

Encoding specific associative memories

Tzu-Ching Esther Lin

A thesis submitted for the higher degree of PhD

Cardiff University, February 2010

UMI Number: U584442

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 U584442

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

DECLARATIONS

DECLARATION

This work has not previously been accepted in substance for any degree and is not concurrently submitted in candidature for any degree.

Signed *Bu Chy Lin* (candidate) Date *10/05/2010*

STATEMENT 1

This thesis is being submitted in partial fulfillment of the requirements for the degree of D.Ed. Psych

Signed *Bu Chy Lin*(candidate) Date *10/05/2010*

STATEMENT 2

This thesis is the result of my own independent work/investigation, except where otherwise stated.

Other sources are acknowledged by explicit references.

Signed *Bu Chy Lin*(candidate) Date *10/05/2010*

STATEMENT 3

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed *Bu Chy Lin*(candidate) Date *10/05/2010*

STATEMENT 4: PREVIOUSLY APPROVED BAR ON ACCESS

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loans **after expiry of a bar on access previously approved by the Graduate Development Committee.**

Signed.....(candidate) Date

Acknowledgements

I would like to express my gratitude to all those who gave me the possibility to complete this thesis. I would like to thank my supervisor Prof. R. C. Honey who helped me with his continued guidance, encouragement, and remarkable patience throughout the process of completing this thesis. His commitment to research and passion are remarkable and inspiring. As a young researcher at the beginning of my career, I am grateful to have the opportunity to work with him.

I would also like to thank Mihaela Iordanova for her assistance with surgery, and thank many lab colleagues, Jeff Lewis, Kerry Housler, Lynn Hill, Clive Mann, Kevin Hotson, Den Simmonds and Dennis Price, who gave me different sorts of support of my research. It is a great pleasure to work with them.

Thanks to my family for their loving considerations and great confidence in me all through these years, even though all they know about my research is that I work with rats. I would also like to thank BBSRC for the financial support in conducting my research.

Last but not least, I thank God.

Abstract

The overarching aim of this thesis was to examine the nature of what is encoded during simple associative learning and configural learning. The experiments used rats as subjects and appetitive conditioning designs to assess the validity of two assumptions that are prevalent within standard accounts of learning in animals. First, animals simply represent the patterns of stimulation that are currently present in the environment. Second, that although the conditions that prevailed during the acquisition of associative knowledge have a marked effect on the strength of the resulting association, the association itself is “blind” with respect to the origin of this influence. The results presented in Chapters 2 and 3 undermine the first assumption by showing that associatively provoked (Experiments 1-3) and short-term traces (Experiments 4-6) can be assimilated into configural representations. The results presented in Chapter 4, from studies involving control rats (Experiment 7) and rats with lesions of the hippocampus (Experiment 8), indicate that animals ordinarily represent the nature of the stimulus trace (immediate or short-term) in the associative structures that are acquired during conditioning. The findings from Chapter 4 are inconsistent with the view that associative learning is blind with respect to the nature of the encoding conditions. Taken together, the novel results presented in this thesis reveal that what is encoded during simple associative learning and configural learning is much richer than has hitherto been realized.

List of publications

The results from group Interval in Experiment 2 and Experiment 6 have been accepted for publication:

Lin, T. E., & Honey, R. C. (in press). Analysis of the content of configural representations: The role of associatively evoked and trace memories. *Journal of Experimental Psychology: Animal Behavior Processes*.

The results have been briefly described in the following book chapter:

Honey, R. C., Close, J., & Lin, T. E. (in press). Acquired distinctiveness and equivalence: A synthesis. In C.J. Mitchell & M.E. Le Pelley (Eds.), *Attention and associative learning: From brain to behaviour*. Oxford: Oxford University Press.

Components of the results from Chapters 2-4 have been presented at conferences:

Honey, R.C., & Lin, T.E. (2010). Encoding specificity in associative learning. Joint meeting of the Experimental Psychology Society and Spanish Experimental Psychology, Granada, Spain.

Lin, T. E., & Honey, R. C. (2010). Are associations blind to the encoding conditions that prevailed during acquisition? Associative Learning Symposium, Greygnog, Wales.

Lin, T. E., & Honey, R. C. (2009). Discriminating stimulus and trace representations in simple conditioning and configural learning. Associative Learning Symposium, Greygnog, Wales.

Lin, T. E., & Honey, R. C. (2008). Sensory preconditioning and configural learning. Associative Learning Symposium, Greygnog, Wales.

List of figures

Figure 1. Each stimulus (A, B, X and Y) consists of four elements (e.g., a_1, a_2, a_3, a_4). The presentation of the AX compound, for example, will activate: the context-independent elements of A and X, that become active irrespective of the presence of other stimuli (i.e., a_1, a_2, x_1, x_2); the context-dependent elements of A and X (i.e., a_x and x_a). Also, the presence of X will inhibit a_4 and the presence of A will inhibit x_4 . Thus, the set of elements activated by each stimulus varies depending on the presence of other stimuli.

Figure 2. The pattern of connections (excitatory = filled triangles, inhibitory = filled circles) formed during a biconditional discrimination (i.e., AX+, AY-, BX-, and BY+). A, B, X, and Y represent input units, AX, AY, BX and BY represent configural units, and the US corresponds to the outcome of the trial (adapted from Pearce, 1994).

Figure 3. A hybrid model for negative patterning (i.e., A+, B+, AB-). The presentation of a compound activates the configural unit, AB, which has an inhibitory connection with the US; whereas the presentation of either A alone or B alone directly activates the US. The presentation of A or B alone is insufficient to activate the AB configural unit.

Figure 4. Experiment 1: Mean discrimination ratios (\pm standard error of the mean, SEM) on test trials in Experiments 1A and 1B.

Figure 5. Experiment 2: Mean discrimination ratios (\pm SEM) for groups Immediate and Interval.

Figure 6. Experiment 3: Mean discrimination ratios (\pm SEM) for test trials conducted in contexts B and D.

Figure 7. Experiment 4: Mean discrimination ratios (\pm SEM) for the 30-second and 90-second simple contextual discriminations.

Figure 8. Experiment 5: Mean discrimination ratios (\pm SEM) for the context present and the context trace discriminations.

Figure 9. Experiment 6: Mean discrimination ratios (\pm SEM) for the context present and the context trace discriminations.

Figure 10. Experiment 7: Mean rate of responding (\pm SEM) during X and Y on the first and the final block of the first-order conditioning stage.

Figure 11. Experiment 7: Mean discrimination ratios (\pm SEM) on the first and the final block of the immediate and short-term trace discriminations during the first-order conditioning stage.

Figure 12. Experiment 7: Mean rates of responding (\pm SEM) during presentations of LI and LT during the second-order conditioning stage.

Figure 13. Experiment 8: Histology. The maximum (grey) and minimum (black) extent of lesions in hippocampal rats. The sections are at specific distances (in mm) from Bregma (top to bottom: -2.4, -3.0, -3.9, -4.7, -5.4).

Figure 14. Experiment 8: Mean rates of responding (\pm SEM) during presentations of stimuli X and Y in groups Sham and HPC during the first-order conditioning stage.

Figure 15. Experiment 8: Mean discrimination ratios (\pm SEM) for the immediate trace discrimination in groups Sham and HPC during the first-order conditioning stage.

Figure 16. Experiment 8: Mean discrimination ratios (\pm SEM) for the short-term trace discrimination in groups Sham and HPC during the first-order conditioning stage.

Figure 17. Experiment 8: The mean number of blocks (\pm SEM) to reach criterion for the immediate and the short-term trace discriminations in groups Sham and HPC.

Figure 18. Experiment 8: Mean rates of responding (\pm SEM) to LI and LT of groups Sham and HPC in the second-order conditioning stage.

List of tables

Table 1: *The design used by Allman and Honey (2006)*

Table 2: *The design of Experiments 1A and 1B*

Table 3: *The design of Experiment 2*

Table 4: *The design of Experiment 3*

Table 5: *The similarity of the training and test patterns in Experiment 3*

Table 6: *The design of Experiment 4*

Table 7: *The design of Experiments 5 and 6*

Table 8: *The within-subjects design used in Experiments 7 and 8*

Table 9: *Stereotaxic coordinates and volume of ibotenic acid for lesions of the
hippocampus*

Table 10: *Simplified second-order conditioning and sensory preconditioning
procedures*

Table 11: *Alternative assessment of the distinction between A1 and A2 states*

Table 12: *Assessment of the nature of latent inhibition*

Table 13: *Extension to an operant procedure*

Contents

	Page
Chapter 1: General Introduction	1
1.1. Associative learning	1
1.2. Configural learning	5
1.3. Modified elemental analyses	8
1.3.1. A unique cue analysis	8
1.3.2. A replaced element model	10
1.4. Configural analysis	16
1.5. A hybrid model	18
1.6. What is represented during configural learning?	20
1.7. A summary of the aims of the thesis	24
Chapter 2: Configural preconditioning	26
2.1. Introduction	26
2.2. Experiments 1A and 1B: Some pilot data	28
2.2.1. Method	29
2.2.2. Results	33
2.2.3. Discussion	35
2.3. Experiment 2: The role of stimulus traces in configural preconditioning I	37
2.3.1. Method	38
2.3.2. Results	41
2.3.3. Discussion	43
2.4. Experiment 3: The role of stimulus traces in configural preconditioning II	45
2.4.1. Method	50
2.4.2. Results and Discussion	51
2.5. General Discussion	53
Chapter 3: Representing traces in simple and configural learning	57
3.1. Introduction	57
3.2. Experiment 4: Representing traces in a simple contextual discrimination	60
3.2.1. Method	61
3.2.2. Results and Discussion	64
3.3. Experiment 5: Representing traces in configural conditioning I	66

3.3.1.	Method	67
3.3.2.	Results and Discussion	69
3.4.	Experiment 6: Representing traces in configural conditioning II	71
3.4.1.	Method	73
3.4.2.	Results and Discussion	74
3.5.	General Discussion	76

Chapter 4: Functional equivalence between trace and associatively provoked memories 81

4.1.	Introduction	81
4.2.	Experiment 7: Assessing encoding specific associations	88
4.2.1.	Method	91
4.2.2.	Result	96
4.2.3.	Discussion	100
4.3.	Experiment 8: Neural dissociation of trace discriminations	104
4.3.1.	Method	107
4.3.2.	Results and Discussion	110
4.3.2.1.	Histology	110
4.3.2.2.	Behavioural results	113
4.3.2.3.	Discussion	120
4.4.	General Discussion	121

Chapter 5: General Discussion 126

5.1.	A brief summary of the new findings	127
5.2.	Theoretical implications	128
5.2.1.	Configural assimilation	128
5.2.2.	Encoding specificity/congruence	131
5.2.3.	The neural bases of encoding specific long-term memories	132
5.3.	Future work	133
5.3.1.	Replication and extension	133
5.3.2.	Distinct associations involving A1 and A2 states	135
5.3.3.	Assessment of the nature of latent inhibition	136
5.3.4.	Extension to an operant procedure	137
5.3.5.	Further dissociations of activity states and trace strength	138

5.3.6.	Assessment of the neural bases of short-term traces	139
5.4.	General summary	141
References		142

Chapter 1

General Introduction

1.1. Associative learning

Understanding the nature of learning and memory in human and nonhuman animals is a core objective of many disciplines (e.g., artificial intelligence, ethology, psychology and neuroscience). One approach to investigating learning and memory is to study simple forms of these processes in nonhuman animals (henceforth animals). One such form is Pavlovian conditioning. In a typical study of Pavlovian conditioning, a neutral stimulus is paired with a stimulus that has some motivational significance to the animal, and as a result of such pairings the presentation of the neutral stimulus comes to provoke some behavioural response. In Pavlov's original studies of salivary conditioning in dogs, the neutral stimulus was an auditory stimulus (e.g., a bell or a tone), the motivationally significant stimulus was food and the conditioned response (CR) was the amount of saliva elicited by the presentation of the auditory stimulus (Pavlov, 1927). One well-established account of such conditioned changes in behaviour is that provided by associative learning theory.

Associative learning theory attempts to provide a simple analysis for how animals learn about and represent their environments. In its basic form, the associative analysis of learning is indeed simple: The environment activates patterns of activity

within the brain and these patterns can become linked to one another, allowing future encounters with one pattern to activate the other through the link (or links) that now exists between them. To use a more recent example, when rats are given a conditioning procedure in which a conditioned stimulus (CS; e.g., a tone) is paired with an unconditioned stimulus (US; e.g., food) they come to show a CR (approaching the site of food delivery) when the tone is presented. This CR is most often assumed to reflect the strengthening of a link between the central representations of the patterns of stimulation activated by the CS and US (e.g., Pearce & Hall, 1980; Rescorla & Wagner, 1972).

Even in its simplest form, the associative account of learning has proven to be extremely powerful, providing explanations for a diverse range of observations. However, when the basic principles of associative learning are placed in a broader theoretical framework the resulting model becomes yet more powerful. For example, Wagner (1981) has provided a theoretical analysis of what he considered to be the standard operating procedures (SOP) of memory. This analysis embodied two principles of associative learning (contiguity and frequency) in a model with different memory states (cf. Atkinson & Shiffrin, 1968), and in which the rules governing performance were made explicit.

Briefly, according to this model presentation of a stimulus results in the representational elements of that stimulus being activated or retrieved from the inactive state (I; in long-term memory) into the primary activity state (A1), and from this state of activation they decay into a secondary activity state (A2) before finally decaying into the inactive state. This model holds that the state (A1 or A2) in which a given stimulus resides can have different consequences for learning and performance. For example, if a CS is in the A1 state it will enter into excitatory associations with a US that is also in the A1 state, but if the CS is in the A2 state then this will not occur. Also, if a stimulus is in the A1 state it can provoke a more marked, or indeed different, response than when it is in the A2 state. Another central feature of this model is that an associatively provoked memory of a stimulus (e.g., the memory of a US provoked by an effective CS) is considered to be in the A2 state. That is, an associatively provoked memory is equivalent to a memory of the same stimulus that has simply decayed into the A2 state from the A1 state. One consequence of these assumptions is that associatively provoked memories cannot enter into simple excitatory associations with other stimuli. There is already some evidence that is inconsistent with this prediction from studies of simple associative learning (for a review, see Hall, 1996; see also Dickinson & Burke, 1996). I will return to this prediction in the context of my rationale for Experiments 1-3 in this thesis. These experiments are presented in Chapter 2 and evaluate whether

associatively activated representations can become part of the (configural) representations formed during configural discrimination. These experiments also examined whether the trace of a stimulus can become part of a configural representation.

A central assumption of Wagner's (1981) model is that although the conditions under which encoding occurs (as characterized in the distribution of elements in the various states, I, A1 and A2) determines the development of excitatory associations, the resulting excitatory associations are themselves blind with respect to the origin of this influence. For example, the strength of an association between the memories of the CS and US can take the same specific value as the result of a limited number of delayed conditioning trials (where the CS co-terminates with the US) or as the result of a greater number of less effective trace conditioning trials (where there is a trace interval between the CS and US). The issue of whether associative memory is path independent, or blind with respect to the conditions under which encoding occurred, is one that will form a central part of this thesis. This issue is one that is theoretically interesting in its own right and it is directly examined in Experiments 4-8 in Chapters 3 and 4. These experiments were motivated by a novel theoretical analysis of the results of Experiments 1-3 - experiments that investigated the content of the representations acquired during configural learning.

Like the Rescorla and Wagner (1972) model, Wagner's (1981) SOP model fails to provide an explicit account for how animals solve configural discriminations. In such discriminations it is the patterns of stimulation, that are critical rather than the individual elements from which they are constructed, that predict the outcome of a trial. The specific issue that was addressed in Chapter 2 is whether associatively activated and trace memories can become assimilated into configural representations. This is clearly an issue that raises a number of problems in the context of Wagner's (1981) model. Before considering the limited evidence that is germane to this issue, it is critical to consider how associative theories of learning and memory have been rendered so that they can account for the ability of animals to acquire configural discrimination.

1.2. Configural learning

A large number of behavioural studies have shown that animals can solve configural discriminations (e.g., Allman & Honey, 2006; Asratyan, 1961, 1965; Honey & Watt, 1998; Wilson & Pearce, 1990) and there have also been a series of studies examining the neural bases of this capacity (e.g., Good & Honey, 1991; Iordanova, Burnett, Aggleton, Good & Honey, 2009; McDonald, Murphy, Guarraci, Gortler, White & Baker, 1997; for a review, see Sutherland & Rudy, 1989). For example, in a

biconditional contextual discrimination, rats are placed in two contexts (A and B; e.g., differently decorated operant chambers) and in each receive separate presentations of two auditory stimuli (X and Y; e.g., a tone and a click). In context A, presentations of X are followed by food and those of Y are not; whereas in context B presentations of X are nonreinforced whereas those of Y are reinforced (i.e., AX+, AY-, BX-, and BY+; + indicates reinforced, - indicates nonreinforced). Under these circumstances, rats will come to show more conditioned responding during presentations of X in A and those of Y in B, than presentations of Y in A and X in B (e.g., Honey & Ward-Robinson, 2002; Saavedra, 1975). A related example is negative patterning where separate presentations of two stimuli compound (A and B) are followed by food, whereas presentation of compound stimulus (AB) is not (i.e., A+, B+, AB-). Animals can also acquire such discriminations, coming to show more responding to the elements (A or B) than to the compound (AB; e.g., Grand & Honey, 2008; Kehoe & Gormezano, 1980; Rescorla, 1972; Woodbury, 1943).

The fact that animals can learn configural discriminations of the kind described in the previous paragraph represents an acute problem for some otherwise powerful and influential models of learning (e.g., Rescorla & Wagner, 1972; Wagner, 1981). These elemental models assume that associative links develop between the stimulus elements presented on a trial and the outcome of that trial. According to this analysis, animals

should be incapable of acquiring the biconditional discrimination described above, because the elements (i.e., A, B, X and Y) presented on each trial type (AX+, AY-, BX-, BY+) are equally often paired with food (+) and no food (-). This observation has prompted three classes of theoretical model: One class supplements traditional elemental analyses of learning (e.g., Rescorla & Wagner, 1972; Wagner, 1981) with the assumption that when stimuli are combined the sets of elements that each stimulus activates differs from those that are activated when the same stimuli are presented alone (e.g., McLaren & Mackintosh, 2002; Wagner & Rescorla, 1972; Wagner, 2003). The second eschews an elemental analysis altogether and replaces it with the view that all learning involves representing patterns of stimulation as configurations (e.g., Pearce, 1994). The final class assumes that there are both elemental and configural learning systems that either interact or operate in parallel (e.g., Kehoe, 1988; Rudy & Sutherland, 1995). I will briefly review these analyses and show how they can each provide an account for how animals acquire configural discriminations. These models share a unifying and simplifying assumption that animals can represent the patterns of stimulation that are physically presented on a given trial. It is this simple assumption that forms the starting point for the research described in this thesis and, in particular, the research described in Chapter 2, where I assess whether or not associatively provoked memories can be assimilated into configural representations.

1.3. Modified elemental analyses

The first set of theoretical analyses to be outlined involved modifications to standard elemental associative models of the type developed by Rescorla and Wagner (1972) and Wagner (1981). These models assume that associative learning involves the formation of direct associations between the elements of stimuli. One of these analyses was specifically developed to explain the ability of animals to acquire configural discriminations whereas the other, more complex, elemental analysis was designed to explain additional results that are inconsistent with simpler elemental models (e.g., the effect of similarity on discrimination learning; e.g., Redhead & Pearce, 1995; and the effects of combining stimuli separately paired with reinforcement; e.g., Aydin & Pearce, 1995).

1.3.1. A unique cue analysis

Wagner and Rescorla (1972) recognized the failure of their recently published model (Rescorla & Wagner, 1972) to explain how animals acquire configural discriminations and were very quick to identify a possible modification that would provide a remedy for it. They argued that each combination of stimuli (e.g., AX) generates a unique cue that can enter into an association with the outcome (e.g. food) of a trial. According to this analysis, each of the four compounds from a configural discrimination (AX, AY, BX,

and BY) generates a unique cue (p, q, r and s, respectively) and the solution to the discrimination involves these unique cues becoming linked to the outcomes of the trials on which they were generated (i.e., $p \rightarrow \text{food}$, $q \rightarrow \text{no food}$, $r \rightarrow \text{no food}$ and $s \rightarrow \text{food}$).

The unique cue analysis is simple, but it fails to provide an analysis of other aspects of associative learning (cf. Aydin & Pearce, 1995; Redhead & Pearce, 1995). Moreover, the analysis is informal and leaves a number of important issues underspecified. For example, is the unique cue generated by the presentation of a compound of A and X, also generated when another stimulus is added (i.e., AXB), or is a different/additional unique cue generated. In an effort to provide a more formal elemental analysis, Wagner (2003) has generated a further model (see also, Brandon & Wagner, 1998; Brandon, Vogel & Wagner, 2003; Wagner & Brandon, 2001) which is in many respects similar to that of McLaren and Mackintosh (2002). Both Wagner (2003) and McLaren and Mackintosh (2002) appeal to the idea that the sets of elements that are activated by a stimulus change as a function of the presence of other stimuli. They also share the views with each other and with previous models (e.g., Wagner, 1981) that (1) what is learnt as the result of a given trial reflects the set of elements that are activated by the environment, and (2) the resulting network of associations among the elements is path independent: To use Wagner's (1981) terminology, while the state of the elements (e.g., I, A1 or A2) influences the development of associations, the associations hold no record

of this influence (aside from the resulting associative strength). Given the degree of overlap between Wagner's and his colleagues theorising and the McLaren and Mackintosh (2002) model, the following discussion will consider only the former.

