05$,

which suggests that participants' sensitivity to the contingencies within the sequence increased with continued exposure during pretraining. Pairwise comparisons between the three levels of the outcome probability variable revealed significant differences between all three: GPH versus GPL trials, $F(1, 19) = 15.64, p < .01$; PPM versus GPL trials, $F(1, 19) = 16.82, p < .01$; GPH versus PPM trials, $F(1, 19) = 7.86, p < .05$.

Accuracy data for the pretraining phase are shown in the bottom panel of Figure 3.3. The same ordinal relationships that are observed in RTs emerge later on in pretraining: Participants are most accurate on GPH trials, least accurate on GPL trials and show intermediate performance on PPM trials. These analyses were subjected to ANOVA, which revealed no effect of outcome probability, $F < 1$, but a significant effect of block, $F(9, 171) = 2.37, p < .05$, indicating a general decline in accuracy towards the end of the pretraining phase. The interaction was not significant, $F < 1$.




Figure 3.3. Data from the pretraining phase of Experiment 5. *Top panel:* RT data for high-probability outcomes following good cues (GPH), low-probability outcomes following good predictor cues (GPL), and medium-probability outcomes following poor predictor cues (PPM). *Bottom panel:* Accuracy data for the same three trial-types.




Figure 3.4. RT data from the training phase of Experiment 5. *Left panel:* RTs to high- (GPH) and low- (GPL) probability outcomes following cues pretrained as good predictor cues, and RTs to high- (PPH) and low- (PPL) probability outcomes following cues pretrained as poor predictor cues. *Right panel:* The data presented as learning scores - RTs on low probability outcomes minus RTs on high probability outcomes, for cues pretrained as good (GP) and poor (PP) predictors.




Figure 3.5. Accuracy data from the training phase of Experiment 5. *Left panel:* Accuracy to high- (GPH) and low- (GPL) probability outcomes following cues pretrained as good predictor cues, and accuracy to high- (PPH) and low- (PPL) probability outcomes following cues pretrained as poor predictor cues. *Right panel:* The data presented as learning scores - accuracy on high probability outcomes minus accuracy on low probability outcomes, for cues pretrained as good (GP) and poor (PP) predictors.

During the training phase, four trial types were of interest: high- and low-probability outcomes following cues that were pre-trained as good predictors (GPH and GPL, respectively) and high- and low-probability outcomes following cues that were pre-trained as poor predictors (PPH and PPL, respectively). Analysis of the training phase followed a similar method to that outlined above for the pretraining phase, however, during the training phase all trials contributed to a trial type average. Consider the sequence 342 and the contingencies given in the right-hand side of Table 3.2. The first transition, 3-4, is of trial type GPH, as cue 3 was pretrained as a good predictor cue and outcome 4 now occurs with a high probability (.8) after cue 3. The next transition, 4-2, is of trial type PPL, since cue 4 was pre-trained as a poor predictor cue and outcome 2 now occurs with a low probability (.1) after cue 4. As any given element in the sequence acts as both the outcome of the preceding transition and the cue for the next transition, every trial in the sequence counts towards a trial-type average.

RT data for the two training blocks are shown in the left-hand panel of Figure 3.4. Overall, learning progressed rapidly for both good and poor predictor contingencies, as indicated by the difference in RTs to high- and low-probability outcomes. These data were subjected to repeated-measures ANOVA, with factors of prior predictiveness (i.e. the predictiveness of the cues during pretraining; good versus poor), outcome probability (high versus low) and block. This revealed a significant main effect of outcome probability, $F(1,19) = 61.62, p < .001$, a marginally significant effect of block, $F(1,19) = 3.27, p = .086$, but no main effect of prior predictiveness, $F(1,19) = 1.28, p = .27$. The interaction of prior predictiveness and block was significant, $F(1,19) = 5.24, p < .05$, which reflects the different pattern of RTs to low probability outcomes for good and poor predictors. No other interactions were significant: prior predictiveness by outcome probability, $F(1,19) = 1.71, p$

= .21; outcome probability by block, $F(1,19) = 2.43$, $p = .14$; nor the three-way interaction, $F(1,19) = 2.84$, $p = .11$.

In the current design, learning about cue–outcome relationships can be assessed by subtracting RTs to high probability outcomes following a particular cue position from RTs to low probability outcomes following the same cue position. This yields a ‘learning score’ which reflects how much participants have learnt about high probability outcomes relative to low probability outcomes, and provides a means of comparing how much has been learnt about good predictors and poor predictors during the training phase, unconfounded from response interference effects which result from pretraining. Given the differential pretraining of good and poor predictor cues, we might expect proactive interference to have a larger detrimental effect on good predictor contingencies than poor predictor contingencies during the training phase. Consider, as an example, a participant whose mapping of elements 1 to 6 (shown in Table 3.2) directly translates onto the stimulus locations 1 to 6 on the screen (e.g. element 1 maps to key X, element 2 to key C, and so on). Suppose that during pretraining this participant has learnt that cue 2 is very likely to be followed by outcome 1 – that is, this participant has learnt to prepare the response of pressing the X key on the trial following cue 2. During training, all cue–outcome contingencies are changed, such that cue location 2 is now usually followed by location 3. This participant’s previously-learnt tendency to prepare an X response following location 2 is therefore now inappropriate, and might be expected to interfere with appropriate responding on this trial (pressing the V key). This can be compared to previously poor predictors, for which participants will not have developed such strong response tendencies during pretraining, and hence for which there will be less interference during the training phase. Consequently, during training, proactive interference could potentially mask any

advantage that might exist for learning the new high probability responses for those cues previously trained as good predictors. However, any proactive interference affecting responding during the training phase will have an equivalent effect on responses made to both high and low probability outcomes. Taking the difference between these two trial types as a measure of learning for the training contingencies therefore allows us to subtract out any influence of proactive interference, such that any difference observed will reflect a difference in the rate of learning about good and poor predictor cues.

While the nonsignificant interaction of prior predictiveness with outcome probability in the ANOVA reported above indicates that this difference between RTs to high and low probability outcomes does not differ significantly between good and poor predictors when assessed across the whole of the training phase, a finer-grained analysis based on learning scores reveals a significant, yet short-lived, influence of prior predictiveness. The difference between RTs on high and low probability outcomes (i.e. the learning score) for both good and poor predictors is plotted in the right-hand panel of Figure 3.4. Pre-planned paired t-tests revealed that the RT difference for good predictors was significantly greater than that for poor predictors in Block 1, $t(19) = 2.26, p < .05$, but that this difference had disappeared by Block 2, $t < 1$. Consistent with both the experimental hypothesis and the results of Experiment 4, the significant difference in Block 1 indicates that learning about good predictor cues was initially at an advantage during the training phase.

Accuracy data for the training phase are shown in the left-hand panel of Figure 3.5. Again, the data show that acquisition progresses rapidly during training, with accuracy on high probability outcomes greater than that on low probability outcomes. These data were again subjected to ANOVA with factors of prior predictiveness, outcome probability, and

block. This revealed a significant effect of outcome probability, $F(1,19) = 27.18$, $p < .001$, but no effect of prior predictiveness, $F < 1$, nor block, $F(1,19) = 2.41$, $p = .137$. There were no significant interactions between any of the factors: outcome probability by prior predictiveness, $F(1,19) = 1.18$, $p = .29$; outcome probability by block, $F(1,19) = 2.45$, $p = .13$; prior predictiveness by block, $F < 1$; nor the three-way interaction, $F < 1$. Learning scores for accuracy data (accuracy on high probability outcomes minus accuracy on low probability outcomes) are shown in the right-hand panel of Figure 3.5. There was no difference between learning scores for good and poor predictors in Block 1, $t < 1$, nor Block 2, $t(19) = 1.22$, $p = .24$.

Discussion

The results of Experiment 5 support the findings of Experiment 4, in that greater learning was observed for cues that were pretrained as good predictors than for cues pretrained as poor predictors. Unlike in Experiment 4, it was clear from the pretraining data in the current experiment that participants had become sensitive to the underlying pretraining sequence. These data therefore corroborate the initial assumption drawn from Experiment 4, that the effect of greater learning about good predictor cues observed in the training data resulted from learning about the differential predictiveness of the different cue types during pretraining.

As in Experiment 4, the effect of prior predictiveness was short lived in the training phase, lasting for only a single block of 150 trials. This was despite the pretraining phase being lengthened from 4 to 10 blocks, which should have provided greater opportunity for participants to learn about differences in predictiveness and hence strengthened any

associability effect. Although the short duration of the prior predictiveness effect is perhaps surprising, it is worth noting that within a single training block each cue-outcome pairing was presented, on average, 20 times. This amount of training is far greater than that used in analogous HCL demonstrations of associability effects (e.g. 4 presentations of each cue-outcome pairing in Le Pelley & McLaren, 2003). We would expect that as training progresses, learning of good predictor cues will reach asymptote, such that participants are unable to express any further learning about these contingencies. At this stage we would expect learning about poor predictor cues to begin to catch up with the learning demonstrated for good predictor cues. In support of this claim, we can see from the RT learning scores (see Figure 3.4) that learning about good predictor cues did not increase between blocks 1 and 2 of the training phase, but learning of poor predictor cues increased dramatically.

One striking difference between the patterns of training data observed in Experiments 4 and 5, was that in the former the effect of prior predictiveness was shown in RTs to high-probability outcomes, whilst in the latter the effect was driven entirely by RTs to the low-probability outcomes. One possible explanation for this difference is that in Experiment 5, participants' RTs were approaching a minimum level, such that participants were unable to respond any faster. Indeed, in Experiment 5 (see Figures 3.3 and 3.4) RTs on GPH trials at the end of pretraining (438ms) were very similar to RTs on GPH trials at the start of training (427ms). Although there was a slight numerical decline in RT across the training phase in Experiment 5, it is likely that this small decline is insufficient to permit RT differences to be observed for responses to high-probability outcomes during training.

It was noted above that proactive interference might have a detrimental effect on learning during the training phase for cues pretrained as good predictors (for which participants will have learnt now-inappropriate cue-outcome mappings during training) as compared to cues pretrained as poor predictors. In fact the RT data shown in Figure 3.4 indicate that good predictors were not at a disadvantage, even in Block 1 of training: RTs following good and poor predictors were in the same range. Note, however, that the data shown in Figure 3.4 are averages across 150 trials in each block, raising the possibility that there has been a rapid recovery from interference as a result of within-block learning. That is, while we might expect a short period at the beginning of the training phase during which interference leads to poorer performance following good predictors, the influence of predictive history on cue processing (facilitating learning for good predictors relative to poor predictors) would over time be expected to counteract and reverse this effect, leading to superior performance for good predictors.

Experiment 6

Experiment 6 aims to provide further support for the prior predictiveness effects shown in Experiments 4 and 5, using a standard six-choice SRT task. The design of the pretraining phase of Experiments 4 and 5 was analogous to previous studies examining the influence of predictive history in animal conditioning and human contingency learning (e.g. Le Pelley & McLaren, 2003), in that certain stimuli acted as cues, whilst others acted as outcomes. It is unclear, however, whether the effect of predictiveness observed in these experiments is limited to this arrangement of designated cue and outcome positions. Experiment 6 sought

to replicate the effect of predictive history using a six-choice SRT task in both phases of the experiment, with all stimulus locations acting as both ‘cues’ and ‘outcomes’ throughout.

Using a six-choice SRT task in both phases of the experiment will also allow the examination of another factor that might potentially have influenced the results of Experiments 4 and 5, namely the change in context between the pretraining and training phases. That is, removing the “outcome” positions following the pretraining phases of Experiments 4 and 5 presumably made it clear to participants that the structure of the task, and the movements of the target, would be different during the training phase. A similar argument applies to most of the previous studies of predictive history effects in human contingency learning (e.g. Bonardi et al., 2005; Le Pelley et al, 2007; Le Pelley & McLaren, 2003), in which there is an explicit change in context between the first phase of the experiment in which predictiveness is established, and the second phase in which the impact of this predictive history on novel learning is assessed. In contrast, in many studies of learned predictiveness effects in animals the same (or very similar) cues and outcomes occur throughout the experiment (e.g. Mackintosh, 1969; Mackintosh, 1973; Holland, 1984). This raises the possibility that effects of predictive history observed in human learning rely on, or are in some way influenced by, the change of context occurring before the critical learning phase. Perhaps, for example, this change in task signals that the cues are now to be involved in different relationships, and therefore leads participants to generalise their previous learning about the predictiveness of the different cues in a way that would not occur if such explicit evidence of a change were not provided. In order to test this suggestion, in Experiment 6 there was no change in context between the two phases of the experiment. That is, both stages involved a six-choice SRT task with the same stimuli, and therefore participants were given no indication that the structure of the task had changed in any way.

While it is at least theoretically possible that a change in context might be required to generate an effect of predictive history on novel learning, it seems unlikely that this would be the case. Studies that have investigated this issue systematically in both animals and humans typically find that a change in context will weaken, rather than enhance, the effect of predictive history (e.g. Lovibond, Preston & Mackintosh, 1984; Nelson & Sanjuan, 2006). Such findings have intuitive plausibility – the greater the difference between the two phases of an experiment, the less likely participants might be to transfer what they have learnt during the first phase to what they are about to learn in the second. To the extent that this applies to the current learning preparation, and other things being equal, we would not expect the removal of an explicit context change in Experiment 6 to weaken the influence of predictive history on novel learning observed in Experiments 4 and 5: if anything the effect might be enhanced.

Although the examination of conscious awareness is not of primary interest in these experiments (see Introduction), in Experiment 6 a basic verbal report questionnaire was given at the end of the experiment in order to provide some measure of the explicit knowledge participants developed during the task. Of particular interest was whether participants had noticed a change between the two stages of the experiment. Considering a higher-order reasoning account of the associability effects reported so far, one might expect that an effect occurs because participants notice a change in the relationships between cue locations at the start of the training phase and then decide to learn selectively as a result of their past experience with the task (i.e. they might focus on cues which have previously been valid sources of target location in the task). It seems, therefore, that particularly compelling evidence for a higher-order reasoning account would be provided by both

participants' awareness of a change in the task contingencies, as well as their ability to locate the point at which this change occurred.

Method

Participants, Apparatus and Stimuli

A new sample of sixteen Cardiff University undergraduates participated for course credit or payment. All apparatus and stimuli were identical to those used in Experiment 2.

Sequence generation

All locations were used as cue and outcome elements in both pretraining and training phases. In describing the generation of the sequence I will use as an example elements 1, 2 and 3 to denote good predictors, and elements 4, 5 and 6 to denote poor predictors. Note, however, that for each participant all six cue elements (three good predictors and three poor predictors) were randomly assigned to the locations 1-6 of the stimulus array.

The left-hand side of Table 3.3 shows the conditional probabilities for the sequence used during the pretraining phase. Positions 1, 2 and 3 were good predictors of their respective outcomes, as all of these positions predicted the location of the target on the following trial with relatively high probability (.9). In contrast, positions 4, 5, and 6 were poor predictors as the position of the target on the following trial could occur in one of two positions with equal probability (.5). The possible outcomes that could follow each of the good predictor cues were always themselves poor predictor cues; each poor predictor cue location acted as a high probability outcome (.9) for one good predictor, and a low

probability outcome (.1) for a different good predictor. Similarly, the possible outcomes that could follow each of the poor predictor cues were themselves always good predictors; each good predictor cue location acted as a medium probability outcome (.5) for two different poor predictor cues. This method of sequence generation doubled the number of presentations of each cue-outcome pairing during pretraining as compared to Experiments 4 and 5, in which half the trials involved transitions from outcome positions (1 and 6) to cue positions (2-5), which were not analysed.

		Trial N (cue) – Pretraining						Trial N (cue) – Training						
		1	2	3	4	5	6	1	2	3	4	5	6	
Trial N+1 (outcome)	1					.5	.5		.1	.8	.1			
	2				.5		.5	.8		.1		.1		
	3				.5	.5		.1	.1					.8
	4	.9		.1					.8			.1	.1	
	5	.1	.9							.1	.8			.1
	6		.1	.9					.1		.1	.8		

Table 3.3. Transition contingencies for the sequences used during the pretraining and training phases in Experiment 6. Numbers 1-6 refer to cues in the stimulus array, although the mapping of numbers in this table to stimulus positions was randomized for each participant. Blank cells indicate zero probability.

The right-hand side of Table 3.3 shows the conditional probabilities for the sequence used during the training phase. Comparing the conditional probabilities for pretraining and training phases in Table 3.3 it is clear that all training transitions were different from those used during pretraining. As in the training phase of Experiment 5, all of the locations now predicted one location with a probability of .8, and two other locations with a probability of .1 each.

Procedure

The procedure was identical to that used in Experiment 5, with the exception that as the training phase also used a 6-choice task, the notice displayed during the rest break following Block 10 in Experiment 5 was omitted. Participants did not receive any indication that the task would change in any way during the experiment. At the end of the 12th block participants answered a series of questions designed to probe their explicit knowledge of the sequence and awareness of the task demands. These were as follows:

1. Did you think any of the locations were presented more often than others? If so, which locations were these?
2. Did you notice any consistent patterns in the movement of the stimulus? If so, could you attempt to tell me what these patterns were?
3. At any point in the experiment did you think the movement of the stimulus might be non-random? If so, in which block did you notice this?

4. If you did notice a pattern, did you feel this pattern changed during the experiment? If so, at what point did it change and in what way?

Participants typed their answers into a text box on the screen.

Results

Trials were excluded on the same basis as in Experiments 4 and 5. One participant produced a mean RT of 1230 milliseconds (median RT of 1094 ms). Given that the majority of this participant's data would have been excluded on the basis of the RT criterion, this participant was excluded from further analysis.

Trials of interest during the pretraining phase were high- and low-probability outcomes following good predictor cues (GPH and GPL respectively) and medium probability outcomes following poor predictor cues (PPM). Since more trials now contributed to each variable, there were no missing data in the current experiment. The RT data for the pretraining phase are shown in the left-hand panel of Figure 3.6. Participants' responses reflected the differing probabilities of the three outcomes: fastest for GPH trials, slowest for GPL trials, and of intermediate speed for PPM trials. These data were subjected to repeated measures ANOVA with factors of outcome probability (high, medium and low) and block, which revealed significant effects of outcome probability, $F(2,28) = 23.10$, $p < .001$, and block, $F(9, 126) = 17.98$, $p < .001$, and a significant interaction, $F(18, 252) = 4.15$, $p < .001$. Pairwise comparisons between the three levels of the outcome probability variable revealed significant differences in RT between all three: GPH versus GPL trials, $F(1, 14) =$

29.45, $p < .001$; PPM versus GPL trials, $F(1, 14) = 23.13$, $p < .001$; GPH versus PPM trials, $F(1, 14) = 11.01$, $p < .01$.

The accuracy data for the pretraining phase are shown in the right-hand panel of Figure 3.6. In line with the findings in RTs, participants were most accurate on GPH trials, least accurate on GPL trials, and of intermediate accuracy for PPM trials. ANOVA revealed a significant effect of outcome probability, $F(2,28) = 16.77$, $p < .001$, and block, $F(9, 126) = 3.33$, $p < .01$, and a significant interaction, $F(18, 252) = 2.89$, $p < .001$. Pairwise comparisons between the three levels of the outcome probability variable revealed significant differences in response accuracy between all three: GPH versus GPL trials, $F(1, 14) = 19.61$, $p < .01$; PPM versus GPL trials, $F(1, 14) = 14.25$, $p < .01$; GPH versus PPM trials, $F(1, 14) = 9.45$, $p < .01$.

It is clear from both the RT and accuracy data that participants were sensitive to the varying levels of predictiveness of the cues. Although Experiment 6 had fewer participants than Experiment 5, the significance levels achieved in most comparisons were greater. This is almost certainly due to an increase in the number of cue-outcome pairings, resulting in greater exposure to the pretraining contingencies, as well as a reduction in the variance due to an increase in the number of sampled trials per data point.




Figure 3.6. Data from the pretraining phase of Experiment 6. *Top panel*: RT data for high-probability outcomes following good cues (GPC), low-probability outcomes following good predictor cues (GPI), and medium-probability outcomes following poor predictor cues (PPM). *Bottom panel*: Accuracy data for the same three trial-types.




Figure 3.7. RT data from the training phase of Experiment 6. *Left panel:* RTs to high- (GPH) and low- (GPL) probability outcomes following cues pretrained as good predictor cues, and RTs to high- (PPH) and low- (PPL) probability outcomes following cues pretrained as poor predictor cues. *Right panel:* The data presented as learning scores - RTs on low probability outcomes minus RTs on high probability outcomes, for cues pretrained as good (GP) and poor (PP) predictors.

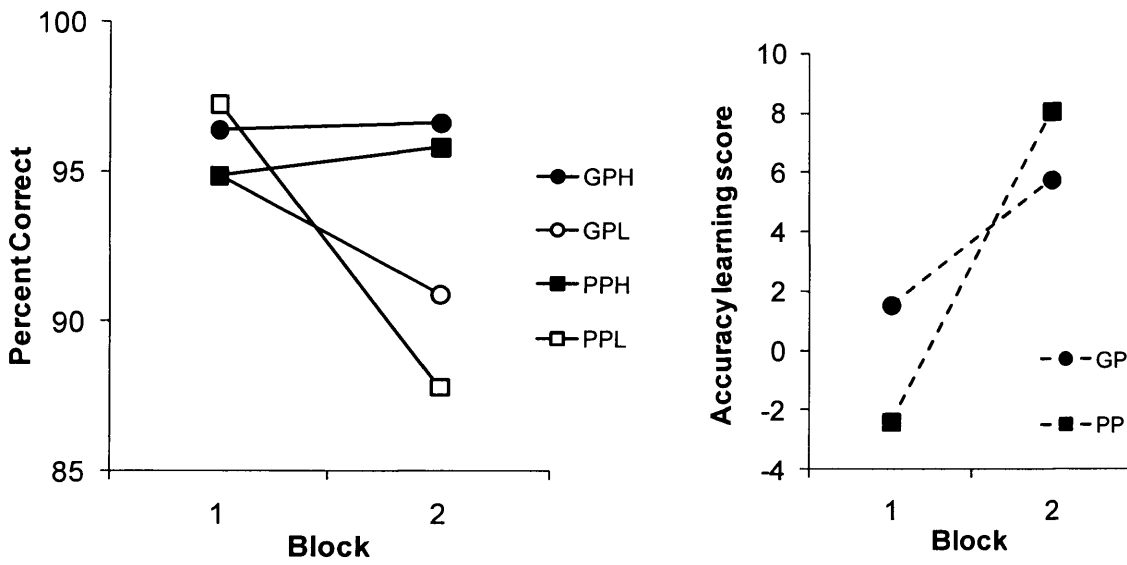


Figure 3.8. Accuracy data from the training phase of Experiment 6. *Left panel:* Accuracy to high- (GPH) and low- (GPL) probability outcomes following cues pretrained as good predictor cues, and accuracy to high- (PPH) and low- (PPL) probability outcomes following cues pretrained as poor predictor cues. *Right panel:* The data presented as learning scores - accuracy on high probability outcomes minus accuracy on low probability outcomes, for cues pretrained as good (GP) and poor (PP) predictors.

The RT data for the training phase were analysed as in Experiment 5, and are presented in Figure 3.7. As previously, trial types GPH and GPL refer to high- and low-probability outcomes, respectively, following cues that were pre-trained as good predictors of outcomes. Similarly, trial types PPH and PPL refer to high- and low-probability outcomes, respectively, following cues that were pre-trained as poor predictors. These data were subjected to repeated-measures ANOVA with factors of prior predictiveness (good versus poor), outcome probability (high versus low) and block. The main effect of outcome probability was significant, $F(1,14) = 10.45$, $p < .01$, indicating faster RTs to high than to low probability outcomes. There was no main effect of prior predictiveness, $F < 1$, nor of block, $F(1,14) = 2.13$, $p = .17$. There was a significant interaction between prior predictiveness and outcome probability, $F(1,14) = 5.60$, $p < .05$, which indicates that the difference in RT between high and low probability outcomes was greater for the good predictor contingencies than for the poor predictor contingencies. The interaction between outcome probability and block was significant, $F(1,14) = 15.86$, $p < .01$, which indicates that overall learning was greater in Block 2 than in Block 1. The interaction between prior predictiveness and block was not significant, $F < 1$, nor was the three-way interaction, $F(1,14) = 1.29$, $p = .28$.

The right-hand panel of Figure 3.7 plots RT data as learning scores (RTs on low probability trials minus RTs on high probability trials) for good and poor predictor contingencies. Planned t-tests revealed that learning scores for good predictor cues were significantly greater than those for poor predictor cues in both Block 1, $t(14) = 2.20$, $p < .05$, and Block 2, $t(14) = 2.15$, $p < .05$.

Figure 3.8 shows the accuracy data for the training phase. These data were subjected to ANOVA with factors of prior predictiveness, outcome probability, and block. This revealed a main effect of outcome probability, $F(1,14) = 5.38$, $p < .05$, which indicates that participants were more accurate on high probability outcomes than on low probability outcomes. There was also a main effect of block, $F(1,14) = 9.47$, $p < .01$, indicating that accuracy decreased from Block 1 to Block 2, driven largely by a decrease in accuracy on low probability outcomes. There was a significant interaction between outcome probability and block, $F(1,14) = 6.49$, $p < .05$, indicating stronger evidence for learning in Block 2 than in Block 1. Prior predictiveness did not exert a significant main effect or interact with any other variable, maximum $F(1,14) = 1.51$, $p = .24$.

The right-hand panel of Figure 3.8 shows the accuracy data as learning scores (accuracy on high probability outcomes minus accuracy on low probability outcomes). There was no difference between these scores for good and poor predictor cues on either Block 1, $t(14) = 1.44$, $p = .17$, or Block 2, $t < 1$.

In response to the first question on the verbal report questionnaire, 9 out of the 15 participants thought that some locations appeared more often than others (in fact all positions were equally frequent). Six participants reported that they believed locations 1 and 6 (the two outermost locations) occurred more frequently than the others.

When asked to report any consistent patterns in the movement of the sequence (Question 2), 12 participants offered specific runs of locations of at least two items (e.g. 43; 13451), whilst one other participant reported that the movement consistently went “inside to outside”. The other two participants were unable to report any consistent transitions. Each run that a participant produced was broken down into single transitions (e.g. 13451 was

broken down to 13; 34; 45; 51) and then each transition was assessed to see if it was a legal transition from either the pretraining or training sequence. For the 12 participants who expressed a specific sequence pattern, the average number of unique transitions reported was 5.1 (min = 1; max = 9; standard deviation = 2.7). Since each participant was free to generate as many transitions as they liked, the number of unique transitions for each variable was converted to a proportion of the total unique transitions generated. Table 3.4 shows the mean proportions for each trial type (along with their standard deviations) and the level expected from chance performance. Chance performance is higher for certain trial types due to the total number of legal transitions used for that trial type (e.g. for the pretraining sequence, there are twice as many PPM transitions than GPH transitions; see Table 3.3).