1.3.2. A replaced elements model

Wagner and his colleagues (Brandon & Wagner, 1998; Brandon, Vogel & Wagner, 2000; Brandon et al., 2003; Wagner & Brandon, 2001; Wagner, 2003) have developed an elemental model that has its origins in stimulus sampling theory (Atkinson & Estes, 1963; Estes, 1959). The key assumption of the model is that the elements activated by a stimulus can change as a function of the presence of other stimuli. This model assumes that each stimulus representation is composed of a set of elements, some of which are activated whenever the stimulus is presented (i.e., context-independent elements), and others whose activation depends on the presence or absence of other stimuli (i.e., context-dependent elements). According to this analysis, the presentation of compound stimulus (e.g., AX) will result in the three types of elements becoming activated or inhibited: *context-independent elements* (i.e., a_i and x_i) where the subscript i denotes independence; *context-dependent elements* (i.e., a_x and x_a) where the subscript x and a denote elements of A and X are determined by the presence of X and A, respectively; and content-dependent elements that are inhibited by the presence of other

stimuli. Thus, some of A's elements will be inhibited by the presence of X (i.e., a_{-x}) and some of X's elements will be inhibited by the presence of A (x_{-a}). The difference between Wagner's (2003) current theory and a previous incarnation (Wagner & Brandon, 2001) is the assumption that the replacement of elements in A when presented in two unrelated contexts (e.g., X and Y), will be statistically independent. According to this analysis, a proportion of the elements of A will be replaced when A is presented with X (i.e., r_x). The proportion of the elements of A that is not replaced, and will occur when A is presented alone or within the AX compound, will be $1-r_x$ (i.e., context-independent elements); and the proportion of the context-independent A elements in the presence of X and Y will be $(1-r_x)(1-r_y)$ (for further details, see Wagner, 2003). The application of this analysis to a biconditional discrimination is relatively simple and is presented below.

One can begin by assuming that each stimulus (A, B, X and Y) is composed of four elements which become active when each stimulus is presented individually: stimulus A is composed of elements a_1, a_2, a_3 and a_4 , stimulus B is composed of b_1, b_2, b_3 and b_4 , stimulus X consists of x_1, x_2, x_3 and x_4 , and stimulus Y consists of y_1, y_2, y_3 and y_4 (see Figure 1). However, the set of elements that is activated when two stimuli are combined (e.g., AX) is not simply: $a_1, a_2, a_3, a_4, x_1, x_2, x_3$ and x_4 . Instead, the set of A's elements that become active varies as a function of the other stimuli with which it is

presented. Figure 1 illustrates the patterns of elements that will become active during each of the compounds in a biconditional discrimination (i.e., AX, AY, BX and BY).

According to the replaced elements model, on both AX and AY trials the context independent elements of A (a_1 and a_2) will become active. However, the activation of A's remaining elements (a_3 and a_4) depends upon the presence of other stimuli – the activation of these elements is context sensitive or context dependent. Consider first a reinforced AX trial. One element of A (e.g., a_4) is inhibited by the presence of X (i.e., context dependent element: $a_{\sim x}$) and is replaced by element a_x which is only activated in the presence of X. In addition, one element of X is inhibited by the presence of A (e.g., x_4) and replaced by x_a . Thus, elements $a_1, a_2, a_3, a_x, x_a, x_3, x_2$ and x_1 are activated and paired with food on AX trials. Now consider a nonreinforced AY trial. The remaining element of A (e.g., a_3) is inhibited by the presence of Y and replaced by element a_y ; and one element of Y (y_3) will be inhibited and replaced by element y_a . Thus, elements $a_1, a_2, a_y, a_4, y_4, y_a, y_2$ and y_1 of compound AY will be paired with no food. The same principles apply to BX and BY trials: On nonreinforced BX trials, the set of elements that become active will be: $b_1, b_2, b_x, b_4, x_4, x_b, x_2$ and x_1 ; whereas, the sets of elements activated on reinforced BY trials will be: $b_1, b_2, b_3, b_y, y_b, y_3, y_2$ and y_1 . Thus, the sets of elements activated on reinforced trials (AX and BY) and nonreinforced trials (AY and BX) contain elements that uniquely predict food and no food, respectively. For

example, AX will activate a_3 , a_x , x_a , x_3 that are not activated on nonreinforced trials involving X (i.e., BX trials) or A (i.e., AY trials).

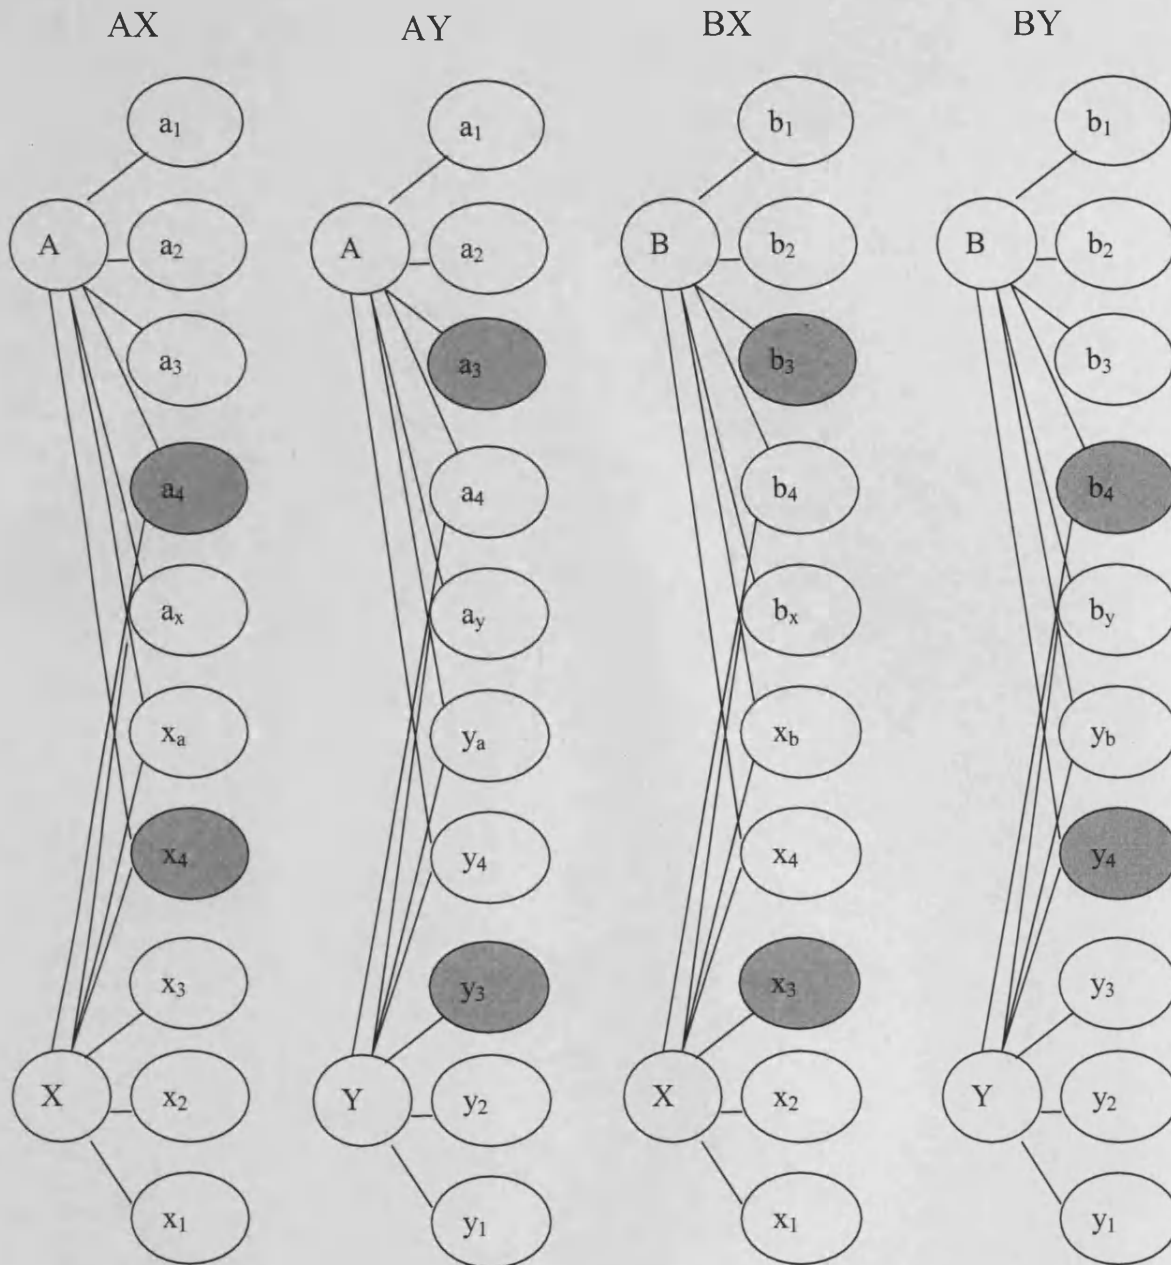


Figure 1. Each stimulus (A, B, X and Y) consists of four elements (e.g., a_1, a_2, a_3, a_4). The presentation of the AX compound, for example, will activate: the context-independent elements of A and X, that become active irrespective of the presence of other stimuli (i.e., a_1, a_2, x_1, x_2); the context-dependent elements of A and X (i.e., a_x and x_a). Also, the presence of X will inhibit a_4 and the presence of A will inhibit x_4 . Thus, the set of elements activated by each stimulus varies depending on the presence of other stimuli.

In Wagner's (1981) original model it was supposed that the proportion of the elements in the different states (i.e., I, A1 and A2) determined, at the level of the stimulus node, learning and performance. In more recent models it has been supposed that the rules governing learning and performance apply to the elements on an individual basis (e.g., Brandon et al., 2003). This view retains the assumption that the patterns of elements that are activated depend upon the patterns of stimulation present in the environment (i.e., AX, AY, BX and BY). Moreover, it does not suppose that the activity states of the elements (A1 or A2) during encoding are represented as a part of long-term associative memory. The new evidence presented in this thesis is directly relevant to both of these central assumptions.

It could be argued that the various elaborations of elemental theories of associative learning that have been summarized above are both complex and ad hoc. There is certainly an uneasy fit between their complexity and the relative simplicity of the results that they seek to explain. In contrast, it might be argued that configural analyses provide a more natural interpretation for many of the same observations. I will now consider one such analysis, that developed by Pearce (1994).

1.4. Configural analyses

In a more radical departure from the prevailing elemental view of associative learning, Pearce (1987, 1994) proposed that each pattern of stimulation becomes linked to a separate configural unit, and it is this unit that becomes linked to the outcome of a trial (e.g., food; see Figure 2 for a summary of the critical features of the model).

Application of this analysis to a biconditional discrimination (i.e., AX+, AY-, BX-, and BY+) is simple. Thus, upon presentation of a pattern of stimulation (e.g., AX) a set of input units will become active (i.e., A and X), and a configural unit will be recruited that can be said to represent the co-occurrence of A and X. This configural unit then becomes linked to the outcome of a trial (in the case under consideration, food) and is activated when AX is re-presented. When a similar compound is presented (e.g., AY) it will elicit generalized conditioned responding through its tendency to activate configural unit AX. This is based on the assumption that compounds AX and AY shared common element, A, and the similarity between AX and AY is derived from the proportion of common elements that they share: $A/AX \times A/AY = \frac{1}{2} \times \frac{1}{2} = \frac{1}{4}$. The new pattern will also recruit a new configural unit because it fails to fully activate the configural unit AX. This new configural unit (AY) will become connected to the outcome unit, but in this case by an inhibitory connection because food is not delivered. Similarly, the configural unit that represents BY will become linked to the food

outcome unit and the configural unit that represents BX will gain an inhibitory link to this outcome unit.

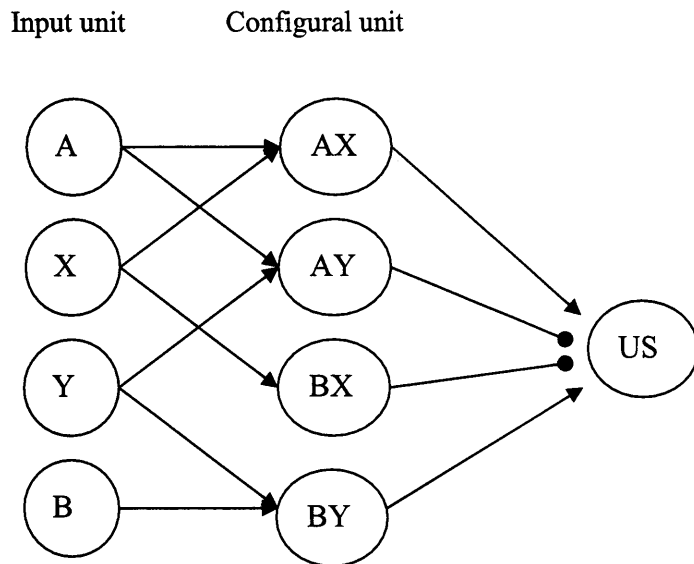


Figure 2. The pattern of connections (excitatory = filled triangles, inhibitory = filled circles) formed during a biconditional discrimination (i.e., AX+, AY-, BX-, and BY+). A, B, X, and Y represent input units, AX, AY, BX and BY represent configural units, and the US corresponds to the outcome of the trial (adapted from Pearce, 1994).

The relative merits of (modified) elemental and configural theories of associative learning remains a matter of considerable debate in the context of behavioural analyses of conditioning phenomena. However, it has been claimed that some neural manipulations result in dissociations between simple discrimination learning (e.g., A+, B-) and discriminations that require configural discrimination (AX+, AY-, BX-, BY+). For example, Good and Honey (1991) showed that lesions of the

hippocampus disrupted a contextual biconditional discrimination (AX+, AY-, BX-, BY+; A and B were contexts, and X and Y were darkness and a clicker), but had no effect on a simple context discrimination (A+, B-). Taken at face value, these results suggest that one might need to appeal to a hybrid associative structure with both elemental and configural processes (for recent evidence, see Iordanova et al., 2009).

One influential theory of hippocampal function, that proposed by Sutherland and Rudy (1989; see also Rudy & Sutherland, 1995), advocates just such a structure.

1.5. A hybrid model

Sutherland and Rudy (1989) suggested that associative learning in animals involves both elemental and configural associations. Ordinarily, both types of association can be formed and will function in parallel. This state of affairs is depicted in Figure 3.

However, without the hippocampus animals are left reliant on elemental associations.

Since this model was first presented, it has been subject to a great deal of empirical scrutiny. It seems clear that whatever one makes of the basic suggestion that elemental and configural processes operate in parallel in animals without hippocampal damage, the suggestion that animals with hippocampal damage are left reliant on elementary associations seems to be too simple to capture the available data. For example, a negative patterning discrimination (i.e., A+, B+, AB-) is considered to be a case where

configural processes are likely to be important; although rats with hippocampal lesions were able to learn this kind of discrimination, they did so less readily than shams (e.g., McDonald et al., 1997; Rudy & Sutherland, 1995). Moreover, the results of other studies seem to be inconsistent with the psychobiological model proposed by Sutherland and Rudy (1989), and more consistent with the view that lesions to the hippocampus disrupt processes allied to Wagner's (1981) model. For example, selective lesions of the hippocampus disrupt both retrieval-generated priming (see Honey, Watt, & Good, 1998; Honey & Good, 2000a, 2000b) and self-generated priming (Honey, Marshall, McGregor, Futter & Good, 2007; Marshall, McGregor, Good, & Honey, 2004). These priming effects have been taken to provide support for the basic tenets of Wagner's (1981) SOP model; and Honey and Good (2000b) have argued that the pattern of results observed in rats with hippocampal lesions suggests that the decay rate from A1 to A2 might be more rapid than in control rats. I will return to this analysis and a prediction that can be derived from it in Chapter 4; where I use a neurological intervention to assess one theoretical analysis for some novel behavioural findings. For the time being, it is sufficient to note that the hybrid model of associative learning remains viable, but like the other models that have been considered (the replaced elements model and the configural model) it leaves learning both tied to stimuli that are present in the environment and path independent.




Figure 3. A hybrid model for negative patterning (i.e., A+, B+, AB-). The presentation of a compound activates the configural unit, AB, which has an inhibitory connection with the US; whereas the presentation of either A alone or B alone directly activates the US. The presentation of A or B alone is insufficient to activate the AB configural unit.

1.6. What is represented during configural learning?

As I have just noted, each of the three theoretical positions that have been outlined share the simplifying assumption that the ability to solve configural discriminations is based solely upon sensitivity to the patterns of stimulation that are present in the environment. However, there is already some evidence that challenges the assumption that configural learning is simply based upon representing combinations of stimulation and linking the resulting representations with the outcome that follows them. In a complex study reported by Allman and Honey (2006), rats received simple exposure to compounds in the preexposure phase and were then given configural discrimination

training in the second phase (see Table 1). During this second phase, when rats were placed in contexts A and B, presentations of X were followed by food and presentations of Y were not (i.e., AX→food and BX→food, AY→no food and BY→no food); and when they were placed in contexts C and D, presentations of X were not followed by food, whereas those of Y were (i.e., CX→no food and DX→no food, CY→food and DY→food). On simple exposure days, rats either received presentations of AB and CD (group Congruent) or AC and BD (group Incongruent). The configural discrimination was acquired more readily in group Congruent than in group Incongruent. The implication of this experiment is simple: Associating the patterns of stimulation (i.e., AX, BX, CX, DX, AY, BY, CY, and DY) with the outcomes that they preceded (i.e., food and no food) was not the sole basis for performance in the configural discrimination.

Table 1: *The design used by Allman and Honey (2006)*

Preexposure	Configural training
	Both groups:
Group Congruent:	AX→food, AY→no food
AB & CD	BX→food, BY→no food
Group Incongruent:	CX→no food, CY→food
AC & BD	DX→no food, DY→food

Note: AB, CD, AC, and BD are hybrid contexts (e.g., involving combinations of wall decoration, odour, temperature and object); X and Y are a tone and a clicker; food and no food indicate trials on which the outcomes of the trials were food and no food.

There are a number of ways to interpret the pattern of results observed in Allman and Honey's (2006) study. One possible analysis is based on the idea that compound preexposure results in configural units being generated for the compounds (i.e., AB and CD for group Congruent; and AC and BD for group Incongruent). These configural units could then be recruited during acquisition of the conditional discrimination. Let us first consider how such a process of recruitment might influence the learning in group Congruent. During the AX trials of the configural learning stage, configural unit AB will be (partially) activated by the presentation of A and the new element X will be assimilated into this configuration. The resulting representation,

ABX, will then become associated with the memory of food. Now, when BX is presented, the ABX representation will be activated and result in rats (correctly) approaching the food well on the reinforced (BX) trial. For group Incongruent, however, the operation of this process of assimilation would interfere with configural learning: Exposure to AC should result in C being encoded as part of the configural representation formed when AX is paired with food; and the resulting representation, ACX, would become active when CX is presented and result in rats (incorrectly) approaching the food well on a nonreinforced (CX) trial.

Another way of interpreting the pattern of results reported by Allman and Honey (2006) relies on the interaction between simple elemental associations and configural associations. Thus, during exposure to compounds AB and CD simple associations might be formed between the elements of the compounds (i.e., A-B and C-D associations). One could envisage these being formed between the input units within Pearce's (1994) model or the hybrid model proposed by Sutherland and Rudy (1989). Now, during configural learning trials, these associations should allow A to associatively provoke the representation of B during AX trials, and allow the representation of B to be assimilated into AX configuration. The resulting ABX representation could then be activated on BX trials and result in rats correctly approach to food well.

The possibilities outlined above do not exhaust the possible ways in which associatively provoked representations might influence configural learning or performance (see Chapter 2). However, the preliminary results reported by Allman and Honey (2006) do suggest a need for models of associative learning to be modified so that events that are not physically present can enter into configural associations (see also, Hall, 1996).

1.7. A summary of the aims of the thesis

The review of the literature presented in this chapter has highlighted a number of areas in which our understanding of simple associative learning in animals is incomplete.

The first area concerns the various modifications to associative theories of learning that allow them to explain configural learning. Each of these modifications shares the simplifying assumption that animals can represent the patterns of stimulation that are physically presented on a given trial. However, there is evidence, from some rather complex studies, that is inconsistent with this view (Allman & Honey, 2006). The first aim of this thesis is to examine in more detail, and using a simpler procedure, whether these results can be replicated and extended. This aim was pursued in Chapter 2 (Experiments 1A, 1B, 2 and 3) that used rats as subjects and appetitive conditioning designs involving contextual cues. The second aim of this thesis was to examine

another assumption that is shared by current associative theories of animal learning; namely, the view that while the encoding conditions influence the acquisition of associative strength, these conditions do not form a part of the associative structure that is acquired. Interest in this assumption was prompted by the results and resulting theoretical analysis presented in Chapter 2. This aim was pursued in Chapter 3 (Experiments 4-6) which again used rats as subjects, contexts as cues and appetitive conditioning procedures. Chapter 4 was concerned with replicating the results presented in Chapter 3, using auditory cues in an appetitive conditioning procedure, and examining a novel theoretical interpretation of them. Briefly, this interpretation relies on the view that the way in which the memories of stimuli are activated (i.e., directly, by association or through a process of trace decay) are encoded as an integral part of associative memory. This view was investigated using both behavioural (Experiment 7) and neural manipulations (Experiment 8).

Chapter 2

Configural preconditioning

2.1. Introduction

In demonstrations of sensory preconditioning (e.g., Brogden, 1939; Rescorla & Cunningham, 1978), after preexposure to a stimulus compound (e.g., AB; a tone and a light, a flavour compound) establishing a conditioned response to one of its components (e.g., A) results in the other component (B) eliciting a conditioned response. The results described by Allman and Honey (2006) suggest that it should be possible to observe a related effect that I will refer to as *configural preconditioning*. Instead of examining whether a simple CR will transfer between the components of a preexposed compound, as in a standard sensory preconditioning paradigm, I examined whether configural learning will transfer between the components of a preexposed compound. To do so, I used the experimental design summarized in Table 2 in which following preexposure to two hybrid contexts, AB and CD, rats received a configural discrimination involving contexts A and C. In context A (e.g., a chamber with spotted walls) presentations of X (e.g., a tone) were followed by food and those of Y (a chain of clicks) were not, whereas in context C (e.g., a chamber with checked walls) presentations of X were nonreinforced whereas those of Y were reinforced (i.e., AX+,

AY-, CX-, and CY+). On the basis of extensive previous research using a similar procedures (e.g., Allman & Honey, 2006; Honey & Ward-Robinson, 2002), I anticipated that rats would acquire the configural discrimination; with rats coming to show more conditioned responding (approaching the site of food delivery) during presentations of AX and CY than during AY and CX. Finally, rats received a test in which an assessment was made of whether what they had learnt in A and C transferred to contexts B and D, respectively. To do this, the tendency of rats to respond to X and Y was measured as a function of whether they were presented in context B or D (i.e., BX, BY, DX, and DY). A configural preconditioning effect would be evident if rats were more likely to show conditioned responding on BX and DY trials than on BY and DX trials.