	GPH _{PT}	GPL _{PT}	PPM _{PT}	GPH _T	GPL _T	PPH _T	PPL _T
Proportion produced	.20	.11	.28	.17	.07	.07	.10
Standard deviation	.33	.15	.30	.24	.11	.10	.15
Chance level	.1	.1	.2	.1	.2	.1	.2

Table 3.4. Quantitative statistics derived from verbal report data from question 2. GPH and GPL refer to high- and low-probability outcomes following a good predictor cue, respectively. PPH, PPM and PPL refer to high-, medium- and low-probability outcomes following a poor predictor cue, respectively. Subscripts PT and T indicate trial types from pretraining and training sequences, respectively.

The modal number of transitions produced for each variable was 0, except for PPM which was .25. As the data were not normally distributed, an analysis was conducted using non-parametric statistics. In order to compare across variables, scores for variables with a

chance level of .2 were divided by 2. Wilcoxon Signed-Ranks tests revealed that there were no differences between the proportion of transitions produced between any of the three pretraining trial types: GPH_{PT} vs. GPL_{PT} , $z = .34$, $N - \text{Ties} = 7$, $p = .74$; GPH_{PT} vs. PPM_{PT} , $z = 0$, $N - \text{Ties} = 11$, $p = 1$; GPL_{PT} vs. PPM_{PT} , $z = .65$, $N - \text{Ties} = 9$, $p = .51$. For training trial types, an equal number of transitions were generated for the high probability and low probability transitions for good predictor cues, GPH_T vs. GPL_T , $z = 1.52$, $N - \text{Ties} = 7$, $p = .13$, and for poor predictor cues, PPH_T vs. PPL_T , $z = .73$, $N - \text{Ties} = 6$, $p = .46$, as well as overall (averaged across GP and PP), $High_T$ vs. Low_T , $z = 1.27$, $N - \text{Ties} = 8$, $p = .21$. There was no difference between good and poor predictor cues in the proportion of extra transitions generated for high than for low probability outcomes, $GPH_T - GPL_T$ vs. $PPH_T - PPL_T$, $z = 1.35$, $N - \text{Ties} = 7$, $p = .18$.

Question 3 asked participants if they had started to believe the movement of the target was non-random, and at what point they noticed this. Six participants either offered no comment, or provided vague answers, such as “some blocks contained non-random bits”. The remaining nine participants provided an estimation of the time-point at which they noticed that the target was moving in a sequenced way. The mean of these estimations was 6.1 (standard deviation of 2.7) indicating that on average participants started to sense the target was moving in a sequenced way at the mid-point of the experiment.

Only 5 participants reported that they felt the sequence changed in some way during the experiment (Question 4). Two of these participants indicated that they felt the sequence changed halfway through the experiment. One indicated that it changed every block, one that it changed before the last four blocks and one that it “changed back to random”, but did

not specify when. These data suggest that the change in the sequence structure for the last two blocks of the experiment went unnoticed by many, if not all the participants.

Discussion

The results of Experiment 6 confirm the general findings of Experiments 4 and 5: the rate at which sequence learning proceeded for a given cue was dependent on the predictive history of that cue. The removal of any explicit change in context between the pretraining and training phases of Experiment 6 indicates that such a change is not necessary for an effect of predictiveness to be observed. Although the verbal report data suggest that the majority of participants were aware that the movement was sequenced, their ability to reproduce the high-probability transitions from either stage was poor. Furthermore, it seems highly unlikely that participants noticed a change in the task sequence during the last two blocks of the experiment, as not a single participant highlighted the first block of training (Block 11) as the block in which they noticed a change in the task structure. Whilst it might be argued that participants would not have kept a running count of the block number, since the questions appeared immediately after Block 12, it seems unlikely that participants would not have been able to report that the task changed during the last two blocks, if this change had been noticed. Although the verbal report data are unlikely to be sensitive enough to allow any strong conclusions to be drawn (see Shanks & St. John, 1994), the data suggest that participants had poor explicit knowledge of the sequenced transitions, but more importantly, that the change in the sequenced material from pretraining to training went unnoticed. This latter finding leads to the conclusion that participants were not strategically applying the

knowledge they had acquired during pretraining to the learning of the training contingencies.

Unlike in Experiment 5, the facilitation in learning during the training phase for those cues pretrained as good predictors was evident for the duration of the training phase. As suggested earlier, it is possible that the longer-lived influence of predictive history in Experiment 6 was a consequence of the lack of an explicit context change as compared to Experiment 5 (cf. Lovibond et al, 1984; Nelson & Sanjuan, 2006) leading to greater transfer between the two stages in Experiment 6. An alternative possibility, however, makes reference to the fact that in Experiment 6 all of the positions in the stimulus array acted as both cues and outcomes. Consequently, participants experienced twice as many presentations of both the good and poor predictor contingencies during the pretraining phase as compared to the pretraining schedule used in Experiment 5 (see Method). It is therefore possible that this difference resulted in the apparently greater effect of predictiveness observed in Experiment 6. On the basis of the current results it is not possible to decide between these two alternatives.

It was noted above that the patterns of data observed in Experiments 4 and 5 were somewhat different: in the former the effect of prior predictiveness on learning during the training phase was observed in responses to high-probability outcomes (see Figure 3.2), whilst in the latter the effect was observed only in responses to low-probability outcomes (see Figure 3.4). In Experiment 6, at least numerically (see Figure 3.7), the effect seems to be present in responses to both high- and low-probability outcomes. In fact, in Block 1 of the training phase, responses to high-probability outcomes following good predictor cues were significantly faster than those following poor predictor cues, $t(14) = 2.21$, $p < .05$,

whilst the difference between responses to low-probability outcomes was not significant, $t(14) = 1.48, p = .16$. The data therefore show a similar pattern to that seen in Experiment 4. A comparison of Figures 3.6 and 3.7 shows that RTs increased between the pretraining (440ms) and training (560ms) phases of Experiment 6. The floor effect I argued was present in the RT data of Experiment 5 is clearly not present in Experiment 6, and therefore there is ample room for an effect to be shown in responses to high-probability outcomes. The current data therefore support the suggestion that floor effects may have masked differences in responses to high-probability outcomes in Experiment 5.

Experiment 7

The results of Experiment 6 revealed a significant effect of predictive history in a task in which the switch between stages was not signalled in any way. As discussed above, the pattern of results was rather different to that seen in Experiment 5: the training effect was observed in RTs to both high- and low-probability outcomes. The primary aim of Experiment 7 was to replicate this pattern of data. I also sought to examine whether it was possible to observe a longer-lived effect of predictive history if the overall rate of learning during the training phase was reduced. It is likely that in Experiments 4-6, the short-lived effect of predictiveness was due to learning about the good predictor cues reaching asymptote, allowing learning about poor predictor cues to catch up towards the end of training. By reducing the chance of high probability outcomes occurring, learning about training contingencies should proceed at a slower rate. In addition, the chance of a high probability outcome occurring in the pretraining phase was increased to .95. Since this will increase the rate of exposure to these contingencies, it should be possible to achieve

equivalent amounts of pretraining (to that of Experiment 6) over a shorter period. To this end, the pretraining phase was reduced to 8 blocks, which I hoped would also reduce the likelihood of participants becoming demotivated in the task.

Method

Participants, Apparatus and Stimuli

A new sample of fourteen Cardiff University students participated for payment of £5. All apparatus and stimuli were identical to those used in Experiment 2.

Sequence generation and Procedure

Sequences were constructed as for Experiment 6, but the probabilities were altered slightly from those shown in Table 3.3. For the pretraining phase, high probability outcomes following good predictor cues now occurred with a probability of .95, and low probability outcomes with a probability of .05. For the training phase, high probability outcomes (for all cues) occurred with a probability of .7, whilst the two low probability outcomes each occurred with a probability of .15. Other aspects of the procedure were as for Experiment 6, with the exception that pretraining was reduced to 8 blocks.

Results

One participant produced a large predictiveness effect during the training phase in the opposite direction to the average data observed in Experiments 4-6: their data showed a greater learning score (RT to low probability outcomes minus RT to high probability

outcomes) for poor predictor cues over good predictor cues. Analysis of the learning scores averaged across both blocks of training revealed that this participant's data lay more than three interquartile ranges below the first quartile and hence were identified as an extreme outlier. This participant's data were therefore excluded from all further analyses. No other participant in this experiment yielded learning scores more than three interquartile ranges above or below the first or third quartile respectively (similar analyses for Experiments 2 to 6 confirmed no outlying learning scores in those studies either⁸).

Trials were excluded on the same basis as in Experiments 3-6. Due to the decreased chance of a low probability outcome occurring after a good predictor cue during pretraining, RT data for this variable were missing on four occasions, and accuracy data were missing on two occasions. These data were replaced using the method described for Experiment 5.

The top panel of Figure 3.9 shows RT data for the pretraining phase, which are consistent with the pattern observed in Experiments 5 and 6. ANOVA revealed significant factors of outcome probability (GPH, GPL, and PPM), $F(2, 24) = 26.09$, $p < .001$, and block, $F(7, 84) = 4.67$, $p < .001$, and a significant interaction, $F(14, 168) = 3.42$, $p < .001$. Pairwise comparisons between the three levels of the outcome probability factor revealed significant differences in RT between each: GPH versus GPL trials, $F(1,12) = 39.36$, $p <$

⁸ For Experiment 4, in which there was an effect of predictiveness in the first block of the training phase, an analysis restricted to this block alone found one outlier. This participant showed a large predictiveness effect in the direction of greater learning about good predictor cues over poor predictor cues. The result of the critical analysis (on Block 1) was significant when this participant was removed. In fact, due to the impact this participant had on the sample variance, the critical t-statistic was larger with this participant removed.

.001; PPM versus GPL, $F(1,12) = 16.39$, $p < .01$; GPH versus PPM trials, $F(1,12) = 15.01$, $p < .01$.

The bottom panel of Figure 3.9 shows accuracy data for the pretraining phase. ANOVA revealed a significant effect of outcome probability, $F(2,24) = 8.52$, $p < .01$, but no effect of block, and no significant interaction, $F_s < 1$. Pairwise comparisons between the three levels of the outcome probability factor revealed a significant difference in accuracy between GPH and GPL trials, $F(1,12) = 9.98$, $p < .01$, and between PPM and GPL trials, $F(1,12) = 7.41$, $p < .05$, but no significant difference between GPH and PPM trials, $F(1,12) = 2.13$, $p = .17$.

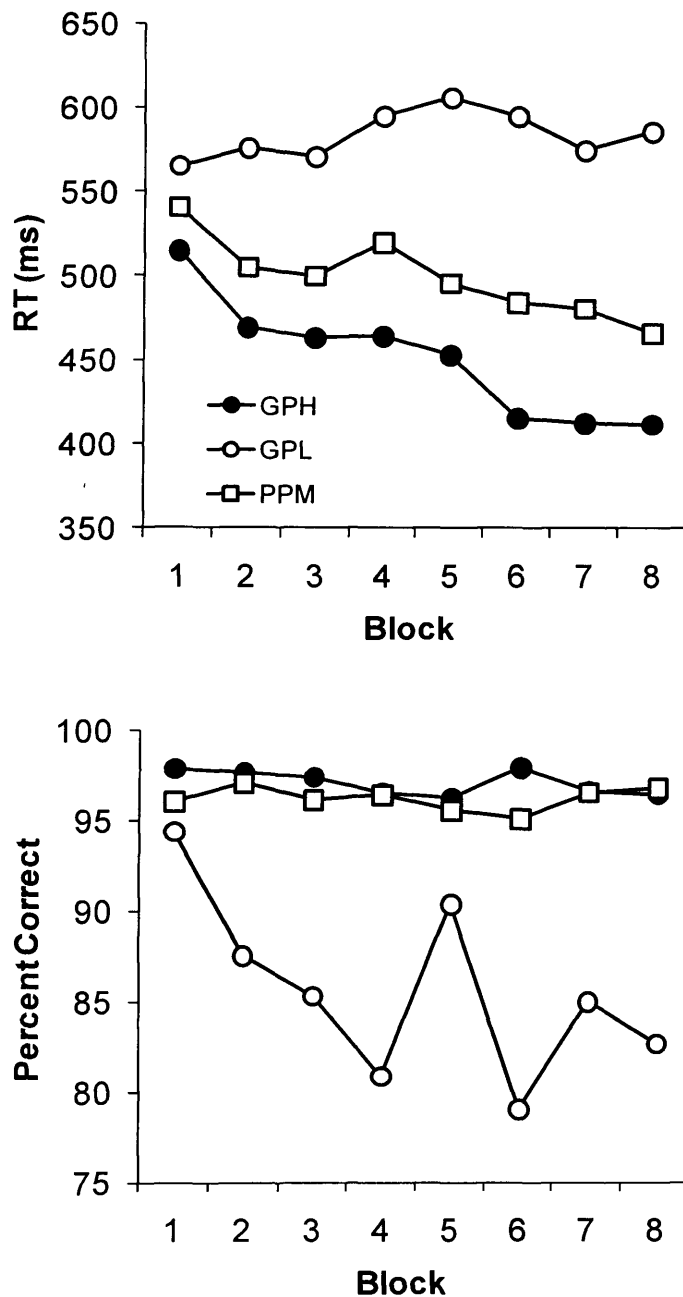


Figure 3.9. Data from the pretraining phase of Experiment 7. *Top panel:* RT data for high-probability outcomes following good cues (GPC), low-probability outcomes following good predictor cues (GPI), and medium-probability outcomes following poor predictor cues (PPM). *Bottom panel:* Accuracy data for the same three trial-types.




Figure 3.10. RT data from the training phase of Experiment 7. *Left panel:* RTs to high- (GPH) and low- (GPL) probability outcomes following cues pretrained as good predictor cues, and RTs to high- (PPH) and low- (PPL) probability outcomes following cues pretrained as poor predictor cues. *Right panel:* The data presented as learning scores - RTs on low probability outcomes minus RTs on high probability outcomes, for cues pretrained as good (GP) and poor (PP) predictors.




Figure 3.11. Accuracy data from the training phase of Experiment 7. *Left panel:* Accuracy to high- (GPH) and low- (GPL) probability outcomes following cues pretrained as good predictor cues, and accuracy to high- (PPH) and low- (PPL) probability outcomes following cues pretrained as poor predictor cues. *Right panel:* The data presented as learning scores - accuracy on high probability outcomes minus accuracy on low probability outcomes, for cues pretrained as good (GP) and poor (PP) predictors.

The RT data from the training phase are shown in Figure 3.10. These data were analysed using ANOVA with factors of prior predictiveness (good versus poor), outcome probability (high versus low) and block. There was a main effect of outcome probability, $F(1, 12) = 23.83, p < .001$, indicating rapid learning of the cue–outcome contingencies during the training phase. There was no main effect of prior predictiveness, nor of block, $F_s < 1$. Unlike in Experiment 5, the prior predictiveness by outcome probability interaction was not significant, $F(1, 12) = 2.00, p = .18$. There was no interaction between prior predictiveness and block, $F < 1$, however, the interaction between outcome probability and block was marginally significant, $F(1, 12) = 4.51, p = .055$, which suggests there was more evidence of learning in Block 2 than in Block 1. The three-way interaction was not significant, $F < 1$.

The right-hand panel of Figure 3.10 plots RT data as learning scores for both good and poor predictor cues. Planned t-tests revealed that learning of good predictor contingencies was significantly better than for poor predictor contingencies in Block 1, $t(12) = 2.30, p < .05$, but not in Block 2, $t < 1$.

Figure 3.11 shows accuracy data for the training phase. These data were again subjected to ANOVA with factors of prior predictiveness, outcome probability and block. There was no main effect of prior predictiveness, $F(1, 12) = 1.72, p = .21$, or block, $F(1, 12) = 1.45, p = .25$, but the main effect of outcome probability approached significance, $F(1, 12) = 4.26, p = .061$, which indicates higher accuracy on high probability outcomes than on low probability outcomes. None of the interaction effects reached significance, all $F_s < 2.75, p_s > .12$. Learning difference scores for these data (accuracy on high probability outcomes

minus accuracy on low probability outcomes) are plotted in the right-hand panel of Figure 3.11. There was no difference between these scores in either block, $t_s < 1$.

Discussion

Experiment 7 replicated the basic effect shown in Experiments 4-6: cues that were pretrained as good predictors were subsequently learnt about more rapidly than were cues pretrained as poor predictors. The advantage was, however, only seen in Block 1 of the training phase, and so Experiment 7 did not provide a perfect replication of the pattern of data shown in Experiment 6. The attempt to produce a longer-lasting effect of prior predictiveness during the training phase (by reducing the likelihood of a high probability outcome occurring) was unsuccessful. It is somewhat surprising that the predictiveness effect observed in Experiment 7 was, if anything, shorter-lived than that observed in Experiment 6. One possibility is that the decreased length of pretraining was successful in keeping motivation at a consistently high level during the task, and that this increase in motivation itself led to faster learning during the training phase, as compared with Experiment 6. Alternatively, it may be that the current participants simply happened to be “faster learners” in this task compared to those tested in Experiment 6. As I argued when discussing the data presented in Chapter 2, one must be cautious about making between-subjects comparisons that might well be influenced by rates of learning, and the same arguments can be applied here too.

One motivation behind the replication of Experiment 6 was as verification of the particular pattern of data observed in Experiment 6 as compared to Experiment 5. Recall that in Experiment 5, the effect of prior predictiveness was driven entirely by a difference in

RTs to low probability outcomes (Figure 3.4), while in Experiment 6 it was driven largely by a difference in RTs to high probability outcomes (Figure 3.7). This difference was ascribed to the presence of a floor effect in RTs (to high probability outcomes) in Experiment 5, but not in Experiment 6. The results of Experiment 7 are consistent with this argument. Average RT to high probability outcomes was 412ms at the end of pretraining (Figure 3.9), whilst average RT to high probability outcomes increased to 544ms at the start of training (Figure 3.10). This increase in RTs between the pretraining and training phases lifts RTs away from any potential floor effect, leaving ample room for a predictiveness effect to emerge in responses to high probability outcomes, and indeed such an effect was observed in the training data.

Summary and discussion of the empirical evidence presented in Chapters 2 and 3

The aim of the research presented in Chapters 2 and 3 was to seek evidence for changes in the associability of cues within an incidental learning task. Experiment 1 adapted an original AGL experiment conducted by Reber (1969), in which participants were given incidental training on letter strings from one grammatical structure, before being transferred to a test phase involving letter strings from a novel grammatical structure. Participants in this condition were able to discriminate between novel grammatical and non-grammatical exemplars during the test phase, whilst participants in a control condition, who received training with non-grammatical stimuli, were not. These results were consistent with the idea of a change in the stimulus associabilities during the training phase.

In Experiment 2, a similar ‘structural transfer’ design was instantiated in an SRT task. In the experimental condition, participants were pretrained on a second-order conditional sequence, which should lead the associabilities of all cues within the task to be high at the end of the pretraining phase. In the control condition, participants were pretrained on pseudo-random transitions, which should lead the associabilities of all cues to be low at the end of the pretraining phase. All participants were then trained on a novel second-order conditional sequence. Despite an initial disadvantage for sequence learning, participants in the experimental condition showed greater learning of the training sequence by the end of the training phase, as compared to the control condition.

Experiment 3 sought to replicate the effect shown in Experiment 2 using simple first-order conditional sequences, again using a between-subjects design with experimental and control conditions. However, this replication failed, and it was suggested that this may have been due to the masking of any associability effect as a result of the rapid learning of first-order conditional sequences during the training phase. Moreover, it was argued that the between-subject variability in the RT measures makes it almost impossible to draw valid conclusions from these data. Although the use of large sample sizes would alleviate the problems seen in Experiment 3, the necessity of using large sample sizes meant it was inefficient to continue using between-subject designs in this research.

In Experiments 4 to 7 the associabilities of the cues were manipulated in within-subject designs: cues within the task acted as either good or poor predictors of the following event. In Experiment 4, dedicated outcomes (with dedicated responses) were used during the pretraining phase – a star and a triangle stimulus, which would appear above the normal SRT array. These outcomes were then removed from the task before the training phase. It

was hoped that the removal of these dedicated outcomes would limit participants' learnt tendencies to initiate responses to these stimuli in the training phase. This method of pretraining is somewhat analogous to the designs used in HCL: certain stimuli in the task function exclusively as cues, whilst others function exclusively as outcomes. During the training phase, cue-outcome contingencies involving cues that were pretrained as good predictors of the dedicated pretraining outcomes were learned about more readily than cues that were pretrained as poor predictors.

Experiment 5 again used dedicated outcomes and responses during the pretraining phase, but these outcomes were standard stimulus positions in a 6-choice SRT task. It was hoped that the removal of the salient, "special" pretraining outcomes used in Experiment 4 would make the contingencies between task cues appear less obvious to participants. Furthermore, Experiment 5 used a method of controlling the impact of proactive interference (generated by the learnt responses from the pretraining phase) on performance in the training phase. As such, it was possible to increase the pretraining phase to over twice the length of that used in Experiment 4, in the hope that this would increase the differentiation in the associabilities of the task cues. In the training phase, the dedicated pretraining outcomes were removed and again greater learning was observed for cues pretrained as good predictors compared to cues pretrained as poor predictors.

One potential problem with the designs of Experiments 4 and 5 is that the change in task between the pretraining and training phases (the removal of the dedicated pretraining stimuli) might act as a signal that the task had changed in some way. It is possible, therefore, that the kind of higher-order reasoning processes that the use of the SRT task was designed to avoid, might in fact be evoked by this salient change. Therefore, in Experiment

6 a standard 6-choice SRT task was used throughout the experiment. The design of the pretraining phase allowed twice as many cue-outcome contingencies to be presented in the same duration of time (as compared to Experiment 5) and, more than likely as a result of this design change, a greater effect of prior predictiveness was observed in Experiment 6. Verbal report data collected after the experiment suggested that participants started to gain some awareness of the sequenced movements of the target mid-way through the experiment. However, participants were unaware of the change in task contingencies for the final two blocks of the task: a finding which I would argue is incompatible with an account of these associability effects in terms of controlled reasoning processes. Experiment 7 provided a replication of the overall pattern observed in Experiment 6, although the effect of prior predictiveness was limited to the first block of training.

The data presented here provide the first empirical evidence from an incidental learning task, that the prior predictive history of a cue modulates the rate at which that cue is learnt about in the future. The data support a number of demonstrations from HCL tasks in which the associability of a cue will change as a result of the prior-predictive history of that cue (e.g. Le Pelley & McLaren, 2003; Le Pelley et al., 2007), and fit within the framework of associative models that allow for cue-associability to vary (e.g. Mackintosh, 1975; Le Pelley, 2004; see Chapter 4). The direction of the predictiveness effect (greater good predictor learning over poor predictor learning during the training phase) is in keeping with the general principles of the Mackintosh (1975) model, in which the associability of the best available predictor of the outcome rises, whilst the associability of poorer predictors will fall.

Although the verbal report data collected in Experiment 6 are unlikely to provide a sensitive measure of participants' explicit knowledge (see Shanks & St. John, 1994), it is clear from these data that participants were unaware of a change in the task contingencies. I would argue therefore, that at the point at which changes in cue associability were assessed (Block 1 of the training phase) it is highly unlikely that participants were applying controlled reasoning in order to learn selectively in the task. One can imagine that in standard HCL tasks (e.g. the food-allergy paradigm) participants might adopt a strategy in which they learn selectively about certain aspects of the task in order to complete the task efficiently. In the SRT task, however, participants perceive and respond to each stimulus consecutively. Since the task simply requires participants to respond to the current location of the target, a strategy of selective learning in this task seems at odds with the fundamental task demands.

I would argue that the associability effects demonstrated here fit most comfortably within the framework of automatic associative learning mechanisms. In the next chapter I shall explore various associative accounts of the results, before going on to examine how changes in cue-associability might be incorporated into current models of sequence learning.

Chapter 4 – Computational simulations of associability effects in sequence learning.

Computational modelling of psychological processes

In Chapter 1 I discussed in detail several associative learning models. Perhaps the main reason why these models have provided such influential theoretical accounts is that they provide a formal mathematical description of the psychological processes they seek to explain. Not only is it therefore very clear how these models function, but moreover each model can be simulated, and verified or falsified as an accurate model of a particular process. As a starting point in examining the mechanisms underlying sequence learning, I shall first test some simple associative learning models with a sample of the sequences used in Chapter 3. Since these models were discussed in detail in Chapter 1, only a general summary shall be provided here, along with details of how the models have been implemented and tested. The failures and successes of the models will be discussed, and will provide a basis for developing a more representative model of sequence learning.

Simulations with the Rescorla-Wagner (1972) model

For over 30 years the Rescorla-Wagner (1972) model has provided a simple associative mechanism for understanding animal and human learning. The essential principle of the model is that learning about the relationship between a cue and an outcome is governed by the extent to which that outcome is well predicted. If an outcome is surprising (i.e. it is not well predicted by the present cues) then large changes in associative strength are made

between those cues and the outcome. Conversely, if an outcome is unsurprising (i.e. it is predicted by some or all of the presented cues) then a relatively small change will be made to the associative strengths of the presented cues.

It has been shown that the Rescorla-Wagner model is mathematically equivalent to the Delta Rule (Widrow & Hoff, 1960) used in many connectionist models (see Sutton & Barto, 1981). Since I shall go on to describe a connectionist architecture in the latter part of this chapter, it seems sensible to keep all description of the models compatible. As such, I shall describe the model as a network of input units (cues; CSs) adjoined to output units (outcomes; USs) by weighted connections (associative strengths). The basic architecture of the single-layer delta-rule model – adapted for sequence learning – is shown in Figure 4.1.




Figure 4.1. A single layer network. The units I₁ to I₆ represent six input units, whilst the units O₁ to O₆ represent six output units

In order to implement the model for the purposes of sequence learning, 6 input (I₁ to I₆) and 6 output units (O₁ to O₆) are used to represent the 6 locations in which the target can appear. For example, the weight between I₁ and O₂ represents the model's level of expectancy for the target to appear in location 2 on trial N+1, given that the target appeared in location 1 on trial N. The model therefore assumes that the sequence contains discrete cue-outcome pairings: if location 2 occurs as an outcome on trial N+1 it acts as the cue on trial N+2.

Before the model was presented with the sequence procedure, all weights within the network were initialised with a random value between 0 and 1. This method of initialization produces a random starting state for the network, akin to the unpredictable prior influences

participants will inevitably bring to the task. Therefore part of the model's learning procedure involves the extinction of these prior influences.

The implementation of the model consisted of two main steps on each trial. Firstly, the input unit representing the location of the target on trial N was activated (set to 1), whilst all other input units were inactivated (set to 0). This will then lead to activation in the layer of output units, with the activation level of each output unit corresponding to the strength of the weight between that output unit and the activated input unit. These activation levels were used to assess the model's ability to predict the location of the target on trial N+1 (the outcome). The simulated 'reaction time' of the model was based on a comparison of the activation of the output unit representing the "correct" outcome (i.e. the location to which the target moved on trial N+1), relative to the summed activation of all output units. This ratio is commonly called the Luce Choice Ratio (hereafter LCR; Luce, 1959):

$$\text{LCR} = \frac{a_{\text{target}}}{\sum_{o=1}^6 a_o} \quad (4.1)$$

where a_o is the activation of output unit o , and a_{target} is the activation of the target output unit (i.e. the output unit relating to the target location on the next trial of the sequence). As the model learns the sequence it will get better at predicting the appropriate output for each input: large weights will develop between that input unit and the target output unit, whilst weights will be small between that input unit and other non-target output units. The higher the LCR, the more the model is committed to a prediction that the target will appear in the correct location as compared to all others. Thus it seems reasonable to assume that LCR is

inversely related to RT, and hence figures showing the performance of all models will be plotted using $(1 - \text{LCR})$.

All weights were then adjusted by the following learning rule:

$$\Delta w_{oi} = \alpha_i \cdot \beta_o \cdot (T_o - a_o) \cdot a_i \quad (4.2)$$

where w_{oi} is the weight between output unit o and input unit i , α_i is a fixed learning rate for input unit i , β_o is a fixed learning rate for output unit o , T_o is the target activation for output unit o (set at .8 if that output is expected, and 0 otherwise), and a_i is the activation of input unit i . Since only the activation of the currently presented input unit was greater than 0, only weights pertaining to this input unit were adjusted. Given the randomization of cues to stimulus positions in the task, α_i and β_o were, on average, the same for all stimuli.

At the start of training (e.g. when $w_{oi} < .8$), the prediction error between the actual output activation and the target output activation will be positive ($T - a_o > 0$), resulting in positive changes for w_{io} . Conversely, for all other weights related to that input unit, the associative strengths will be equal to or greater than the target activation (i.e. $T = 0$). As a result, any positive weights between these input units and non-target output units will be reduced ($0 - a_o < 0$), resulting in less expectation for these outcomes following this cue on future trials.

Since the largest training phase effect was observed in Experiment 6, the design of this experiment was used in all the model simulations presented in this chapter (the RT data of Experiment 6 are presented in Figure 4.2). Each simulation contained 15 simulated participants trained with the exact sequences used in Experiment 6. Data from the 15 simulated subjects were averaged to produce the model's performance for each parameter

set. In the Rescorla-Wagner model the values of α and β are fixed for the duration of the experimental procedure. In the current simulations these two constant parameters were combined into one constant learning rate parameter, which was manipulated between simulation runs. This was the only free parameter in the Rescorla-Wagner model simulations.

In general, sequence learning tasks involve a large number of presentations of each cue-outcome contingency. For instance, in Experiment 6, each high-probability contingency (GPH; see Table 3.3) occurred, on average, 22.5 times in each pretraining block. However, during the pretraining phase of Experiment 6, learning emerged gradually, such that there was an interaction between outcome probability and block during this phase (i.e. there was greater evidence for learning at the end of the pretraining phase than at the start). In order to find a rate of learning which produced a similar pattern of performance in the simulation results, a range of learning rate parameters were tested: .0001; .001; .01; .02; .05; .07; .1; .3; .5; .7. Learning rates of .01 and lower failed to produce asymptotic learning during the pretraining phase, whilst rates of .1 and higher led to asymptotic learning within the first two blocks of the pretraining phase. Learning rates of .02, .05 and .07 produced a more accurate reflection of the pretraining performance seen in Experiment 6 (see Figure 3.6). The results of an example simulation, using a learning rate of .05, are shown in Figure 4.3.




Figure 4.2. RT data from Experiment 6. *Left panel:* Pretraining phase data – RTs to high-probability outcomes following good cues (GPC), low-probability outcomes following good predictor cues (GPI), and medium-probability outcomes following poor predictor cues (PPM). *Right panel:* Training phase data – RTs to high- (GPH) and low- (GPL) probability outcomes following cues pretrained as good predictor cues, and RTs to high- (PPH) and low- (PPL) probability outcomes following cues pretrained as poor predictor cues.




Figure 4.3. An example of the simulation results produced by the Rescorla-Wagner model trained on the sequences used in Experiment 6. Learning rate was set at .05. See Figure 4.2 for a description of the trial-types. *Top panel:* Pretraining phase performance. *Bottom panel:* Training phase performance.

The example shown in Figure 4.3 is typical of the pattern of results produced by the Rescorla-Wagner model across the range of learning rates examined. Following successful learning of the pretraining contingencies, learning of the training phase contingencies commences, but proceeds at the same rate for outcomes following good and poor predictor cues⁹. Thus, unlike in the empirical data, there is no effect of prior predictiveness on new learning.

These findings are unsurprising, given a basic assumption of the Rescorla-Wagner model is that learning is ‘path-independent’. That is, new learning about a cue is independent of the associative history of that cue; the cue-specific learning rate (α) remains constant across training. Of course, this is not to say that the associations pertaining to a cue will not influence how associations develop in the future, rather that the means by which that association was accrued will not affect learning in the Rescorla-Wagner model. It is worth noting here that one would expect a certain amount of proactive interference (as a result of pretraining) to obstruct the learning of the training phase contingencies. In fact, the Rescorla-Wagner model consistently predicted slightly less learning about good predictor cues than poor predictor cues in the training phase. However, as a result of the rapid extinction of pretraining associations in the Rescorla-Wagner model, the effect of proactive interference is short-lived, producing only a minute impairment on new learning (see Figure 4.3).

⁹ As in Chapter 2, for any discussion relating to the training phase procedure, the terms *good* and *poor* will be used to describe how these sets of cues were pretrained, despite the fact that both sets of cues were equally predictive of their respective outcomes during the training phase.

Simulations with the Mackintosh (1975) model

Unlike in the Rescorla-Wagner model, in the Mackintosh (1975) model, changes in associative strength between a cue and an outcome are influenced by the predictive history of that cue: the rate at which associative learning occurs for a given cue is influenced not only by the current associative strengths pertaining to that cue, but also by the process by which that associative strength was accrued. Mackintosh (1975; see also, Krushcke, 2001) suggested that the rate at which learning proceeds for a cue is determined by the extent to which that cue has been a consistent signal of reinforcement in the past, relative to all other currently presented cues. Specifically, Mackintosh suggested that the associability of a cue (α) should increase if a cue is a better predictor of the current outcome than are all other presented cues, whilst the value of α should decrease if a cue is a poorer predictor of the current outcome than are all other presented cues. This mechanism for changes in α reflects how simultaneously presented cues are assumed to compete with each other for limited attentional resources, i.e. it involves a direct comparison of the relative predictiveness of several simultaneously presented cues with respect to the same outcome. Hence this model is clearly designed with multiple cue presentation procedures in mind. In order to assess whether a variant of the Mackintosh model can account for the empirical findings shown in Chapter 3, it is necessary to examine how the rules governing changes in α can be modified to accommodate a learning procedure that uses a single-cue presentation mode.

Relative predictiveness (with respect to context) as a determinant of cue-associability

It seems unlikely that changes in associability occur as a result of a direct comparison between cues in the SRT task (as would be envisaged by the original instantiations of the

Mackintosh and Kruschke models), since only a single cue is presented at a time in the task. However, one possibility is that cue validity is assessed relative to the experimental context. In fact, this notion of cue-competition with the context was put forward by Mackintosh (1975) as a means by which changes in α could occur during single cue conditioning (e.g. latent inhibition; learned irrelevance). Mackintosh (1975) suggested that under such circumstances the cue can be thought to form a compound stimulus with the experimental context. In the case of learned irrelevance (see Chapter 1), for example, since the task cue is no better a predictor of the occurrence of reinforcement than the experimental context, a decline in its associability will occur.

The model shown in Figure 4.1 was amended to include an additional input unit that represented the context. The weights linking this input unit to the output units therefore represent the context-outcome associations. In order to implement changes in these context-outcome associations, on each trial the model was presented with a compound stimulus of the cue stimulus and the experimental context. Following each trial of the sequence procedure, the context stimulus was presented to the model alone, which simulated the period between trials in which the context was present in the absence of any outcome.

On each trial the accuracy of the model's prediction was assessed using the Luce Choice Ratio (Equation 4.1). Associative strengths were adjusted using the following learning rule:

$$\Delta w_{i_o} = \alpha_i \cdot \theta \cdot (T_o - a_o) \cdot a_i \quad (4.3)$$

where α_i represents the learning rate associated with input unit i , and θ represents a fixed learning rate parameter. This equation is essentially identical to the learning rule for the Rescorla-Wagner model (Equation 4.2). However, the critical difference is that now,

following changes in associative strength, the α value for each presented cue (i.e. the ‘task cue’ and the ‘context cue’) was adjusted. The associability of a cue increased if that cue was a better predictor of the outcome than other presented cues, whilst the associability of a cue decreased if that cue was a poorer predictor of the outcome than other presented cues. The magnitude of the change in the associability was determined by the relative predictiveness of the two cues:

$$\Delta\alpha = \phi \cdot \left(|T_o - w_{xo}| - |T_o - w_{io}| \right) \quad (4.4)$$

where w_{xo} reflects the associative strength between the competing cue and the outcome (when calculating the associability of the task cue, the “competing cue” is the context cue; when calculating the associability of the context cue, the “competing cue” is the task cue). ϕ is a fixed learning rate parameter determining the general rate at which changes in associability occur. Alpha values were restricted to the range from .05 to 1. The lower limit prevents a cue from becoming “frozen out” of the learning process entirely, whilst the upper limit reflects an assumption that there is some limit on the amount of attention that can be devoted to a cue.

All other aspects of the simulations were the same as those conducted with the Rescorla-Wagner model. The simulations examined the effect of three parameters on the model’s performance: the learning rate governing changes in associative strength (θ), the learning rate governing changes in associability (ϕ), and the starting value of α . Values of .001, .01, .05, .1, .3, .5, and .7 were used for learning rate parameters θ and ϕ . Values of .1, .3, .5, and .7 were used as the starting values of α . The combination of these parameter values resulted in a total of 196 parameter sets.

Each simulation result was first analysed on the basis of its performance with the pretraining contingencies. Successful pretraining was assumed to have occurred when the average LCR on GPH trials was greater than the average LCR on PPM trials by a margin of at least .2, in the last block of the pretraining phase. A total of 37 simulations failed to meet this criterion, consisting mainly of those simulations employing the lowest value of θ (.001). For the remaining 159 simulations, the effect of prior predictiveness during the training phase was calculated using the difference in average LCRs for high- and low-probability outcomes, across the two blocks of the training phase:

$$(GPH - GPL) - (PPH - PPL) \quad (4.5)$$

Positive values on this measure indicate greater learning about good predictor cues, whilst negative scores reflect greater learning about poor predictor cues. Only 13 (8.2%) simulations produced a positive effect of greater learning about cues pretrained as good predictors, and these effects were all of a small magnitude (mean = .01, standard deviation = .009). In contrast, 144 (90.6%) simulations produced a small negative effect (mean = -.02, standard deviation = .005), whilst the remaining 2 simulations produced a zero score (to 4 decimal places).

For many of the simulations (80.5%) the model parameters led to asymptotic values of α for both good and poor predictor cues by the end of the pretraining phase (both $\alpha > .95$). Any simulation that fails to produce a clear differentiation in the α values of good and poor predictor cues at the end of the pretraining phase would not be expected to result in an effect of prior predictiveness during the training phase. Indeed, of the 144 simulations producing a negative effect, 87.5% produced asymptotic alpha values at the end of the pretraining phase. Effectively these simulations would function in the same manner as the

Rescorla-Wagner model during the training phase, since the associabilities of all cues would be approximately equal to 1.

Of the 13 parameter sets showing a positive training effect, only 1 produced asymptotic values of α at the end of the pretraining phase. Those 12 parameter sets producing sub-asymptotic α values were facilitated by the use of a low learning rate affecting changes in associability ($\phi \leq .05$), and in the most successful simulations, a low starting value for α (e.g. .03 or .01). Figure 4.4 shows the model's performance with the parameter set that produced the largest effect of prior predictiveness during the training phase. The combination of a low starting value of α and a low ϕ , ensured that the associabilities of all cues remained low (below .25) across the pretraining phase. Since good predictor cues are far better at predicting the next target location than the context is, the α values for good predictor cues increase across the pretraining phase. The α values for poor predictor cues also rise steadily across pretraining, since these cues are also better predictors of the target location than the context is. However, since the difference in predictiveness of the poor predictor cues and the experimental context is less than that between the good predictor cues and the context, a small differentiation in α can be seen as the pretraining phase progressed. This in turn led to a small, positive effect of prior predictiveness during the training phase (.032 – see Equation 4.5).




Figure 4.4. An example simulation with the Mackintosh model, using an associability mechanism governed by relative predictiveness ($\theta = .1$; $\phi = .001$; starting $\alpha = .1$). See Figure 4.2 for an explanation of trial types. *Top left:* Model performance (1 – LCR) during the pretraining phase. *Top right:* Mean alpha (α) value for cues pretrained as good (GP) and poor predictors (PP) during the pretraining phase. *Bottom left:* Model performance during the training phase, with novel cue-outcome pairings. *Bottom right:* The training data expressed as learning scores (low-probability outcome minus high-probability outcome).

Despite a wide parameter search, the model was only able to produce an effect of prior predictiveness within a limited parameter range (in 8.2% of the analysed simulations). The main cause of the model's poor performance lies in the use of the context as a competing stimulus for changes in cue-associability. Learning trials in which the context occurs in the absence of any outcome, ensure that the associability of the context declines rapidly during the pretraining phase. Furthermore, the context is paired equally with all outcomes during the pretraining phase, and so only weak associations will develop between the context and the six target locations. Thus, even for those poor predictor cues, the relative predictiveness of the task cues is always far greater than that of the competing, context cue during the pretraining phase. This inevitably leads to either asymptotic values of α for the task cues, or values which fail to differentiate sufficiently to produce a prior predictiveness effect of any reasonable magnitude during novel learning in the training phase.

Relative predictiveness (across trials) as a determinant of cue-associability

The issues with the previously-described instantiation of the Mackintosh model arise from its use of relative predictiveness to determine α – since both GP and PP cues are relatively better predictors of the outcome than is the context, assessing the predictiveness of the task cues with respect to the context will mean that only small differentiations in α will occur between the two. An alternative approach, which might overcome this problem, would instead base α on a comparison of the relative predictiveness of cues *across trials*, i.e. to compare the predictiveness of task cue X on a particular trial to the predictiveness of other task cues on previous trials. Quite how such a mechanism would be implemented is unclear, however. That is, while current attentional models (e.g. Mackintosh, 1975; Kruschke 2001)

suggest that changes in cue-associability are determined by the relative ability of the currently presented cues to predict a particular outcome, it is unclear exactly how this cue-comparison process would operate when conducted across trials. As an example, take an SRT procedure (e.g. like that used in Experiment 6) in which a good predictor cue, 1, accurately predicts the outcome with which it is consistently paired, 2, whilst a poor predictor cue, 3, consistently predicts the occurrence of two different outcome locations, 4 and 5. Whilst cue 1 is a better predictor of outcome 2 than cue 3 is, it is also clear that cue 1 is a poorer predictor of outcome 4, than cue 3 is. Thus, the problem lies in the notion of cue-predictiveness: in a mechanism in which predictiveness is compared across trials, it seems necessary for predictiveness to be defined more abstractly as the extent to which a cue predicts outcomes in general, rather than any one specific outcome. However, even by adopting an abstract notion of predictiveness, further complications arise when one considers which cues should be involved in a comparison across trials. For example, a comparison between two good predictor cues would fail to find one cue more predictive than another. This problem could be circumvented by assuming that the good predictor cues are compared to poor predictor cues, though such a comparison assumes that the participants must already have knowledge of the predictiveness of the cues, and thus the circularity of this process makes it somewhat unlikely. Although it is possible that there are other means by which a comparison across trials could occur (i.e. by comparing the predictiveness of the current cue with an average of the predictiveness of previously presented cues), this method seems to demand several additional assumptions about the attentional mechanism. Thus it seems logical to reject this method in favour of a more parsimonious process.

Absolute predictiveness as a determinant of cue-associability

A simple means of overcoming the issues described above, is to allow α to reflect the absolute predictiveness of that cue, as determined by the prediction error generated by that cue. Thus, a cue which consistently predicts the outcomes with which it is paired should have a high associability, whilst cues which are less perfect predictors should have a lower associability.

A winner-takes-all mechanism was used to assess which outcome was predicted by the model on each trial. The associability of a cue increased if the cue predicted the correct outcome, whilst the associability of the cue decreased if an incorrect prediction was made. The magnitude of associability change on a trial was governed by the prediction error of the model, in this case using the LCR.

$$\text{If correct then } \alpha_i = \alpha_i + (\text{LCR})^N \quad 4.6$$

$$\text{If incorrect then } \alpha_i = \alpha_i - (\text{LCR})^N$$

where N is a free parameter used to reduce the magnitude of changes in associability. The model contained three free parameters: the learning rate affecting changes in associative strength, θ (using values of .001, .01, .05, .1, .3, .5, and .7), the starting value of α (using values of .1, .3, .5, and .7), and the exponent in the adjustments to associability, N (using values of 1, 2, and 4). The combination of these values produced a parameter space of 84 simulations.

Performance was evaluated in the same way as for the previous instantiation of the model. Of the 84 simulations, 22 failed to show satisfactory learning of the pretraining contingencies, and again this was mainly those simulations that used the lowest learning rate

parameter of .01 or lower. Of the remaining 62 simulations, 31 (50%) produced a positive training phase effect of greater learning about cues pretrained as good predictors, whilst 27 (43.5%) simulations produced a negative effect, and 4 produced a zero difference (to 4 decimal places).

Of the 27 simulations producing a negative training phase effect, 59.3% of the parameter sets led to asymptotic values in α at the end of the pretraining phase. It seems somewhat peculiar that the rules used to adjust α (Equation 4.6) could produce asymptotic values for poor predictor cues, since for these cues the model's prediction will be incorrect on half of all trials during this phase. As a result one would expect the associability of a poor predictor cue to decrease as often as it increases. However, whilst this is true, since LCR values will be greater for correct than incorrect predictions, positive changes in associability will always be greater than negative changes. This inequality is especially pronounced during the early stages of learning. Unsurprisingly therefore, asymptotic values in α for poor predictor cues were often the result of a high learning rate (e.g. $\theta > .3$), as shown by a highly significant positive correlation between the learning rate (θ) and the average value of α for poor predictor cues at the end of the pretraining phase, $r(62) = .61$, $p < .001$.

None of the 31 simulations producing a positive training phase effect had asymptotic α values for both good and poor predictor cues. The average size of the training phase effect was .046 (standard deviation = .028). The largest positive training phase effects were produced by simulations with low learning rates (e.g. $\theta < .1$), and in the majority of instances with low starting values of α . Moreover, when the combination of parameters led to a clear differentiation in the α values for good and poor predictor cues by the end of the pretraining phase, large effects of prior predictiveness were observed in the training phase

data. In support of this conclusion, there was a significant positive correlation between the difference in α values at the end of the pretraining phase (average α for GP cues minus average α for PP cues) and the size of the training phase effect of prior predictiveness, $r(31) = .46, p < .01$. An example simulation is shown in Figure 4.5.




Figure 4.5. An example simulation with the Mackintosh model, using an associability mechanism governed by absolute predictiveness ($\theta = .1$; starting $\alpha = .3$; $N = 2$). See Figure 4.2 for an explanation of trial types. *Top left*: Model performance (1 - LCR) during the pretraining phase. *Top right*: Mean alpha (α) value for cues pretrained as good (GP) and poor predictors (PP) during the pretraining phase. *Bottom left*: Model performance during the training phase, with novel cue-outcome pairings. *Bottom right*: The training data expressed as learning scores (low-probability outcome minus high-probability outcome).

Interim summary

Simulations with a connectionist implementation of the Rescorla-Wagner (1972) model showed that it was unable to capture the prior predictiveness effect shown throughout Chapter 3 (specifically that shown in Experiment 6). It is clear from the simulations based on a version of the Mackintosh (1975) model that this failure is most likely due to the use of a fixed learning rate parameter governing changes in associative strength. When the associability of each cue was allowed to vary to reflect that cue's predictiveness within the task, simulations were found that matched the general effect of greater learning about cues pretrained as good predictors of outcomes. Changes in associability were best handled by a mechanism using the absolute predictiveness of the cue (i.e. as determined by the prediction error on the trial), rather than by one based on a cue's relative predictiveness (i.e. as determined by its prediction error relative to the context).

A common characteristic of the successful Mackintosh simulations was the use of a low starting level for α . Since the α values for both good and poor predictor cues would rise over the course of pretraining (see above), a low starting value prevented α from reaching asymptotic values and as a result becoming unable to modulate learning effectively. Although this method of allowing α to rise from a low starting level provides a simple computational mechanism, it is worth noting that it is perhaps not completely in keeping with the empirical evidence on changes in cue-processing. For instance, any attentional account of latent inhibition and learned irrelevance (see Chapter 1) would seem to require α to start high, before subsequently declining as a result of pretraining. Whilst latent inhibition and learned irrelevance are perhaps obvious examples of a need for a high starting value of α , it has been suggested that perhaps all empirical phenomena that provide evidence for a change in cue-processing might be accounted for by a decrease in α (Le Pelley, 2004).

The adjustments required to bring the current mechanism in line with the empirical phenomena from animal conditioning are fairly simple. For example, one could employ separate learning rate parameters governing changes in α , which could be set to ensure that decreases in α were of a greater magnitude than increases in α . Indeed, the use of different learning rates controlling increases and decreases in associative strength has precedent in models of animal conditioning (e.g. Pearce & Hall, 1980). Is it particularly troublesome that the most successful simulations presented here used a low starting value of α ? The current simulations offer an effective means by which the predictive history of a cue can influence future processing of that cue. Given the fact that any changes to a model (e.g. additional learning rates), are likely to increase the complexity of the model and possibly the number of free parameters used, one could argue that the current method is favourable as a simple model of the empirical phenomena it aims to account for.

The remainder of the chapter shall focus on a widely studied model of sequence learning, namely the Simple Recurrent Network (SRN), and examines whether this model also requires a mechanism for changes in cue-processing in order to model the current data. Whilst it is clear that a version of the Mackintosh (1975) model is able to model the empirical data satisfactorily, the aim here is to produce a general-purpose model of sequence learning that will also be able to simulate a variety of effects in the sequence learning literature – effects that are beyond the scope of the relatively simple associative models presented thus far.

Evidence for the acquisition of remote contingencies in sequence learning

Single layer associative learning models, such as the Rescorla-Wagner (1972) and Mackintosh (1975) models, were primarily designed to apply to experimental settings in which each cue-outcome pairing occurs as a discrete learning trial, presented in isolation from other cue-outcome pairings. In a sequence learning experiment, however, the location of the target can often be contingent on the sequence of target locations on a number of preceding trials. For example, Lewicki, Hill, and Bizot (1988; see also Lewicki, 1986; Lewicki, Czyzewska, & Hoffman, 1987, for similar results in non-SRT tasks) used a four-choice SRT task, in which the screen was divided into quadrants and the target would appear in one of these four locations on each trial. Each sequence block contained 48 segments of 5 trials. These 5-trial segments were designed such that the first two target locations were randomly chosen, whilst the last three target locations were dependent on the location of the target on the preceding two trials. For instance, if the movement of the target during the first two locations was vertical (e.g. from bottom left to top left) then the movement between the second and third trial would be diagonal (i.e. from top left to bottom right). The same rules were then applied to trials four and five. In order to anticipate the correct target location on sequenced trials, participants had to encode the relationship between the previous two trials. Lewicki et al. (1988) found that participants' RTs increased considerably when these rules were altered later in the experiment, providing evidence that participants had learnt two elements of the sequence context.

In a six-choice SRT task, Cleeremans and McClelland (1991; see also Remillard & Clarke, 2001) used a probabilistic grammar to generate sequenced blocks (see Figure 4.6). The letters S to X refer to the six stimulus positions on the screen. Sequences are generated by starting at node 0 and then traversing between nodes of the grammar (an example of a

five element sequence is V-X-Q-X-P-S). Where two possible routes are available, a route was chosen at random, and the sequence looped between the two instances of node 0, such that sequences could be of infinite length. On occasion, a sequenced location was replaced with a random, unsequenced location.




Figure 4.6. The artificial grammar used by Cleeremans and McClelland (1991).

Cleeremans and McClelland replicated the common sequence learning effect that participants were faster to respond on sequenced than unsequenced trials. Furthermore, the authors examined in more detail the complex temporal contingencies that participants had become sensitive to during training with sequences generated from this grammar. For instance, consider transitions XTV, PTV and QTV, noting that each sequence run has a different initial element but the same final two elements. In addition, note that the TV transition in XTV is generated by a different part of the grammatical structure than the TV

transition in PTV and QTV. All three transitions predict, with a .5 probability, the occurrence of X on the next trial. However, they make different predictions with regard to the other legal transition. Thus, in order for a participant to respond optimally after each sub-sequence, the 3rd order contingency must be learned. If participants encode the entire three item context in each case, then location S will be anticipated after the sequence XTV, whilst T will be anticipated after PTV and QTV. Cleeremans and McClelland showed that participants were faster to respond to predictable transitions (e.g. S after XTV; T after PTV; T after QTV) than unpredictable transitions (e.g. T after XTV; S after PTV; S after QTV). In a subsequent experiment, Cleeremans and McClelland examined whether participants were able to learn 4th order contingencies. However, even after 62,000 trials of the SRT task there was no evidence for any learning of this kind.

To summarise, the data from several studies show that training on the SRT task can lead to the acquisition of higher-order contingencies. Responses to the target on a given trial are primed not only by the location of the target on the previous trial, but also by several preceding trials (i.e. a context of up to three trials). Any model of sequence learning must, therefore, be able to learn such high-level contingencies within the sequence and make context-relevant responses.

The Simple Recurrent Network

The general architecture of the Simple Recurrent Network (SRN) was first developed by Elman (1990), who demonstrated its capacity for processing sequential material. However, work by Cleeremans and colleagues (Cleeremans, Servan-Schreiber, & McClelland, 1989; Cleeremans & McClelland, 1991; Cleeremans, 1993) established the current name, and more importantly began to test the model on findings from the field of implicit learning.

The SRN was developed within the Parallel Distributed Processing (PDP) framework, a computational method that has since become synonymous with ‘Connectionism’. In the mid-eighties the work of the PDP research group – led primarily by Jay McClelland and David Rumelhart – laid down the fundamental principles of modern connectionism, which led to the development of long-standing models of human behaviour. The basic principle of the PDP or connectionist framework is that fundamental cognitive processes can be conceptualized as the parallel activation and interaction of a network of computational units. A multi-layer network is shown in Figure 4.7. The network receives input via the ‘Input units’ (I_1 to I_6) in the form of an input vector of activation. Activation on the input units tends to represent the physical properties of relevant stimuli within the environment (though the detail of exactly how such stimulation is perceived is rarely considered a concern of the network).




Figure 4.7. *Left panel:* A two-layer network. I_1 to I_6 represent input units. H_1 to H_3 represent hidden units. O_1 to O_6 represent output units. Solid lines represent weights between units. *Right panel:* The same network (input and output units and respective weights have been omitted for clarity) with the additional context layer feature of the SRN. C_1 to C_N represent context units. The 'copy' arrow represents the process of copying the activation values of the hidden units to the context units at the end of each time step.