As mentioned in Chapter 1, this configural preconditioning effect would be beyond the scope of accounts of configural learning which suppose that only those stimuli that are present on a trial are encoded as a part of the configural representation of that trial (e.g., Kehoe, 1988; Pearce, 1994; Wagner & Rescorla, 1972; Wagner, 2003).

As will become evident, the results of the experiments presented in Chapter 2 (Experiments 1A, 1B, 2 and 3) suggest that the content of configural representations is much richer than these conventional analyses of configural learning have assumed.

Indeed, taken together, the results have far reaching implications for our understanding

of associative learning more generally which are directly investigated in Chapters 3 and

4.

Table 2: *The design of Experiments 1A and 1B*

Preexposure	Configural training	Test
AB	AX→food, AY→no food	BX versus BY
CD	CX→no food, CY→food	DX versus DY

Note: AB and CD denote hybrid contexts (e.g., a spotted chamber with a cool floor and a checked context with a warm floor); X and Y denote a tone and a clicker; food and no food indicate the outcomes of the trials.

2.2. Experiments 1A and 1B: Some pilot data

Experiments 1A and 1B both used variants of the experimental design summarized in

Table 2 and described above. The principal difference between the experiments was the

number of days on which rats received preexposure (Experiment 1A = 8 days;

Experiment 1B = 16 days) and the number of days of configural training (Experiment

1A = 6 days; Experiment 1B = 8 days).

2.2.1. Method

Subjects

A total of sixty-four naïve male Lister Hooded rats (supplied by Harlan Olac, Oxon, UK) were used in Experiment 1A ($n = 32$) and 1B ($n = 32$). The rats were housed in pairs in a colony room that was illuminated between the hours of 8 a.m. and 8 p.m. Behavioural training began at, approximately, 9 a.m. on each day. The rats received a restricted amount of food every day (supplied by Harlan Tekland, Bicester, Oxfordshire, England) in order to maintain them at 80% of their ad-lib weight ($M = 382$ g; range = 350-444 g).

Apparatus

Two sets of four operant chambers (Test chamber CI-410; Campden Instruments Ltd., Loughborough, England) were used. Each set was arranged in 2×2 array and was located in different experimental rooms. Each chamber (24.5 cm wide \times 23 cm deep \times 21 cm high) was positioned within a sound-attenuating box and had three aluminium walls and ceiling. The front wall was constructed from transparent plastic wall and served as the door of the chamber. There was a food well in the left hand aluminium wall into which 45-mg of food pellets (supplied by P. J. Noyes, Lancaster, NH) could be delivered. A top-hinged transparent plastic flap guarded access to this food well. Food-well entries were automatically recorded when the top-hinged magazine flap was

pushed approximately 3 mm. Each of the chambers was illuminated by a 3-W light bulb, positioned in the centre of the ceiling panels.

The hybrid contexts (AB and CD) used during preexposure were constructed in the following manner. The top left chamber and the bottom right chamber were decorated with spotted wallpaper (diameter: 15 mm; centre-to-centre distance: 25 mm) that was mounted behind transparent plastic panels. The top right and bottom left chambers were decorated with black and white checked wallpaper (30 mm × 30 mm squares) that was also mounted behind plastic panels. The aluminium floors in the top chambers were warmed to 28 °C, whereas the floors of the bottom chambers were cooled to 10 °C (for further details, see Ward-Robinson & Honey, 2000).

During configural training, all four chambers in both rooms had standard 16-bar grid floor (stainless steel bars, diameter 0.47 cm, spacing from bar centre to bar centre, 0.93 cm) and the chambers remained decorated with the same wallpaper as during preexposure. Two auditory stimuli (X and Y) were used during the configural learning stage: A 2-kHz tone and a 10-Hz clicker, produced by an audio generator. These stimuli were presented at an intensity of approximately 78 dB from a speaker located in the ceiling of the chamber. A computer controlled the apparatus and recorded food well entries. During the critical test, the wallpaper was removed and the standard floors were

replaced with aluminium floors. The floors of the top left and bottom right chambers were cooled and the floors of the remaining chambers were warmed.

Procedure

Preexposure. On each of 8 days (Experiment 1A) or 16 days (Experiment 1B), rats received two 10-min sessions, that were separated by approximately 1 min, in the two hybrid contexts (AB and CD). On the first day of preexposure, half of the rats were exposed in context AB (e.g., spotted+cool) in the first session and context CD (e.g., checked+warm context) in the second session. This arrangement was reversed for the remaining rats. The reader should be alerted to the fact that the assignment of contexts (to A and C) and floors (to B and D) was not fully counterbalanced; all rats received the compounds described above (i.e., spotted+cool and checked+warm). As it turns out, this fact is unlikely to have generated the pattern of test results (see Section 2.2.3). The order in which the pairs of hybrid contexts were presented alternated across days (e.g., Day 1: AB-CD; Day 2: CD-AB; Day 3: AB-CD; Day 4: CD-AB and so on). No responses were recorded during this stage.

Magazine training and configural learning training. On the day after the final day of preexposure, all rats were trained to collect food pellets from food well in undecorated operant chambers with standard grid floors. In the first 20-min session, the flap in front of the food well was taped open, allowing the rat unimpeded access to 20

food pellets that were delivered on a variable time 60-s schedule (VT 60). During the second session, the flap was returned to its normal resting state, and 20 pellets were delivered according to the same VT 60 schedule. On the following 6 days (Experiment 1A) or 8 days (Experiment 1B), rats were given 2 sessions of configural learning involving the visual contexts (A and C) and auditory stimuli (X and Y). The sheet floors were replaced with standard grid floors during this stage. When rats were placed in context A, they received 10-s presentations of X that were followed by food and 10-s presentations of Y were followed by no food (i.e., AX→food and AY→no food); and when they were placed in context C, they received nonreinforced presentations of X and reinforced presentations of Y (i.e., CX→no food and CY→food). There were 10 presentations of X and Y in each session and the inter-trial interval (ITI) was 30 s. Two mirror-imaged sequences were used to present the stimuli. The sequences had the constraint that no more than two trials of the same type occurred in succession. The use of these sequences alternated across days.

Tests. On the test day, the decoration in the boxes was removed and the grid floors were replaced with the sheet aluminium floors. All of the rats were placed in contexts B and D (i.e., cool and warm floors) where they received four 10-s nonreinforced presentation of X and Y (i.e., they received BX, BY, DX and DY). There was a 30-min interval between placement in one context and placement in the

other. Half of the rats received BX and BY trials in the first session, and DX and DY trials in the second session, and for the remaining rats this arrangement was reversed. The sequence in which X and Y were presented was XYYYXYYX for half of rats, and YXXYXXYYX for the remainder, and the ITI was 30 s. For a given rat, the same designated sequence was used in both test sessions.

Discrimination ratios were used to assess configural learning during training and test. These ratios took the following form during training: rate of responding during reinforced trials (i.e., AX and CY) divided by the rate of responding during both reinforced and nonreinforced trials (i.e., AX, CY, AY and CX). When this measure is used, scores above .50 indicate that discrimination training has been successful. During the test, the ratios took the following form: rate of responding during BX and DY divided by the rate of responding during all test trials (i.e., BX, DY, BY and DX). In this case, scores above .50 indicate that preexposure to AB and CD has allowed configural learning (where AX→food, AY→no food, CY→food, CX→no food) to transfer to the test patterns (BX, BY, DY and DX, respectively).

2.2.2. Results

Configural discrimination training. Over the course of the 6 days of training, rats in both groups acquired the discrimination (consecutive 2-day blocks for Experiment

1A: .51, .57, and .58; and Experiment 1B: .51, .56, .56, and .57). An independent analysis of variance (ANOVA) conducted on the results from each experiment confirmed that there was a significant effect of block in Experiment 1A and 1B, $F(2, 62) = 10.76, p < .0001$, and $F(3, 93) = 7.44, p < .05$, respectively. One-sample t tests revealed that the rats' discrimination ratios were significantly above chance on the final block in Experiment 1A, $t(31) = 6.23, p < .001$, and in Experiment 1B, $t(31) = 6.24, p < .001$. The mean rates of responding (in responses per minute, rpm) on the nonreinforced trials during the final block, were 10.80 responses per minute (rpm) for rats in Experiment 1A and 13.76 rpm for those in Experiment 1B.

Test. Figure 4 depicts the mean discrimination ratios (pooled across contexts B and D) for the two replications. Inspection of the left panel of this figure indicates that rats in Experiment 1A showed no evidence of transfer of configural learning from A to B (or from C to D). In fact, if anything, the mean score was slightly below .50. Similarly, in Experiment 1B, there was no indication that configural learning had transferred from A and C to B and D, respectively: the scores were close to .50. One sample t -tests confirmed that the ratio scores in neither experiment differed from chance (i.e., .50; Experiment 1A: $t(31) = -1.13, p > .05$; Experiment 1B: $t(31) = .19, p > .05$).

A separate analysis was also conducted on the raw rates of responding during the test. Again, paired-sample t tests confirmed that the rates of responding during the

test trials with BX and DY (Experiment 1A: $M = 7.13$ rpm; Experiment 1B: $M = 11.50$ rpm) did not differ from the rates of responding during BY and DX (Experiment 1A: $M = 8.58$ rpm; Experiment 1B: $M = 11.75$ rpm; largest $t(31) = -1.37, p > .05$).




Figure 4. Experiment 1: Mean discrimination ratios (\pm standard error of the mean, SEM) on test trials in Experiments 1A and 1B.

2.2.3. Discussion

There was no evidence of a configural preconditioning effect in either Experiments 1A or 1B. These failures to observe this effect occurred in spite of the facts that (1) the preexposure stage was very similar to Allman and Honey (2006; but see below), and (2) the rats acquired the configural discrimination involving contexts A and C. One

difference between the procedures used here and those in Allman and Honey (2006) was that in the latter rats received training in four contexts (i.e., A, B, C and D) over the course of, approximately, an hour. This meant that there was often a long interval (i.e., on average, thirty minutes) between placement in one context and another (e.g., A and C). It is not immediately apparent why the introduction of this difference should have compromised our ability to detect a configural preconditioning effect. However, this clear difference in procedure between Experiments 1A and 1B and those of Allman and Honey (2006) encouraged to examine this possibility in Experiment 2.

One thing that needs to be re-visited, at this point, is the fact that the counterbalancing of the combination of contexts was limited: the spotted context was always paired with cool floor and checked context was always paired with warm floor in Experiments 1A and 1B. It is not clear why this fact could have contributed to the failure to observe a configural preconditioning effect. In Allman and Honey (2006) the designation of contexts, floors, object and odour to the various roles was counterbalanced, and the effect that they observed was a general one that was not restricted to a subset of the counterbalancing. In any case, it is clearly important to counterbalance the experiment more fully, and in Experiment 2 the assignment of contexts (to A and C) and of floors (to B and D) was fully counterbalanced.

2.3. Experiment 2: The role of stimulus traces in configural preconditioning I

The experimental design used in Experiment 2 was identical to that employed in Experiments 1A and 1B with the exception that after the preexposure stage, rats were separated into two groups (see Table 3). The groups differed solely in the way that the configural training sessions were arranged. For rats in group Immediate, the configural training sessions in one context (e.g., A) immediately preceded those in the other context (e.g., C). For rats in group Interval, the configural training sessions in A and C were separated by an interval of two hours. So, the treatment given to rats in group Immediate was the same as that given to rats in Experiments 1A and 1B; whereas the treatment given to rats in group Interval was more akin to those given to rats in the experiments described by Allman and Honey (2006). As it has already been noted, it is not immediately apparent why, at a theoretical level, the introduction of this long interval would influence the likelihood of observing a configural preconditioning effect. However, if the introduction of this interval does make a difference, a configural preconditioning effect should be observed in group Interval but not in group Immediate (cf. Experiments 1A and 1B).

Table 3: *The design of Experiment 2*

Preexposure	Configural training	Test
AB	AX→food, AY→no food Immediate/Interval	BX versus BY
CD	CX→no food, CY→food	DX versus DY

Note: AB and CD denote hybrid contexts (e.g., a spotted chamber with a cool floor and a checked context with a warm floor); X and Y are a tone and a clicker; food and no food indicate the outcomes of the trials. For rats in group Immediate, the two daily configural training sessions involving A and C immediately followed one another (A then C, and C then A on alternate days), whereas for rats in group Interval, a two hour interval separated the training sessions within a day.

2.3.1. Method

Subjects and apparatus

Thirty-two naïve male Lister Hooded rats (supplied by Harlan Olac, Oxon, UK) were used in Experiment 2. The rats were housed in the same way as in Experiments 1A and 1B and were maintained at 80% of their ad-lib weight ($M = 382$ g; range = 350-444 g).

The apparatus was the same as in Experiments 1A and 1B.

Procedure

The training procedure for Experiment 2 was similar to Experiment 1B, with the notable exception that the design was counterbalanced and that the interval between configural training sessions was manipulated. On each of 16 preexposure days, rats received two 10-min sessions that were separated by approximately 1 min in the two hybrid contexts (AB and CD). For half of the rats, spotted+cool and checked+warm served as AB and CD; and for the remaining rats, spotted+warm and checked+cool served as AB and CD. For both subgroups, the identity of the hybrid contexts that served as AB and CD was counterbalanced. On the first day of preexposure, half of the rats were exposed in context AB (e.g., spotted+cool) in the first session and context CD (e.g., checked+warm context) in the second session. This arrangement was reversed for the remaining rats. The order in which the pairs of hybrid contexts were presented alternated across days (e.g., Day 1: AB-CD; Day 2: CD-AB; Day 3: AB-CD; Day 4: CD-AB). No responses were recorded during this stage.

After magazine training, rats were then randomly assigned to groups Immediate and Interval for the following 8 days of configural discrimination training. As in the Experiment 1B, rats were given 2 sessions of configural learning involving the visual contexts (A and C) and auditory stimuli (X and Y). In context A, rats received presentations of X were followed by food and presentations of Y were followed by no

food; and when in context C, they received nonreinforced presentations of X and reinforced presentations of Y (i.e., AX→food, AY→no food, CX→no food and CY→food). For group Immediate, placement in one context within a day was immediately followed by placement in the other context, and for group Interval, there was a 2-hr interval between the two sessions. The times of day at which the two groups received their sessions was matched; with half of the rats in group Immediate receiving their sessions at the same time of day as the first session for group Interval, and the remainder receiving their sessions at the same time of day as the second session for group Immediate. The procedure used for the test day was identical to that of Experiment 1B. However, analysis of the results was restricted to the first three 10-s nonreinforced presentation of X and Y (i.e., they received BX, BY, DX and DY) because the level of responding on the fourth extinction trials was very low. There was a 30-min interval between placement in one context and placement in the other. Half of rats received BX and BY trials in the first session and DX and DY trials in the second session, and the remaining rats this arrangement was reversed. The sequence in which these trials were presented was XYYXYX for half of rats, and YXXYXY for the remainder, and the ITI was 30 s. For a given rat, the same designated sequence was used in both test sessions.

2.3.2. Results

Configural discrimination training. Over the course of the 8 days of training, rats in both groups acquired the discrimination (consecutive 2-day blocks for group Immediate: .52, .56, .60 and .57; and for group Interval: .50, .57, .60, and .58). An analysis of variance (ANOVA) with group and block as factors revealed that there was a significant effect of block, $F(3, 90) = 9.49, p < .0001$, but no effect of group and no interaction between these factors, $F_s < 1$. One-sample t tests revealed that the discrimination ratios were significantly above chance on the final block in both group Immediate, $t(15) = 4.28, p < .001$, and in group Interval, $t(15) = 3.80, p < .01$. The rates of responding on nonreinforced trials during the final block, with means of 12.34 rpm for the group Immediate and 11.34 rpm for group Interval, did not differ significantly, $F < 1$.

Test. Figure 5 depicts the mean discrimination ratios (pooled across contexts B and D) for the two groups. Inspection of this figure indicates that rats in group Immediate showed no transfer of configural learning from A to B (or from C to D): the scores from this group were, on average, slightly below .50. In contrast, rats in group Interval did show a configural preconditioning effect: their scores were above .50. ANOVA confirmed that the ratios for group Interval were significantly higher than those for group Immediate, $F(1, 30) = 5.04, p < .05$. One sample t tests revealed that

the scores from both groups did not differ significantly from .50 (Immediate: $t(15) = -1.28, p > .05$; Interval: $t(15) = 1.97, p > .05$). However, ANOVA conducted on the raw rates of responding from which the ratios were derived revealed that there was an interaction between group and trial type (BX+DY versus BY+DX), $F(1, 30) = 4.96, p < .05$, but no effect of either factor ($F_s < 1$). Analysis of simple main effects showed that responding on the BX+DY trials ($M = 13.75$ rpm) was higher than on BY+DX trials ($M = 9.75$ rpm) in group Interval, $F(1, 15) = 4.85, p < .05$, but that this was not the case in group Immediate (reinforced trials, $M = 12.25$ rpm; nonreinforced trials, $M = 14.75$ rpm; $F(1, 15) = 1.20, p > .05$).




Figure 5. Experiment 2: Mean discrimination ratios (\pm SEM) for groups Immediate and Interval.

2.3.3. Discussion

The results of Experiment 2 demonstrate that after preexposure to two hybrid contexts (AB and CD), a configural discrimination acquired in contexts A and C transfers to contexts B and D. This configural preconditioning effect reveals that configural learning does not solely involve the pattern of stimulation that is currently impinging on an animal and are consistent with the results of Allman and Honey (2006). One interpretation for this configural preconditioning effect is that following preexposure to AB, B is assimilated into the configural representations acquired during AX→food and AY→no food trials. This assimilation process might be produced by virtue of the presentation of A associatively activating a representation of B during the two types of trial (AX→food and AY→no food); or because the configural representation (or representations) of AB formed during preexposure are recruited during the later discrimination learning trials (see Allman & Honey, 2006). In either case, the resulting configural representations (i.e., ABX→food and ABY→no food) could mediate generalization to BX and BY during testing. This analysis is straightforward, but it provides no account of the fact that the configural preconditioning effect was evident when there was an interval between training sessions involving contexts A and C (as it was for group Interval), but was not observed when there was no such interval (as was the case in group Immediate). This finding requires further theoretical analysis.

One potential explanation for the influence of the interval between configural discrimination trials in A and C that was observed in Experiment 2 relies on the simple notion that the trace of one context (e.g., A) might remain active when training trials in the other context (e.g., C) are taking place. Thus, if configural training trials in A are followed by training trials in C, then the trace of A will be present during the training trials in C (i.e., $CX \rightarrow \text{no food}$ and $CY \rightarrow \text{food}$); and when training trials in C are followed by A, the trace of C will be present during the training trials in A ($AX \rightarrow \text{food}$ and $AY \rightarrow \text{no food}$). This will mean that representations that include both contexts will be acquired (i.e., ACX and ACY) and be activated both on trials on which food occurs (i.e., $AX \rightarrow \text{food}$ and $CY \rightarrow \text{food}$) and on trials on which no food occurs ($AY \rightarrow \text{no food}$ and $CX \rightarrow \text{no food}$). The fact that each of these representations will have been paired with both food and no food will provide an additional source of generalization to the test compounds (e.g., BX and BY), which might well obscure the likelihood of observing a configural preconditioning effect. However, a straightforward implication of this analysis is that acquisition of the configural discrimination should have been impaired in group Immediate relative to group Interval in Experiment 2, and there was no sign of such an effect.

What is needed, therefore, is an analysis of the influence of a trace that accommodates the facts that configural learning proceeded equally readily in groups

Immediate and Interval, but one that predicts that the transfer of this learning to the test trials is less evident in group Immediate than in group Interval. One such analysis, presented in detail below, is based on two propositions: First, animals can learn different things about the immediate trace of a stimulus, occasioned by the presentation of the stimulus (denoted by an uppercase A), and the short-term trace of the same stimulus (i.e., denoted by a lowercase a); second, the trace of a stimulus is equivalent to the associatively activated representation of the same stimulus (i.e., a; cf. Wagner, 1981). As will become evident, application of the resulting, somewhat complex analysis, generates a straightforward prediction that was assessed in Experiment 3.

2.4. Experiment 3: The role of stimulus traces in configural preconditioning II

I proceed by applying the two ideas outlined above to the restricted case in which after exposure to AB and CD, configural training trials in A (i.e., AX→food and AY→no food) are always immediately followed by trials in C (i.e., CX→no food and CY→food), and then rats are tested with BX, BY, DX and DY (see Table 4). A configural preconditioning effect would be reflected in BX and DY provoking greater conditioned responding than BY and DX. Following exposure to AB and CD, the training trials in A will allow the associatively activated representation of B (i.e., b) to be assimilated into what is learnt as a consequence of those trials (i.e., AbX→food and

AbY→no food; see Table 5). When rats are then immediately placed in context C, the trace of A (i.e., a) will remain active and become assimilated into the configural representations acquired in C, together with the associatively activated representation of D (i.e., d). The resulting representations will be linked to the outcomes that they precede (i.e., aCdX→no food and aCdY→food). If we assume that the immediate trace of A (i.e., A) and the short-term trace of A (i.e., a) are discriminable, then the presence of a during conditioning trials in C need not markedly disrupt the acquisition of the configural discrimination.