The behaviour of the network is manifest through the pattern of activation produced on the output units (O_1 to O_6). Usually a decision rule is employed in order to decide the network's response on the basis of the output activation (see below). The output units receive activation directly from a set of hidden units. These hidden units themselves receive input directly from the input units, and modulate the level of activation that is passed forward to the output layer. Therefore, the hidden units do not represent a feature of the physical environment – they are not an analogue of stimulus input or response output. Instead, these hidden units reflect the internal representation of the network and allow a nonlinear recoding of the input activation, allowing the network to solve complex problems in which the input and output patterns of activation are less similar (e.g. the 'exclusive-or', and other 'parity' problems, see Rumelhart, Hinton & Williams, 1986). The number of hidden units used in a network is often set as a free parameter, with more units generally permitting a greater degree of flexibility within the model.

In the same manner as for the single-layer models described previously, weights within the network provide a means of modulating the level of activation passed between units. On each trial the activation on the output units (the network's prediction) is compared to the actual outcome on that trial. Error values across the output units are then used to adjust weights within the network by a process of 'back-propagation of error' (see below; Rumelhart, Hinton & Williams, 1986).

What sets the SRN apart from earlier PDP models is its use of a 'context loop', as shown in the right-hand panel of Figure 4.7. On each trial the pattern of activation across the hidden units is copied to the context units. The context units then act as additional input units by sending this pattern of activation back into the hidden units on the next time step. This recurrent feature allows the SRN to retain a level of internal activation across several

trials. Therefore, the activation on the context units is an amalgamation of the continuously changing pattern of activation across the hidden units. This provides the network with a history of the previous target locations in the sequence and ensures that target locations from all previous trials have an impact on the network's prediction on the current trial.

The current implementation of the SRN functioned in much the same way as that described by Elman (1990). On each trial, the activation of one input unit was set to 1 (the location of the target on the current trial) and all other input units were set to 0. The activation was then fed forward to the hidden units by calculating the sum of all products of input activations and their respective connection strengths with each hidden unit. Thus, for hidden unit h :

$$in_h = B_h + \sum_{i=1}^6 w_{hi} \cdot a_i \quad 4.7$$

where in_h is the input for hidden unit h ; B_h is the input from a continually active bias unit associated with hidden unit h ; w_{hi} is the weight of the connection between hidden unit h and input unit i ; a_i is the activation of input unit i . Note that input units here refer to all units providing input to the hidden units, which therefore includes input from the context units. This input is then transformed into an activation value for hidden unit h , by the 'logistic activation function' given in Rumelhart, Hinton and Williams (1986):

$$a_h = \frac{1}{1 + e^{-in_h}} \quad 4.8$$

The input to, and activation of, the output units is calculated in much the same way as that of the hidden units, such that for output unit o :

$$\text{in}_o = \sum_{h=1}^H w_{oh} \bullet a_h \quad 4.9$$

$$a_o = \frac{1}{1 + e^{-\text{in}_o}} \quad 4.10$$

where in_o is the input for output unit o ; w_{oh} is the weighted connection between output unit o and hidden unit h ; a_h is the activation of hidden unit h ; and H is the total number of hidden units in the network.

When applied to sequence learning, the target output on each trial is the next element in the sequence. As in the simulations with single-layer models, the accuracy of the model in predicting the next element in the sequence was calculated as the activation of the target output unit divided by the total activation of all output units; the Luce Choice Ratio (LCR).

Target values for ‘active’ and ‘inactive’ stimulus positions were set at .9 and .1, respectively. Although only one target stimulus position (output unit) is ‘active’ at any one time, values of 1 and 0 cannot be reached without infinitely large weights in a network with non-linear activation functions, and so effectively cannot be achieved (Rumelhart, Hinton & Williams, 1986).

As is the case for single-layer networks, following each response made by the network, an error signal is calculated which reflects the discrepancy between the network output and the desired ‘target’ output (i.e. the error in the network’s prediction):

$$\delta_o = (t_o - a_o) \bullet (1 - a_o) \bullet a_o \quad 4.11$$

where δ_o refers to the error on output o .

In the case of multi-layer networks, this error is ‘back-propagated’ through the network to update the weights between each layer of units. In order to adjust weights

leading to the hidden layer, the network must be able to accurately attribute the error generated on the output units to the hidden units, such that those hidden units producing the greatest error on the output units undergo the largest weight changes:

$$\delta_h = \left(\sum_{o=1}^{o=6} \delta_o \bullet w_{oh} \right) \bullet (1 - a_h) \bullet a_h \quad 4.12$$

where δ_h refers to the error on hidden unit h.

These errors are then used to update the weights and biases in the network using the ‘generalised delta rule’ (Rumelhart, Hinton & Williams, 1986):

$$\Delta w_{oh} = \theta_{oh} \bullet \delta_o \bullet a_h \quad 4.13$$

$$\Delta w_{hi} = \theta_{hi} \bullet \delta_h \bullet a_i \quad 4.14$$

$$\Delta w_{hc} = \theta_{hc} \bullet \delta_h \bullet a_c \quad 4.15$$

$$\Delta B_h = \theta_b \bullet \delta_h \quad 4.16$$

where θ denotes a learning rate parameter, a_c is the activation of context unit c, and w_{hc} is the weighted connection between context unit c and hidden unit h. Note that in this implementation of the model, independent learning rate parameters were used for the three sets of weights, and for the hidden unit biases (θ_{oh} , θ_{hi} , θ_{hc} , and θ_b). Given that the aim of these simulations was simply to establish whether a “standard” SRN model (that is, an SRN that does not contain an additional associability parameter) could feasibly capture the effects of predictive history observed empirically, it seemed sensible to specify as few constraints on the model as possible. Hence if it is possible, by whatever means, for this model to account for the observed effects of predictive history, then this approach of providing the greatest possible flexibility gives the best chance of detecting this ability.

The implementation of different learning rate parameters also allows for an examination of a specific hypothesis regarding the ability of multi-layered networks to model associability effects. In their discussion of an intradimensional shift effect in rats, Oswald et al. (2001) suggested that ‘attention-like’ effects could be modelled using a multi-layer network, without the need of a specific attentional component, provided one assumes that the weights between input and hidden units are subject to slower changes in strength than weights between hidden and output units.

Consider how this restraint on the network would impact on the performance of the network when applied to the current experimental procedures. In the case of good predictor cues, after consistent reinforced training on trials GPH during the pretraining phase, weights between good predictor input units and certain hidden units will be strong, and weights between these hidden units and respective output units will also be strong. However, in the case of poor predictors, for which the outcome on trial PPM is less consistent, the weights between input and hidden units, and hidden and output units, will be less strong by the end of the pretraining phase. That is to say, after pretraining the network will develop stronger associative connections between those input units representing good predictor cues and their respective outcomes, than between input units representing poor predictor cues and their respective outcomes. When the network is then presented with the training phase sequences, it is able to benefit from these existing associative pathways, in particular by using existing connections between input and hidden units for the good predictor cues. If the network maintains existing strong connections between these input units and hidden units, whilst simultaneously changing weights between hidden and output units rapidly, a ‘rewiring’ of the associative connection to new output units could feasibly be achieved, thereby facilitating rapid acquisition of the new training phase contingencies (see Figure 4.8). Since

poor predictor cues will not have these strong existing connections between input and hidden units, this rewiring process will not be so effective. By using independent learning rates for each set of weights within the network, this hypothesis can be tested directly. That is, on the basis of a 'rewiring' hypothesis, better performance in the SRN would be expected when the learning rate governing weight changes in the hidden to output layer is higher than that governing weight changes in the input to hidden layer.




Figure 4.8. An illustration of how the associative connections between hidden and output units might be ‘rewired’ in the SRN. The high contingency between the good predictor cue (GP) and outcome 1 (O_1) during the pretraining phase will lead to strong connections between these units, mediated by a hidden unit H_1 . During the subsequent training phase GP is paired with O_2 . If the learning rate for input-hidden connections is low, while that for hidden-output connections is high, this change in contingency can be learnt by a rewiring: the strong associative connection between GP and H_1 will remain largely intact (as extinction of this connection will be slow), while H_1 will rapidly develops a new connection to O_2 . Relatively low contingencies between the poor predictor cue (PP) and O_3 and O_5 during pretraining will produce weaker associations to the hidden layer, and hence PP will be less able to benefit from this rewiring process.

Simulations with the standard SRN

In searching for model configurations which would produce the pattern of data observed in Experiment 6, wide ranges of learning rate parameters (.01, .05, .1, .3, .5 or .7), and numbers of hidden units (3, 15, 30, or 50) were examined. Given that each of the four independent learning rates (see Equations 4.13 to 4.16) within the network could take one of six values, the simulation space produced 5,184 networks. As was the case for the single-layer network simulations, for each parameter set, the model was trained using the exact sequences given to the 15 experimental participants of Experiment 6. For each simulated participant, all weights within the network were initialised with random values between -.5 and .5.

Initial pilot simulations highlighted that the output generated by the SRN is far less stable than that observed for the single-layer models. Amongst other factors, this is likely to be due to the complexity of the network and the fact that the recurrent loop will, to a certain extent, preserve the activation caused by randomly initialized weights. In order to compensate for this variability in the network's predictions, the performance of each parameter set was taken to be the average performance across 100 separate simulated experiments. Thus, each parameter set was effectively run with 1,500 simulated participants.

Each parameter set was first examined on its performance with the pretraining contingencies. If the network produced the ordinal relationship shown in the empirical data (i.e. higher Luce Ratio for GPH than for PPM, and higher Luce Ratio for PPM than for GPL), then it was considered to have successfully learnt the pretraining contingencies. In total 237 of the 5,184 parameter sets (4.5%) failed to meet this criterion and were excluded from further analysis. The parameter sets that failed tended to be those with only 3 hidden

units, or low learning rate values (e.g. $\leq .01$) for weight changes between the input and hidden layer, and between the hidden and output layer.

For each of the remaining 4,947 sets, learning scores (LCR on high probability outcomes minus LCR on low probability outcomes) for good and poor predictor cues were compared across the two blocks of the training phase. This revealed that 1,707 parameter sets mirrored the empirical data, showing greater learning about good predictor cues than poor predictors (i.e. a score of greater than zero using Equation 4.5).

Although this seems to suggest that the SRN is potentially able to produce the observed effect of prior predictiveness in the training phase, many of these simulations resulted in a weak effect. In order to examine the performance of the model more thoroughly, a t-test was conducted between learning scores for good and poor predictor cues during the training phase, for each simulated experiment with 15 participants. The mean t-statistic was then calculated from the 100 simulated experiments for each parameter set. This revealed that 395 parameter sets (8% of the total number of parameter sets producing the correct pretraining ordinal relationships) produced a mean t-value of greater than 1, whilst only 50 parameter sets (1%) produced a t-value of greater than 2. Given the critical t-value for a paired test with 15 participants is 2.145, the performance of the standard SRN is severely limited in terms of the range of parameters which produce a robust training phase effect.

Further analysis explored the means by which the SRN was able to produce the training phase effect of prior predictiveness in this minority of “successful” simulations. It was argued earlier that a rewiring process might produce a learning advantage for GP cues over PP cues, and that this process would be most likely to operate when the learning rate

for hidden-output weights (θ_{oh}) is higher than that for input-hidden weights (θ_{hi}). Consistent with this hypothesis, for those simulations producing a t-value of greater than 1, there was a significant positive correlation between the size of the t-value, and the difference in these learning rates (i.e. $\theta_{oh} - \theta_{hi}$), $r(395) = .40$, $p < .001$. This correlation is illustrated in Figure 4.9, which plots the t-value for the training phase effect against the difference in learning rates ($\theta_{oh} - \theta_{hi}$), for all parameter sets that produced a t-value of greater than 1.




Figure 4.9. The magnitude of training phase effect as a function of the combination of learning rates used, for simulations producing a t-value of greater than 1 in the direction of greater learning about cues pretrained as good predictors. Positive values on the learning rate difference scale are simulations which used a higher learning rate for weight changes between hidden and output units, than for weight changes between input and hidden units.

Every parameter set that produced a t-value of greater than 2 had a positive difference between these two learning rate parameters (i.e. θ_{oh} was greater than θ_{hi}). That is, all of the strongest effects of associability produced by the standard SRN used learning rates that would lead to faster changes to the hidden-output weights, and slower changes to the input-hidden weights.

In addition to the pattern observed in the learning rate parameters, all of the parameter sets producing a t-value of greater than 1 used 30 or more hidden units, whilst all of those producing a t-value of greater than 2 used 50 hidden units. This finding is also in-keeping with a rewiring hypothesis. With more hidden units in the network, each hidden unit is more likely to act as a ‘unique gateway unit’, linking a particular input unit to a particular output unit. Conversely, with fewer hidden units, each hidden unit is more likely to act as a ‘common gateway unit’, mediating activation from several strong connections between input and output units. In a network with common gateway units, appropriate new links between hidden and output units will be difficult to establish, since weight changes involving a particular hidden unit will occur not only on trials in which that outcome occurs, but also on trials on which that outcome does not occur; activation of the hidden unit will occur for several different contingencies, some of which will promote that weight between the hidden and output unit, and some of which will not. On the other hand, in a network with unique gateway units, activation of the hidden unit will only occur for a particular input-output pairing, which will therefore lead to rapid positive changes to the new weight between hidden and output units.

An example simulation, using the parameters that produced the most robust prior predictiveness effect, is shown in Figure 4.10. It is clear from the example that the effect of prior predictiveness is not particularly strong, although by Block 2 of the training phase

there is greater learning (performance on low-probability outcomes minus performance on high probability outcomes) about good predictor cues, over poor predictor cues. One striking feature of the example in Figure 4.10 is that the model's responses to outcomes following poor predictor cues are overall far less accurate than its responses to outcomes following good predictor cues. This main effect of cue (i.e. GPH + GPL vs. PPH + PPL) was indeed highly significant for all 100 simulations using this example parameter set, all $t_{s(14)} > 3.49$, $p_s < .01$, however, no main effect of cue was observed in any of the training phase results from Experiments 4 to 7. Specifically, in Experiment 6 participants were fastest to GPH trials, slowest on GPL trials, and were of intermediate speed on PPH and PPL trials.




Figure 4.10. An example simulation with the standard SRN ($\theta_{hi} = .3$; $\theta_{oh} = .5$; $\theta_{ch} = .01$; $\theta_b = .3$; 50 hidden units). See Figure 4.2 for an explanation of the trial types. *Left panel*: Model performance (1 - LCR) during the pretraining phase. *Right panel*: Model performance during the training phase, with novel cue-outcome pairings.

In order to assess the ability of the SRN to produce the correct ordinal relationships between the training phase trial-types, the empirical data from Experiment 6 (see Figure 4.2) were examined in more detail. Only some of the RTs to the four trial types were found to be significantly different: collapsing across the two blocks of the training phase, it was found that RTs on GPH outcomes were significantly faster than those on both PPL and GPL outcomes, whilst RTs on PPH outcomes were significantly faster than those on GPL outcomes, all $t(14)s > 2.62$, $ps < .05$. Therefore a parameter set was assessed as having produced the correct ordinal prediction overall if it made the correct ordinal predictions with respect to these differences considered individually (i.e. $GPH < PPL$, $GPH < GPL$, and $PPH < GPL$). Of the 395 parameter sets producing a t-value of greater than 1, a total of 213 produced the correct ordinal relationships shown empirically (53.9%, or 4.7% of the analysed parameter sets). Whilst for the 50 parameter sets showing a robust effect of $t > 2$, just 3 sets produced the correct ordinal relationships shown empirically (6%, or 0.1% of the analysed parameter sets).

To summarise, simulations were conducted with a standard implementation of the SRN, using a wide parameter search, with the results of each parameter set averaged across 100 separate simulations. The model was able to produce a prior predictiveness effect of sorts within a limited parameter range, and an analysis of these parameter values suggests that the SRN is able to achieve the best results by a process of rapidly rewiring connections between hidden and output units. However, of those simulations that produced a robust effect ($t > 2$), the vast majority (94%) were unable to accurately produce the correct ordinal relationship between the trial types seen empirically in Experiment 6.

Simulations with the Alpha SRN

It was clear from the simulations with the Rescorla-Wagner and Mackintosh models, that when the associability of a cue is allowed to vary as a function of its predictive history, a transfer effect of associability on new learning is more readily observed. It seems logical therefore, that an SRN which incorporates a cue-processing mechanism should also be better-equipped to capture the pattern of results observed in Experiment 6.

Since only input-to-hidden weights are directly connected to cues (input units) within the SRN, in order for a cue's associability to directly modulate the amount of learning that accrues to it, it is logical to allow associability to modulate weight changes only at this level. The learning rule for weight change between input and hidden units in the standard SRN (Equation 4.14) was modified to incorporate associability as follows:

$$\Delta w_{hi} = \alpha_i \bullet LR_{hi} \bullet \delta_h \bullet a_i \quad 4.17$$

where α_i is the associability of cue i . Recall that simulations with the Mackintosh model were far more successful when the equation for adjusting the associabilities of cues was based on the absolute predictiveness of that cue, rather than by means of relative predictiveness (compared to the experimental context). A mechanism based on absolute predictiveness was therefore also used in the 'Alpha SRN':

$$\begin{aligned} \text{If correct then } \alpha_i &= \alpha_i + (LCR)^4 \\ \text{If incorrect then } \alpha_i &= \alpha_i - (LCR)^4 \end{aligned} \quad 4.18$$

As for Equation 4.6, these rules specify that the associability of a cue increases if that cue predicts the correct outcome, and decreases if it predicts the incorrect outcome.

Values of α were allowed to vary between lower and upper limits of .1 and 1. Pilot simulations with this model found that raising the LCR to the fourth power results in small changes in α on each trial, ensuring a gradual approach to these limits. In addition to the parameters used in the simulations with the standard SRN, in the Alpha SRN simulations, various starting values of α (.1, .2, .3, .4, and .5) were also examined, yielding a parameter space of 25,920 sets. All other procedural and analytical aspects of these simulations were as for the standard SRN.

Of the 25,920 parameter sets, 2,168 failed to produce the ordinal pattern of results observed in the pretraining phase of Experiment 5, and were therefore excluded from further analysis. Of the remaining 23,752 parameter sets, a total of 12,144 produced a predictiveness effect of greater learning about good predictors than poor predictors in the training phase data. As in the simulations with the standard SRN, an average t-value was calculated from the 100 simulated experiments for each parameter set. A total of 4,467 (18.8%) simulations produced a t-value of greater than 1 in the direction of greater learning about good predictor cues, over twice the proportion produced by the standard SRN. The Alpha SRN was also far better at producing robust effects in this direction: 1,142 (4.8%) parameter sets produced an effect with a t of greater than 2, and 220 (0.9%) with $t > 3$.

Mean t-value	Proportion of parameter sets showing empirical predictiveness effect.		Proportion showing empirical ordinal pattern of data.	
	Standard	Alpha	Standard	Alpha
t > 1	8.0%	18.8%	4.3%	15.2%
t > 2	1.0%	4.8%	0.1%	4.3%
t > 3	0	0.9%	0	0.9%

Table 4.1. Simulation results for the Standard SRN and the Alpha SRN broken down by the robustness of the produced effect. Proportions are of the total number of parameter sets producing the ordinal pattern of data observed in the pretraining phase of Experiment 6 (SRN: 4,947; Alpha-SRN: 23,752). t-values are the result of a comparison of the learning scores in the training phase (LCR for high probability outcomes minus LCR for low probability outcomes) for good and poor predictor cues, across the 15 participants of each simulated experiment. Mean t values are an average of the 100 simulated experiments for each parameter set.

Table 4.1 shows a summary of simulation results from the standard SRN and the Alpha SRN. It is clear that many more parameter sets resulted in robust effects in the Alpha SRN simulations than in the standard SRN simulations. Moreover, the Alpha SRN produces far more simulation results matching the ordinal relationships between the four trial-types, as observed in Experiments 6 (see above). Whilst it was the case for the standard SRN that the most robust prior predictiveness effects rarely reproduced the correct ordinal relationships between the training phase trial types, for the Alpha SRN, the converse was true. In fact, for the 220 parameter sets producing a t-value of greater than 3 – the strongest effects observed with the Alpha SRN – all 220 parameter sets produced the correct ordinal relationships between the trial types.

Figure 4.11 shows an example simulation result from the Alpha SRN. As was the case in the simulation results from the Mackintosh model, using a measure of absolute predictiveness to drive changes in associability ensures that the associabilities of poor predictor cues rise more gradually across the pretraining phase (see the discussion of the Mackintosh model for a discussion of why associability increases at all for poor predictor cues). By the end of the pretraining phase there is a clear differentiation in the associabilities of good and poor predictor cues. This differentiation in the associabilities clearly leads to facilitation in the learning about novel contingencies pertaining to good predictor cues, over those pertaining to poor predictor cues.

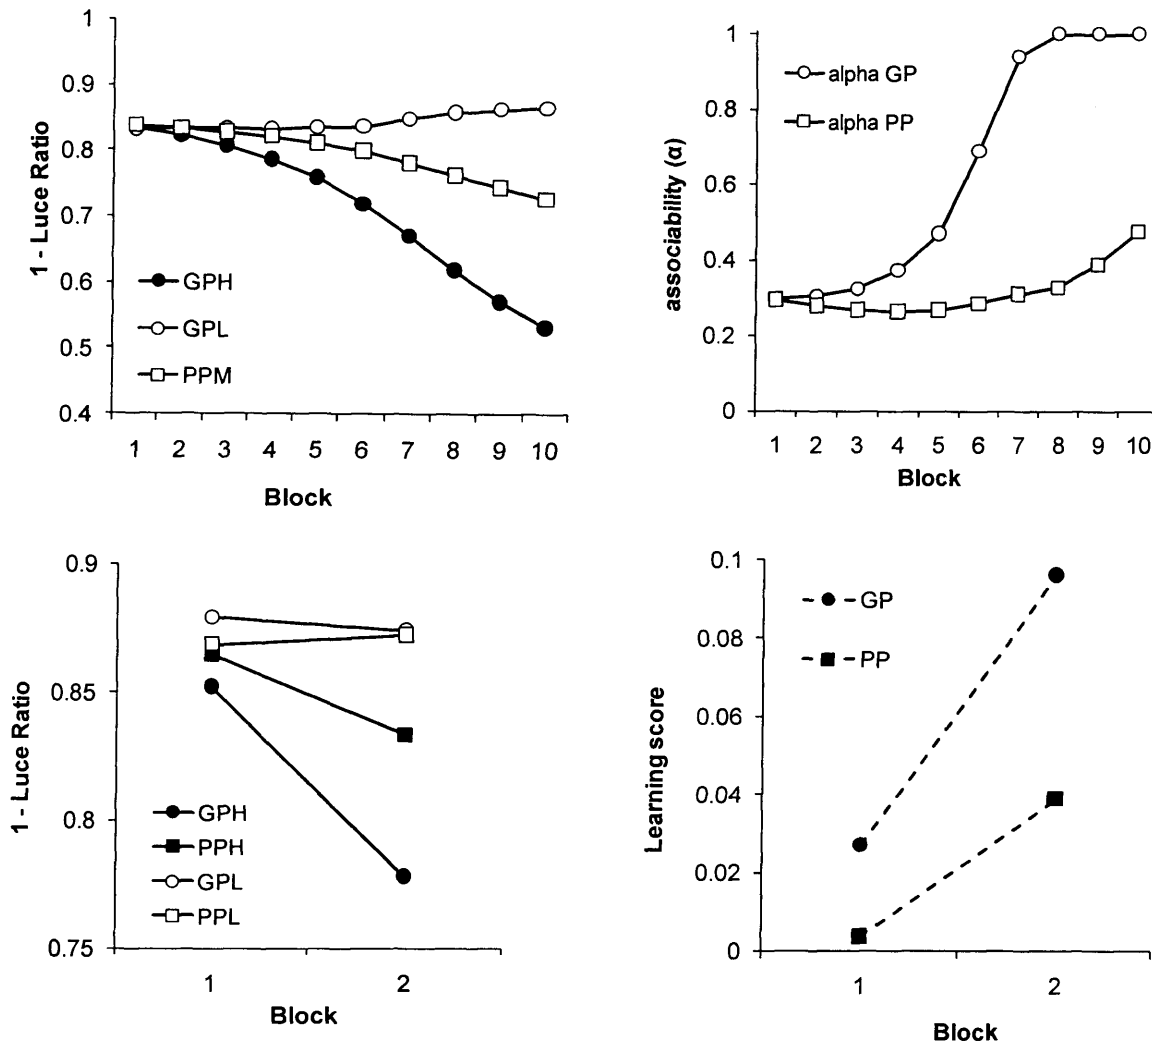


Figure 4.11. An example simulation with the Alpha SRN ($\theta_{hi} = .7$; $\theta_{oh} = .3$; $\theta_{ch} = .1$; $\theta_b = .1$; starting value of $\alpha = .3$; 15 hidden units). See Figure 4.2 for an explanation of trial types.

Top Left panel: Model performance (1 - LCR) during the pretraining phase. *Top Right panel:* Average associability value for good predictor (GP) and poor predictor (PP) cues during the pretraining phase. *Bottom left panel:* Model performance during the training phase, with novel cue-outcome pairings. *Bottom right panel:* The training phase data presented as learning scores (performance on low-probability outcomes minus performance on high probability outcomes) for cues pretrained as good (GP) and poor (PP) predictors.

Although the Alpha SRN performs considerably better than the standard SRN, the model fails to predict greater learning about good predictors than poor predictors during the training phase for almost half of the parameter sets tested. Further analysis revealed that these failures are largely a consequence of undifferentiated α values at the end of the pretraining phase. As was the case with the Mackintosh simulations employing an associability mechanism based on absolute predictiveness, there was a significant positive correlation between the difference in α at the end of the pretraining phase (α for good predictor cues minus α for poor predictor cues) and the size of the training phase effect, $r(23752) = .57, p < .001$. This correlation is shown in Figure 4.12. This difference is particularly striking if one compares those parameter sets showing an effect in the direction of greater learning about good predictors than poor predictors (the 4,467 sets with $t > 1$), and those showing an effect in the opposite direction (the 983 sets with $t < -1$). For the former, the average difference between the mean α value for good predictor and poor predictor cues at the end of Stage 1 was .34, whilst for the latter it was only .05. Thus the simulations that failed to show the observed training phase effect tended to be those in which α for good predictors failed to rise above those of poor predictors during the pretraining phase.