Table 4: *The design of Experiment 3*

Preexposure	Configural training	Test
AB	AX→food, AY→no food	BX versus BY
	↓	
CD	CX→no food, CY→food	DX versus DY

Note: AB and CD denote hybrid contexts (e.g., a spotted chamber with a cool floor, and a checked context with a warm floor); X and Y are a tone and a clicker; and ↓ indicates that the daily configural training sessions in A always immediately preceded those in C; food and no food indicate the outcomes of the trials.

Table 5. *The similarity of the training and test patterns in Experiment 3*

Training patterns	Test patterns			
	aBX	aBY	cDX	cDY
AbX (→ food)	2	1	1	0
AbY (→ no food)	1	2	0	1
aCdX- (→ no food)	2	1	2	1
aCdY+ (→ food)	1	2	1	2

Note: Uppercase letters indicate that the immediate trace of the stimulus is active and lowercase letters indicate that a short-term trace of a stimulus is active; that can be produced by the recent presentation of a stimulus or by an associate of the same stimulus. The greater the similarity between the training and test patterns the more likely it is that the test pattern will evoke the outcome that was associated with the training pattern (see text for details).

The question of interest is how the rats will respond when tested with BX, BY, DX and DY. We will first assume that both BX and BY will activate a representation of A (i.e., a) and create patterns aBX and aBY, respectively; and similarly that DX and DY will activate a C and create patterns cDX and cDY. We will assume that these test patterns will activate the training patterns to the extent to which they (1) share components with the training patterns, and (2) these components are either trace congruent (i.e., immediate trace during training and test; or short-term trace during training and test) or trace incongruent (immediate trace during training and short-term

trace during test or vice versa). In particular, a matching test component will be given the score of 1 if it is trace congruent and .50 if it is trace incongruent. Using this scheme, the pattern activated by BX during the test (i.e., aBX) is most similar to one training pattern that was associated with food (AbX: score = $0.50+0.50+1.0 = 2$) and one associated with no food (i.e., aCdX: score = $1.0+0+0+1.0 = 2$). Similarly, the pattern activated by BY during the test (i.e., aBY) is most similar to one training pattern associated with food (aCdY; score = $1.0+0+1.0 = 2$) and one associated with no food (AbY; score = $0.50+0.50+1.0 = 2.0$). Thus, there are no clear grounds for predicting that the presentation of BX or BY will provoke different amounts of responding on the basis of the similarity between the components of the patterns that they activate and those activated by the original training patterns.

The state of affairs is different when the same analysis is applied to the presentation of DX and DY. The pattern activated by DY during the test (i.e., cDY) is most similar to a training pattern that was associated with food (aCdY: score = $0+0.50+0.50+1.0 = 2$), and the pattern activated by DX (cDX) is most similar to a training pattern that was associated with no food (i.e., aCdX: score = $1.0+0+0+1.0 = 2$; see Table 5). Now, there are clear grounds for predicting that DY will elicit more

responding than will DX during the test; that is, there are grounds for predicting a configural preconditioning effect involving C and D.¹

The prediction that configural learning is more likely to transfer between contexts C and D than between A and B, is based on the restricted scenario in which training trials in context A are always followed by trials in context C within a day.

However, this was not the case for group Immediate in Experiment 2. The fact that the order in which rats were placed in contexts A and C alternated across days, means that the effects of configural training in group Immediate will, on average, be equivalent when considering the transfer of configural learning from A to B and from C to D.

There are, therefore, no clear grounds for predicting transfer of configural learning under these circumstances. In contrast, once the possibility of the trace of the contexts being assimilated is removed (by the introduction of an interval), there are clear grounds for anticipating a configural preconditioning effect: configural training will result in AbX and CdY, but not AbY and CdX, becoming linked to food. Now, when BX, for example, is presented at test it will produce a pattern of activity, aBX, that is

¹ The relative weighting given to a match in stimulus identity alone versus a match in both identity and the nature of the trace (immediate or short term) does not influence the pattern of predictions for the DX and DY trials. However, if the weighting given to a matching trace is reduced (matching identity = 0.50, and matching identity+trace = 1) then BX is predicted to elicit greater responding than BY (i.e., a configural preconditioning effect); but, if the weighting given to a matching state is increased (matching identity = 0.50, and matching identity+trace = 1) then BY is predicted to elicit greater responding than BX.

more similar to a configuration that was linked to food (i.e., AbX) than to any other configuration.

In Experiment 3, I assessed the prediction, outlined above, that the order in which the configural training takes place in contexts A and C will influence the pattern of test results observed. To do so, the design outlined in Table 4 was employed in which, within a day, training trials in context A always immediately preceded those in context C. The question of interest was whether there would be greater transfer of configural learning from context C to context D, than from context A to context B. This prediction is based upon the suggestion that the trace of a stimulus can become assimilated into a configural representation (e.g., aCX→no food) that differs from the one recruited by the direct application of the same stimulus (e.g., AX→food).

2.4.1. Method

Subject and apparatus

Sixty-three naïve male Lister Hooded rats from the same supplier as in Experiments 1A and 1B were used: thirty-two rats were run in one replication and thirty-one rats were run in the other replication. The rats were housed in the same way as in Experiments 1A and 1B and were maintained at 80% of their ad-lib weights ($M = 329$ g; range = 286-392 g). The apparatus was the same as in Experiments 1A and 1B.

Procedure.

The training procedure was almost identical to Experiment 2 with the important exception that context A was always presented immediately before context C during configural discrimination training. In Experiment 3, rats received two test days that were arranged in the same way as Experiments 1A and 1B with the order of context placements reversing across days; for half of the rats the order was BD on day 1 and DB on day 2, and for the remainder it was DB on day 1 and BD on day 2. In each test session they received four 10-s nonreinforced presentation of X and Y in a counterbalanced sequence (XYYXYYXY for half of the rats and YXXYXXYYX for the remainder). Other details of the experiment that have not been mentioned were the same as Experiments 1A and 1B.

2.4.2. Results and Discussion

Configural discrimination training. There was a marked and significant increase in configural discrimination ratios over the four consecutive 2-day blocks of training ($M_s = .52, .56, .57, .57$; $F(3, 186) = 12.31, p < .0001$), and a one-sample t test revealed that the discrimination ratios were significantly above chance on the final block, $t(62) = 7.29, p < .0001$. The rate of responding during the nonreinforced trials on the final block was 14.63 rpm. The discrimination ratios on the final block for context A ($M = .55$) and

context C ($M = .58$) did not differ from one another, $t(62) = -1.64, p > .05$. One-sample t test revealed that the discrimination ratios were significantly above chance in context A, $t(62) = 3.41, p < .01$, and context C, $t(62) = 6.89, p < .001$.

Test. Inspection of Figure 6 reveals that the discrimination ratios were above chance (i.e., .50) in context D, but not in context B. An ANOVA confirmed that there was a difference between the discrimination ratios in contexts B and D, $F(1, 62) = 4.13, p < .05$, and one-sample t tests showed that the discrimination ratios for context D were significantly above chance, $t(62) = 2.37, p < .05$, but that those for context B were not, $t(62) = -0.83, p > .05$. The mean rates of responding to the nonreinforced trials were 17.96 and 14.61 rpm for contexts B and D, respectively; and these did not differ significantly, $t(62) = 1.45, p > .05$. These results are consistent with the theoretical analysis presented in the introduction of Experiment 3.




Figure 6. Experiment 3: Mean discrimination ratios (\pm SEM) for test trials conducted in contexts B and D.

2.5. General discussion

The results of Chapter 2 show that a configural preconditioning effect can be observed and this, in itself, is a novel finding: After exposure to AB and CD, a configural discrimination involving A and C transferred to B and D, respectively. To be more specific, after exposure to AB and CD, rats learnt that when they were in context A presentations of X would be reinforced and those of Y would be nonreinforced, whereas when they were in C presentations of X would be nonreinforced and those of Y would be reinforced. During subsequent testing, rats were more likely to respond when X was presented in B and Y was presented in D, than when X was presented in D and Y was presented in B. These results are analogous to the findings reported by Allman and Honey (2006), and indicate that configural learning involves more than simply

representing the patterns of stimulation that are presented on a given trial (cf. Pearce, 1994; Wagner, 2003). Perhaps of greater theoretical interest than the fact that such a configural preconditioning effect can be observed, are the conditions under which it is observed. These conditions prompt a number of intriguing conclusions. Briefly, that the short-term trace of a stimulus can enter into an association that differs for the immediate trace of the same stimulus; and that the short-term trace of a stimulus is treated as equivalent (or similar) to the associatively activated memory of the same stimulus. Direct evidence for these claims will be sought in the remaining empirical chapters of this thesis (Chapters 3 and 4). However, I should now consider various explanations for the effect of primary immediate interest, configural preconditioning.

There are two general types of account for configural preconditioning that mirror those that have been put forward for sensory preconditioning (cf. Rescorla & Cunningham, 1978; Ward-Robinson & Hall, 1996). One account proposes that *during testing*, B provokes the memory of A and D provokes that of C, and these associatively provokes memories combine with X and Y to recreate the trained configurations (AX, AY, CX and CY) during the test. The second account appeals to a process of configural assimilation in which the exposure to AB and CD allows B and D (i.e., the associatively provoked representations of A and C) to become part of the configural representations acquired during the configural discrimination (cf. Hall, 1996; Holland, 1981). The

supplementary observations from Experiment 2, that the configural preconditioning effect depended on the conditions that obtained when the configural discrimination was acquired, is more consistent with the second class of account. Thus, the configural preconditioning effect was only apparent when there was an interval of several hours between the training trials in contexts A and C; and when there was no interval between the training trials in context A and C there was no transfer of configural learning to B and D.

It is a relatively simple step to extend the configural assimilation account outlined above to explain the pattern of results observed in Experiment 2. This step involves the assumption that when training sessions involving A and C occur in close temporal proximity, the short-term traces of the two contexts (denoted as a and c) can become part of the configural representations of the stimuli that are physically present during configural training (see Section 2.4.). Moreover, it must be supposed that the short-term trace of a given stimulus can enter into an association that differs from the immediate trace of the same stimulus. That is, the conditions under which encoding occurs (involving either a short-term or an immediate trace) become represented, in some way, as part of the associative structures acquired during configural learning. Although the evidence from Experiments 1-3 supports this analysis, it does so only indirectly. The focus of the remaining empirical work and theoretical analysis in this

thesis now focuses on whether associative memory is encoding specific, and if it is then what is the basis of this specificity.

Chapter 3

Representing traces in simple and configural learning

3.1. Introduction

Chapter 2 provides clear evidence of a configural preconditioning effect. This effect is not anticipated by many models of associative learning (e.g., Pearce, 1994; Rudy & Sutherland, 1995; Wagner, 2003) and suggests that associatively activated representations can be assimilated into configural representations. The theoretical analysis that was developed to explain the conditions under which this effect was observed was based upon two propositions: First, that animals can learn different things about the memory activated by the immediate trace of a stimulus (e.g., A), and a short-term trace of the same stimulus (i.e., a); and second, that the short-term trace of a stimulus is equivalent to the associatively activated representation of that stimulus (i.e., a). The principal aim of Chapters 3 and 4 is to assess the validity of these propositions. In particular, Chapter 3 examines whether rats can learn that the immediate trace of a stimulus predicts one outcome and the short-term trace of the same stimulus predicts a different outcome. This issue was assessed in both a simple contextual conditioning procedure (Experiment 4) and in a configural learning procedure involving contexts (Experiments 5 and 6). Chapter 4 then assesses whether associatively activated and trace representations are functionally equivalent (Experiments 7 and 8).

It is well established that a trace conditioning procedure, where there is an interval between the CS and US, results in weaker conditioned responding than when there is no such interval (i.e., during delayed conditioning; e.g., Beylin, Gandhi, Wood, Talk, Matzel & Shors, 2001; Honey & Hall, 1992; Kamin, 1965; Pavlov, 1927; Revusky, 1968). In Chapter 3, variants of a trace conditioning procedure were used to assess whether rats can learn that the memory immediately activated by a stimulus (or the recent presentation of a stimulus) predicts one outcome, and the short-term trace of the same stimulus predicts a different outcome. The experiments involved simple discrimination learning (Experiment 4) and configural learning (Experiments 5 and 6). However, rather than being concerned with responding during the stimuli, in both cases I examined responding during trace intervals that followed the stimuli but preceded the outcomes of the trials. For example, in the simple discrimination procedure used in Experiment 4, rats were placed in contexts A and B and were then moved to a third context C after an interval of zero or sixty second. When the rats had been placed in context A and were then immediately moved to C, food was delivered, but when there was a sixty-second interval, no food was delivered. In contrast, when the rats had been placed in context B and there was a sixty-second interval before they were placed in C food was delivered, but when there was no interval, no food was delivered (see Table 6). In terms of the theoretical analysis outlined in Chapter 2, the rats should be capable of

solving this discrimination by forming the following four associations: $A \rightarrow \text{food}$, $a \rightarrow \text{no food}$, $B \rightarrow \text{no food}$, and $b \rightarrow \text{food}$. However, the fact that rats can solve such a discrimination does not provide unique support for such an analysis. The procedures outlined in Table 6 are directly analogous to those that have been employed to investigate the timing of conditioned responding (e.g., Desmond & Moore, 1991; see for a review, Gallistel, 1990; Gallistel & Gibbon, 2000). Theoretical analyses of such discriminations have been proposed that rely on rats being capable of representing the time at which events (such as reinforcers) will be delivered. However, it will become evident, over the course of Chapters 3 and 4, that it is possible to discriminate between a theoretical analysis based upon a novel modification to a model of associative learning (Wagner, 1981), and an analysis based upon an explicit process of timing (e.g., Barnet, Arnold, & Miller, 1991; Cole, Barnet, & Miller, 1995; Desmond & Moore, 1988; Miller & Barnet, 1993; Schreurs & Westbrook, 1982).

Table 6: *The design of Experiment 4*

Simple Contextual Conditioning
30-second discrimination: A-30s→food, B-30s→no food
90-second discrimination: A-90s→no food, B-90s→food

Note: A and B denote contexts (spotted or checked); 30s and 90s denote the trace interval between presentation of the contexts and the outcomes of the trials (food or no food). Conditioned responding was assessed in the 30s period immediately preceding the delivery of the outcomes when they were in a third context, C, in which these outcomes were delivered (see text for details).

3.2. Experiment 4: Representing traces in a simple contextual discrimination

Rats were placed in two contexts, A and B, and received what will be referred to as a 30-second discrimination and a 90-second discrimination (see Table 6). In the 30-second discrimination, rats were placed in context A for 160 s and were then immediately moved to a third context, C, for another 75 s where they received presentations of food during the final 45 s (i.e., the rats were placed in C for 30 s prior to the delivery of food). The rats were also placed in context B for the same amount of time, after which they were moved to context C where they received no food (i.e., A-30s→food and B-30s→no food). In the 90-second discrimination, after exposure to

contexts A and B, rats were placed in home cage for 60 s before they were placed in context C (for 75 s) where they received food (during the final 45 s) if they had recently been placed in B and no food if they had been placed in A (i.e., A-90s→no food and B-90s→food). The rate of food well entries in the first 30 s period in C (when no food was presented) was used as the index of whether rats had acquired the discrimination: I anticipated that rat would be more likely to approach the food well on trials on which A had been recently presented than when B had been recently presented (i.e., A-30s→food and B-30s→no food), and that they would be more likely to approach the food well when B had been presented remotely than when A had been presented remotely (i.e., B-90s→food and A-90s→no food).

3.2.1. Method

Subjects and apparatus

Sixteen naïve male Lister Hooded rats (supplied by Harlan Olac, Oxon, UK) were used in Experiment 4. The rats were housed in the same way as in Experiments 1A and 1B and were maintained at 80% of their ad-lib weight ($M = 351$ g, range = 323-370 g).

The apparatus was identical to that used in Experiments 1A and 1B with the exception that there were only four operant chambers in a single experimental room. The upper two chambers were decorated with either spotted (A) or checked (B) wallpaper that

was mounted behind transparent plastic panels. The lower two chambers were undecorated standard operant chambers. These chambers served as context C. A computer controlled the apparatus and recorded food well entries.

Procedure

Discrimination training. After magazine training (see details described in Experiments 1A and 1B), rats were given the 30-second and 90-second discriminations on alternate days for the following 40 days (i.e., 20 days of training for each discrimination). On each day, half of the rats were given the 30-second discrimination and the remainder were given the 90-second discrimination. Thus, half of the rats received the 30-second discrimination on odd days and received the 90-second discrimination on even days, and for the remainder this arrangement was reversed. In the 30-second discrimination, a given rat was placed in context A (e.g., spotted context) for 160 s and then immediately placed in context C (i.e., a blank box) for 75 s where it receives 10 food pellets on a fixed-time (FT) 5-s schedule during the final 45 s of the session. The rat was then removed from context C. After 40 min, the rat was placed in context B (e.g., checked context) for 160 s and immediately placed in C for 75 s where it received no food. In the 90-second discrimination, after exposure to context A for 160 s, the rat was moved to home cage for 60 s and then placed in C for 75 s where it received no food; whereas, after exposure to context B for 160 s, the rat was moved to home cage for 60 s

and was then placed in C where it received 10 food pellets on a FT 5-s schedule during the final 45 s of the session. Thus, a rat was given A-30s→food and B-30s→no food trials in the 30-second discrimination, and A-90s→no food and B-90s→food trials in the 90-second discrimination. In both discriminations, half of the rats were given context A in the first training session and context B in the second training session and the remainder received the reverse arrangement. The identities of the contexts that served as A and B were fully counterbalanced. For half of the rats, the spotted environment served as context A and checked environment served as context B, and for the remainder this arrangement was reversed.

Behavioural measures. The rate of food well responding during the first 30 s in context C was recorded. Discrimination ratios were used to assess the acquisition of the contextual discriminations. These ratios took the following form: rate of responding during the first 30 s period on the reinforced trials (e.g., A) divided by the rate of responding during the first 30 s period on both reinforced and nonreinforced trials (A+B). When this measure is used, a score above .50 indicates that discrimination training has been successful.

3.2.2. Results and Discussion

Figure 7 depicts the mean discrimination ratios for both the 30-second and 90-second discriminations over the course of 40 days of training. Inspection of this figure reveals that rats acquired both the 30-second and 90-second discriminations. ANOVA with discrimination (30-second or 90-second) and block as factors revealed an effect of block, $F(4, 60) = 2.91, p < .05$, but neither the effect of discrimination nor the interaction between these factors was significant, both $F_s < 1$. One-sample t tests revealed that discrimination ratios on the final block of both 30-second or 90-second discriminations were significantly above chance, $t(15) = 3.67, p < .01$, and $t(15) = 2.48, p < .05$, respectively. The rate of responding on the nonreinforced trials on the final block of the 30-second discrimination ($M = 8.71$ rpm) and of the 90-second discrimination ($M = 7.44$ rpm) did not differ significantly, $t(15) = 0.66, p > .05$.

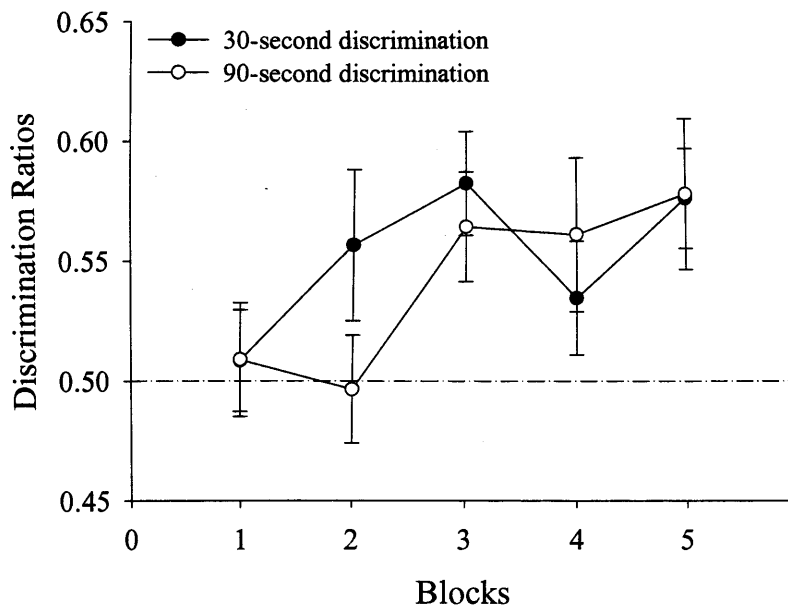


Figure 7. Experiment 4: Mean discrimination ratios (\pm SEM) for the 30-second and 90-second simple contextual discriminations.

The results of Experiment 4 are consistent with the suggestion that rats can acquire a simple contextual discrimination in which the immediate and short-term traces of the same stimulus become associated with different outcomes. Thus, the directly activated or immediate trace of context A became associated with food, whereas the short-term trace of A became associated with no food. In contrast, the immediate trace of B became associated with no food, whereas the short-term trace of B became associated with food. The results of Experiment 4 are, therefore, consistent with the interpretation offered for Experiments 1-3 - which relied on rats being able to learn distinct associations involving the immediate and short-term traces of a given stimulus. However, the analysis of Experiments 1-3 assumed that these distinct traces (or some

correlate of them) could become parts of different configural associations. In

Experiments 5 and 6 this assumption was directly assessed.

3.3. Experiment 5: Representing traces in configural conditioning I.

The design of Experiment 5 is summarised in Table 7. All rats were given two configural discriminations, one that involved contextual stimuli that were physically present (i.e., the context present condition) and another that involved the memory traces of the same contexts (i.e., the context trace condition). In the context present discrimination, when the rats were placed in context A, they received presentations of X followed by food and those of Y were followed by no food. In contrast, when they were placed in context B, they were given presentations of X that were followed by no food and those of Y were followed by food. In the context trace discrimination, rats were again placed in contexts A and B, and then were immediately moved to a third context, C, to receive presentations of X and Y. When the rats had previously been placed in A, presentations of X were paired with no food and those of Y were paired with food, whereas when the rats had previously been placed in B, presentations of X were paired with no food and those of Y were paired with food. The question of interest was whether rats can acquire the context present and trace configural discriminations,

thereby providing support for the suggestion that immediate trace (e.g., A) and short-term trace (e.g., a) of stimulus A can enter into independent (configural) associations.

Table 7: *The design of Experiments 5 and 6*

Configural training
Context present discrimination:
AX→food, AY→no food; BX→no food, BY→food
Context trace discrimination:
A-X→no food, A-Y→food; B-X→food, B-Y→no food

Note: A and B denote contexts (black and white), and X and Y denote auditory stimuli (tone and clicker); – indicates that there is an interval between presentations of the contexts and the delivery of X and Y in context C; food and no food indicate the outcomes of the training trials.

3.3.1. Method

Subjects and apparatus

Sixteen naïve male Lister Hooded rats (supplied by Harlan Olac, Oxon, UK) were used in Experiment 5. The rats were housed in the same way as in Experiments 1A and 1B and were maintained at 80% of their ad-lib weights ($M = 321$ g; range = 293-357 g).

The apparatus was the same as in Experiment 4 with the exception that one of the

chambers was decorated with black walls (upper left chamber) and the other with white walls (upper right chamber; see Honey & Watt, 1999).

Procedure

Following magazine training (using the procedure from Experiments 1A and 1B), rats received 32 days of configural training involving two types of conditional discrimination: context present and context trace. In the context present discrimination, a given rat was placed in context A for 160 s where it received presentations of X that were followed by food and those of Y were followed by no food; and when the same rat was placed in context B, presentations of X were followed by no food and those of Y were followed by food (i.e., AX→food, AY→no food, BX→no food, and BY→food). In the context trace discrimination, the same rat was placed in context A for 160 s and immediately moved to context C, where it received nonreinforced presentations of X and reinforced presentations of Y; after exposure to context B for 160 s, the same rat was moved immediately to context C, where it was given reinforced presentations of X and nonreinforced presentations of Y (i.e., A-X→no food, A-Y→food, B-X→food, B-Y→no food).

In the context present sessions (in A and B) and context trace sessions (in C), there were two 10-s presentations of X and Y with an ITI of 30 s. The sequences in which the stimuli were presented alternated across days, between XYYX and YXXY.

The two sessions from the context present discrimination occurred at one time of day, with an interval of 40 min between them; and the two context trace sessions occurred at a different time of day, with an interval of 40 min between them. Within a day, there was a 40 min interval between the second session of one type of discrimination and the first session of the second type of discrimination; and the order in which the sessions involving A and B were presented for both types of discrimination was consistent (A-B for half of the rats and B-A for the remainder). The order in which the context present and context trace discriminations were presented was counterbalanced across rats and alternated across days.

Configural learning was again assessed using a discrimination ratio where the rate of responding during the reinforced stimulus presentations was divided by the combined rate of responding for both reinforced and nonreinforced presentations. With this ratio, scores above .50 indicates that responding is greater during the reinforced stimulus than the nonreinforced stimulus.

3.3.2. Results and Discussion

Figure 8 depicts the mean discrimination ratios for both discriminations over the course of 32 days training. The results are presented in consecutive, four-session blocks.

Inspection of Figure 8 suggests that rats acquired both the context present and context

trace discriminations. ANOVA with discrimination and block as factors revealed that there was an effect of block, $F(7, 105) = 3.45, p < .001$, but neither the effect of discrimination nor the interaction between these factors was significant, $F_s < 1$. One-sample t tests revealed that the discrimination ratios on the final block of training for both context present and context trace discriminations were significantly above chance, $t(15) = 3.87, p < .01$ and $t(15) = 2.50, p < .05$. A paired-sample t test revealed that the rates of responding during the nonreinforced trials on the final block of discrimination training in the context present discrimination ($M = 6.33$ rpm) and of the context trace discrimination ($M = 8.88$ rpm) were significantly different, $t(15) = -2.23, p < .05$. This finding compromises any direct comparison of the rate at which the two discriminations were acquired; but given the fact that this was not the main purpose of Experiment 5, this is not a basis for great concern. The results of Experiment 5 are of primary importance because they suggest that the immediate and short-term trace of the same stimulus (e.g., A and a) can enter into distinct configural associations. The broader implications of these results and alternative interpretations of them will be considered after the presentation of the final experiment in this chapter, Experiment 6.

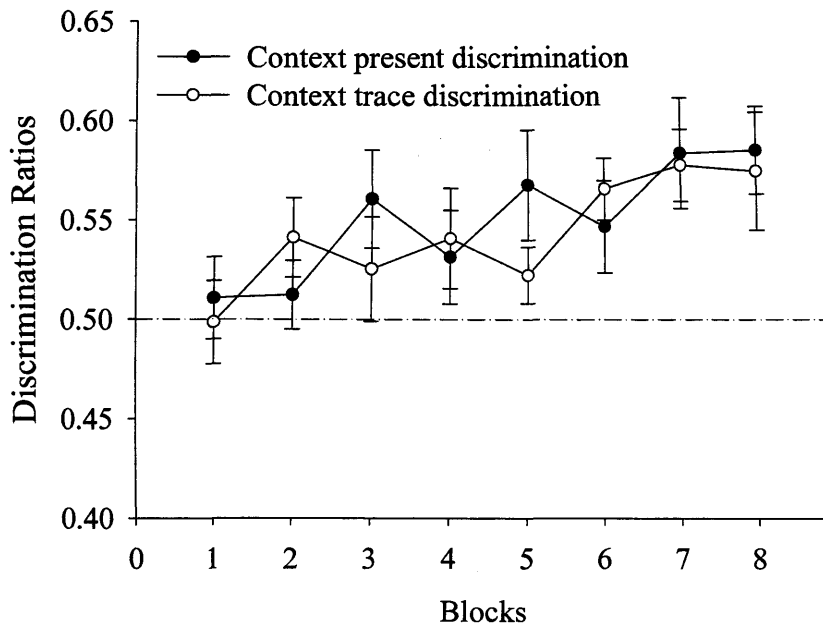


Figure 8. Experiment 5: Mean discrimination ratios (\pm SEM) for the context present and the context trace discriminations.

3.4. Experiment 6: Representing traces in configural conditioning II

The results of Experiment 5 suggested rats can acquire distinct configural associations involving the memory immediately activated by a stimulus and the short-term trace of the same stimulus. This suggestion is both novel and theoretically controversial - controversial in the sense that it is not anticipated by current theories of associative learning (e.g., Pearce, 1994; Wagner, 1981; but see, Brandon et al., 2003; Cole et al., 1995). Therefore, before accepting that rats are capable of this form of discrimination, alternative theoretically uninteresting interpretations should be considered; even if they are implausible. One such explanation for the results of Experiment 5 (that also applies to Experiment 4) relies on the fact that contexts A and B were in different locations.

This fact means that when rats entered context C from contexts A and B, different cues might have been available to them that could serve as a basis for the "trace" discrimination. For example, the visual cues (or handling cues) that the rat could perceive upon exiting context A (let us call them V1) might differ from those that they could perceive upon exiting context B (V2). If this was the case, then V1 and V2 (in Experiment 4) or V1 and V2 in conjunction with X and Y (in Experiment 5) could serve as a basis for discrimination learning in these experiments. One way to avoid this problem would be to replace A and B with punctate stimuli (e.g., two auditory stimuli) presented in the same experimental chamber. This was the approach taken in Chapter 4, where the theoretical basis for the type of discrimination observed in Experiment 4 was also further assessed. Experiment 6, however, used the same class of stimuli as in Experiments 4 and 5 (i.e., contexts), but modified the design so that the cues associated with movement from contexts A and B to context C were the same. To do so, the location in which contexts A and B were presented was always the same. The design of Experiment 6 was the same as Experiment 5, with the notable exceptions that contexts A and B were presented in the same location and were created by spotted and checked wallpapers rather than black and white wallpapers.

3.4.1. Method

Subjects and apparatus

Sixteen naïve Lister Hooded rats (supplied by Harlan Olac, Oxon, UK) were used in Experiment 6. The rats were housed in the same manner as in Experiments 1A and 1B and maintained at 80% of their ad lib weights ($M = 328$ g; range = 304-352 g). The apparatus was the same as that used in Experiment 5 with the exception that spotted and checked wallpaper served as contexts A and B in Experiment 6. Also, the box in which contexts A and B was housed was the same for a given rat (e.g., left); and the box housing context C was immediately below that housing contexts A and B. This was achieved by moving the decorated chambers from one box to another, and was so arranged to ensure that it was the immediate and short-term traces of A and B were the sole basis for the two types of discrimination.

Procedure

Following magazine training (using the procedure from Experiments 1A and 1B), rats received 36 days of configural training involving two types of conditional discrimination: context present and context trace. The procedure was identical to Experiment 5 with the exception that for the first 20 days of training rats were immediately placed in context C after being removed from contexts A or B; that is, the interval was nominally 0; and for the following 16 days, after being placed in contexts

A and B, rats were placed in a holding cage for 60 s and then placed in context C to receive presentations of auditory stimuli, X and Y. This change was introduced to reduce the likelihood that the immediate trace would remain active during the context trace discrimination.

Configural learning was assessed using a discrimination ratio. In addition to these ratio scores, the rates of responding during the first presentations of X and Y in each context were recorded on the final four days of training (i.e., the final block). This allowed me to confirm that discriminative performance was being controlled by the contexts and their traces rather than the reinforcement contingencies that were in force (i.e., whether food was presented after the tone or the click during a given session).

3.4.2. Results and Discussion

Figure 9 depicts the mean discrimination ratios for both the context present and context trace discriminations over the 36 days of training. Inspection of this figure indicates that rats acquired both the context present and the context trace discriminations - with the ratios for both increasing across the 4-day blocks of training. Unlike in Experiment 5, there was a clear indication that the context present discrimination was acquired more readily than the context trace discrimination. One plausible reason for this difference between the results of Experiments 5 and 6 is that in the former, but not the latter

experiment, rats could have used cues other than the traces of contexts A and B to solve the discriminations. In any case, statistical analysis confirmed the accuracy of the previous description of the results of Experiment 6.

ANOVA confirmed that there was an effect of block, $F(8, 120) = 3.64, p < .001$, and an effect of discrimination, $F(1, 15) = 44.04, p < .001$, but the interaction between these factors was not significant, $F(8, 120) = 1.02, p > .05$. One-sample t tests revealed that the discrimination ratios on the terminal three blocks of training were significantly above chance (.50) in context present discrimination, $t(15) = 8.15, p < .001$, and in the context trace discrimination, $t(15) = 2.18, p < .05$. The rates of responding on nonreinforced trials during these final blocks of discrimination training, in the context present ($M = 10.1$ rpm) and context trace ($M = 9.1$ rpm) discriminations, did not differ significantly, $t(15) = 1.24, p > .05$. Over the final four days of training (i.e., the final block of training), the mean discrimination ratios for the first trial responding for both the context present ($M = .64$) and context trace ($M = .58$) discriminations were different from chance, $t(15) = 4.13, p < .05$, and $t(15) = 2.30, p < .05$, respectively. The mean rate of responding of the first trial on the nonreinforced trials during the final block of discrimination training in the context present ($M = 10.22$ rpm) and context trace ($M = 8.25$ rpm) discriminations did not differ significantly, $t(15) = 1.44, p > .05$. These results thereby confirm the reliability of those from Experiment 5, and suggest that the

immediate and short-term traces of a given stimulus can enter into distinct configural associations. It should be acknowledged that there are several ways in which this general form of analysis could be implemented, and these should now be considered.

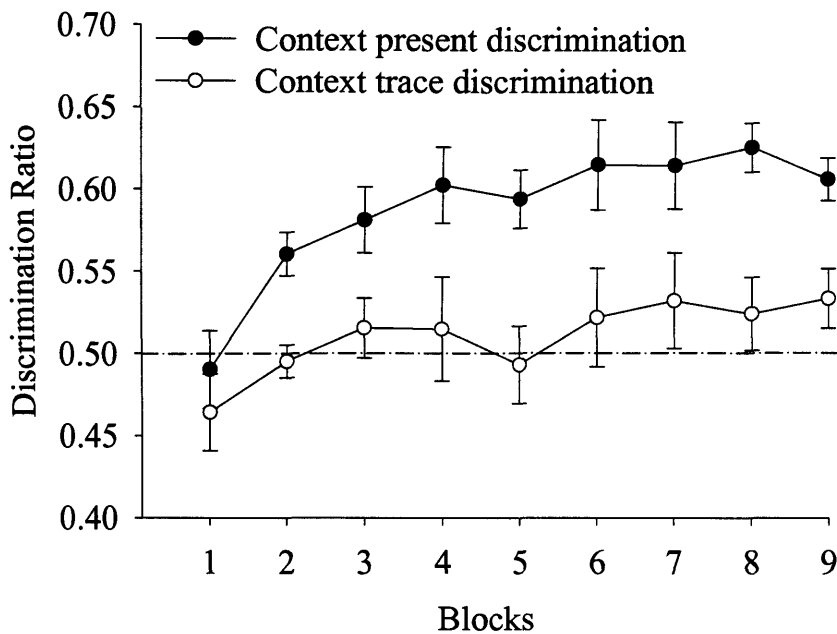


Figure 9. Experiment 6: Mean discrimination ratios (\pm SEM) for the context present and the context trace discriminations.

3.5. General discussion

The results of Experiments 4-6 suggest that the immediate trace (e.g., A) and the short-term trace (e.g., a) of the same stimulus can enter into distinct simple associations (Experiment 4) and configural associations (Experiments 5 and 6). For example, in Experiment 4, rats were more likely to approach the food well during the immediate trace of context A (i.e., A) than during the equivalent trace of B (i.e., B), but were more

likely to approach the food well during the short-term trace of B (i.e., b) than during the short-term trace of A (i.e., a). Similarly, in Experiment 6, rats were more likely to approach the food well during configurations involving the immediate trace of context A (i.e., AX) than during configurations involving the equivalent trace of B (i.e., BX), but were more likely to approach the food well during configurations involving the short-term trace of B (i.e., bX) than during configurations involving the equivalent trace of A (i.e., aX). In Experiments 5 and 6, rats acquired configural associations that involved the immediate trace of a stimulus (e.g., AX→food and AY→no food) and the short-term trace of the same stimulus (e.g., aX→no food, aY→ food). The pressing issue that now emerges is how the difference between the immediate trace of a stimulus (e.g., A) and the short-term trace of the same stimulus (i.e., a) should be conceived. There are three plausible alternatives that will be discussed: one based on temporal encoding (e.g., Cole et al., 1995), another based upon C-SOP (Brandon et al., 2003), and, finally, one based on a novel modification to Wagner's SOP that was offered in Chapter 2.

One possible explanation is to assume that what I have been referring to as immediate and short-term traces actually corresponds to temporal information about when a US will occur relative to the CS (or configuration): where the immediate trace corresponds to a short (or no) interval and the short-term trace corresponds to a long

interval. For example, rats might encode the temporal interval between the offset of CS and the onset of US during training (Barnet et al., 1991; Cole et al., 1995; Miller & Barnet, 1993, see also Desmond & Moore, 1988). That is, rats not only encode the relationship between CS and US, but also encoded the interval at which the US would be delivered. This temporal information (e.g., "60 seconds after context A") could then be used to generate appropriate performance in both a simple discrimination, and in conjunction with other cues (e.g., X and Y) in a configural discrimination.

Alternatively, the results of the current chapter can also be interpreted in terms of C-SOP (Brandon et al., 2003). It will be remembered from Chapter 1, that the basic assumptions of C-SOP are similar to Wagner's (1981) SOP model. The principal difference between the models that allows C-SOP to explain the results of this chapter is that the elements activated by a given stimulus vary in a consistent fashion across its duration and, presumably, during a trace interval that follows the offset of that stimulus. Thus, those elements that are activated by the presence (or recent presentation) of a context might be quite different than those that are active some time later. This state of affairs would allow the type of discriminations described in Experiments 4-6 to be solved on the basis of the different populations of elements that were active on immediate and short-term trace trials involving the same nominal stimulus.

The final alternative analysis is based upon a novel modification to Wagner's (1981) SOP. Wagner's (1981) SOP assumes that the mnemonic activity states (i.e., A1 and A2) influence the course of excitatory (and inhibitory) learning and also performance. The novel modification that is proposed is to allow that the state of the CS representation (A1 or A2) during encoding to become a component of long-term associative knowledge. Thus, in an analogous way to the suggestion that associative learning is context specific (e.g., Hall & Honey, 1989) what is being proposed is that associative learning is also "memory state" dependent. This analysis can be illustrated with reference to the results of Experiment 4. Here, when a stimulus is present or has recently been so, animals encode an association between the A1 state of that stimulus and the outcome, but when there is a trace interval then the A2 state of the same stimulus is encoded as part of the association with the outcome. This would allow the presentation of the same stimulus (e.g., A) to enter into distinct associations (e.g., A in A1→food and A in A2→no food) that would subsequently become evident in performance when the memories of stimulus A entered either the A1 or A2 states. This theoretical analysis could be considered analogous to the depth, or more immediately relevant, congruence of processing idea developed by Craik and Tulving (1975) in the context of human memory. The three classes of explanation presented above represent the focus of interest in Chapter 4.

In summary: The results presented in Chapter 3 provide general support for the interpretation of the conditions under which a configural preconditioning effect was observed in Chapter 2. That is, using similar procedures and stimuli, it was shown that the immediate and short-term traces of the same stimulus can enter into opposing simple and configural associations. Three plausible theoretical analyses were identified to explain these observations: the temporal encoding (e.g., Cole et al., 1995), C-SOP (Brandon et al., 2003), and a novel modification to SOP (Wagner, 1981). The experimental designs used in Chapter 3 (i.e., Experiments 4-6) do not allow a choice to be made between these three accounts identified above. The aim of the experiments presented in my final experimental chapter, Chapter 4, was to enable such a choice to be made.

Chapter 4

Functional equivalence of trace and associatively provoked memories

4.1. Introduction

One traditional and widespread view of memory processes holds that sensory inputs activated by a given item or stimulus are encoded temporarily into short-term memory stores (with limited capacity) and then transferred to long-term memory stores via a process of rehearsal (e.g., Atkinson & Shiffrin, 1968). According to this view, memory can be conveniently separated into a set of processes (encoding, consolidation and retrieval) that operate on stimulus traces distinguished by their longevity (immediate, short-term and long-term memory traces). A common view is that while the nature of the stimulus traces (immediate, short-term or indeed retrieved long-term traces) present during the acquisition of new long-term knowledge can have a profound influence on long-term memory formation, the resulting long-term traces are themselves *blind* with respect to the origin of this influence. This view is one that is common to many theories of associative learning in animals (e.g., Mackintosh, 1975; Pearce, 1994; Pearce & Hall, 1980, Rescorla & Wagner, 1972; Sutherland & Rudy, 1989; Wagner, 1981). To place this view in a concrete context, we can use the acquisition of associative knowledge during standard delayed conditioning and trace conditioning. It

is well established that the development of Pavlovian conditioned responding occurs more readily when the CS (e.g., a tone) immediately precedes or co-terminates with the US (e.g., food), than when there is a trace interval between the CS and US. The standard interpretation of this effect rests on the idea that the presentation of the tone activates its corresponding memory, and this immediate trace of the CS decays into a short-term trace once the tone is turned off. It is assumed that the immediate trace of the tone is more effective than its short-term equivalent in engaging the processes of learning. For example, Wagner (1981) assumes that during trace conditioning the memory of the CS is less likely to be in a form (i.e., in the A1 state) that engenders the development of an excitatory association. However, in general terms, the fact that different encoding conditions (immediate trace versus short-term trace) preceded the delivery of food is only represented in long-term memory by the fact that the associative bond is weak. To put it bluntly, a long-term association brought about by many trace conditioning trials is held to be equivalent to an association that has been brought about by a few delayed conditioning trials. In both cases, upon presentation of the tone the memory of food will only be weakly activated. As stated in Chapter 1, this assumption of *path independence* is one that is central to theories of animal learning and memory (e.g., Rescorla & Wagner, 1972; Wagner, 1981; 2003; McLaren & Mackintosh, 2002; Pearce & Hall, 1980).

In contrast to the views considered in the previous paragraph, the results of Chapters 2 and 3, however, suggest that the long-term associative memories of rats are not blind to the encoding conditions that were present during simple and configural learning. For example, rats were able to learn that the immediate trace of a stimulus predicts one outcome whereas the short-term trace of the same stimulus predicts a different outcome (see Experiments 4-6). An account of these findings was offered in the General Discussion to Chapter 2 that was based on a modified form of Wagner's (1981) SOP model. In particular, it was assumed that the nodal activity (A1 or A2) that was present prior to the delivery of an outcome (food or no food) could become represented as part of long-term associative memory. However, in the General Discussion of Chapter 3 it was also acknowledged that there were alternative interpretations of the results of Experiments 4-6. Namely, I considered explanations based upon the *C-SOP* model (Brandon et al., 2003) and the temporal coding hypothesis (Barnet et al., 1991; Cole et al., 1995; Miller & Barnet, 1993; see also, Desmond & Moore, 1988). The results presented in Chapter 3 did not allow these alternatives to be discriminated from one another. The aim of the Chapter 4, therefore, was to investigate different predictions made by these three accounts.

Part of the basis for supposing that a stimulus can provoke different types of mnemonic activity (A1 and A2) rather than activation having a continuous function (i.e.,

trace strength) is the observation that the immediate presentation of a stimulus (e.g., footshock in rats; see Blanchard & Blanchard, 1969; for other examples, see also Brandon et al., 2003; Wagner, 1981) can provoke a different response (i.e., heightened activity) than the trace of the same stimulus (i.e., freezing). It is difficult to imagine that differences in simple trace strength could provide a basis for these opposing response topographies. Rather these findings do seem to be more consistent with Wagner's (1981) assumption that the immediate presentation of a stimulus provokes one type of activity (i.e., A1) whereas the trace of the same stimulus provokes a different type of activity (i.e., A2). A critical further assumption of the model is that the form of activity that is provoked by association is the secondary, A2 state. For example, when an effective CS is presented it will provoke the A2 state in the US representation. That is, Wagner's model supposes that a memory retrieved by association is equivalent to the decayed form of a memory. If we take the latter assumption in conjunction with the novel modification to Wagner's theory outlined above (concerning A1 and A2 being encoded in the long-term association), a simple prediction can be generated. I will now consider this prediction in detail, contrasting it with the predictions from the alternative accounts (i.e., Brandon et al., 2003; Cole et al., 1995) as it forms the basis of the remainder of this thesis.