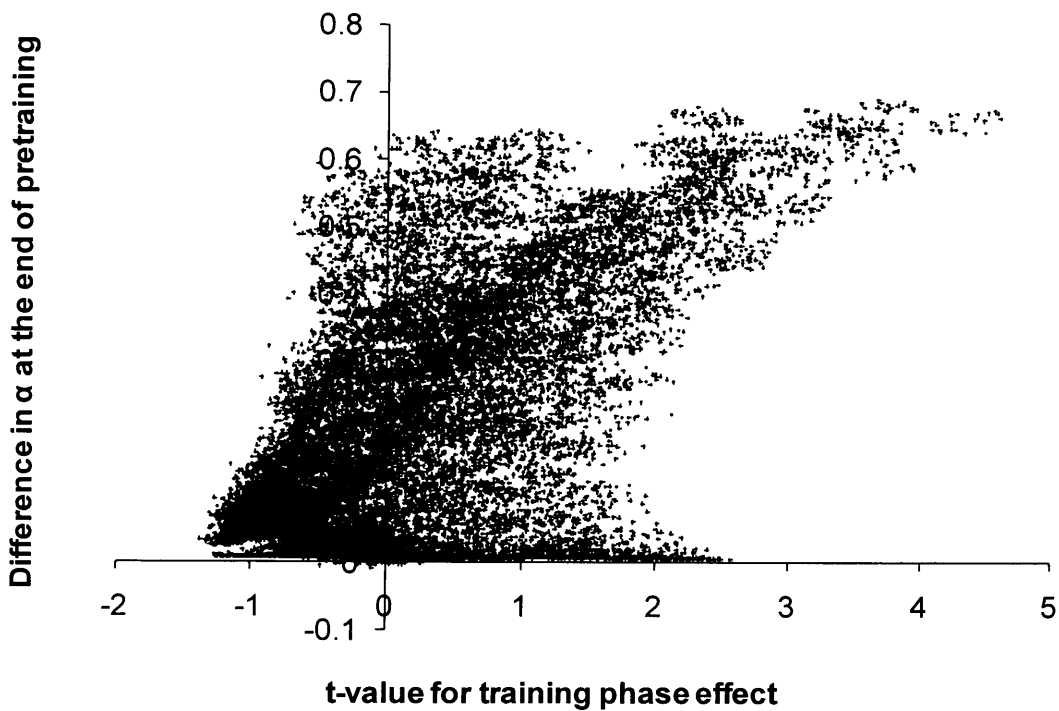


Figure 4.12. The magnitude of training phase effect as a function of the difference in α for good and poor predictor cues at the end of pretraining, for simulations with the Alpha SRN. Positive values on the α difference scale reflect a greater average α for good predictor cues over the average for poor predictor cues. Positive t values reflect greater learning about cues pretrained as good predictors over cues pretrained as poor predictors in the training phase.

There are several reasons why the α values for good and poor predictor cues might fail to differentiate during the pretraining phase, depending on the particular parameter set involved. For example, the learning rates might be sufficiently low that learning about the differential predictiveness of the two classes of cues is too weak to drive differences in their α values. In support of this, the magnitude of the learning effect during the pretraining phase (i.e. LCR on GPH trials minus LCR on PPM trials) was positively correlated with the average of the learning rate parameters used on the three sets of weights, $r(23752) = .562$, $p < .001$. In addition, parameter sets featuring very high learning rates will lead the α values for both good and poor predictor cues to rise to ceiling levels by the end of pretraining phase, thus compressing the difference between them. In support of this, there was a significant positive correlation between the average learning rate and the value of alpha at the end of the pretraining phase for both good predictor cues, $r(23752) = .445$, $p < .001$, and poor predictor cues, $r(23752) = .476$, $p < .001$.

Table 4.2 shows, for each starting value of α , the number of parameter sets showing an effect of greater learning of cues pretrained as good predictors with $t > 1$ in the training data, and of those the number that also predict the correct ordinal pattern between the trial types in the training phase (see above). It is clear that those parameter sets using high starting values of α are generally less successful in reproducing the empirical results of Experiment 6. There was a significant negative correlation between the starting value of α and the difference in α for good and poor predictor cues at the end of the pretraining phase, $r(23752) = -.15$, $p < .001$. This suggests that higher starting values of α were more likely to produce undifferentiated values of α at the end of pretraining, which presumably contributed to the poorer performance of these parameter sets in accounting for the prior predictiveness effect observed in the empirical training data.

Training phase results			
Starting value of α	Number (N) of parameter sets with correct pretraining ordinal predictions	N showing greater learning for good predictors with $t > 1$	N showing correct ordinal predictions when $t > 1$
.1	4,606	1036 (22.5%)	844 (18.3%)
.2	4,683	1041 (22.2%)	851 (18.2%)
.3	4,768	1006 (21.1%)	836 (17.5%)
.4	4,832	877 (18.1%)	719 (14.9%)
.5	4,862	507 (10.4%)	350 (7.2%)

Table 4.2. Simulation results for the Alpha SRN broken down by starting value of α . Values in parentheses show number of parameter sets as a proportion of the total number of parameter sets producing a correct pretraining ordinal prediction for each starting value of α (N). For a description of how t-values were calculated see Table 4.1.

In general, provided the parameters permit a clear differentiation in the α values for good and poor predictors by the end of the pretraining phase, the model will tend to predict faster learning for good predictor cues than poor predictor cues during the training phase. Equation 4.17 indicates that, once α values have differentiated, the learning rate for input–hidden unit connections (θ_{hi}) will play a key role in determining the magnitude of the preferential learning of good predictor cues over poor predictor cues. If θ_{hi} is high, then any differences in the associabilities of the cues will lead to a significant difference in the rate of learning about cues; if θ_{hi} is very low, then even quite large differences in α will only have a weak overall impact on learning. Consistent with this suggestion, an analysis of those parameter sets producing a differentiation in α values between good and poor predictors of greater than .1 (at the end of the pretraining phase), revealed a significant positive correlation between θ_{hi} and the magnitude of the learning advantage for good predictors during the training phase, $r(13236) = .542, p < .001$.

Summary of the simulation results

The purpose of the research presented in this chapter was to examine possible associative mechanisms governing changes in cue-processing, and to evaluate how successful mechanisms could be incorporated into a model of sequence learning. Firstly, several single-layer models were examined. A connectionist implementation of the Rescorla-Wagner model (Rescorla & Wagner, 1972) – one of the most extensively examined models of associative learning – was incapable of modelling the effects shown in Chapter 3. Simulations with a version of the Mackintosh (1975) model were far more successful, and provide a clear demonstration that the effects shown in Chapter 3 fit comfortably within an

associative learning model that permits changes in the cue-specific learning rate (α). A variety of mechanisms for permitting changes in α were considered; simulation results showed that determining α on the basis of the cue's absolute predictiveness (the extent to which the cue exclusively predicts the current outcome) provides a simple and effective mechanism.

An extensive simulation programme was conducted with the Simple Recurrent Network (SRN) – a connectionist architecture that has offered an accurate model of sequence learning phenomena. In an attempt to give the model the best chance of reproducing the empirical data, it was implemented in such a way as to permit a high degree of flexibility. Despite this flexibility, the SRN's ability to produce the empirical data was severely limited. The simulations reproducing the most robust effects tended to be those that attempted to 'rewire' the associative connections acquired during the pretraining phase in order to facilitate the learning of the training phase contingencies. Whilst this is, in principle, one means by which effects of prior predictiveness might be produced, the ordinal pattern of data that emerged from this process was generally quite different to that observed in Experiment 6.

Simulations were then conducted with a modified version of the SRN (the Alpha SRN), in which the learning rate (α) used to modulate weights between input and hidden units was allowed to vary across training. On the basis of the successful simulations with the Mackintosh model, α was determined by the absolute predictiveness of the cue. When the simulation parameters enabled a differentiation in the values of α for good and poor predictor cues by the end of the pretraining phase, the model readily produced an effect of prior predictiveness in the training phase. The proportion of parameter sets producing robust effects in the same direction as the observed data was four times that seen in the standard

SRN. Moreover, in each of the simulations that produced the strongest effects of prior predictiveness, the Alpha SRN also reproduced the ordinal pattern of data observed empirically.

Chapter 5 – Examining the attentional component to sequence learning by measuring eye gaze.

The relationship between associability and attention

The simulation results presented in Chapter 4 demonstrated that associative learning mechanisms incorporating changes in cue-specific learning rate parameters are able to capture the patterns of data observed in Chapter 3. Mackintosh (1975) proposed that associability reflects the attention that is paid to a cue: when the associabilities of certain cues increase relative to the associabilities of other cues, it is due to a change in how attention is being distributed between those cues. In essence, the Mackintosh model states that selective attention acts to filter out cues which are unhelpful in predicting events in the environment, in order to allow more processing resources to be devoted to the most useful stimuli. Although this has long been a fundamental assumption of models incorporating changes in associability (e.g. Mackintosh, 1975; Kruschke, 2001; Le Pelley, 2004) there have been few attempts to measure directly whether changes in associability reflect changes in attention.

One widely used measure of attention is overt eye gaze. Introspectively at least, we know there is a tendency to look towards the same spatial location as where our attention is directed. Of course, there are exceptions to this rule: it is possible to fixate on an object whilst attending to another location in space (e.g. Posner, 1980). Nevertheless, it seems reasonable to assume that a measure of overt eye gaze will correlate highly with attention and there is a wealth of support for the use of eye gaze as a measure of attention, much of which comes from research on reading (for a review see Rayner, 1998).

Support from more relevant research was provided by Rehder and Hoffman (2005a; see also 2005b), who used eye-tracking equipment to measure participants' eye gaze in a category learning task. Many theories of category learning have suggested that attention will be directed towards relevant dimensions that are diagnostic of category membership (e.g. Shepard, Hovland & Jenkins, 1961; Medin & Schaffer, 1978; Nosofsky, 1984). Yet, as I have suggested is the case for models of HCL and animal conditioning, demonstrations of such changes within the field of categorization had previously been limited to indirect measures of attention (e.g. generalization to new exemplars). Rehder and Hoffman (2005a) presented participants with stimuli consisting of three feature dimensions, and asked them to categorise these stimuli as belonging to one of two sets, providing feedback as to the correct category membership on each trial. The complexity of the category membership rule was manipulated, such that in one condition correct categorization of the stimuli required information from only one dimension, in another condition categorization required two dimensions, and in another all three. Categorization was easiest when a single dimension was diagnostic, whilst it was most difficult when all three dimensions were integral to classification decisions. In line with models of categorisation and attention (e.g. Shepard et al, 1961; Medin & Schaffer, 1978; Nosofsky, 1984), Rehder and Hoffman found that as participants learnt the categorisation rules, they fixated on the same number of dimensions as was required to solve the category discrimination. This result suggests that attention is directed towards those dimensions that are relevant to category membership, ignoring irrelevant dimensions as a consequence.

Krushke, Kappenman and Hetrick (2005; see also Wills, Lavric, Croft & Hodgson, 2007) used eye tracking to examine attentional changes resulting from blocking and highlighting preparations. In the highlighting procedure (see Chapter 1 for an explanation of

the blocking procedure) participants were first presented with a compound stimulus that predicts an outcome: AB – 1. In a second phase, participants were presented with an additional compound that predicts a new outcome: AC – 2, along with further AB – 1 presentations. In a subsequent causal rating phase, when presented with cue A alone, participants were more likely to rate it as a cause of outcome 1, than as a cause of outcome 2 (presumably reflecting the higher base rate for outcome 1). More importantly, participants were more likely to rate compound BC as causing outcome 2, than outcome 1; a result which indicates that the association between cue C and outcome 2 was stronger than the association between cue B and outcome 1. This is somewhat surprising, given that cue B had been paired more often with outcome 1 than cue C had with outcome 2. The highlighting effect can, however, be accommodated by models of selective attention. For instance, the Mackintosh (1975) model suggests that when compound AC is presented, cue A is a relatively poor predictor of this outcome (due to its previous pairing with outcome 1), and therefore selective attention leads to greater learning about cue C, to the detriment of learning about the contingency between cue A and outcome 2. Therefore, learning about cue C proceeds relatively unimpaired, whilst learning is more equally divided between cue B and cue A (overshadowing occurs).

Kruschke et al. (2005) found evidence in support of both blocking and highlighting effects in participants' causal ratings. Furthermore, attentional accounts of these results were bolstered by eye gaze duration data. In the case of blocking (A+ training, followed by AB+ and CD+), duration of eye gaze was significantly longer for control cues (D in CD) than for blocked cues (B in AB). These data suggest that learning about cue B is blocked because participants fail to devote attention to the blocked cue. Similarly in the case of highlighting, participants devoted more time to looking at cue C in compound AC, than they did for cue

B in compound AB. The results provide further support for the suggestion that the effect of highlighting is driven by selective attention.

To summarise, in several experimental conditions in which attentional processes have been assumed to play a prominent role, eye gaze has been shown to correlate highly with the expected patterns of attentional change. These data therefore offer strong support for the premise that eye gaze provides an accurate measure of attention.

Eye gaze as a measure of perceptual learning

It was mentioned briefly in Chapter 1, that some debate has arisen over the nature of the associations acquired during sequence learning. In very general terms the debate focuses on the extent to which the acquired knowledge can be best characterized as a series of motor responses, or a series of expected spatial locations. In the case of pure motor learning, participants are assumed to develop associations between contingent responses: response X is always followed by response Y. On the other hand, in the case of pure perceptual learning, participants are assumed to learn a sequence of spatial locations: after the stimulus has appeared in location X it next appears in location Y. By this account, sequence learning demonstrates an ability to shift attention towards an expected location in space, such that a rapid response can be made if the target appears in that position.

Of course, these two accounts of sequence learning are not mutually exclusive: it is possible that sequence learning is driven by both motoric and perceptual learning systems. I described briefly in Chapter 1 a study by Mayr (1996) which sought to examine the contribution of motoric and perceptual systems to sequence learning. In this study,

participants were given an SRT task in which the target could appear as one of 4 objects (i.e. black square, white square, black circle, white circle) and in one of 4 locations; responses were made to object identity only (e.g. black circle – C key). Participants were trained incidentally on both a sequence of target objects and a sequence of target locations, and these two sequences were independent of one another so target location could not be predicted on the basis of the target object sequence, and vice versa. Learning of each sequence was tested by separately replacing each sequence with random transitions. For example, learning of the object sequence was tested independently from learning of the location sequence by replacing the object sequence with random transitions, while maintaining the location sequence. Mayr found that RTs increased when each sequence was replaced with random transitions, which suggests that participants had learnt both a perceptual sequence of locations and a sequence of motor responses to the target object. The results conflict with those of a similar study by Willingham, Nissen and Bullemer (1989), who found evidence for only motor learning. Mayr argued that perceptual learning was not observed by Willingham et al. because they used stimulus positions which were too close together, thus rendering eye-movements too small for either perceptual learning to develop or to be measured.

Remillard (2003) provided further support for perceptual learning in an SRT task in which participants were trained with a single sequence. In this task, one of two bigrams (“XO” or “OX”) was presented in each of the six positions of the SRT task. One of the positions was then signalled as the target (a line appeared below it) and the participant had to detect which bigram appeared in that position (i.e. XO – left key; OX – right key). On each trial, the bigrams were pseudo-randomly assigned to locations, such that there was no consistent sequence of responses. However, the movement of the target position was

sequenced; successful learning of the target location sequence would result in anticipatory movements of attention (or eye gaze) towards the expected target location. Remillard hypothesised that this anticipation of the target location will then afford more processing time for the bigram in that position, and hence the correct response to that stimulus will be prepared and executed rapidly. In a series of these tasks Remillard demonstrated that perceptual learning can occur independently of motor learning in the SRT task for simple (e.g. first-order) but not complex sequences (e.g. second-order).

While studies such as those described above implicate perceptual learning in the SRT, it is not a particularly robust finding. For instance Deroost and Soetens (2006; see also, Rüsseler, Münte, & Rösler, 2002; Willingham, Nissen, & Bullemer, 1989) found evidence for perceptual learning in a between-subjects design with simple sequences, but not when these sequences were complex. However, when these complex sequences were used in a replication of Mayr's (1996) within-subjects experiment, a perceptual learning effect was observed. Deroost and Soetens suggested that perceptual learning only emerges in the absence of motor learning for simple, first-order sequences. More complex, second-order conditional sequences can be learnt perceptually, but this learning needs to proceed concurrently with motor learning.

Several studies have examined whether sequence learning can occur if the target sequence is simply observed and not responded to (e.g. Howard, Mutter, & Howard, 2002; Kelly & Burton, 2001). However, it seems under these conditions that explicit knowledge of the sequence is likely to develop, especially when simple sequences are used. The results therefore offer less convincing demonstrations of implicit perceptual learning than those offered with more complex sequences (e.g. Mayr, 1996), or with probabilistic sequence generation (e.g. Remillard, 2003).

Three studies have attempted to look at perceptual components to sequence learning by measuring eye gaze. Marcus, Karatekin, and Markiewicz (2006) trained participants in one of four conditions. Participants were either trained with a 10-item deterministic sequence, or with random target movement, and were asked to either observe or respond to the movement of the target. Eye gaze data suggested that anticipations of the next target location (i.e. when eye gaze moved to the target location prior to stimulus onset) were far greater for sequence-trained than for random-trained participants, but equally frequent in the observe and respond sequence learning conditions. These data suggest that eye gaze provides a sensitive measure of perceptual sequence learning, which can occur independently of motor performance.

Albouy et al. (2006) used a variant of the SRT task in which the sequence of target movements was not self-paced, but instead followed a fixed, regular timing. Despite the absence of motoric responses, Albouy et al. observed anticipatory saccadic eye movements to targets during sequence learning, and 'saccadic reaction times' increased when this sequence was altered. Data from a generation task suggested that very few participants had become consciously aware of the sequence, which Albouy et al. suggested was due to use of a dual-task procedure.

Kinder, Rolf and Kliegl (2008) used a variant of the SRT task in which target movement was controlled by the participants' eye movements. That is, participants were required to look at the target in order to invoke the next trial in the sequence. Kinder et al. argued that since the appropriate target responses were identical to response locations, stimulus-response associations in the task were already fully established in participants before the task began. In a 4-choice task, using SOC sequences, participants showed decreasing saccadic response times during sequenced training. When the sequence was

replaced with a novel sequence, saccadic response times increased significantly. Kinder et al. suggested that these results provide strong evidence for a perceptual basis to sequence learning, when stimulus-response associations are presumably fully established.

Rationale for the current experiment

The work presented in this chapter examines two main theoretical issues. Firstly, the results of the computational simulations (see Chapter 4) suggested that the sequence learning effects shown in Chapter 3 are in keeping with changes in cue-associability. I have described how models incorporating changes in cue-associability (e.g. Kruschke, 2001; Mackintosh, 1975) draw parallels between associability and attention. Thus, the question is: to what extent are changes in cue-associability in the SRT task a result of changes in the attention paid to certain cues within the task? Modifications made to the SRN to allow changes in cue-associability involved a measure of absolute, rather than relative predictiveness (see Chapter 4). This choice of mechanism reflects the fact that the SRT task involves the sequential presentation of individual stimuli. In this sense, it is difficult to imagine how a mechanism based on relative predictiveness would modulate learning in this task; it is unclear why the attention paid to cues pretrained as good predictors will limit the attention to cues pretrained as poor predictors. Nevertheless, Experiment 8 uses a modification of the design used in Experiment 5 to examine any potential changes in the attention paid to cues in the SRT task.

Secondly, as in the studies detailed above (e.g. Marcus et al., 2006; Albouy et al., 2006; Kinder et al., 2008) measuring eye gaze allows for an examination of perceptual learning in the SRT task. The current study is most closely related to the respond condition

used in the Marcus et al. (2006) study, in the sense that it employs a self-paced standard SRT task with manual responses. However, manipulating cue-predictiveness within a probabilistic sequence allows an examination of whether perceptual sequence learning reflects the different sequence contingencies, as shown in the RT data. That is, unlike in the Marcus et al. study which measured anticipations in sequenced versus unsequenced blocks, the current study will also examine whether anticipatory eye movements reflect the range of contingencies within a probabilistic sequence (e.g. for high, medium, and low probability outcomes). To the extent that eye gaze offers a sensitive measure of sequence learning, it should also be possible to examine whether the prior predictiveness effect seen in Experiment 5 will be mirrored in the eye gaze data. Marcus et al. (2006) also asked participants to track the target closely with their eyes: an instruction that is seldom given in the standard SRT task. As such, the current study also offers a more natural examination of eye gaze in the SRT task.

Experiment 8

Experiment 8 employed a similar design to that used in Experiment 5, but with some important changes. During the pretraining phase of Experiment 5, four stimulus positions functioned as cues, whilst two other positions functioned as outcome positions. The dedicated outcome positions were then removed from the screen for the training phase. This general approach was used again in the Experiment 8, although now the four cue positions were placed in the top half of the screen, whilst the two outcome positions were positioned in the bottom half of the screen (see Figure 5.1).

Two of the four cue positions functioned as good predictors of their respective outcomes during pretraining, whilst two functioned as poor predictors. The sequence contingencies are shown in Table 5.1. In order to keep the design compatible with that of Experiment 5, the four cue positions are numbered 2-5, whilst the two outcome locations are numbered 1 and 6. Each good predictor cue (locations 2 and 3 in Table 5.1) consistently predicted the appearance of the target in one of the outcome locations (with a probability of .9), whilst the two poor predictor cues (locations 4 and 5 in Table 5.1) were equally likely to be followed by the target appearing in either outcome position on the next trial. After the target had appeared in an outcome position, there was an equal chance of it appearing in any one of the cue positions on the next trial.




Figure 5.1. *Top panel:* A screenshot of the stimulus positions used in Experiment 8. The top four positions acted as cues (left to right: cues 1 to 4), whilst the two bottom positions acted as outcomes (left to right: outcomes 1 and 2). *Bottom panel:* A schematic of the 9-point calibration arrangement.

		Trial N (cue) – Pretraining						Trial N (cue) – Training						
		1	2	3	4	5	6	1	2	3	4	5	6	
Trial N+1 (outcome)	1		.9	.1	.5	.5								
	2	.25					.25			.1	.8	.1		
	3	.25					.25	.8		.1	.1			
	4	.25					.25	.1	.1		.8			
	5	.25					.25	.1	.8	.1				
	6		.1	.9	.5	.5								

Table 5.1. Transition contingencies for the sequences used during the pretraining and training phases in Experiment 8. Numbers 2-5 refer to the cue positions in the top half of the screen (see Figure 5.1), whilst numbers 1 and 6 refer to the outcome positions in the bottom half of screen. Outcome positions 1 and 6 did not appear in the training phase of the experiment. Blank cells indicate zero probability.

Unlike in previous experiments, the assignment of good and poor predictor cues was not randomised for each participant. Instead, the two good and two poor predictor cues were always positioned on the same side of the screen as each other. By setting the good and poor predictor cues on opposite sides of the screen it was possible to examine biases in attention. Consider the transition between outcome and cue positions. After the target appeared in an outcome position there was an equal chance of it appearing in any one of the four cue positions. As such, if attention is distributed evenly between the four cue position, eye gaze should, on average, be positioned in the horizontal centre of the screen during this transition. To put it another way, the most efficient strategy for preparing to respond to the target appearing in a cue position is to position gaze directly in the horizontal mid-point of the screen, between cues 2 and 3. Any deviation away from this mid-point would suggest that participants are devoting more attention towards one half of the screen than the other. Therefore, by placing the good and poor predictor cues on opposite sides of the screen it is possible to measure changes in attention towards these cues. If the associability effects observed in Experiments 4-7 are as a result of changes in attention to good predictor cues, given the arrangement of cues in the current experiment, attention would be expected to move towards the side of the screen on which good predictor cues are located¹⁰.

¹⁰ Although such a manipulation would be possible with the standard horizontal arrangement of stimulus positions (i.e. that used in Experiment 5), one can imagine that the large distances between outcome and cue positions used in Experiment 8 should accentuate any shift in attention that might be observed in the standard arrangement (for a similar argument see Mayr, 1996).

Method

Participants, Apparatus and Stimuli

Eighteen Cardiff University students participated for payment. A screenshot of the arrangement of stimulus positions is shown in Figure 5.1. The four positions in the top half of the screen acted as cue positions in the task and were responded to using the keys P, O, T, and R. The two positions in the bottom half of the screen acted as outcome positions in the task and were responded to using the keys F8 and F5. The keyboard was turned 180° so that responses to keys F5 and F8 were made with the right and left thumbs respectively.

Participants were asked to respond to keys P and O with the middle and index fingers of the left hand respectively, and to keys T and R with the index and middle fingers of the right hand respectively. The arrangement of keys closely resembled the arrangement of the stimulus positions. Error signals were issued by computer speakers.

The experiment was conducted using a Tobii 1750 Eye Tracker (Tobii technology) – a monitor-mounted eye tracker capable of recording eye gaze at a resolution of 50Hz (i.e. a recording every 20 milliseconds). The tracker functions by emitting infra-red light towards the eye and recording changes in the pattern of reflection. It offers a non-intrusive method of eye tracking and is able to compensate for small head movements.

Sequence generation

The sequence-generation procedure for the pretraining phase was identical to that used in Experiment 5 and the sequence contingencies are presented in Table 5.1. During this

phase, two cue locations were good predictors of their respective outcomes, whilst two cue locations were poor predictors of their respective outcomes. Good and poor predictor cues were always positioned on the same side of the screen. However, the assignment of cues to positions was counterbalanced, such that for half of the participants, good predictor cues were in positions 1 and 2, and for the other half good predictor cues were in positions 3 and 4. Counterbalancing across participants in this way ensures that any potential bias in eye gaze cannot simply be due to a general, between-subjects bias towards one side of the screen over the other. Within the set of good predictor cues, the contingencies between cues and outcomes were also counterbalanced. For example, for participants whose good predictor cues were on the right, half had cue 1 predicting outcome 1, whilst the other half had cue 1 predicting outcome 2.

The sequence generation procedure for the training phase was similar to that used in Experiment 5, with slightly different pairings between locations. Since the good and poor predictor locations were always presented on opposite sides (i.e. there was less randomisation in the assignment of design elements to cue locations), the contingencies used in Experiment 5 would have led to the salient sequence “2-3-4-5” being used as a run of high-probability transitions for all participants. In the current design, the somewhat less salient sequence “2-3-5-4” was used instead. Note that the average distance between a cue and its high-probability outcome was equivalent for all cues (1.5 locations), regardless of whether the good predictor cues occupied the left- or the right-hand side of the screen.

Procedure

Participants were given the same instructions as in Experiments 5, with additional information about the calibration procedure for the eye tracker. Each participant was asked

to keep their head as still as possible during the experiment. The eye tracker was then calibrated with a 9-point calibration procedure (see Figure 5.1). A new calibration test was conducted when needed, or no more than 3 blocks after the previous calibration. All participants received a calibration before the last two blocks of the experiment (the training phase). Participants were tested individually and the experimenter remained in the testing room throughout the experiment. The addition of several calibrations throughout the experiment increased the time required to run the experiment. As such, the pretraining phase was shortened to eight blocks (compared to ten in Experiment 5) in order to keep the length of the study to around 45 minutes.

Before the training stage commenced (the last two blocks of the experiment) the two outcome positions were removed from the screen. The experimenter informed the participant of this change and that the target stimulus would now only appear in the remaining four positions on the screen. They were also told that they should continue to use the same four keys to respond to these positions.

Results

Data were excluded on the same basis as in Experiment 5. During the eye gaze recording process the Tobii eye tracker is inevitably unable to detect the location of gaze on occasion. Often this is due to eye blinks, or because the eye tracker has difficulty recording for participants with low eyelids or long eye lashes. For each participant, the proportion of missing data for each eye was calculated and the data from the best eye was chosen for further analysis. For one participant, the proportion of missing data was above 40% for each eye, and therefore this participant was excluded from further analysis. The average

proportion of errors on the best eye for the remaining seventeen participants was 3.5% (standard deviation of 4.4%).