Imagine first that rats receive the same form of training that was given to those in Experiment 4. That is, they receive trials where there is a short interval (30 s) between the offset of the stimulus and the delivery of the outcome (i.e., A-30s→food) and trials with another stimulus where there is a longer interval (90 s) following the offset of stimulus and the outcome (i.e., B-90s→food). It is assumed that this arrangement allowed the A1 state of stimulus A (i.e., uppercase A) to be represented as a part of the information encoded in the long-term A→food association, and also allows the A2 state of stimulus B (i.e., lowercase b) to be encoded as a part of the long-term B→food association. That is, the long-term associations that are formed are coloured by the specific mnemonic activity that was present at encoding (i.e., A→food and b→food). Now imagine that stimulus A and B are used as second-order reinforcers for two visual stimuli, light immediate (LI) and light trace (LT). That is, the rats now receive pairings of LI with A and of LT with B (see Table 8). This stage should allow LI to activate the A2 state of stimulus A (i.e., a) and LT to activate the A2 state of stimulus B (i.e., b). If the first stage of training has resulted in the encoding of the specific associations described above (i.e., A→food and b→food), then LT should provoke more second-order responding than LI: LT will evoke the A2 state of stimulus B (b) which is linked to food, whereas LI will evoke the A2 state of stimulus A (a)

which is not linked to food. This highly counterintuitive prediction is not made by alternative accounts.

Without the novel modification to Wagner's (1981) model introduced above, SOP predicts that A should be a more effective second-order reinforcer than B; and therefore, LI should elicit more responding than LT. This is because the association between A and food should be stronger than that between B and food. The prediction that A should be a more effective second-order reinforcer than B also follows from C-SOP (Brandon et al., 2003) and the temporal coding hypothesis (Cole et al., 1995). C-SOP makes this prediction because, as a result of LI→A pairings, LI should provoke activity in the elements activated by the onset of A, and these would also have been likely to be contiguous with the delivery of food during the first stage of training (i.e., those activated when A is first presented). In contrast, LT will, as a function of second-order conditioning, provoke activity in elements that were not contiguous with food (i.e., those activated when B is first presented). The temporal coding hypothesis makes the same prediction, because whereas LI has, as a result of second-order conditioning, become linked to a stimulus that predicts that food will arrive shortly (i.e., A), LT has become linked to B, that predicts that food will not be delivered until much later on.

Table 8: *The within-subjects design used in Experiments 7 and 8*

First-order conditioning	Second-order conditioning
Immediate trace trials:	
X-10s→food, Y-10s→no food	LI→X
Short-term trace trials:	
X-40s→no food, Y-40s→food	LT→Y

Note: X and Y were auditory stimuli (tone or clicker); 10s (immediate trace) and 40s (short-term trace) were the intervals between the offset of X and Y and the outcomes of the trials (i.e., food or no food). LI and LT denote visual stimuli (left light or right light).

The accuracy of the differential predictions made by the novel modification to SOP, and both the original SOP, C-SOP and the temporal coding hypothesis were assessed in Experiments 7 and 8 using a second-order conditioning procedure outlined in Table 8. In this case, A and B were auditory stimuli and LI and LT were visual stimuli. The use of auditory and visual stimuli allowed me to adopt widely used second-order conditioning procedures that would not have been possible with the contexts used in Experiments 4-6. Moreover, it allowed the second-order conditioning procedure to be fully automated. Experiment 7 was a simple behavioural study, whereas Experiment 8 introduced a neural manipulation. The neural intervention

involved pre-training lesions targeted at the hippocampus. The rationale for this intervention will be presented in greater detail in the introduction to Experiment 8. For the present purposes, and as outlined in Chapter 1, it is sufficient to note that there is evidence showing that lesions to the hippocampal formation disrupt behavioural effects indicative of the presence of A1 and A2 activity states (e.g., associative priming; see Honey & Good, 2000a b, Honey, Watt & Good, 1998; nonassociative priming: Honey et al., 2007; Marshall et al., 2004). There is also complementary evidence that has been interpreted as indicating that the hippocampus is involved in temporal processing and short-term mnemonic processes (see Olton, Becker & Handelmann, 1979; Rawlins, 1985; Shors, 2004; Solomon, van der Schaaf, Weisz, & Thompson, 1986; but see, Kyd, Pearce, Haselgrove, Amin, & Aggleton, 2007). It is, therefore, of considerable general interest to ascertain whether hippocampal lesions disrupt the mnemonic processes assessed in the design outlined in Table 8.

4.2. Experiment 7: Assessing encoding specific associations

The first-order conditioning phase of Experiment 7 was similar to the training given to rats in Experiment 4 with the exception that auditory stimuli, rather than contexts, were used as X and Y. Rats received two types of discrimination that were intermixed during training sessions: An immediate trace discrimination and a short-term trace

discrimination. The immediate trace discrimination trials consisted of presentations of X that were followed by food after an interval of 10 s (i.e., X-10s→food) and presentations of Y that were followed by no food after an interval of 10 s (i.e., Y-10s→no food). Food well responding during the 10-s intervals that immediately followed X and Y allowed a measure of this simple discrimination to be assessed under comparable conditions to how learning on the short-term trace trials was assessed. That is, responding was assessed in the immediate traces generated by X and Y. The short-term trace trials consisted of presentations of X that were followed by no food after an interval of 40 s (i.e., X-40s→no food) and presentations of Y that were followed by food after an interval of 40 s (i.e., Y-40s→food). As with the immediate trace trials, development of the discrimination was assessed by recording responding during the 10-s periods that preceded the outcomes (food and no food). These specific intervals were chosen on the basis of the fact that similar intervals can result in systematic effects on discrimination learning (e.g., Honey & Hall, 1992). It was assumed that rats would acquire this discrimination: that is, they would show greater responding immediately after X than immediately after Y, and that the reverse would be the case when responding was assessed when period of 30-s had elapsed after the presentation of X and Y. According to the modified SOP analysis, the basis for this pattern of results is that rats will learn that the A1 state of X predicts food (i.e., X→food) and the A1 state

of Y does not (i.e., $Y \rightarrow \text{no food}$); and that the A2 state of Y predicts food (i.e., $y \rightarrow \text{food}$) and the A2 state of X does not (i.e., $x \rightarrow \text{no food}$). In addition, the A2 state of X does not predict the occurrence of food, because it is evoked after food on $X-10s \rightarrow \text{food}$ and food does not occur $X-40s \rightarrow \text{no food}$. Similarly, the A1 state of Y does not predict food because on the reinforced trial the A1 state would have decayed (into the A2 state) by the time food arrives, and food is not presented on the other type of trial on which the A1 state of Y would be active.

Of course, a more standard interpretation of the acquisition of such a "temporal" discrimination is that rats could use temporal information to guide their behaviour, and one example of this claim comes in the form of the temporal coding hypothesis (e.g., Barnet et al., 1991; Cole et al., 1995). Alternatively, they could base their behaviour on elements of X and Y whose activation is systematically correlated with elapsed time since the presentation of these stimuli (e.g., Brandon et al., 2003). The second stage allowed these accounts to be discriminated from one another, as mentioned in the Introduction to Chapter 4. In this stage, X and Y were used as second-order reinforcers for two visual stimuli, LI and LT: LI immediately preceded the delivery of X and LT immediately preceded Y. To recap, the modified version of SOP predicts that LT will elicit more second order conditioned responding than LI, whereas

the remaining theories (Brandon et al., 2003; Cole et al., 1995) predict the opposite outcome.

4.2.1. Method

Subjects

Sixteen naïve Lister Hooded rats were used. The rats were maintained in the same manner as Experiments 1A and 1B ($M = 335$ g; range = 313-362 g). The housing conditions were same as those in Experiments 1A and 1B.

Apparatus

Eight operant chambers (Test chamber 80004-D001; Campden Instruments Ltd., Loughborough, England), arranged in 4×2 array, were used. Each chamber (30.5 cm wide \times 26 cm deep \times 20 cm high) was positioned within a sound-attenuating box and had two aluminium side walls, a transparent perspex back wall and transparent perspex ceiling. The front wall was also constructed from transparent perspex and served as the door to the chamber. There was a food well (4.5 cm wide \times 3 cm deep \times 4 cm high) in a central position at the base of the left hand aluminium wall into which 45-mg food pellets (supplied by P. J. Noyes, Lancaster, NH) were delivered. A top-hinged transparent flap guarded access to this food well, and food-well entries were automatically recorded when the top-hinged magazine flap was pushed into the well by,

approximately, 1mm. A 3-W light bulb, with a white plastic cover, positioned centrally and at 13.5 cm above the floor, illuminated the chamber. Two 10-s auditory stimuli were used during the first-order conditioning stage and served as X and Y: a 2-kHz tone and a 10-Hz clicker. These stimuli, presented at an intensity of approximately 78 dB, were produced by an internal audio generator through a speaker located above the ceiling of the chamber. Two additional visual stimuli were used during second-order conditioning and served as LI and LT: illumination of covered 3-W jewel lights that were located on the left- and right-hand sides of the left aluminium wall that contained the food well. These lights were two of the three lights located on this wall. The central wall light, which was not illuminated during the experiment, was mounted 13.5 cm above the floor and was positioned over the food well. The lights used during second-order conditioning were mounted at the same height above the floor as the central light, but were displaced 9.2 cm to the left and right of the central light. These lights were both constantly illuminated throughout their 10-s durations. A 19-bar grid floor (stainless steel bars, diameter 0.47 cm, spacing from bar centre to bar centre, 1.07 cm) served as the floor of the chamber, beneath which was a tray that was lined with absorbent paper. A computer (Mark II Control Unit) controlled the apparatus, operated the program (using Behavioural Net Controller Control 1.0) and recorded food well

entries (all equipment and software was supplied by Campden Instruments Ltd., Loughborough, England).

Procedure

Discrimination training. After two days of magazine training (using the procedure described in Experiments 1A and 1B), there followed 56 days of discrimination training. On each day, rats received four types of trial: X-10s→food and Y-10s→no food (i.e., immediate trace trials), and X-40s→no food, Y-40s→food (i.e., short-term trace trials; see Table 8). On X-10s→food trials, stimulus X (e.g., a tone) was presented and followed by the delivery of two food pellets after an interval of 10 s. Y-10s→no food trials were the same as X-10s→food trials with the exceptions that stimulus Y (e.g., a clicker) replaced X and no food was delivered. On Y-40s→food trials, stimulus Y was presented and followed by the delivery of two food pellets after an interval of 40 s. X-40s→no food trials were the same as Y-40s→food trials with the exception that stimulus X replaced Y and no food was presented. There were five trials of each type per session that were presented in a pseudo-random order with the constraint that there were no more than two trials of the same type in succession. The inter-trial interval (ITI), as measured from the offset of the designated outcome (food or no food) and the onset of the next auditory stimulus, was 2 min.

On the next two days, in the morning session rats received refresher first-order conditioning trials. These trials were arranged in exactly the same way as on the preceding days. In the afternoon sessions, rats received second-order conditioning trials that were of two types: LI→X and LT→Y (see Table 8). On LI→X trials, the presentation of LI (e.g., the left light) immediately preceded the presentation of X, and on LT→Y trials, the presentation of LT (e.g., the right light) immediately preceded that of Y. For half of the rats, the tone served as X and the clicker served as Y, and for the remainder this arrangement was reversed. For half of the rats from each of the subgroups created by the previous counterbalancing operation, the left light served as LI and the right light served as LT, and for the remainder this arrangement was reversed. In each second-order conditioning session there were 6 trials of each type, and the order in which they were presented alternated and was counterbalanced: for half of the rats the alternating sequence was started with LI, and for the rest it began with LT. The ITI, as measured between the offset of LI or LT and the onset of X or Y, was also 2 min. A given rat received the same sequence on both days of second-order conditioning.

Behavioural measures. In order to assess the development of the immediate and short-term trace discriminations, ratios were used. These ratios were calculated using food-well responses during the 10-s trace periods before the outcome of the trial (food or no food). For the immediate trace discrimination (i.e., X-10s→food and Y-10s→no

food), the rate of responding during the 10-s period immediately following X (and preceding food) was divided by the combined rate of responding during the 10-s periods following both stimuli. For the short-term trace discrimination (i.e., X-40s→no food and Y-40s→food), the rate of responding during the final 10-s period of the 40-s interval that followed Y (and preceded food) was divided by the combined rate of responding during this period and the equivalent period after X. Using these ratios, scores of above .50 indicate that rats had acquired the discrimination. In fact, to assess the acquisition of the discriminations a comparison between the first and final seven-day block of discrimination training was used. The scores from the intervening blocks (immediate trace discrimination; blocks 2 = .56, 3 = .56, 4 = .53, 5 = .54, 6 = .53, 7 = .55; short-term trace discrimination; blocks 2 = .49, 3 = .49, 4 = .53, 5 = .55, 6 = .52, 7 = .53) took values between the first and final blocks (see Figure 11). However, there was a great deal of individual variation from one block to the next in the intervening blocks. Responding during X and Y during the same blocks was assessed in order to confirm that the trace intervals that were used produced a conventional trace conditioning deficit (i.e., with stimulus X eliciting greater responding than stimulus Y). During second-order conditioning, the rates of food-well responding during LI and LT were recorded.



4.2.2. Results

First-order conditioning. The rate of responding *during* presentations of auditory stimuli, X and Y, on the first block (the first seven days) and the final block (the final seven days) was assessed to evaluate whether or not the current procedures produced a standard trace conditioning deficit (see Figure 10). Inspection of Figure 10 reveals that the levels of responding were more marked to stimulus X than Y, and there was some indication that the overall levels of responding increased from the first to the final block. This description of the results was broadly supported by an ANOVA with stimulus (X versus Y) and block (first block versus final block) as factors. This analysis confirmed that there was an effect of stimulus, $F(1, 15) = 12.81, p < .01$, however, neither the effect of block nor the interaction between these factors were significant, $F(1, 15) = 3.53, p > .05$, and $F(1, 15) = 3.03, p > .05$, respectively.




Figure 10. Experiment 7: Mean rate of responding (\pm SEM) during X and Y on the first and the final block of the first-order conditioning stage.

The discrimination ratios used to gauge acquisition of the immediate trace and short-term trace discriminations, over 56 days of first-order conditioning training, are shown in Figure 11 for the first and final blocks. Inspection of this figure suggests that rats acquired both the immediate trace discrimination (left-hand side of Figure 11) and the short-term trace discrimination (right-hand side of Figure 11). ANOVA with factors of discriminate type (i.e., immediate trace versus short-term trace) and block (i.e., first versus final) revealed that there was an effect on block, $F(1, 15) = 4.72, p < .05$, but neither the effect of discrimination nor the interaction between these factors was significant, $F_s < 1$. A one-sample t test confirmed that the mean discrimination ratio on the final block of training ($M = .55$) was significantly above chance, $t(15) = 2.59, p$

< .05. During the first block of training, the rates of responding on the nonreinforced trials from the immediate trace and short-term trace discriminations (immediate trace: $M = 7.24$ rpm and short-term trace: $M = 6.43$ rpm) did not differ significantly, $t(15) = 1.43, p > .05$. Also, during the final block of training, the rates of responding on nonreinforced trials for the immediate trace and short-term trace discriminations (immediate trace: $M = 6.64$ rpm and short-term trace: $M = 5.69$ rpm) did not differ significantly, $t(15) = 2.07, p > .05$.

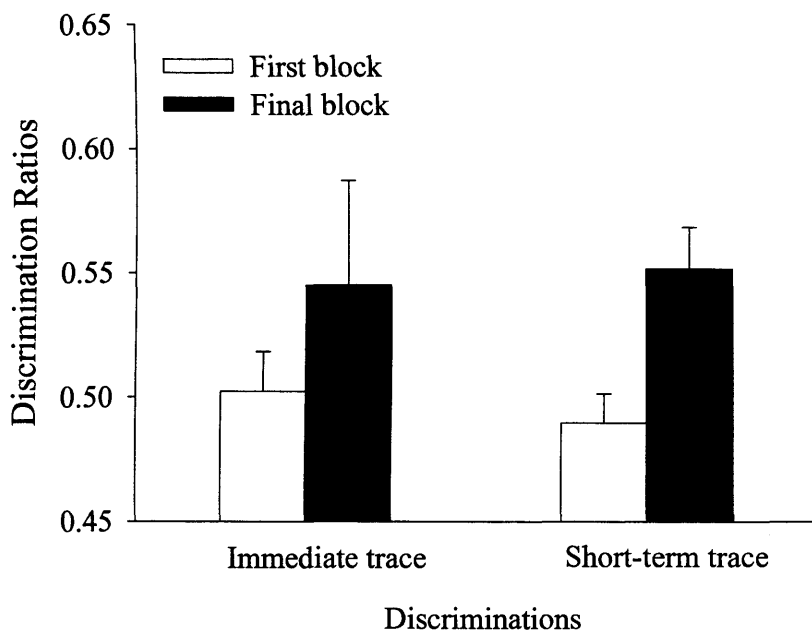


Figure 11. Experiment 7: Mean discrimination ratios (\pm SEM) on the first and the final block of the immediate and short-term trace discriminations during the first-order conditioning stage.

Second-order conditioning. After first-order conditioning stage, stimulus X and stimulus Y served as second-order reinforcers for LI and LT, respectively. The mean

rates of responding during X and Y on refresher trials still showed some indication of trace conditioning deficit ($X = 7.73$ rpm and $Y = 6.26$ rpm); however, the difference was no longer statistically significant $t(15) = 1.79, p > .05$. The ratios for the immediate trace and short-term trace discriminations were also similar to those during the final block of first-order training (immediate trace: $M = .58$ and short-term trace: $M = .56$). Paired-sample t test revealed that there was no difference between the ratios from the two discriminations, $t(15) = .27, p > .05$; and a one-sample t test confirmed that the overall or pooled discrimination ratios ($M = .57$) were significantly above chance, $t(15) = 2.27, p < .05$. The rates of responding during the nonreinforced trials for immediate trace, with a mean of 9.15 rpm, and short-term trace discrimination, with a mean of 7.95 rpm, did not differ significantly, $t(15) = 1.18, p > .05$.

Figure 12 depicts mean rates of responding during presentations of LI and LT. Inspection of this figure shows that the level of responding to LT was greater than that to LI. A paired-sample t test confirmed that LT provoked more responding than LI, $t(15) = -2.17, p < .05$.




Figure 12. Experiment 7: Mean rates of responding (\pm SEM) during presentations of LI and LT during the second-order conditioning stage.

4.2.3. Discussion

The results from the first-order conditioning stage of Experiment 7 (involving auditory cues) are similar to those observed in Experiment 4 (using contextual cues; see Chapter 3). In both experiments, rats learnt, in some way, that one stimulus (X) would be followed by food after a short interval, whereas another stimulus (Y) would not; and that once a longer interval had elapsed since stimulus Y, food would be delivered, whereas this was not the case after stimulus X. In the discussion of Chapter 3 (Section 3.5.) and the Introduction to Chapter 4 (Section 4.1.) several accounts of this finding

were identified: One based on C-SOP (Brandon et al., 2003), another based on temporal coding (e.g., Cole et al., 1995), and a final explanation based upon a novel modification to Wagner's (1981) SOP. The predictions of these theoretical analyses were evaluated in the second stage of Experiment 7 in which X and Y served as the second-order reinforcers for LI and LT, respectively. During this stage, LT came to elicit greater responding than LI. This is a counterintuitive finding, at a general level, because the second-order reinforcer for LT (i.e., Y) elicited less responding than the second-order conditioned for LI (i.e., X; see also Cole et al., 1995). It is also inconsistent with predictions derived from C-SOP and the temporal coding hypothesis; both of which predict that LI should come to elicit greater responding than LT. The results of Experiment 7 were, however, predicted by a modification to SOP that involved the suggestion that the state in which CSs are active at the point of US delivery (A1 in the case of X and A2 in the case of Y) become part of the long-term association involving those CSs. Once this idea is coupled with the suggestion that during second-order conditioning LI and LT will come to provoke X and Y into their respective A2 states, then second-order conditioning should be more effective with a CS that has been established when a relatively long interval has elapsed between CS and US (as it had in the case of Y).

The findings from Experiment 7, while inconsistent with the temporal coding hypothesis, are broadly consistent with the pattern of results observed by Cole et al. (1995). Cole et al. (1995) demonstrated that after first-order conditioning with either delayed conditioning trials (i.e., $X \rightarrow US$) or trace conditioning trials (i.e., X -interval $\rightarrow US$), rats received "backward" second-order conditioning trials with A (i.e., $X \rightarrow A$). Second-order conditioning to A was then assessed in the absence of X, and was found to be more marked in the group that had received trace than delayed first-order conditioning. The modification to SOP that provides an account for the results of Experiment 7 can also explain the results reported by Cole et al. (1995). Thus, following trace conditioning it can be assumed that the A2 state of X became associated with the US. Now, when X precedes A it is plausible to assume that it is the A2 state of X than is paired with A; that is, A will be paired with the stimulus (i.e., x) that was associated with the US during first-order conditioning. This will not be the case for those rats that received delayed conditioning trials for whom it is the A1 state of X (i.e., X) that became associated with the US, but for whom the A2 state of X (i.e., x) was paired with A. The novel modification of SOP provides an account for both the results of Experiment 7 and for those reported by Cole et al. (1995). Now, I will proceed to provide an alternative account for the results of Experiment 7, and to suggest one way

in which this account can be distinguished from the explanation based on my modification to SOP.

The alternative explanation is based upon the traditional notion of trace strength and decay (e.g., Hull, 1943). The first assumption that one needs to make is that the strength (or intensity) of the trace at the point of US delivery becomes a cue that is associated with that US. For example, according to this account, during Experiment 7 a strong trace of X will be associated with food whereas a relatively weak trace of Y will become associated with food. This will allow the immediate and short-term trace discriminations to be acquired, but it could also provide an explanation for the results from the second-order conditioning stage. Thus, when LI is paired with X and when LT is paired with Y, respectively, LT might activate a trace of Y that is like the weak trace that was associated with food; but that LI will activate a trace of X than is much weaker than the trace that was associated with food during the previous stage. Unlike the analysis based on a modification to SOP, the analysis based on trace strength relies on the specific choice of parameters resulting in the trace interval following Y generating a trace strength that happened to match that activated by LT. While this is possible, it certainly might be considered to be implausible. However, I chose to conduct a further experiment to distinguish between these alternatives. The manipulation used was a neurological one, but was one for which there was a very clear theoretical basis.

The general rationale for Experiment 8 was simple. Imagine that there is a manipulation (behavioural or neural) that results in more rapid decay between the A1 and A2 states. On the basis of my analysis for how rats solve the immediate and short-term trace discriminations, this manipulation should have a clear-cut effect: It should disrupt acquisition of the immediate trace discrimination, but not the short-term trace discrimination. However, if one simply assumed that a given trace decays in a continuous fashion across an interval, then any manipulation that influenced the immediate trace discrimination should also disrupt the short-term trace discrimination. These different predictions were assessed in Experiment 8 where the manipulation of interest was a selective lesion of the hippocampus.

4.3. Experiment 8: the role of hippocampal formation on trace discrimination

Olton et al. (1979) suggested that hippocampus is essential for tasks that require working memory - tasks that require the memory trace of one stimulus to be maintained in order to direct future behaviour (e.g., choosing the correct arm to visit in a radial arm maze after a delay). In a related vein, it has been shown that the hippocampus disrupts behavioural effects that seem to be best interpreted in terms of the mnemonic activity states (A1 and A2), and their consequences for performance, that are central to Wagner's (1981) SOP model. For example, Honey and Good (2000a) investigated the

influence of associative primes on the orienting response (OR) in rats. In brief, during training rats received presentations of a tone (e.g., X) that signalled the illumination of two constant lights (e.g., V1) and presentations of a clicker (e.g., Y) that signalled two pulsed lights (e.g. V2; i.e., $X \rightarrow V1$, $Y \rightarrow V2$). During the test, the presentation of X was followed by the illumination of V1 and V2 simultaneously and rats were more likely to orient to V2 than V1. This result suggests that a primed light (here V1) is less likely to elicit an OR than an unprimed or unexpected light (V2 in this case). This observation is consistent with Wagner's (1981) claim that when the representation of a stimulus is in the A2 state it prevents the presentation of the light from provoking A1 activity and thereby a marked OR. In rats with hippocampal lesions made prior to behavioural training, this effect was not observed. Instead, V1 (the primed light) elicited greater responding than V2 (the unprimed light). The same influence of hippocampal lesions has also been observed in cases of, so-called, self-generated priming (e.g., Honey et al., 2007; Marshall et al., 2004). Honey and Good (2000b, p. 203) suggested that these effects would be predicted by Wagner's SOP if there was rapid decay from A1 to A2. Briefly, under these conditions, the fact that there would be more elements active on a primed trial (albeit in the A2 state) should result in greater orienting than on an unprimed trials when fewer elements would be active (for a more detailed exposition of this argument, see Honey & Good, 2000b, p. 202).

On the basis of Honey and Good's (2000b) specific suggestion, lesions of the hippocampus should disrupt the immediate discrimination while leaving the short-term discrimination relatively uninfluenced. In the immediate discrimination, the A1 state of X should rapidly decay and the A2 state should remain active after the presentation of food; neither the A1 nor the A2 state of X will be a good predictor of food. Whereas, in the short-term discrimination, the A2 state of Y will be present prior to the delivery of food; and the rats should learn that the A2 state of Y predicts food whereas the A2 state of X predicts no food. It should not be possible to observe this type of dissociation on the basis of the view that trace strength is a continuous process; because a manipulation that disrupted a strong trace of a given stimulus (e.g., by reducing its strength) should have a similar effect on a weaker trace of the same stimulus.

The predictions detailed in the previous paragraph were assessed in Experiment 8. There were two groups of rats, one group had received hippocampal lesions made using ibotenic acid prior to behavioural training and the other group had received sham lesions. All rats then received the same form of discrimination given to rats in the first stage of Experiment 7. Following this training, rats received second-order conditioning in the same way as Experiment 7. This second stage of training allowed (1) an assessment to be made of the reliability of the results of Experiment 7 (in group Sham), and (2) an evaluation of the whether rats with lesions targeted at the hippocampus

acquired the immediate trace and short-term trace discriminations in the same way as rats in group Sham.

4.3.1. Method

Subjects, apparatus, and procedure

Thirty one naïve Lister Hooded rats were used in Experiment 8. Sixteen rats received sham operations (group Sham) and fifteen rats received ibotenic acid lesions of hippocampus (group HPC). Following a minimum two weeks of postoperative recovery, rats were gradually reduced to 80% of their ad lib weights ($M = 356$ g, range = 315-406 g). The housing conditions were the same as in Experiments 1A and 1B. The apparatus and the behavioural procedure that were used in this experiment were identical to Experiment 7, with the exception that in Experiment 8 there were 90 sessions of training in Stage 1. These sessions were combined to make 6 consecutive 15-session blocks for the purpose of statistical analysis.

Discrimination ratios were again used to assess acquisition of the immediate and short-term discriminations. However, in order to contrast the two groups of rats, I also used the number of blocks of training that it took the rats to acquire the discriminations to a criterion. The discrimination criterion that was used was .54. This level was achieved in both discriminations and by both groups in the majority of cases (Sham:

immediate discrimination = 11 rats and short-term discrimination = 8 rats; HPC:

immediate discrimination = 8 rats and short-term discrimination = 11 rats). Rats that did not attain this (albeit modest) criterion were given a score of 7, which represents the first block on which they could have achieved the criterion had training continued.

Surgery and histology

The surgical procedure was closely modelled on Marshall et al. (2004). Briefly, rats were first anaesthetized with Isoflurane and then placed in a stereotaxic frame (David Kopf Instruments, Tujunga, CA). After scalp incision, the bone overlying the area of neocortex directly above the hippocampus was removed, and injections of ibotenic acid (Biosearch Technologies, San Rafael, CA; dissolved in phosphate-buffered saline [pH 7.4] to provide a solution with a concentration of 63 mM) were made with 2- μ l Hamilton syringe mounted on the stereotaxic frame. Table 9 shows the coordinates and volume of infusions for rats in group HPC. Injections of 0.05-0.10 μ l were made at 28 sites with a KD Scientific electronic pump (Model 5000; Boston, MA) at a rate of 0.05 μ l/min. After each injection, the needle was left in position for 2 min to allow diffusion of the ibotenic acid and to limit the spread of the drug into overlying cortical areas. Sham-operated rats received an identical treatment with the exception that dura was perforated with a 25-gauge Microlance3 needle (Becton Dickinson, Drogheda, Ireland) and no fluid was infused.

Table 9. *Stereotaxic coordinates and volume of ibotenic acid for lesions of the hippocampus*

	AP	ML	DV	Volume (μ l)
From bregma:	-5.4	\pm 4.2	-3.9	0.10
		\pm 5.0	-6.1	0.08
	-4.7	\pm 4.0	-5.3	0.08
			-4.5	0.09
		-7.2	0.10	
		-3.5	0.05	
		\pm 4.5	-6.5	0.05
		\pm 2.2	-3.0	0.10
	-3.9	\pm 3.5	-1.8	0.10
			-2.7	0.10
		\pm 1.4	-3.0	0.10
		-2.1	0.10	
	-3.1	\pm 3.0	-2.7	0.10
		\pm 1.0	-3.0	0.05

Note: AP, ML and DV indicate the coordinates in relation to bregma from anterior to posterior (AP), from medial to lateral (ML) and from dorsal to ventral (DV).

Following the behavioural procedures, all rats received a lethal overdose of sodium pentobarbitone (Euthatal). The rats were then transcardially perfused, first with 0.9% saline and then with 10.0% formal-saline. Their brains were first extracted and postfixed for 24 hr, and then transferred to phosphate-buffered (0.1 M) 30.0% sucrose solution in which they remained for a further 24 hrs. Subsequently, all brains were frozen in a -20 °C cryostat and sectioned coronally. The 40- μ m sections were collected on gelatine-coated slides, left to dry at room temperature over 24 hrs and then stained

with cresyl violet. The sections were examined under a microscope and histological borders of hippocampal lesions were verified with reference to the boundaries defined by Paxinos and Watson (1998).

4.3.2. Results and Discussion

4.3.2.1 Histology

Figure 13 depicts a series of coronal sections of hippocampal formation adapted from Paxinos and Watson (1998), and shows the maximum (grey) and minimum (black) extent of cell loss for rats in group HPC. One rat was excluded from behavioural analysis due to extensive sparing of the dorsal and ventral parts of CA1, CA3, dentate gyrus as well as the ventral subiculum. Of the remaining 14 rats, 9 had extensive cell loss in the dorsal but less in the ventral part of hippocampal formation. The damaged areas in these 9 rats included CA1, CA2, CA3, dorsal subiculum and dentate gyrus including polymorph and granular dentate gyrus. These 9 animals also sustained fimbria damage. Two of the 14 rats had limited damage of dorsolateral CA1, CA2 and CA3, but left polymorph and granular dentate gyrus intact. Three of the 14 rats had more limited damage on pyramidal cells, radiatum layers and dorsal CA2, sparing most of dorsal and ventral CA1 and CA3. In 13 rats, there was limited damage to the posterior part of primary and secondary motor cortex; the remaining rat had damage to the

ventral part of the primary and secondary motor cortex. Also, in 13 rats there was damage to the anterior part of primary visual cortex, mediomedial and mediolateral secondary visual cortex. For the remaining rat only the ventral parts of these areas were damaged. Importantly, the amount of damage to the areas adjacent to the hippocampus was not correlated with the behavioural effects of interest.




Figure 13. Experiment 8: Histology. The maximum (grey) and minimum (black) extent of lesions in hippocampal rats. The sections are at specific distances (in mm) from Bregma (top to bottom: -2.4, -3.0, -3.9, -4.7, -5.4).

4.3.2.2. Behavioural results

First-order conditioning. Figure 14 depicts mean rates of responding during presentations of stimulus X and stimulus Y for both groups. Inspection of this figure suggests that there was a trace conditioning deficit in both groups, with stimulus X eliciting greater responding than stimulus Y. ANOVA with group (i.e., Sham versus HPC), stimulus (i.e., X versus Y) and block as factors revealed that there was an effect of block, $F(5, 140) = 8.54, p < .001$, but neither the effect of group nor stimulus was significant, $F < 1$ and $F(1, 28) = 1.98, p > .05$, respectively. However, the interaction between stimulus and block was significant, $F(5, 140) = 4.57, p < .01$, but there was no three-way interaction, $F < 1$. Analysis of simple main effects showed that the effect of stimulus was significant on block 6, $F(1, 28) = 4.85, p < .05$. It should be noted that the greater variability in group HPC was the consequence of the behaviour of a single rat who responded at approximately four times the rate of the mean of the remaining rats in group HPC.

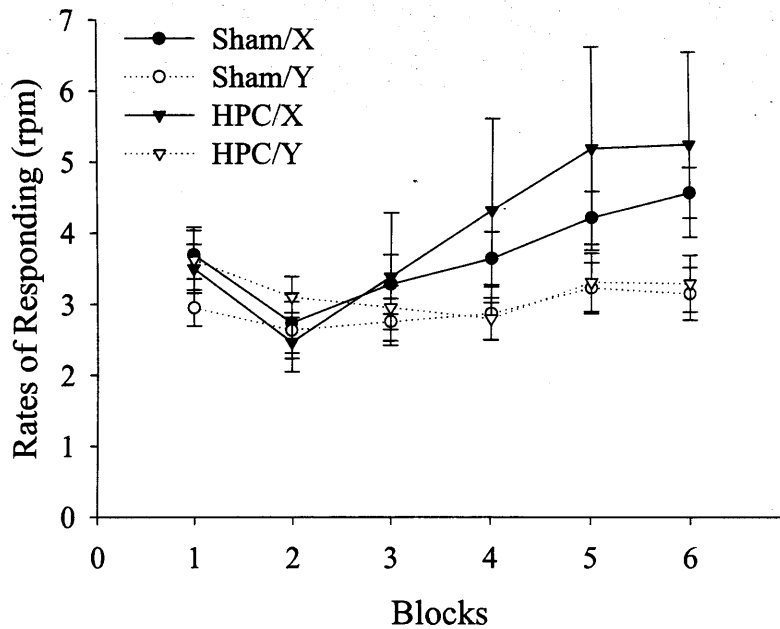


Figure 14. Experiment 8: Mean rates of responding (\pm SEM) during presentations of stimuli X and Y in groups Sham and HPC during the first-order conditioning stage.

Figures 15 and 16 show the acquisition of immediate trace and short-term trace discriminations for groups Sham and HPC. Comparison of Figures 15 and 16 suggests that in group Sham the immediate trace discrimination was acquired more effectively than the short-term discrimination. In contrast, in group HPC this difference was not observed; if anything the short-term trace discrimination was acquired more rapidly than the immediate trace discrimination. Also, comparison of the groups suggests that there was a selective impact of hippocampal lesions on the immediate trace

discrimination. ANOVA with group (Sham versus HPC), discrimination (immediate trace versus short-term trace) and block confirmed that there was an effect of block, $F(5, 140) = 3.85, p < .05$, but neither the effects of group nor discrimination were significant, $F_s < 1$. The three-way interaction between these factors was not significant, $F < 1$. However, the interaction between group and discrimination was significant, $F(1, 28) = 4.25, p < .05$. Analysis of simple main effects revealed that the difference between the overall discrimination scores approached the conventional level of statistical significance in group Sham, $F(1, 28) = 3.80, p = .06$, but not in group HPC, $F = 1$.

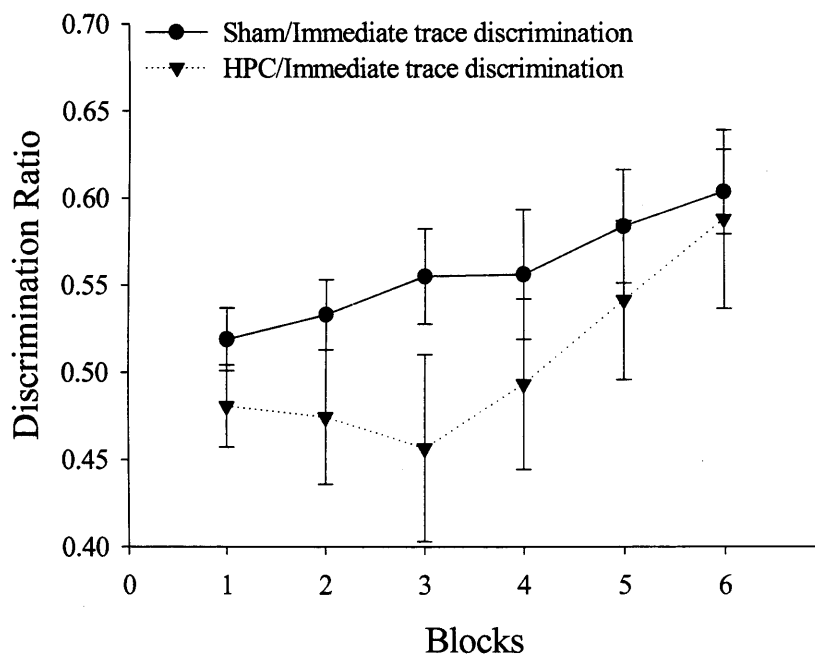


Figure 15. Experiment 8: Mean discrimination ratios (\pm SEM) for the immediate trace discrimination in groups Sham and HPC during the first-order conditioning stage.




Figure 16. Experiment 8: Mean discrimination ratios (\pm SEM) for the short-term trace discrimination in groups Sham and HPC during the first-order conditioning stage.

ANOVA conducted on the rates of responding on nonreinforced trials during the final block of discrimination training, with group and discrimination as factors (group Sham: immediate $M = 8.90$ rpm, and short-term $M = 6.20$ rpm; group HPC: immediate $M = 9.11$ rpm, and short-term $M = 6.35$ rpm) showed that there was an effect of discrimination, $F(1, 28) = 14.44, p < .05$, but neither the effect of group nor the interaction between these factors was significant, $F_s < 1$.

Figure 17 depicts the mean number of blocks that the rats required to reach criterion (i.e., .54). Inspection of the left-hand side of this figure suggests that rats in group Sham required fewer blocks to reach criterion in the immediate discrimination

than those in group HPC. In contrast, inspection of the right-hand side of the same figure suggests that there was no between-difference in the number of blocks to reach criterion in the short-term discrimination. Although, ANOVA with group and discrimination as factors revealed that the effects of group, discrimination or the interaction between these factors were not significant, largest $F(1, 28) = 1.90, p > .05$; Kruskal-Wallis tests showed that the difference between groups Sham and HPC was significant during the immediate trace discrimination, $H(2) = 3.96, p < .05$, but there was no such difference on the short-term trace discrimination, $H(2) = .01, p > .05$.

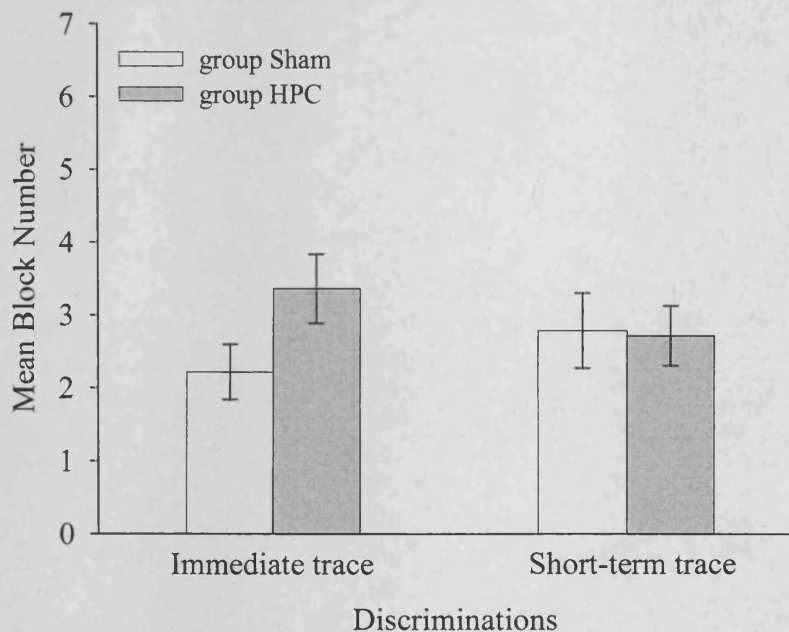


Figure 17. Experiment 8: The mean number of blocks (\pm SEM) to reach criterion for the immediate and the short-term trace discriminations in groups Sham and HPC.

Second-order conditioning. During the refresher trials there was a trace conditioning deficit in both groups Sham and HPC. The rates of responding during presentations of X (group Sham: $M = 6.20$ rpm, and group HPC: $M = 6.50$ rpm) were greater than those of Y (group Sham: $M = 4.47$ rpm and group HPC: $M = 4.00$ rpm). ANOVA with group and stimulus as factors confirmed that there was a main effect of stimulus, $F(1, 28) = 12.77, p < .001$, but neither the effect of group nor the interaction between these factors was significant, $F_s < 1$. The ratios for the immediate and short-term trace discriminations in group Sham (immediate trace: $M = .66$ and short-term trace: $M = .50$) and for group HPC (immediate trace: $M = .58$ and short-term trace: $M = .54$) were similar to the final block of first-order conditioning training. ANOVA with factors of group and discrimination revealed that there was no effect of discrimination or group, and no interaction between these factors, $F_s < 1$. One-sample t tests revealed that the overall discrimination ratios for the immediate and short-term trace conditions in each group (group Sham: $M = .58$ and group HPC: $M = .56$) were significantly above chance, $t(15) = 1.83, p < .05$ and $t(13) = 1.49, p < .05$, respectively. The mean rates of responding during the nonreinforced trials of group Sham ($M = 9.15$ rpm) and group HPC ($M = 7.95$ rpm) were not significantly different, $t(15) = 1.18, p > .05$.

Figure 18 depicts the mean response rates (in rpm) to two visual stimuli, LI and LT, in groups Sham and HPC during the second-order conditioning phase. Analysis of

the results on the second-order conditioning training was restricted to the first eleven 10-s nonreinforced presentation of LI and LT, because the level of responding on the final trial of each type was very low and variable². Examination of this figure shows that LT elicited more responding than LI in both groups, and that the level of responding was generally higher in group HPC than in group Sham. ANOVA with group (groups Sham and HPC) and stimulus (LI and LT) as factors revealed significant main effects of group, $F(1, 28) = 4.45, p < .05$, and stimulus, $F(1, 28) = 5.73, p < .05$, but there was no interaction between these two factors, $F < 1$.

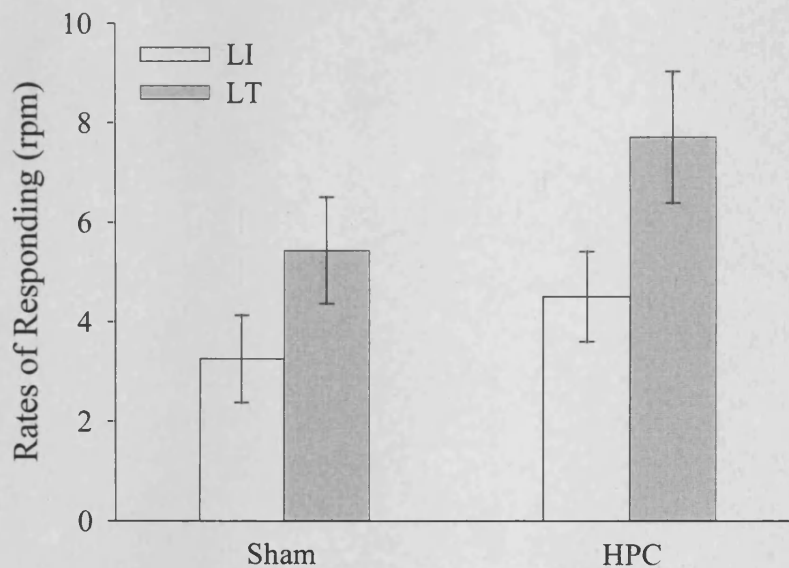


Figure 18. Experiment 8: Mean rates of responding (\pm SEM) to LI and LT of groups Sham and HPC in the second-order conditioning stage.