A full analysis of the accuracy data was conducted, but will not be reported here for two reasons. Firstly, accuracy was consistently high for participants in this experiment: mean accuracy was 97.6% (standard deviation of 1.5%). This high level of accuracy was most likely due to participants being aware of the presence of the experimenter in the testing room, who was able to hear the error signals from the computer speakers. Such high levels of accuracy will inevitably limit the chance of observing statistical differences due to ceiling effects. Secondly, the findings of the statistical analyses matched those reported in Experiment 5: there were no significant differences in accuracy to the three trial-types during pretraining; there was a main effect of outcome probability during the training phase, but no effect of prior predictiveness, nor any significant interactions.

Cue-outcome contingency learning during the pretraining phase

Missing RT data for the low-probability outcomes were replaced using the technique described in Experiment 5. The RT data for outcome stimuli during the pretraining phase are shown in the top panel of Figure 5.2. The overall pattern of results is similar to that seen in Experiments 5-7: participants were fastest on high-probability outcomes following good predictor cues (GPH), slowest on low-probability outcomes following good predictor cues (GPL), and were of intermediate speed on medium-probability outcomes following poor predictor cues (PPM). A repeated measures ANOVA revealed a significant effect of outcome probability, $F(2, 32) = 27.40, p < .001$, a significant effect of block, $F(7, 112) = 8.39, p < .001$, and a significant interaction, $F(14, 224) = 1.98, p < .05$. Pairwise

comparisons between the three levels of the outcome-probability variable revealed significant differences between all three trial types: GPH vs. GPL, $F(1, 16) = 35.86$, $p < .001$; GPH vs. PPM, $F(1, 16) = 15.02$, $p < .01$; GPL vs. PPM, $F(1, 16) = 20.73$, $p < .001$.

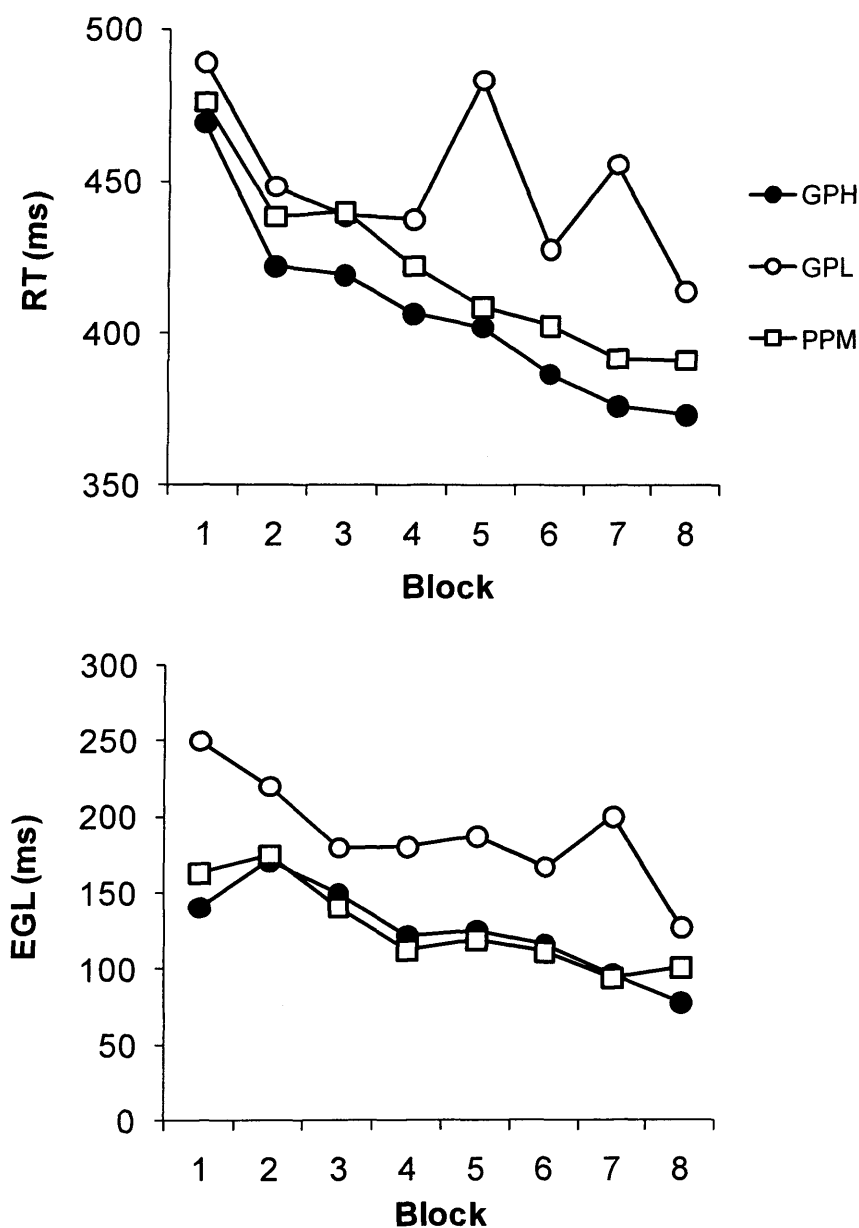


Figure 5.2. Data from the pretraining phase of Experiment 8. *Top panel*: RTs to high-probability outcomes following good cues (GPH), low-probability outcomes following good predictor cues (GPL), and medium-probability outcomes following poor predictor cues (PPM). *Bottom panel*: Eye gaze latencies (EGLs) to the same trial types.

In addition to the measure of learning provided by RT, it is also possible to conduct a complementary analysis with the eye gaze data, by measuring the time taken for the participants' gaze to reach the target location square (see Figure 5.1) after stimulus onset. However, on average, participants' eye gaze reached the target location on only 64.8% of trials. There was also a considerable amount of variation across participants in the percentage of trials on which eye gaze reached the target location: highest – 97.6%; lowest – 5.7%; standard deviation – 30.5%. Thus, it was impossible to include certain participants in the following statistical analyses due to a lack of data on the EGL measure. Five participants were excluded due to eye gaze failing to reach the target on more than 50% of all trials. For the remaining twelve participants, missing EGL data were replaced using the same technique as for RTs (i.e. an average of the adjacent data).

The bottom panel of Figure 5.2 shows the eye gaze latencies (EGLs) to outcome stimuli during the pretraining phase. The EGLs to outcomes were subjected to repeated measures ANOVA with factors of outcome-probability and block. This revealed a main effect of outcome probability, $F(2, 22) = 19.22, p < .001$, indicating that EGLs were affected by the predictability of the outcome. There was also a main effect of block, $F(7, 77) = 3.34, p < .01$, indicating that EGLs were shorter at the end of the pretraining phase than they were at the start. There was no interaction between outcome probability and block, $F < 1$. Pairwise comparisons between the three levels of the outcome-probability variable revealed that EGLs were significantly shorter for high-probability outcomes following good predictor cues than for low-probability outcomes following good predictor cues (GPH vs. GPL), $F(1, 11) = 27.09, p < .001$. EGLs were also significantly shorter for medium-probability outcomes following poor predictor cues than for low probability outcomes following good predictor cues (PPM vs. GPL), $F(1, 11) = 18.79, p < .01$. However, there was no difference

in EGLs between high-probability outcomes following good predictor cues and medium probability outcomes following poor predictor cues (GPH vs. PPM), $F < 1$.

Anticipatory eye-movements towards outcome locations during the pretraining phase

It was clear from the EGL data that some participants responded to the target without first fixating on the target location. This strategy of relying on peripheral vision to detect target location led to many instances of missing data in the EGL measure. One method of compensating for this strategy, whilst simultaneously providing a measure of target anticipations, is to measure the location of eye gaze at various stages before and after stimulus onset. Following the appearance of the target in a poor predictor cue location, it is impossible for participants to accurately predict in which outcome location the target will appear. Hence participants must wait for the target to appear in one of the outcome locations before they can initiate an eye movement towards that location. In this case, the time to make an appropriate eye movement towards the outcome location will reflect the time taken to perceive the target, and to plan and execute an eye movement. In contrast, following the appearance of a target in a good predictor cue location, provided participants have learnt the contingency between that cue and the predicted outcome, participants could feasibly make anticipatory movements towards the expected outcome location. Hence one would expect to see selective eye movements towards the correct outcome sooner after target onset on trials following good predictors than on trials following poor predictors (or possibly even prior to target onset following good predictors).

The difference in the horizontal coordinate of eye gaze between trials on which the target appeared in the left and right outcome positions (i.e. eye gaze position on right

outcome trials minus eye gaze position on left outcome trials) was used as a measure of the location of eye gaze. A positive score indicates that participants made an appropriate differential movement to the target in an outcome location: when the target appears in the outcome location on the right hand side of the screen the horizontal pixel coordinate is greater than when the target appears on the left hand side of the screen. A zero score on this measure reflects that no differential movement in eye gaze has been made to the two outcomes. It is clear that by 200 ms after the onset of the target in an outcome location, participants are making appropriate eye movements towards the side of the screen on which the target is located. However, during the early stages of the target onset, and even prior to target onset, these appropriate movements are only made following the target appearing in a good predictor location on the previous trial. That is, only after the target has appeared in a good predictor location are participants able to make early, or even anticipatory, eye movements towards the correct outcome location.




Figure 5.3. The difference in the horizontal coordinate of eye gaze between trials on which the target appeared in the left and right outcome positions (eye gaze position on right outcome trials minus eye gaze position on left outcome trials), as a function of the cue that preceded the outcome and also the time relative to stimulus onset. Positive values indicate movements towards the appropriate outcome. *Top panel:* Blocks 1 to 8. *Bottom panel:* Blocks 5 to 8.

The pixel difference data (plotted in Figure 5.3) were subjected to repeated measures ANOVA with factors of predictiveness of the cue (good and poor) and time (-100 to 0 ms; 0 to 100 ms; 100 to 200 ms; 200 to 300 ms). This revealed a significant main effect of cue, $F(1, 16) = 13.52, p < .01$, indicating that the difference in eye gaze position between trials on which the outcome occurred on the left and the right, was greater when the outcome appeared after a good predictor cue than when it appeared after a poor predictor cue. The main effect of time was also significant, $F(3, 48) = 42.25, p < .001$, indicating that differential movements were smaller immediately before or after the stimulus onset, and much greater once the stimulus had been on screen for some time (e.g. after 200 ms). The interaction between cue and time was not significant, $F(3, 48) = 1.42, p = .25$. In order to assess the time course of anticipatory eye movements following good and poor predictor cues, tests of simple main effects were conducted on each level of the time variable. This revealed significant differences in all three bins during the post-stimulus onset period (0 to 100 ms, 100 to 200 ms, and 200 to 300 ms), all $F_s(1, 16) > 10.91, p_s < .01$, however the difference during the pre-stimulus period (-100 to 0 ms) did not quite reach significance, $F(1, 16) = 3.59, p = .076$. One would expect that anticipatory eye movements towards the correct outcome location would only be made once participants had learnt the contingencies between cues and outcomes, thus we would expect these pixel differences to be more pronounced in the latter half of the pretraining phase. An analysis restricted to the last four blocks of the pretraining phase produced a similar pattern of statistical results (see bottom panel of Figure 5.3), but importantly the pre-stimulus-onset anticipatory movements were also significantly greater following good predictor cues than following poor predictor cues, $F(1, 16) = 7.39, p < .05$.

In order to assess the point at which participants started to make eye movements towards the target location, one sampled t-tests against zero were calculated for the data shown in Figure 5.3. For the data following good predictor cues (black bars), all means were significantly different from zero for the post-stimulus onset period (0 ms onwards), all $t(16) > 3.45$, $p < .01$, however the analysis with the data from the pre-stimulus onset period (-100 ms to 0 ms) was only marginally significant, $t(16) = 2.03$, $p = .060$. For the data following poor predictor cues (white bars), the data were not significantly different from zero for the first two bins (-100 to 0 and 0 to 100 ms), both $t(16) < 1.29$, $p > .22$. However, the data from 100 ms onwards were significantly greater than zero, both $t(16) > 2.57$, $p < .05$. An analysis restricted to the last four blocks of the pretraining phase found the same pattern of results, with the addition of a significant difference from zero for the pre-stimulus onset period (-100 to 0 ms) for outcomes following good predictor cues, $t(16) = 2.52$, $p < .05$. Thus, for predictable outcomes following good predictor cues, participants will make anticipatory movements towards the expected outcome, whilst when it is impossible to predict the outcome on the next trial, eye movements towards the target location occur at around 150 ms after target onset.

Changes in attention to cues during the pretraining phase

It was hypothesised that the effect of prior predictiveness observed in the training phase data of Experiments 4-7 might have resulted from participants learning to increase the allocation of attention to good predictor cues over poor predictor cues as pretraining progresses. It was reasoned that biases in attention could be assessed during the pretraining period by analysing eye gaze during the RSI period immediately following the target

appearing in an outcome location. Specifically, since all of the outcome-to-cue contingencies were equal (i.e. there was a probability of .25 of each cue occurring after the target has appeared in a particular outcome location; see Table 5.1), following the target appearing in an outcome location, biases in eye gaze towards one set of cues over another would indicate changes in the allocation of attention during the task.

As a starting point for assessing any attentional bias, the RTs and EGLs to cue locations were analysed and are shown in Figure 5.4. It might be expected that, if more attention is devoted to good predictor cues over poor predictor cues, RTs and EGLs to good predictor cue locations should be faster than those to poor predictor cue locations. The RT data were subjected to a repeated measures ANOVA with factors of block and cue (good versus poor), which revealed an effect of block, $F(7, 112) = 7.96$, $p < .001$, indicating that participants' speed in responding to the cue positions increased across training. There was, however, no effect of cue, nor an interaction between cue and block, both $F_s < 1$. As discussed above for the analysis of EGLs to the outcome stimuli, 5 participants were removed as their eye gaze failed to reach the target on more than 50% of the trials. The EGLs from the remaining 12 participants were subjected to an ANOVA, which found no main effects, nor any interaction between cue and block, all $F_s < 1$.

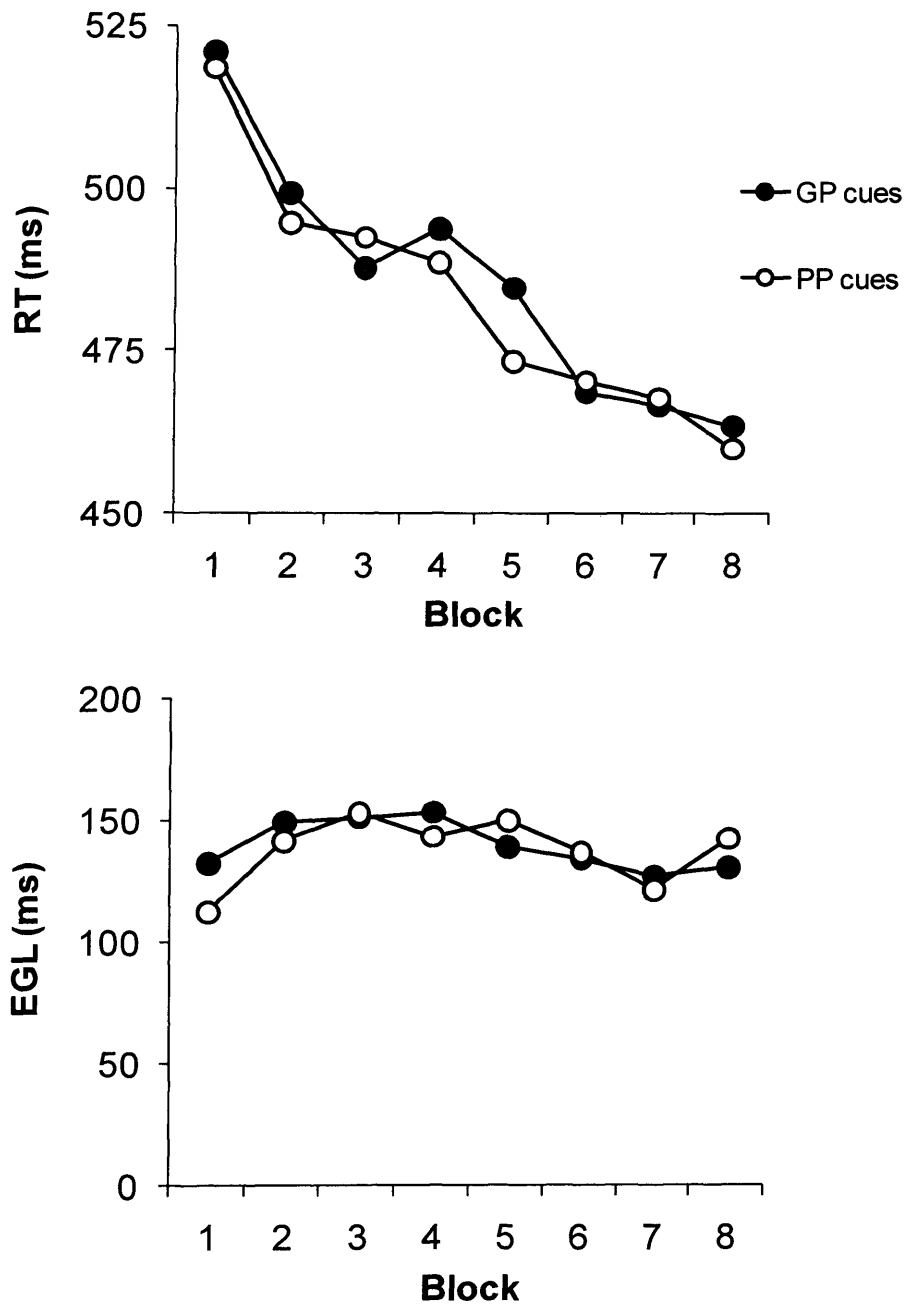


Figure 5.4. Data from the pretraining phase of Experiment 8. *Top panel*: RTs to cue locations pretrained as good-predictors of outcomes (GP) and to cue locations pretrained as poor predictors of outcomes (PP). *Bottom panel*: Eye gaze latencies (EGLs) to the same cue locations.

Although there was no evidence to suggest an attentional bias towards one set of cues over the other in the RT or EGL data, it is possible that attentional effects were quite subtle, and that the RT and EGL measures were insufficiently sensitive to detect a small change in the eye gaze brought about by a shift in attention. In order to further examine any possible attentional changes during the pretraining phase, the eye gaze positional data were analysed during the entire response-stimulus interval (RSI) following a response to a target in an outcome location.

Since the side on which the good predictor cues were positioned was counterbalanced across participants, the horizontal coordinates of the eye gaze data were standardised so that for all participants the data corresponded to the good predictor cues being located on the left hand side of the screen. The left hand panel of Figure 5.5 plots the mean horizontal coordinate of eye gaze during the RSI period preceding each trial on which the target appeared in a cue location, across the 8 blocks of the pretraining phase. The intersection of the vertical axis on the graph is anchored to the mean value from the first block of pretraining. The graph therefore plots the change in attention away from a baseline position early on in the pretraining phase, when the impact of contingency learning on attention should be minimal¹¹.

¹¹ An alternative method of analysis would be to assess changes in eye gaze away from the middle of the screen (by comparing data to the most central pixel). However, this analysis would fail to compensate for any general bias in the participants' eye gaze as a result of uncontrolled factors (e.g. due to calibration inaccuracies, stimulation from the background environment, etc.). Whilst it is possible to minimize the impact of these factors by counterbalancing the stimuli (see Method), only the employed method of using a baseline at Block 1 will allow an assessment of the development of bias on an individual participant level.

The trend in the data suggests that as pretraining progresses, eye gaze tends to move towards the good predictor cues. However, a one-way ANOVA found no main effect of block, $F(7, 112) = 1.54$, $p = .16$. Furthermore, a related samples t-test comparing the horizontal coordinate in Block 1 with an average of the horizontal coordinates for Blocks 2 to 8, failed to find a significant change in eye gaze position, $t(16) = 1.15$, $p = .27$. Further related samples t-tests, comparing the eye gaze during Block 1 with every other block, revealed a significant difference in Block 5, $t(16) = 2.16$, $p < .05$, but not in any other block, all $t(16) < 1.33$, $p > .20$. The difference in Block 5 is not significant when the α value is Bonferroni adjusted to compensate for multiple comparisons, an appropriate adjustment given there was no a priori reason for expecting a difference only in Block 5.




Figure 5.5. *Left panel:* Horizontal eye gaze data during the RSI following outcome stimuli, as a function of pretraining block. Negative values reflect a bias in eye gaze towards cues pretrained as good predictors, whilst positive values reflect a bias towards cues pretrained as poor predictors. *Right panel:* The mean standard deviation of horizontal eye gaze data, during the same period, as a function of pretraining block.

The left panel of Figure 5.5 indicates that eye gaze initially moved towards good predictor cues during the pretraining phase, but towards the end of pretraining eye gaze seemed to retreat back towards the anchor point of Block 1. One possibility is that, in general, participants developed a strategy for fixating in the middle of the screen as the pretraining phase progressed. As was suggested by the analysis of the eye gaze latency data, many participants began to use their peripheral vision in order to detect the target (in fact at the end of the experiment, several participants verbally reported that they had adopted this behaviour during the task). It is therefore possible that the movement back towards the centre point during blocks 6 to 8 is due to a general reduction in the horizontal eye gaze movement as the experiment progressed. This interpretation of the data was supported by an analysis of the horizontal gaze data during the outcome-to-cue RSI. The right hand panel of Figure 5.5 shows the mean standard deviation of this horizontal gaze data across the 8 blocks of the pretraining period for the outcome-to-cue RSI. A one-way repeated measures ANOVA revealed a significant effect of block, $F(7, 112) = 2.82, p < .05$, with a significant linear trend, $F(1, 16) = 15.52, p < .01$. This finding suggests that the horizontal movement in eye gaze during the outcome-cue RSI became more restricted as the pretraining phase progressed.

One possibility that follows from the above analysis is that as eye gaze movement became more restricted during the course of pretraining, there was less scope for observing changes in attentional bias. In other words, any bias in eye gaze resulting from a shift of attentional resources is likely to be counteracted later on in the pretraining phase by the overall reduction in eye gaze variation. On this basis, all other things being equal, one would predict that those participants who show the largest reduction in the variance of their eye gaze (e.g. those participants who develop a strategy to detect targets using their

peripheral vision) would be the least likely to show any attentional bias. Conversely, those participants for whom eye gaze is sufficiently varied (e.g. those participants who are still detecting targets foveally) would be most likely to show a larger attentional change, if indeed such a behavioural change occurred. This prediction was supported by a significant negative correlation between the reduction in variance across pretraining (standard deviation in Block 1 minus the standard deviation in Block 8) and the change in attention during pretraining (horizontal coordinate in Block 1 minus the horizontal coordinate in Block 8), $r(17) = -.51, p < .05$. Thus, those participants showing the greatest reduction in variance for eye gaze movements during the outcome-cue RSI period, also showed the smallest change in attentional bias towards good predictor cues during the same period. On the other hand, when movements in eye gaze were sufficiently varied in general, changes in attention were more pronounced, and these changes tended to be directed towards cues that were good predictors of outcomes.

Cue-outcome contingency learning during the training phase

The left hand panel of Figure 5.6 shows the RT data for the training phase. Learning proceeded rapidly for both good and poor predictor cues during the early part of the training phase, as shown by the difference in RTs to high and low probability outcomes. These data were subjected to an ANOVA with factors of outcome probability (high vs. low) and prior predictiveness of cue (good vs. poor), and block. This revealed a significant effect of outcome probability, $F(1, 16) = 57.90, p < .001$, indicating strong learning of the training contingencies. There was no effect of prior predictiveness of cue, nor an effect of block, both $F_s < 1$. The prior predictiveness by block interaction was significant, $F(1, 16) = 5.83, p$

< .05, which is likely to be driven by the different pattern of RTs across the training phase in the low-probability outcomes: RTs increase for low-probability outcomes following cues pretrained as poor predictors, whilst, if anything, there is a decrease in RTs for low-probability outcomes following cues pretrained as good predictors. The outcome-probability by block interaction was also significant, $F(1, 16) = 6.97, p < .05$, indicating greater learning in Block 2 than in Block 1. There was no prior predictiveness by outcome-probability interaction, $F(1, 16) = 1.83, p = .20$, nor a significant three-way interaction, $F < 1$.

Learning scores (RTs on low-probability outcomes minus RTs on high probability outcomes) are shown in the right-hand panel of Figure 4.6. Whilst there was a numerical trend towards greater learning about outcomes following good predictor cues over poor predictor cues in Block 1, this difference failed to reach a conventional level of significance, $t(16) = 1.59, p = .13$. The difference in Block 2 was also not significant, $t < 1$. The bottom panel of Figure 5.6 shows the RT data from the first block of the training phase split into three 50-trial sub-blocks. The pattern of data suggests that there was a short lived effect of prior predictiveness across the first 100 trials of the training phase. An analysis restricted to the first 100 trials found a marginally significant interaction between outcome probability and the prior predictiveness of cue, $F(1, 16) = 3.45, p = .082$.

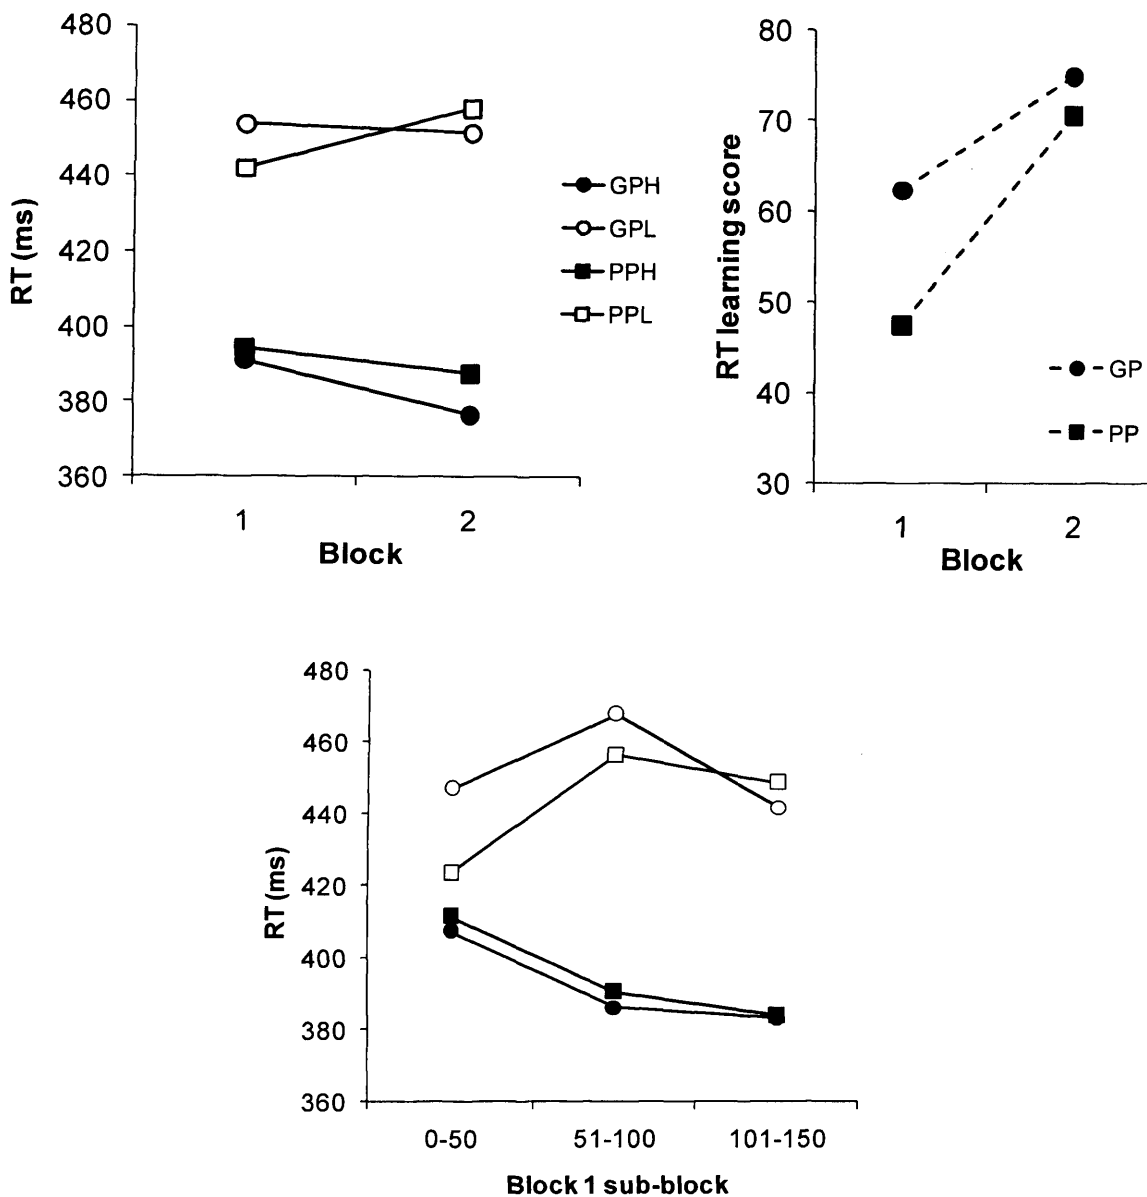


Figure 5.6. RT data from the training phase of Experiment 8. *Top-left panel:* RTs to high- (GPH) and low- (GPL) probability outcomes following good predictor cues, and RTs to high- (PPH) and low- (PPL) probability outcomes following poor predictor cues. *Top-right panel:* The data presented as learning scores - RTs on low probability outcomes minus RTs on high probability outcomes, for good (GP) and poor (PP) predictor cues. *Bottom panel:* RT data from the first block analysed as three 50-trial sub-blocks.

The left hand panel of Figure 5.7 shows the EGL data for the training phase. On the whole, the proportion of trials on which participants' eye gaze reached the target location was similar to that observed for the pretraining phase: 67.84% (highest = 97.3%; lowest = 4.0%; standard deviation = 28.5%). Five participants had less than 50% of trials in which eye gaze reached the target location; the data from these participants were removed from all subsequent analyses. These five participants were the same five removed from the corresponding analysis on the pretraining data.¹²

¹² Four of these participants provided enough data to be analysed in a full ANOVA. Although the data contained far more variance than that shown in Figure 4.7, the results of the key statistical tests were the same: a marginally significant interaction between prior-predictiveness of cue and outcome probability, and a greater learning score for cues pretrained as good predictors in Block 1.




Figure 5.7. Eye gaze latency (EGL) data from the training phase of Experiment 8. *Left panel:* EGLs to high- (GPH) and low- (GPL) probability outcomes following good predictor cues, and EGL to high- (PPH) and low- (PPL) probability outcomes following poor predictor cues. *Right panel:* The data presented as learning scores - EGLs on low probability outcomes minus EGLs on high probability outcomes, for good (GP) and poor (PP) predictor cues.

Figure 5.7 shows that there was a clear difference in latency between high- and low-probability outcomes throughout the training phase, and this difference was, at least during the first block, greater for outcomes following good predictor cues than for outcomes following poor predictor cues. A repeated measures ANOVA with factors of outcome probability, prior predictiveness of cue and block, revealed a significant main effect of outcome probability, $F(1, 11) = 13.09$, $p < .01$, indicating shorter EGL to high-probability outcomes than to low-probability outcomes. There was no main effect of block, $F < 1$, nor a significant main effect of prior predictiveness, $F(1, 11) = 2.03$, $p = .18$. There was no interaction between outcome probability and block, $F < 1$, but the interaction between the prior predictiveness of cue and block was significant, $F(1, 11) = 10.96$, $p < .01$, which indicates that latency decreased across blocks for outcomes following cues pretrained as good predictors, whilst (on the whole) it increased across blocks for outcomes following cues pretrained as poor predictors (this effect seems to be driven largely by the contrasting pattern of data in the latencies to low-probability outcomes). Crucially, the interaction between outcome probability and prior predictiveness approached significance, $F(1, 11) = 4.61$, $p = .055$, with the overall trend suggesting that the difference in EGL to high- and low-probability outcomes was greater following good predictor cues than following poor predictor cues. The three-way interaction was also significant, $F(1, 11) = 5.06$, $p < .05$, which reflects the fact that the difference in EGL to high- and low-probability outcomes decreased across the training phase following good predictor cues, whilst it increased across the training phase following poor predictor cues.

EGL learning scores (EGL to low-probability outcomes minus EGL to high-probability outcomes) are shown on the right-hand side of Figure 4.7. Analysis of the EGL data using related-samples t-tests found a significant advantage for learning about cues

pretrained as good predictors in Block 1, $t(11) = 3.57$, $p < .01$, but no difference in Block 2, $t < 1$.¹³

Discussion

Experiment 8 examined two key theoretical issues concerning sequence learning. Firstly, to what extent does sequence learning reflect the learning of a series of perceptual locations? Eye-tracking offers a particularly useful tool in addressing this question as it allows the measurement of anticipatory eye movements towards expected target locations. There was a considerable amount of evidence from the pretraining phase that participants learnt a sequence of eye movements towards targets: participants were faster to fixate on high-probability outcomes than on low-probability outcomes. Perhaps the best evidence, however, came from the eye gaze position data from the pretraining data (see Figure 5.3). This measure took the horizontal coordinate of the eye gaze data in order to assess differential movement to the two outcomes at the bottom of the screen. A positive score on this measure reflects an appropriate movement towards the outcome on the next trial. When the outcome location on the next trial was unpredictable (i.e. following a poor predictor cue), eye movements to the appropriate location were evident at around 150 ms after stimulus onset. However, when target location could be predicted (i.e. following a good predictor cue) evidence from the second half of the pretraining phase suggested that participants made anticipatory movements towards the appropriate target location in order to

¹³ Although there was a trend of greater learning scores in Block 1 than Block 2 for cues pretrained as good predictors, this difference was not significant, $t(11) = 1.46$, $p = .17$.

process and detect an outcome occurring in that location more rapidly. These findings provide strong evidence in favour of a perceptual component to sequence learning.

The second theoretical issue of interest was the extent to which the effect of prior predictiveness on new learning – observed in Experiments 4 to 7 – was due to a bias in the allocation of resources towards learning about good predictor cues. To the extent that attention correlates with overt eye movements, eye gaze should provide a measure of changes in attention during the task. The results, however, were not terribly conclusive. To a certain extent, there seems to be an overall trend for attention to be allocated towards good predictor cues during the pretraining phase (see Figure 5.5), however there was scant statistical evidence to support this trend: only a significant change in eye gaze between Block 1 and Block 5, and only then at an α level that was not adjusted for multiple comparisons. The overall pattern of data suggested a movement away from the centre of the screen for the first five blocks, followed by a return to the centre point towards the end of the pretraining phase.

One possible reason for this pattern of data is that a conflicting process might be acting to reduce the extent of the attentional changes towards the end of pretraining. An analysis of the variation in eye movements supported this interpretation of the data: there was a strong tendency for the horizontal movement in eye gaze, in general, to reduce as the pretraining phase progressed (see Figure 5.5). This reduction in the horizontal variation in eye gaze is likely to have had a significant impact on the measure of attention employed in the task, counteracting any movement away from the initial starting point that might have occurred as a result of changes in attention. Indeed, correlational analysis suggested that the

largest changes in ‘attention’ were shown in those participants who showed the smallest reduction in their eye movement across the training phase.

It is difficult to be sure of exactly what led to this reduced variation in eye gaze during the pretraining phase. One possibility is that with extended practice on the task, participants realise that they are able to accurately respond to the target location without fixating on it, and therefore become more reliant on their peripheral vision for target detection. On the basis of this suggestion, it is difficult to rule out that an effect of prior predictiveness in these tasks results from a shift in attention, but that such shifts are made covertly and will therefore go undetected by a measure of eye gaze.

It is important, however, to consider the possibility that the effect of prior predictiveness observed in the training data of Experiments 4-7 did not result from a change in attention towards good predictor cues during pretraining. An alternative explanation of the effect of prior predictiveness in sequence learning is that predictiveness modulates the associability of the cue, but not how attention is allocated. This definition sees changes in associability as reflecting changes in the cue-specific learning rate, but is not committed to a manifestation of associability in terms of changes in attention.

It is, however, difficult to decide between these two possibilities, especially given the evidence for an effect of prior predictiveness in the current task was not as strong as that seen in Experiment 5. Although the RT data produced the same trend as that seen in Experiments 5-7, there was no significant facilitation for learning about cues pretrained as good predictors, even in the early part of the training phase. However, stronger evidence was provided by the same analysis conducted with the EGL data. By this measure greater learning was observed for cues pretrained as good predictors in the first block of the training

phase. The data seem to suggest therefore, that EGL provides a sensitive measure of sequence learning, and is able to tease apart small differences in learning that the RT measurement cannot.

One caveat to this however, is that the EGL measure seemed to be less sensitive than the RT measure with respect to differences between trial types in the pretraining phase (see Figure 5.2, e.g. GPH vs. PPM). However, an analysis conducted with eye gaze position (see Figure 5.3) did find differences in eye movements towards outcomes following good and poor predictor cues during the pretraining phase, which suggests that the ‘insensitivity’ of the EGL measure during the pretraining phase was possibly due to the particular arrangement of target locations in the pretraining procedure used here, rather than a problem with the sensitivity of this measure for sequence learning per se.

To conclude, the data presented here support a number of studies (e.g. Mayr, 1996; Remillard; 2003) which have suggested that one component of sequence learning is the acquisition of a series of perceptual locations. When the contingencies in the task permit them, anticipatory eye movements towards target locations will be made, otherwise eye movements will be initiated at around 150 milliseconds following stimulus onset. On many occasions, participants were able to detect targets without first fixating on them. This finding suggests that an instruction to follow the target closely (e.g. Marcus et al., 2006) will, in many cases, lead participants to behave differently than they perhaps would do in the absence of such instruction.

The evidence for changes in the allocation of attention towards certain cues in the task was not particularly strong. A conservative conclusion would therefore be that the most likely mechanism driving changes in cue-associability during sequence learning is one

concerned with changes in the cue-specific learning rate, than the allocation of selective attention (e.g. Kruschke, 2001; Mackintosh, 1975). However, it seems that the current task might not have been a particularly sensitive test of what might well be small differences in eye gaze. The evidence from the RT data in the training phase, suggests that the associability effect on new learning was quite small in the current experiment, and this could well be an indication of why a more robust change in attention was not observed during the pretraining phase. In addition, this failure was perhaps compounded by the reduction in the pretraining length from 8 to 10 blocks. Nevertheless, an associability effect on new learning was found in the EGLs to outcomes during the training phase. These data therefore suggest that EGL might be a particularly sensitive measure for detecting small effects in sequence learning, such as those seen in Experiments 4-7. In addition, the absence of an effect in the RT measure suggests that the associability effects seen in Experiments 4-7 might well be driven primarily by changes in the perceptual, rather than the motoric, component of sequence learning.

Chapter 6 – Conclusions

It seems trivial to state that learning is an essential property of human cognition. Learning is, after all, the means by which an organism becomes sensitized to the relationships between elements of the environment, and is therefore essential for survival. It has been over 100 years since Pavlov first presented his empirical methods for examining the means by which such associations develop in animals and the field has since seen an explosion in the examination of the processes underlying such learning. Our understanding of learning has benefited hugely from the controlled examination of the conditions and procedures affecting the rate at which associative learning occurs. This thesis has focussed on one important factor affecting associative learning: the role that the predictive history of a cue plays in the rate at which conditioning occurs for that cue.

In Chapter 1 I provided a brief review of findings from studies on animal conditioning that have provided evidence for changes in cue-associability. Recent studies have also suggested that similar changes can be observed in human learning, offering evidence to support the suggestion that common mechanisms underlie these effects in both human and non-human animals. The empirical data presented in this thesis provide further evidence for changes in cue-processing in humans, and extend our knowledge of the conditions under which these effects can occur.

Although previous findings in human learning have been couched in the terms of associative learning theory (e.g. Lochman & Wills, 2003; Le Pelley & McLaren, 2003; Le Pelley et al., 2007; Bonardi et al., 2005; Suret & McLaren, 2003), it has also been argued that learning in these studies could be controlled by higher-order reasoning processes (see De Houwer et al., 2005; Mitchell et al., in press). The current experiments used incidental

learning tasks in an attempt to limit the application of controlled reasoning processes to learning. The work presented here provides the first examination of changes in cue-associability in such tasks. Hence, the main aim of the empirical work was to provide unequivocal evidence of associability processes under conditions of incidental learning.

Experiments 1-3, presented in Chapter 2, attempted to manipulate the associabilities of the stimuli on a between-subjects basis, by training one condition on structured material and one condition on unstructured material. Changes in associability were then assessed by comparing the performance of the two groups on a subsequent transfer task. Although the data from both an AGL (Experiment 1) and an SRT task (Experiment 2) provided preliminary support for changes in cue-associability, these data also seem open to a number of alternative explanations (perhaps most notably in terms of between-group differences in motivation), all of which seemed to be inherent features of between-subjects designs.

Experiments 4-7, presented in Chapter 3, minimised these factors by seeking evidence for changes in cue-associability in within-subjects designs. Participants were given an SRT task in which some cues were pretrained as good predictors of the next target location, whilst some were pretrained as relatively poor predictors. In a subsequent training phase, all cues were paired with new outcomes. Evidence for changes in the associabilities of the cues during pretraining was assessed by analysing the rate of novel learning about these cues during the training phase at the sub-sequence level. A consistent finding across these within-subjects designs was that learning about good predictor cues proceeded at a faster rate than that for poor predictor cues. These effects of prior predictiveness were frequently short lived in the current SRT tasks, with only Experiment 6 showing an effect beyond the first block of the training phase.

In Experiment 6, a post-experimental questionnaire was used to assess the extent to which knowledge of the sequence became consciously accessible. These data suggested that most participants became aware that the target moved in a sequenced fashion and that, numerically at least, participants identified more high-probability transitions than low-probability transitions. However, a statistical analysis of these numerical differences failed to reveal any conclusive evidence of conscious awareness, though it should be noted that this verbal report test is unlikely to offer the sensitivity required to elicit evidence of what is likely to be weak conscious knowledge (cf. Shanks & St. John, 1994). Moreover, there was no evidence to suggest that participants were aware that the task contingencies changed at the end of the procedure, suggesting that the observed associability effects were not the product of controlled learning processes.

Examining conscious awareness

To what extent would the current experiments have benefited from the use of more sensitive measures of conscious knowledge, such as the generation task? There seem to be several issues with the application of such a test to the current designs. Firstly, it is not easy to see how any such test could be applied in such a way as to generate meaningful data. The general aim of these tests would be to assess whether the effect of prior predictiveness on novel learning in the training phase depended on consciously-available knowledge. In order to make this assessment, it is necessary to obtain a measure of the extent to which participants are aware of the pretraining cue–outcome relationships, and a measure of the effect of prior predictiveness on learning during the training phase. One approach would be to test participants' awareness at the end of the pretraining phase. However, given that a test

of awareness will invariably involve informing participants of the structured nature of the trial sequence, this would clearly alter their approach to the task in the training phase – any learning occurring after the awareness test would no longer be incidental. Consequently this approach would not reveal whether awareness of the pretraining contingencies is necessary for an effect of prior predictiveness to occur in the training phase.

Alternatively, participants' awareness of the pretraining contingencies could be tested after participants have completed the training phase (e.g. the approach taken in Experiment 6 using a post-experimental questionnaire). However, this procedure will inevitably result in any measure of awareness becoming contaminated by new learning in the training phase, since participants must extinguish responding appropriate to the pretraining phase in order to learn the new contingencies in the training phase.

A final option would be to test awareness of the pretraining contingencies using a separate participant sample. Although this would provide a measure of the general level of awareness after the pretraining phase, it would not allow for an assessment of the extent to which the predictiveness effects observed in the training phase rely on participants' awareness. In summary, even supposing that an unambiguous test of awareness were available, it seems that there is no satisfactory way of assessing the extent to which awareness of the pretraining information is necessary for an effect of prior predictiveness to occur.

It should be reemphasised, however, that the issue of whether participants become consciously aware of the task contingencies is not the issue of central importance here. Indeed, one view of human memory is that the knowledge expressed on 'implicit' and 'explicit' tests is based on the same underlying knowledge store; whenever one observes

evidence of knowledge on an ‘implicit’ measure, one should also expect to observe sensitivity on ‘explicit’ measures (e.g. Kinder & Shanks, 2001).

It is also important to note that an observation of explicit knowledge does not entail that this knowledge has necessarily resulted from the operation of inferential reasoning processes. That is, the issue of awareness is orthogonal to associative accounts (Shanks, 2007) – there is no reason why people should *necessarily* remain unaware of associative links that have formed. In contrast, proponents of controlled, higher-order reasoning accounts are clear in stating that the learning resulting from such processes will always be accompanied by conscious awareness (Mitchell et al., in press). In relation to the current experiments, if participants were aware of the predictive relationships involving the cues of the SRT tasks, then they should necessarily be aware of a change in those relationships; this assumption is necessary if a reasoning account is to explain the learning of novel information about these cues in the training phase. And yet the verbal report data from Experiment 6 suggest that participants were unaware of a change in the task contingencies, and therefore suggest that an inferential account of the current results is unlikely.

Predictive history and the partial reinforcement extinction effect

It is worth noting a similarity between many of the current designs and those used to examine the partial reinforcement extinction effect (PREE). The PREE refers to the finding that extinction of conditioned responding proceeds more slowly for stimuli conditioned on a partial reinforcement schedule than those conditioned on a continuous reinforcement schedule, and is well-established in both animals (e.g. Haselgrove, Aydin, & Pearce, 2004;

Leonard, 1975; Pearce, Redhead & Aydin, 1997) and humans (Leonard, 1975; Pittenger & Pavlik, 1989). The most influential accounts of the PREE view it as depending on the extent of the mismatch between the effective context during conditioning and extinction (e.g. Amsel, 1967; Capaldi, 1967; Pearce et al., 1997). For example, Pearce et al. (1997) suggested that following continuous reinforcement, the onset of extinction results in a change of 'internal context' for the animal, brought about by the conflict arising from non-reinforced trials: the animal experiences a period of frustration. The generalization decrement (see Mackintosh, 1974, pp.440-443) caused by this change of internal context will lead to a reduction in conditioned responding, hence facilitating extinction. During partial reinforcement, however, conditioning proceeds whilst the animal is experiencing the internal conflict brought about by the non-reinforced trials. Consequently, the generalization decrement resulting from the change to extinction will be reduced, such that conditioned responding will persist for longer.

In some sense, the SRT procedures used in the current experiments are similar to the pretraining schedules commonly used to examine the PREE: the pretraining of poor predictor cues resembles a partial reinforcement schedule (with cues intermittently followed by one of two equally-likely outcomes), whilst the pretraining of good predictor cues resembles a continuous reinforcement schedule (with cues followed somewhat consistently by a single outcome). In this sense, given the evidence from PREE, we might expect the rate of extinction for the pretraining associations pertaining to good predictor cues to be faster than the rate of extinction for pretraining associations pertaining to poor predictor cues. These different rates of extinction could, therefore, potentially account for the advantage for learning about the good predictors observed during the training phase. However, it seems unlikely that the results of the current experiments could be explained by an account of the

PREE such as that advanced above, since in Experiments 4-7, the predictiveness of the cues was manipulated on a within-subjects basis. Thus training with both good and poor predictor cues occurred within a context of conflict arising from non-reinforced trials, and so any change in context experienced at the onset of the training phase would have been the same for both classes of cues. Hence any reduction in responding caused by generalisation decrement will apply equally to good and poor predictors and consequently no PREE should occur. Indeed, while many prior between-subjects studies have reported a PREE in both animals and humans (e.g. Haselgrove, et al. 2004; Leonard, 1975; Pearce, et al. 1997; Pittenger & Pavlik, 1989), within-subjects studies have typically found either no PREE at all or a reversed effect, with more rapid extinction for partially reinforced cues (e.g. Amsel, 1967; Crawford, Steirn, & Pavlik, 1985; Pearce et al., 1997; but see Rescorla, 1999).

Although a discussion of the present results in the context of the PREE is certainly of value, perhaps the most important aspect of this similarity is that the current results may have implications for our understanding of the mechanisms driving the PREE. Indeed, it has been suggested that one mechanism contributing to the PREE might be the different changes in cue-associability resulting from partial and continuous reinforcement schedules (Haselgrove et al. 2004).

Mechanisms for changes in cue-associability

In Chapter 4, several associative learning models were assessed on their ability to reproduce the observed empirical pattern of results from Experiment 6. Simulations with the Rescorla-Wagner (1972) model – in which cue-associability was fixed – failed to produce an

associability effect on new learning. However, simulations with the Mackintosh (1975) model – in which cue-associability was allowed to vary as a function of a cue’s absolute predictiveness – more readily produced a pattern of data similar to that seen empirically. Thus, simulations with these simple models of associative learning provided support for a mechanism incorporating changes in cue-associability.

Simulations were then conducted with the SRN, a well established model of sequence learning. Since the SRN is a multi-layered connectionist model, the complexity of the model makes it difficult to know intuitively whether it will be able to reproduce a given learning effect. Therefore, simulations were run with a standard implementation of the model without a mechanism for changes in cue-processing. These simulations also allowed for the examination of a theoretically-possible mechanism by which multi-layer networks could produce associability-like effects; namely a rapid rewiring of the associations between hidden and output units, while associations between input and hidden units remain largely intact (see Oswald et al. 2001). It was reasoned that this mechanism would benefit from a combination of two parameter settings: a high learning rate affecting weight changes on the hidden to output layer, and a low learning rate affecting the weight changes on the input to hidden layer. Consistent with this suggestion, the most robust effects of prior predictiveness produced by the SRN occurred when the learning rate parameters were set in this way. However, simulations that benefited from this rewiring mechanism also tended to produce inaccurate predictions for the ordinal relationships between training phase variables. Thus, the model was severely limited in terms of the parameter values that could accurately reproduce the observed pattern of data.

A new set of simulations with the SRN examined whether these parameter limitations could be alleviated by the inclusion of a mechanism incorporating changes in cue-associability. In line with the results of the simulations with the Mackintosh (1975) model, changes in associability reflected the absolute predictiveness of the cue. This modification greatly improved the ability of the resulting model to anticipate robust effects of prior predictiveness on novel learning in line with the empirical results. Furthermore, the parameter values producing the strongest effects also readily produced the correct ordinal relationships seen in the empirical data. Thus, it would seem that the simulations with the 'Alpha SRN' suggest that models of sequence learning ought to allow for changes in the cue-specific learning rate in order to capture the empirically observed influence of prior predictiveness on new learning.

Associability and attention

Current models incorporating changes in cue-processing have drawn parallels between changes in cue-associability and changes in attention (e.g. Mackintosh, 1975; Kruschke, 1996; 2001; Le Pelley, 2004). Since there exists a strong correlation between eye-movements and overt changes in attention (e.g. Hoffman, 1998; Deubel & Schneider, 1996), eye-tracking equipment can potentially provide a useful tool in the examination of such attentional changes. Indeed, recent findings in humans have suggested that attention is allocated selectively to predictive events of a task (e.g. Rehder & Hoffman, 2005a; 2005b; Kruschke et al. 2005; Wills et al. 2007) and that such changes in attention may act to modulate future learning (Beesley & Le Pelley, in prep).

Experiment 8 examined eye gaze data in combination with a design manipulating cue-associability. It was hypothesised that the allocation of attention to target locations might be biased as a function of the predictiveness of cues appearing in these locations. However, the results were inconclusive. Whilst there was a general trend consistent with the possibility of an attentional bias towards good predictor cue locations, there was insufficient statistical evidence to enable any strong conclusions to be drawn. Several factors may have contributed to this failure, though perhaps the most likely explanation is that many participants developed a tendency to limit the variation in their eye gaze as the experiment progressed. This general reduction in the variation of eye gaze is likely to reduce the magnitude of any attentional bias. In support of this argument, those participants who showed sufficiently varied movement in eye gaze also tended to show the largest attentional bias to good predictor cues.

One possibility that develops from this account is that as training progressed, participants became more proficient in detecting targets with their peripheral vision, whilst keeping the eyes fixated on a central area of the screen. Given the evidence that attention can be dissociated from eye gaze when the eyes are not moving (e.g. Posner, 1980), it is possible that changes in attention may have occurred covertly in this task. However, this tentative conclusion should be treated with caution, especially given the attentional effects on new learning observed in Experiment 8 were not as strong as those seen in the previous demonstrations (i.e. there was no significant effect of prior predictiveness on new learning in the training phase RT data).

Alternatively, if one assumes that the procedure did act to modulate the associability of the task cues (i.e. given the evidence for an effect of prior predictiveness on new learning

in the training phase EGL data), one might take the lack of evidence for attentional changes to indicate that cue predictiveness in the SRT task modulates the rate of learning for that cue, but not the allocation of attentional resources to learning about it per se. Such a conclusion is not unprecedented in the literature, given the theoretical accounts proposed for the reinforcer-specificity of associability effects (e.g. Dickinson & Mackintosh, 1979).

Finally, by analysing eye-gaze it was possible to explore the extent to which learning in the SRT task is facilitated by shifts in eye gaze to expected locations. Indeed, there was strong evidence to support a perceptual learning component to SRT performance: during the latter half of the pretraining phase, participants made anticipatory movements towards predictable target locations, with movements towards unpredictable target locations occurring around 150 ms later. Thus the data from Experiment 8 are consistent with previous findings in providing evidence for a perceptual component to sequence learning (e.g. Mayr, 1996; Remillard, 2001).

Future directions and final remarks

The current empirical and computational work provides the first examination of associability effects in these tasks, and could therefore act as a base for several theoretically interesting future projects. In respect to the current designs, it would perhaps be profitable to examine the procedure of Experiment 8 in more detail. I argued in this section that increasing the spatial distance between cues should increase the chances of detecting attentional changes, though the data suggest that this manipulation may well have inadvertently caused many participants to reduce their variation in eye gaze (see above). A

replication using the standard arrangement of target positions could alleviate the problems. Alternatively, one might induce greater variation in eye gaze by using a secondary target property and a task in which target fixation is necessary (e.g. respond when the target is a circle; inhibit response when the target is a square).

It may also be interesting to investigate the impact of a resource-demanding dual-task procedure on the observed effects of associability. Given the evidence that the learning of first-order contingencies is unimpaired under dual-task conditions (e.g. Cohen et al. 1990), one might expect that the current associability effects might still be observed, and therefore lend further support to the claim that these effects are the product of automatic associative mechanisms.

Finally, the present results go some way towards bridging the gap between research on HCL and animal conditioning, and research in the field of implicit learning. The amalgamation of theories from these fields, shown here in the application to models of sequence learning, provides further evidence that common mechanisms might underlie human and animal behaviour.

Appendix

Letter strings used in Experiment 1.

GA	GB	NG₁	NG₂
MMRV	VVTR	RXRV	TMRX
RVRV	VVTRM	MXRT	MVRV
RTXVX	TMXXM	VMTXT	VTTM
MRVRV	VTMXM	RVMVT	RRXRV
MRTVX	TRMTR	RTMVR	MXVXT
RTTXV	TRMXM	TVVMX	TMMTV
RTTVX	VTMTR	MXMRT	RRVXM
MMRVX	VTMTRM	TRTRTX	VRMVR
RTXR VX	TRRMTR	RXVXRR	TTRTXT
RTVMRV	TRRMXM	MTMTVM	RMTTVT
MRVR VX	VTMXXM	VMRXVV	VXXTVT
RTTXVX	TMXMTR	XVTVTR	MTXR XV
MRTTVX	VTVVTR	VRRMTX	VRVTRT
RTXVMRV	VTMXTR	TTMVTXR	TMVRVMX
MRTXR VX	TRMTRM	XR XVVRM	XVRTMTM
MRTVMRV	TRMXXTR	MRRXVTX	XXTRVXR
RTTXR VX	TMXXMTR	TVMVVTX	XRRTMXV
RTTVMRV	TMXTRXM	XMTXTVR	MVXR TMV
MRTTXVX	VTRRMTR	VRMTVMT	RVTVXTR
MRTTXRV	VTMXTRM	XMRMXVR	XTVRXXR

References

- Albouy, G., Ruby, P., Phillips, C., Luxen, A., Peigneux, P., & Maquet, P. (2006). Implicit oculomotor sequence learning in humans: Time course of offline processing. *Brain Research, 1090*, 163–171.
- Allan, L. G. (1980). A note on measurements of contingency between two binary variables in judgment tasks. *Bulletin of the Psychonomic Society, 15*, 147-149.
- Alloy, L. B., & Abramson, L. Y. (1979). Judgment of contingency in depressed in nondepressed students: Sadder but wiser? *Journal of Experimental Psychology: General, 108*, 441-485.
- Amsel, A. (1967). Partial reinforcement effects on vigor and persistence. In K. W. Spence & J. T. Spence (Eds), *The Psychology of Learning and Motivation* (Vol. 1, pp. 1-65). New York: Academic Press.
- Aydin, A., & Pearce, J. M. (1994). Prototype effects in categorization by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 20*, 264–277.
- Baker, A. G. (1976). Learned irrelevance and learned helplessness: Rats learn that stimuli, reinforcers, and responses are uncorrelated. *Journal of Experimental Psychology: Animal Behavior Processes, 2*, 130–141.
- Baker, A. G., & Mackintosh, N. J. (1977). Excitatory and inhibitory conditioning following uncorrelated presentations of CS and UCS. *Animal Learning & Behavior, 5*, 315–319.
- Beesley, T. & Le Pelley, M. E. (in prep.). The impact of blocking on cue-associability, attention and novel cue-outcome learning.
- Bennett, C. H., Wills, S. J., Oakeshott, S. M., & Mackintosh, N. J. (2000). Is the context specificity of latent inhibition a sufficient explanation of learned irrelevance? *Quarterly Journal of Experimental Psychology, 53B*, 239–254.

- Bonardi, C., Graham, S., Hall, G., & Mitchell, C. (2005). Acquired distinctiveness and equivalence in human discrimination learning: Evidence for an attentional process. *Psychonomic Bulletin & Review*, *12*, 88-92.
- Bonardi, C., & Ong, S. Y. (2003). Learned irrelevance: A contemporary overview. *Quarterly Journal of Experimental Psychology*, *56B*, 80–89.
- Brewer, W. F. (1974). There is no convincing evidence for operant or classical conditioning in adult humans. In W.B. Weimer & D.S. Palermo (Eds.), *Cognition and the symbolic processes* (pp. 1–42). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Brooks, D. N., & Baddetey, A. D. (1976). What can amnesic patients learn? *Neuropsychologia*, *14*, 111-122.
- Bush, R. R., & Mosteller, F. (1951). A mathematical model for simple learning. *Psychological Review*, *58*, 313–323.
- Capaldi, E. J. (1967). A sequential hypothesis of instrumental learning. In K. W. Spence & J. T. Spence (Eds), *The Psychology of Learning and Motivation* (Vol. 1, pp. 67-156). New York: Academic Press.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.
- Cheng, P. W., & Holyoak, K. J. (1995). Complex adaptive systems as intuitive statisticians: Causality, contingency, and prediction. In J. A. Meyer & H. Roitblat (Eds.), *Comparative approaches to cognition* (pp. 271-302). Cambridge, MA: MIT Press.
- Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology*, *58*, 545-567.
- Cheng, P. W., & Novick, L. R. (1992). Covariation in natural causal induction. *Psychological Review*, *99*, 365-382.
- Cleeremans, A. (1993). *Mechanisms of Implicit Learning: Connectionist Models of Sequence Processing*. Cambridge, MA: MIT Press.

- Cleeremans, A., Destrebecqz, A., & Boyer, M. (1998). Implicit learning: News from the front. *Trends in Cognitive Sciences*, 2, 406-416.
- Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, 120, 235-253.
- Cleeremans, A., Servan-Schreiber, D., & McClelland, J. L. (1989). Finite state automata and simple recurrent networks. *Neural Computation*, 1, 372-381.
- Cock, J. J., Berry, D. C., & Buchner, A. (2002). Negative priming and sequence learning. *European Journal of Cognitive Psychology*, 14, 27-48.
- 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.
- Cohen, N. J., & Squire, L. R. (1980). Preserved learning and retention of pattern-analyzing skill in amnesia: Dissociation of knowing how and knowing that. *Science*, 210, 207-210.
- Crawford, L. L., Steirn, J. N., & Pavlik, W. B. (1985). Within- and between- subjects partial reinforcement effects with an autoshaped response using Japanese quail (*Coturnix coturnix japonica*). *Animal Learning and Behavior*, 13, 85-92.
- Curran, T. (1995). On the neural mechanisms of sequence learning [On-line]. *Psyche*, 2. Available: <http://psyche.cs.monash.edu.au/v2/psyche-2-12-curran.html>
- De Houwer, J. (2002). Forward blocking depends on retrospective inferences about the presence of the blocked cue during the elemental phase. *Memory & Cognition*, 30, 24-33.
- De Houwer, J., Beckers, T., & Glautier, S. (2002). Outcome and cue properties modulate blocking. *Quarterly Journal of Experimental Psychology*, 55A, 965-985.
- De Houwer, J., Beckers, T., & Vandorpe, S. (2005). Evidence for the role of higher order reasoning processes in cue competition and other learning phenomena. *Learning & Behavior*, 33, 239-249.

- Dienes, Z. (1993). Computational models of implicit learning. In D. C. Berry & Z. Dienes. *Implicit learning: theoretical and empirical issues*. Hove: Erlbaum.
- Delamater, A. R. (1998). Associative mediational processes in the acquired equivalence and distinctiveness of cues. *Journal of Experimental Psychology: Animal Behavior Processes*, *24*, 467-482.
- Deroost, N., & Soetens, E. (2006). The role of response selection in sequence learning. *Quarterly Journal of Experimental Psychology*, *59*, 449-456.
- Destrebecqz, A., & Cleeremans, A. (2001). Can sequence learning be implicit? New evidence with the process dissociation procedure. *Psychonomic Bulletin and Review*, *8*, 343-350.
- Destrebecqz, A., & Cleeremans, A. (2003). Temporal effects in sequence learning. In L. Jimenez (Ed.), *Attention and implicit learning* (pp.181-213). Amsterdam: John Benjamins.
- Deubel, H., & Schneider, W. X. (1996). Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vision Research*, *36*, 1827-1837.
- Dickinson, A. (2001). The 28th Bartlett memorial lecture: Causal learning: An associative analysis. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology*, *54B*, 3-25.
- Dickinson, A., & Mackintosh, N. J. (1979). Reinforcer specificity in the enhancement of conditioning by posttrial surprise. *Journal of Experimental Psychology: Animal Behavior Processes*, *5*, 162-177.
- Dickinson, A., Shanks, D. R., & Evenden, J. L. (1984). Judgement of act-outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology*, *36A*, 29-50.
- Dienes, Z., Broadbent, D. E., & Berry, D. (1991). Implicit and explicit knowledge bases in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 875-887.

- Dulany, D. E., Carlson, R. A., & Dewey, G. I. (1984). A case of syntactical learning and judgment: how conscious and how abstract? *Journal of Experimental Psychology: General*, *113*, 541–555.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*, 179-211.
- Estes, W. K. (1950). Toward a statistical theory of learning. *Psychological Review*, *57*, 94-107.
- Frensch, P. A., Lin, J., & Buchner, A. (1998). Learning versus behavioral expression of the learned: The effects of a secondary tone-counting task on implicit learning in the serial reaction task. *Psychological Research*, *61*, 83–98.
- Fu, Q., Fu, X., & Dienes, Z. (2008). Implicit sequence learning and conscious awareness. *Consciousness and Cognition*, *17*, 185-202.
- George, D. N., & Pearce, J. M. (1999). Acquired distinctiveness is controlled by stimulus relevance not correlation with reward. *Journal of Experimental Psychology: Animal Behavior Processes*, *25*, 363–373.
- Gomez, R. L., & Schvaneveldt, R. W. (1994). What is learned from artificial grammar? Transfer tests of simple association. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 396–410.
- Graf, P., & Schacter, D. L. (1985). Implicit and explicit memory for new associations in normal and amnesic subjects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*, 501-518.
- Hall, G. (1991). *Perceptual and Associative Learning*. Clarendon Press, Oxford, England,.
- Hall, G., & Pearce, J.M. (1979). Latent inhibition of a CS during CS-US pairings. *Journal of Experimental Psychology: Animal Behavior Processes*, *5*, 31-42.
- Hall, G., & Pearce, J. M. (1982). Restoring the associability of a pre-exposed CS by a surprising event. *Quarterly Journal of Experimental Psychology*, *34B*, 127–140.

- Haselgrove, M. Aydin, A., & Pearce, J. M. (2004). A partial reinforcement extinction effect despite equal rates of reinforcement during Pavlovian conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *30*, 240-250.
- Hiroto, D. S., & Seligman, M. E. P. (1975). Generality of learned helplessness in man. *Journal of Personality and Social Psychology*, *31*, 311-327.
- Hoffman, J.E. (1998), Visual attention and eye movements. In H. Pashler (Ed.), *Attention*. London: University College London Press, 119-154.
- Holland, P. C. (1984). Unblocking in Pavlovian appetitive conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *10*, 476-497.
- 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*, 1029–1039.
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*, 513–541.
- Jiménez, L., & Méndez, C. (1999). Which attention is needed for implicit sequence learning? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 236–259.
- Jiménez, L., & Méndez, C. (2001). Implicit sequence learning with competing explicit cues. *Quarterly Journal of Experimental Psychology*, *54A*, 345–369.
- Jiménez, L., Méndez, C., & Cleeremans, A. (1996). Direct and indirect measures of sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 948–969.
- Kamin, L. J. (1969). Predictability, surprise, attention and conditioning. In B. A. Campbell & R.M. Church (Eds.). *Punishment and aversive behavior*. New York: Appleton-Century-Crofts.

- Kaye, H., & Pearce, J. M. (1984). The strength of the orienting response during Pavlovian conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *10*, 90–109.
- Kelly, S. W., & Burton, M. A. (2001). Learning complex sequences: No role for observation? *Psychological Research*, *65*, 15–23.
- Kehoe, E. J., Horne, A. J., Horne, P. S., & Macrae, M. (1994). Summation and configuration between and within sensory modalities in classical conditioning of the rabbit. *Animal Learning & Behavior*, *22*, 19–26.
- Kinder, A., Rolfs, M., & Kliegl, R. (2008). Sequence learning at optimal stimulus-response mapping: Evidence from a serial reaction time task. *Quarterly Journal of Experimental Psychology*, *61*, 203–209.
- Kinder, A., & Shanks, D. R. (2001). Amnesia and the declarative/nondeclarative distinction: A recurrent network model of classification, recognition, and repetition priming. *Journal of Cognitive Neuroscience*, *13*, 648–669.
- Knowlton, B. J., & Squire, L. R. (1994). The information acquired during artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 79–91.
- Kremer, E. F. (1978). The Rescorla–Wagner model: Losses of associative strength in compound conditioned stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, *4*, 22–36.
- Kruschke, J. K. (1996). Dimensional relevance shifts in category learning. *Connection Science*, *8*, 201–223.
- Kruschke, J. K. (2001). Towards a unified model of attention in associative learning. *Journal of Mathematical Psychology*, *45*, 812–863.
- Kruschke, J. K., Kappenman, E. S., & Hetrick, W. P. (2005). Eye gaze and individual differences consistent with learned attention in associative blocking and highlighting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 830–845.

- Le Pelley, M. E., Beesley, T., & Suret, M. B. (2007). Blocking of human causal learning involves learned changes in stimulus processing. *Quarterly Journal of Experimental Psychology*, *60*, 1468 – 1476.
- Le Pelley, M. E., & McLaren, I. P. L. (2003). Learned associability and associative change in human causal learning. *Quarterly Journal of Experimental Psychology*, *56B*, 68–79.
- Le Pelley, M. E., Oakeshott, S. M., & McLaren, I. P. L. (2005). Blocking and unblocking in human causal learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *31*, 56–70.
- Lewicki, P. (1986). Processing information about covariations that cannot be articulated. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*, 135-146.
- Lewicki, P., Czyzewska, M., & Hoffman, H. (1987). Uncousicous acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 523-530.
- Lewicki, P., Hill, T., & Bizot, E. (1988). Acquisition of procedural knowledge about a pattern of stimuli that cannot be articulated. *Cognitive Psychology*, *20*, 24-37.
- Leonard, D. (1975). Partial reinforcement effects in classical aversive conditioning in rats and human beings. *Journal of Comparative and Physiological Psychology*, *88*, 596-608.
- Lochmann, T., & Wills, A. J. (2003). Predictive history in an allergy prediction task. In F. Schmalhofer, R. M. Young, & G. Katz (Eds.), *Proceedings of EuroCogSci 03: The European Conference of the Cognitive Science Society* (pp. 217–222). Mahwah, NJ: Erlbaum.
- 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*, 97–106.

- Lovibond, P. F., Been, S. L., Mitchell, C. J., Bouton, M. E., & Frohart, R. (2003). Forward and backward blocking of causal judgment is enhanced by additivity of effect magnitude. *Memory & Cognition*, *31*, 133-142.
- Lovibond, P. F., Preston, G. C., & Mackintosh, N. J. (1984). Context specificity of conditioning and latent inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*, *10*, 360-375.
- Luce, R. D. (1959). *Individual Choice Behavior: A Theoretical Analysis*. New York: Wiley.
- Lubow, R. E. and Moore, A. U. (1959). Latent inhibition: The effect of non-reinforced preexposure to the conditioned stimulus, *Journal of Comparative and Physiological Psychology*, *52*, 415-419.
- Lubow, R. E. (1989). *Latent Inhibition and Conditioned Attention Theory*. Cambridge University Press, New York.
- Mackintosh, N. J. (1963). Extinction of a discrimination habit as a function of overtraining. *Journal of Comparative and Physiological Psychology*, *56*, 842-847.
- Mackintosh, N. J. (1969). Further analysis of the overtraining reversal effect. *Journal of Comparative and Physiological Psychology*, *67*, 1-18.
- Mackintosh, N. J. (1973). Stimulus selection: Learning to ignore stimuli that predict no change in reinforcement. In R. A. Hinde & J. S. Hinde (Eds.), *Constraints on learning* (pp. 75-96). London: Academic Press.
- Mackintosh, N. J. (1974). *The psychology of animal learning*. London: Academic Press.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*, 276-298.
- Mackintosh, N. J., & Little, L. (1969). Intradimensional and extradimensional shift learning by pigeons. *Psychonomic Science*, *14*, 5-6.

- Mackintosh, N. J., & Turner, C. (1971). Blocking as a function of novelty of CS and predictability of UCS. *Quarterly Journal of Experimental Psychology*, *23*, 359–366.
- Marcus, D., Karatekin, C., & Markiewicz, S. W. (2006). Oculomotor evidence of sequence learning on the Serial Reaction Time (SRT) Task. *Memory and Cognition*, *34*, 420–432.
- Mathews, R. C., Buss, R. R., Stanley, W. B., Blanchard-Fields, F., Cho, J., & Druhan, B. (1989). The role of implicit and explicit processes in learning from examples. A synergistic effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 1083–1100.
- Mayr, U. (1996). Spatial attention and implicit sequence learning: Evidence for independent learning of spatial and nonspatial sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 350–364.
- McLaren, I. P. L., & Suret, M. (2000). Transfer along a continuum: Differentiation or association? *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238.
- Merikle, P. M. (1994). On the futility of attempting to demonstrate null awareness. Comment on Shanks & St. John. *Behavioural and Brain Sciences*, *17*, 412.
- Miller, R. R., Barnet, R. C., & Grahame, N. J. (1995). Assessment of the Rescorla–Wagner model. *Psychological Bulletin*, *117*, 363–386.
- Milner, B. (1962). Les troubles de la mémoire accompagnant des lésions hippocampiques bilatérales. In *Physiologie de l'hippocampe* (pp. 257–272). Paris: Centre National de la Recherche Scientifique. English translation: B. Milner & S. Glickman (Eds.). Princeton: Van Nostrand, 1965 (pp. 97–111).
- Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (in press). The propositional nature of human associative learning. *Behavioral and Brain Sciences*.

- Mitchell, C. J., & Lovibond, P. F. (2002). Backward and forward blocking in human electrodermal conditioning: Blocking requires an assumption of outcome additivity. *Quarterly Journal of Experimental Psychology*, *55B*, 311–329.
- Mitchell, C. J., Lovibond, P. F., Minard, E., & Lavis, Y. (2006). Forward blocking in human learning sometimes reflects the failure to encode a cue-outcome relationship. *Quarterly Journal of Experimental Psychology*, *59*, 830-44.
- Nelson, J. B., & Sanjuan, M. C. (2006). A context-specific latent inhibition effect in a human conditioned suppression task. *Quarterly Journal of Experimental Psychology*, *59*, 1003-1020.
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, *19*, 1-32.
- Nissen, M. J., Willingham, D., & Hartman, M. (1989). Explicit and implicit remembering: When is learning preserved in amnesia? *Neuropsychologia*, *27*, 341–352.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 104–114.
- Oswald, C. J. P., Yee, B. K., Rawlins, J. N. P., Bannerman, D. B., Good, M., & Honey, R. C. (2001). Involvement of the entorhinal cortex in a process of attentional modulation: Evidence from a novel variant of an IDS/EDS procedure. *Behavioral Neuroscience*, *115*, 841–849.
- Pavlov, I. P. (1927). *Conditioned reflexes: an investigation of the physiological activity of the cerebral cortex*. Translated and edited by G. V. Anrep. London: Oxford University Press.
- Pearce, J. M., & Hall, G. (1979). Loss of associability by a compound stimulus comprising excitatory and inhibitory elements. *Journal of Experimental Psychology: Animal Behavior Processes*, *5*, 19-30.

- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian conditioning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87, 532–552.
- Pearce, J. M., Kaye, H., & Hall, G. (1982). Predictive accuracy and stimulus associability. In M. L. Commons, R. J. Herrnstein, & A. R. Wagner (Eds.), *Quantitative analyses of behavior: Acquisition* (Vol. 3). Cambridge, MA: Ballinger.
- Pearce, J. M., Redhead, E. S., & Aydin, A. (1997). Partial reinforcement in appetitive Pavlovian conditioning with rats. *Quarterly Journal of Experimental Psychology*, 50B, 274–294.
- 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, 785–800
- Perruchet, P., & Pacteau, C. (1990). Synthetic grammar learning: implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General*, 119, 264–275.
- Perruchet, P., Vinter, A., Pacteau, C., & Gallego, J. (2002). The formation of structurally relevant units in artificial grammar learning. *Quarterly Journal of Experimental Psychology*, 55A, 485–503.
- Pittenger, D. J., & Pavlik, W. B. (1989). Resistance to Extinction in Humans: Analysis of the generalized Partial Reinforcement Effect. *Learning and Motivation*, 20, 60-72.
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32, 3-25.
- Rayner, K. (1998). Eye movements in reading and information processing: Twenty years of research. *Psychological Bulletin*, 124, 372–422.
- Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, 77, 317-327.

- Reber, A. S. (1969). Transfer of syntactic structure in synthetic languages. *Journal of Experimental Psychology*, *81*, 115-119.
- Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, *118*, 219–235.
- Reber, A. S., & Allen, R. (1978). Analogy and abstraction strategies in synthetic grammar learning: A functionalist interpretation. *Cognition*, *6*, 189-221.
- Reber, A. S., & Lewis, S. (1977). Toward a theory of implicit learning: The analysis of the form and structure of a body of tacit knowledge. *Cognition*, *5*, 333-361.
- Redington, M., & Chater, N. (1996). Transfer in artificial grammar learning: A re-evaluation. *Journal of Experimental Psychology: General*, *125*, 123–138.
- Redington, M., & Chater, N. (2002). Knowledge Representation and Transfer in Artificial Grammar Learning. In R. French and A. Cleeremans (Eds.), *Implicit Learning and Consciousness* (pp. 121-143). Hove, East Sussex: 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*, 585–594.
- Rehder, B. & Hoffman, A.B. (2005a). Eyetracking and selective attention in category learning. *Cognitive Psychology*, *51*, 1-41.
- Rehder, B., & Hoffman, A. B. (2005b). Thirty-something categorization results explained: Selective attention, eyetracking, and models of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 811-829.
- Reid, L. S. (1953). The development of noncontinuity behavior through continuity learning. *Journal of Experimental Psychology*, *46*, 107–112.
- Reingold, E. M. & Merikle, P. M. (1988). Using direct and indirect measures to study perception without awareness. *Perception & Psychophysics*, *44*, 563-575.

- Remillard, G. (2003). Pure perceptual-based sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 581–597.
- Remillard, G., & Clark, J. M. (2001). Implicit learning of first-, second-, and third-order transition probabilities. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 483–498.
- Rescorla, R. A., & Coldwell, S. E. (1995). Summation in autoshaping. *Animal Learning and Behavior*, *23*, 314-326.
- Revusky, S. (1971). The role of interference in association over a delay. In W. K. Honig & P. H. R. James (Eds.). *Animal memory* (pp.155-213). New York: Academic Press.
- Rescorla, R. A. (1999). Within-subject partial reinforcement extinction effect in autoshaping. *Quarterly Journal of Experimental Psychology*, *52B*, 75-87.
- Rescorla, R. A., & Coldwell, S. E. (1995). Summation in autoshaping. *Animal Learning and Behavior*, *23*, 314-326.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.
- Rowland, L. A., & Shanks, D. R. (2006). Attention modulates the learning of multiple contingencies. *Psychonomic Bulletin & Review*, *13*, 643-648.
- Rumelhart, D.E., Hinton, G.E., & McClelland, J.L. (1986). Learning internal representations by error propagation. In D.E. Rumelhart, & J.L. McClelland (Eds.) *Parallel Distributed Processing: Explorations in the microstructure of cognition. Volume 1: Foundations*. Cambridge, MA: MIT Press.
- Rüsseler, J., Münte, T. F., & Rösler, F. (2002). Influence of stimulus distance in implicit learning of spatial and nonspatial event sequences. *Perceptual and Motor Skills*, *95*, 973–987.

- Sedek, G., & Kofta, M. (1990). When cognitive exertion does not yield cognitive gain: toward an informational explanation of learned helplessness. *Journal of Personality and Social Psychology*, *58*, 729-743.
- Servan-Schreiber, D., & Anderson, J. R. (1990). Learning artificial grammars with competitive chunking. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 592-608.
- Shanks, D. R. (1985). Forward and backward blocking in human contingency judgement. *Quarterly Journal of Experimental Psychology*, *37B*, 1-21.
- Shanks, D. R. (2005). Implicit learning. In K. Lamberts & R. L. Goldstone (Eds.), *Handbook of cognition* (pp. 202-220). London: Sage.
- Shanks, D. R. (2007). Associationism and cognition: Human contingency learning at 25. *Quarterly Journal of Experimental Psychology*, *60*, 291-309.
- Shanks, D. R., & Darby, R. J. (1998). Feature- and rule-based generalization in human associative learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *24*, 405-415.
- 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*, 1435-1451.
- Shanks, D. R., Johnstone, T., & Staggs, L. (1997). Abstraction processes in artificial grammar learning. *Quarterly Journal of Experimental Psychology*, *50A*, 216-252.
- Shanks, D. R., Rowland, L. A., & Ranger, M. S. (2005). Attentional load and implicit sequence learning. *Psychological Research*, *69*, 369-382.
- Shanks, D. R., & St John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, *17*, 367-447.
- Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs*, *75*, 517.

- 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*, 598-616.
- Stadler, M. A. (1995). Role of attention in sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 674-685.
- Stanley, W. B., Mathews, R. C., Buss, R. R., & Kotler-Cope, S. (1989). Insight without awareness: On the interaction of verbalization, instruction and practice in a simulated process control task. *Quarterly Journal of Experimental Psychology*, *41A*, 553-577.
- Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory*, *82*, 171-177.
- Suret, M. B., & McLaren, I. P. L. (2003). Representation and discrimination on an artificial dimension. *Quarterly Journal of Experimental Psychology*, *56B*, 30-42.
- Suret, M.B. & McLaren, I.P.L. (2005). Elemental Representation and Associability: An Integrated Model. In Andy J. Wills (Ed.). *New Directions in Human Associative Learning*: LEA.
- Sutherland, N. S., & Mackintosh, N. J. (1971). *Mechanisms of animal discrimination learning*. New York: Academic Press.
- Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: expectation and prediction. *Psychological Review*, *88*. 135-170.
- Thorndike, E. L. (1911). *Animal Intelligence: Experimental studies*, New York: Macmillan.
- Whitlow, J. W., & Wagner, A. R. (1972). Negative patterning in classical conditioning: summation of response tendencies to isolable and configural components, *Psychonomic Science*, *27*. 299-301.
- Widrow, G., & Hoff, M. E. (1960). *Adaptive switching circuits*. *Institute of Radio Engineers, Western Electronic Show and Convention, Convention Record*. 96-104.

- Wilkinson, L., & Shanks, D. R. (2004). Intentional control and implicit sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 354–369.
- Willingham, D. B., Nissen, M. J., & Bullemer, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 1047-1060.
- Wills, A. J., Lavric, A., Croft, G. S., Hodgson, T. L. (2007). Predictive Learning, Prediction Errors, and Attention: Evidence from Event-related Potentials and Eye Tracking. *Journal of Cognitive Neuroscience*, *19*, 843–854.