² ANOVA with group and stimulus (LI and LT) as factors that was conducted on all 12 trials revealed an effect of group, $F(1, 28) = 5.38, p < .05$, and an effect of stimulus that approached the conventional level of statistical significance, $F(1, 28) = 3.86, p = .059$. There was no interaction between these two factors, $F < 1$.

4.3.2.3. Discussion

The results from group Sham in Experiment 8 replicated, broadly speaking, the pattern of results observed in Experiment 7. First, stimulus X elicited greater responding than stimulus Y (i.e., there was a trace conditioning deficit). Second, a discrimination involving a short trace interval was acquired more rapidly than one involving a longer interval; however, unlike in Experiment 7 the short-term trace discrimination was not apparent (when analyzed in isolation) by the end of training. Finally, and most importantly, stimulus Y supported more second-order conditioning (to LT) than did stimulus X (to LI). The pattern of results in group HPC during X and Y was similar to that seen in group Sham – stimulus Y elicited more responding than stimulus X (cf. Kyd et al., 2007). However, the pattern of responding during the trace intervals themselves differed between the two groups: group HPC acquired the immediate trace discrimination less rapidly than group Sham, but there was no difference between the groups in the short-term trace discrimination. In the final test, in group HPC there was more second-order conditioning to LT than LI. This finding suggests that by the end of training in stage 1 the rats in the two groups had acquired the discrimination in the same manner. The fact that there was a higher overall rate of responding during second-order conditioning in group HPC than group Sham might simply reflect the fact that the two visual stimuli (LI and LT) elicited less orienting in the former than the latter group (cf.

Oswald, Yee, Bannerman, Rawlins, Good & Honey, 2002): Orienting to the lights would, therefore, be less likely to compete with food-well responses in group HPC than in group Sham. It should be acknowledged, however, that the difference in overall levels of responding during the second-order conditioning stage makes it somewhat difficult to assess the difference in responding to LI and LT between the two groups. That is, it should be recognized that there is a scaling effect.

There are several issues that remain to be discussed: If it is assumed that rats in group HPC have greater difficulty in associating the immediate trace of A with food (e.g., because the immediate trace rapidly decays), then what is the basis of the trace conditioning deficit that appears to be equally marked in group HPC as it is in group Sham. This issue will be addressed in the General Discussion that immediately follows. The general theoretical implications of the findings of Chapter 4 will be more fully discussed in Chapter 5.

4.4. General discussion

The series of experiments in Chapter 3, demonstrated that rats can acquire discriminations in which the time elapsed since the presentation of two contexts predicts whether food or no food will be delivered. These experiments join those demonstrating that rats are sensitive to time (see, for example, Meck, 2005) and

temporal arrangements of conditioning procedures (e.g., Cole et al., 1995; Desmond & Moore, 1988; Kamin, 1965). For example, in one type of discrimination, a stimulus A signalled that food would be delivered after a relatively short period and stimulus B signalled that no food would be delivered after the same amount of time; whereas in another simultaneously acquired discrimination, A signalled no food after a longer period of time and B signalled food after the same amount of time (see Experiment 4). Monitoring the levels of responding during the intervals after A and B indicated that the rats had acquired this discrimination: responding was more vigorous in thirty-second period after A than the same period after B, and was more vigorous in the sixty-to-ninety period after B than during the equivalent period after A. The fact that rats can acquire these discriminations can be understood in terms of several theoretical accounts: the modified version of SOP that I offered in Chapter 2, C-SOP, and temporal coding analysis (e.g., Brandon et al., 2003; Cole et al., 1995; see Section 4.1. for further details). The results reported in Chapter 3 did not allow one to make choice between these three classes of account, but those from Chapter 4 do.

The experiments reported in Chapter 4 (Experiments 7 and 8) used more conventional procedures, including the nature of the stimuli and temporal intervals (c.f., Honey & Hall, 1992). In Experiments 7 and 8, rats first received immediate trace conditioning trials with one stimulus (i.e., X-10s→food) and short-term conditioning

trials with another stimulus (i.e., Y-40s→food). Then they were given second-order conditioning trials in which two lights, LI and LT, were paired with the two auditory stimuli, X and Y, respectively. According to the modified SOP model, the first stage of training should result in the A1 state of X becoming linked to food, and the A2 state of Y becoming linked to food. If it is suppose that during second-order conditioning trials LI and LT come to evoke the A2 states of X and Y, respectively, then a simple prediction follows: Because LT evokes the memory of Y in the same state as it was paired with food whereas LI does not, then LT should elicit more responding than LI. The results of Experiments 7 and 8 support this prediction. As I have already illustrated (see Section 4.1), this pattern of results is neither predicted by C-SOP (Brandon et al., 2003) or the temporal coding hypothesis (Cole et al., 1995; see also, Barnet et al., 1991).

A further alternative analysis was developed in the discussion of Experiment 7 (Section 4.2.3.). This analysis was based on the idea that rats could use the strength of a trace to predict food. For example, a strong trace of X might be represented and associated with food; whereas a weaker trace of Y combined with the conditioning context might become associated with food. The first-order temporal discrimination would then be conceptually equivalent to the following discrimination: A→food, aC→no food, B→no food, and bC→food. According to this analysis any manipulation that reduced the strength of the traces of A and B should also influence the weaker

traces. In Experiment 8, however, rats with lesions of the hippocampus showed an impairment in learning the immediate trace discrimination (presumed to be equivalent to: A→food and B→no food), but not the short-term trace discrimination (equivalent to: bC→food and aC→no food). This selective influence of hippocampal lesions, or indeed any manipulation, is inconsistent with an analysis in terms of trace strength. Instead this dissociation is consistent with the idea that there are distinct activity states that can be separately influenced. In fact, there is already some evidence that is consistent with this suggestion from quite different procedures (e.g., Honey & Good, 2000b; Marshall et al., 2004).

There is one unresolved issue that I now need to return - in spite of the fact that there is no obvious satisfactory resolution of it. It is the fact that rats with hippocampal lesions showed a deficit in learning the discrimination that was based upon immediate traces, while showing a trace conditioning deficit that was equivalent to that shown by control rats. One explanation for this complex pattern of dissociations (i.e., involving both behavioural effects and lesions), is based on Honey and Good's (2000b) claim that lesions to the hippocampus results in more rapid decay from A1 to A2 state. If this is the case, then on X- 10s→food trials while some of the elements of X might be in A1 others might be in A2, and both sets of elements might become linked to food (cf. Brandon et al., 2003). The fact that some of X's elements are paired with food when

they are in the A1 state, whereas none of Y's elements would have been paired with food in the A1 state, might be a sufficient basis upon which to observe trace conditioning deficit. It might also be worth remembering that the trace conditioning deficit was only significant on the final block of training, and it was only on this block that rats in group HPC had acquired the immediate trace discrimination. Under these conditions, it is difficult to be certain about the basis for the trace conditioning deficit. This ambiguity concerning the effect of hippocampal lesions is, among the issues identified in Chapter 5, one that I should like to resolve in the future. One possibility would be to use a further pair of auditory cues in place of LI and LT.

In summary: Chapter 4 demonstrated two important empirical findings: one is that the short-term trace of a stimulus is functionally equivalent to an associatively provoked representation (Experiment 7); and the second is that the acquisition of a discrimination involving in the immediate traces of two stimuli can be dissociated from the acquisition of a discrimination based on the short-term traces of the same stimuli (Experiment 8). In Chapter 5, I will discuss in greater detail of the theoretical implications of Experiments 1-8.

Chapter 5

General Discussion

The overarching aim of the research reported in this thesis was to increase our understanding of the representational content of associative learning. Recent theoretical debate has focussed on the nature of the contributions made by elemental and configural processes to a variety of phenomena (e.g., McLaren & Mackintosh, 2000, 2002, Pearce, 1994; Wagner, 2003). It was noted, in the introduction to this thesis, that several theories of associative learning share two important and related assumptions: First, animals represent the patterns of stimulation that are physically presented on a given trial; second, while the mnemonic encoding conditions influence the acquisition of associative strength, these conditions do not form a part of the associative structure that is acquired (cf. Wagner, 1981). The experiments presented in Chapter 2 cast considerable doubt on the veracity of the first assumption, and those presented in Chapters 3 and 4, undermine the second assumption. I will now briefly review the principal findings from this thesis, before considering their general theoretical implications, and then describing the research that I would like to conduct in the future.

5.1. A brief summary of the new findings

There are four principal novel findings reported in this thesis. First, both the associatively provoked (Experiments 2 and 3) and trace memories of a stimulus (Experiments 5 and 6) can be assimilated into configural representations. That is, the content of configural representations extends to cues that are not physically present. Second, details of the conditioning trials correlated with the interval between the CS and US are represented as a component of the associative structure (Experiments 4-8). Thus, associative learning is not blind to the encoding conditions that obtained during conditioning. Third, the associatively activated memory of a stimulus (during second-order conditioning) is treated as similar or equivalent to the memory of the short-term trace of the same stimulus that was previously associated with food during trace conditioning (Experiments 7 and 8). Finally, the hippocampus has a selective role in the maintenance of the immediate trace of a stimulus, but not the secondary or short-term trace of the same stimulus (Experiment 8). It is now time to consider the detailed theoretical implications of these findings for both configural and elemental analyses of associative learning, and for our understanding of hippocampal function.

5.2. Theoretical implications

5.2.1. Configural assimilation

Pearce (1994) suggested that each time that an animal is confronted with a novel pattern of stimulation a new configural unit is recruited and this configural unit becomes associated with the outcome that is presented. If the same pattern should be re-presented then this configural unit would become fully active, but if the pattern differs then a new unit will be recruited. It is in this sense that activation of a hidden unit can be said to represent a previously experienced pattern. This analysis has been extended to provide an account of sensory preconditioning (see Pearce, 2002). For example, during exposure to a pattern, AB, a hidden unit (AB) might be recruited and bi-directional links might form between input unit A and the hidden unit (i.e., $A \rightarrow AB$ and $AB \rightarrow A$), and B and the same hidden unit (i.e., $B \rightarrow AB$ and $AB \rightarrow B$). Now, when A is paired with a US both input units A and B will become active, which will mean that no new hidden unit is recruited when A is paired with the US; and this will mean that B will be capable of eliciting responding at test because it will (partially) activate the AB unit. This form of analysis does not apply very readily to the configural preconditioning effects demonstrated in Chapter 2 (Experiments 2 and 3). In these experiments, new stimuli (X and Y) were added to those that were preexposed (i.e., A and C). This fact should ensure that new configural units are recruited. However, if

one were to relax this constraint, then X (or Y) might be assimilated into the AB configural unit (to form AbX) and provide a basis for generalization to the test compounds (e.g., aBX).

A more radical departure from configural analyses developed by Pearce (1987, 1994, 2002) is to allow preexposure to a compound to result in the formation of direct, input unit-to-input unit links (e.g., $A \rightarrow B$ and $B \rightarrow A$). These links would allow the associatively provoked memory of B, for example, to become active on $AX \rightarrow \text{food}$ trials. Again, AbX could mediate generalization to aBX at test. There is nothing in the results presented in this thesis that allows a choice to be made between these alternatives. However, recent research on the neural mechanisms that underlie different forms of sensory preconditioning, suggests the need to allow preexposure to patterns of stimulation to result in the formation of both configural representations and more direct links between the components of a pattern (see Iordanova, Burnett, Aggleton, Good & Honey, 2009).

The results of Experiments 5 and 6 show that the trace of stimuli (e.g., a and b) can become assimilated into configural representations (i.e., $aX \rightarrow \text{no food}$ and $bX \rightarrow \text{food}$) that are distinct from the configural representations directly activated by stimuli (i.e., $AX \rightarrow \text{food}$ and $BX \rightarrow \text{no food}$). This finding seems to pose another challenge for configural analyses of associative learning: If the trace of a stimulus

simply results in less activity in the input unit than does the direct application of the same stimulus, then it is not immediately clear how the configural discrimination involving a, b, A and B is solved (e.g., Pearce, 1987, 1994, 2002). One possibility relies on the following idea: as activity in a given input unit declines (e.g., from A to a) a different hidden unit becomes recruited through the increasing activation of the context input unit (C), which, in turn, is brought about by a reduction in the inhibition from input unit A to input unit C.

The analysis presented in the immediately preceding paragraph might be deemed implausible; and it is not clear how it deals with the results of Experiments 7 and 8. Why do the rats treat the trace of a stimulus conditioned during first-order conditioning in Experiments 7 and 8 (i.e., $C_y \rightarrow \text{food}$) as similar to the test configuration activated as the result of second-order conditioning (presumably L2y). The presence of L2 should result in the recruitment of a different hidden unit - one that is not associated with food. For this reason, it might be more parsimonious to assume that the trace memory of a stimulus (e.g., a) is simply equivalent to the associatively provoked memory of the same stimulus (i.e., a; cf. Wagner, 1981); and to allow this memory to enter into an association (elemental or configural) that is distinct from that involving the directly activated memory of the same stimulus (i.e., A). This view is considered in more detail in the next section.

5.2.2. Encoding specificity/congruence

Wagner's SOP model is the dominant elemental analysis of associative learning and memory. One important component of this model, perhaps its central component, is the distinction between different states of activity of a memory or node (A1 and A2). It is this distinction that allows the model to provide an analysis for the different "priming" phenomena (i.e., habituation and sensitization) that are a feature of both simple exposure and associative learning (e.g., Donegan, 1981; see also, Honey & Good, 2000b). The distinction between these two states also provides a way of explaining the transition between different forms of responding: The A1 state unconditionally activates one response (e.g., hyperactivity), whereas the A2 state activates a quite different response (e.g., freezing; Blanchard & Blanchard, 1969). Moreover, the identity between the decayed form of a memory for a given stimulus and the associatively provoked form of the same stimulus allows the model to anticipate the observed form of the conditioned response. The results of Experiments 7 and 8 suggest that not only can the A1 and A2 states unconditionally support different behaviours, but they can come to evoke different behaviours as the results of training. Presented in this way, the results of Experiments 7 and 8 seem to support an entirely parsimonious extension to the analysis of associative learning presented by Wagner (1981). At the same time, they reinforce the view that the associatively retrieved form of a memory

(i.e., A2) should be equated with the memory of a stimulus that has decayed. This analysis is reminiscent of the encoding specificity or congruence effects identified in humans (Craik & Tulving, 1975): when the encoding state in which learning occurs matches the state in which testing occurs retrieval is better than when these states do not match. Much closer to the research reported in this thesis, however, this idea provides one way of interpreting trace conditioning deficits (e.g., Pavlov, 1927) of the kind observed in Experiments 7 and 8. In standard, delayed conditioning the memory of the CS is in the A1 state when it is paired with food and when it is tested, but in trace conditioning the memory of the CS is in A2 when it is paired with the US and A1 when it is tested.

5.2.3. The neural bases of encoding specific long-term memories

A neural manipulation was used to distinguish between explanations for the encoding specificity/congruence effects based upon the modification to SOP described above and one based upon simple trace strength. The effect of this manipulation (lesions targeted at the hippocampus) - a selective disruption to a discrimination based on the immediate traces of stimuli - was more consistent with the account based upon modified SOP than that based upon trace strength. These results are consistent with the suggestion that the hippocampus supports the maintenance of A1, but not A2; a specific suggestion that

was based on the finding that associative priming effects are aberrant in rats with hippocampal lesions (Honey & Good, 2000a, 2000b), but that was foreshadowed by Olton et al. (1979; see also, Rawlins, 1985). One interesting issue that remains is whether such a deficit could explain other aspects of hippocampal function (cf. Iordanova et al., 2009) and another involves the locus of the short-term traces of stimuli. I will return to these issues in the next section that concerns the research questions that are important to pursue in the future.

5.3. Future work

5.3.1. Replication and extension

The idea that animals equate the trace memory of a stimulus with the associatively activated memory of the stimulus was assessed in Experiments 7 and 8 in a procedure that required many weeks of training. The theoretical analysis that I have developed suggests that such a complex training procedure should not be required in order to observe the effect of interest. Instead the procedure described in Table 11 could be used in which the first stage simply involved standard conditioning with one stimulus (X) and trace conditioning with a second stimulus (Y). After this training, which should only take a matter of days, second-order conditioning based upon Y should be more effective than that based upon X (see upper panel of Table 10). The theoretical

analysis that I have presented could be assessed further by reversing the order of the two stages, so that the procedure changes from one of second-order conditioning to one of sensory preconditioning (see lower panel of Table 10). The prediction is that L2 will provoke greater responding at test than L1, because L2 will evoke the A2 state of Y which was paired with food during Stage 2.

Table 10: *Simplified second-order conditioning and sensory preconditioning procedure.*

Stage 1	Stage 2
X-10s→food	L1→X
Y-40s→food	L2→Y

Stage 1	Stage 2	Stage 3
L1→X	X-10s→food	L1
L2→Y	Y-40s→food	L2

Note: X and Y denote auditory stimuli (tone or clicker); L1 and L2 denote visual stimuli (left light or right light); 10s and 40s indicate the trace interval between the offset of CS and the delivery of food.

5.3.2. Distinct associations involving A1 and A2

The theoretical analysis outlined above supposes that the A1 and A2 states of a given stimulus can enter into distinct associations. The evidence which supports this contention comes from the fact that animals can learn discriminations involving the immediate and short-term traces of stimuli (Experiments 4-8). Another way to assess the suggestion is summarized in Table 11. In Stage 1, rats (or indeed pigeons; cf. Jenkins & Moore, 1973) receive B→A pairings, and then in Stage 2 receive A-10s→food and A-40s→water trials. According to the modification of SOP model, the second stage training shall allow the immediate trace of A (i.e., A1) to become associated with food and the short-term trace of A (i.e., A2) to become associated with water. If the presentation of B during the test provoke the A2 state of A, then conditioned responses indicative of water rather than food should be observed (cf. Holland, 1983). Alternatively, two types of food (e.g., pellets and sucrose) could be substituted for food and water, and the effects of devaluing either one of the foods on test performance to B would be assessed. In this instance, the prediction is that devaluing the food presented 10 seconds after A should have a greater effect on performance to B than should devaluing the food presented 40 seconds after A.

Table 11: *Alternative assessment of distinction between A1 and A2 states*

Stage 1	Stage 2	Test
B→A	A-10s→food A-40s →water	B

Notes: A and B denote a tone and a clicker; 10s and 40s indicate the interval between the presentations of A and the outcomes of the trials (food and water).

5.3.3. Assessment of the nature of latent inhibition

The theoretical analysis advocated in this thesis has implications for our understanding of latent inhibition: When a novel stimulus is paired with food the modified SOP models supposes that it will be the A1 state of the CS that becomes associated with the US; whereas, when a preexposed stimulus is conditioned it will be the A2 state of the CS that becomes associated with the US. The fact that latent inhibition occurs might then reflect a difference in the learning rate about A1 and A2 or the fact that when two novel stimuli are paired they are both in the same state whereas when a familiar and a novel stimulus are paired they are in different states (for a recent analysis of the effect of similarity on conditioning, see Grand, Close, Hale & Honey, 2007). One prediction that follows from this analysis is that a preexposed CS should be a more effective second-order reinforcer than should a novel CS. The design of such experiment is

summarized in Table 12. The rats first receive preexposure to stimulus A and then receive pairings of A with food and a novel stimulus with food prior to second-order conditioning with L1 and L2. The prediction is that L1 will elicit greater responding than L2 because L1 will evoke a memory of A in the same state as it was conditioned (i.e., A2).

Table 12: *Assessment of the nature of latent inhibition*

Stage 1	Stage 2	Stage 3
A	A→food	L1→A
	B→food	L2→B

Notes: A and B denote auditory stimuli, tone and clicker; L1 and L2 denote a left light and a right light; food indicates the outcome of the trials.

5.3.4. Extension to an operant procedure

Thus far all of the effects that I have described use Pavlovian conditioning procedures. It would clearly be of interest to establish whether the same effects can be found in an operant conditioning procedure. The design for one such experiment is summarized in Table 13. Initially rats receive pairings of a discriminative stimulus with each of two responses (i.e., A→Left response and A→Right response) while responding on an instrumental baseline for an irrelevant food type. The rats would then receive pairings

of the Left response with the immediate delivery of food and pairings of the right response with the delayed delivery of food (i.e., Left response→food and Right-----→food). In the test, stimulus A would be presented and the expectation would be that rats would be more likely to produce the right response than the left response.

Table 13: *Extension to an operant procedure*

Stage 1	Stage 2	Test
A→Left response	Left response→food	A
A→Right response	Right response-----→food	

Notes: A denotes an auditory stimulus, ----- denotes a trace interval between the response and food.

5.3.5. Further dissociations of activity states and trace strength

It has been argued that the effects of hippocampal lesions demonstrated in Experiment 8 constitute a challenge to the idea that rats are using simple trace strength to predict when food will arrive. There are a number of more direct ways to assess this suggestion. For example, one could contrast the effects of a trace interval (observed in Experiments 7 and 8) with directly reducing the intensity of one of the stimuli (i.e., B; see Table 13). If this manipulation simply reduces the proportion of elements that enter into the A1

state, then the effects of this manipulation should be quite different from those of introducing a trace interval. First, in control rats, second-order conditioning should now be more effective to L1 than L2; simply because A has a stronger association with food than does B, and L1 and L2 will evoke memories of both A and B in a state (i.e., A2) that differs from the state in which they were paired with food in Stage 1 (i.e., A1).

Table 13: *Dissociations of activity states and trace strength*

Stage 1	Stage 2	Test
A→food	L1→A	L1
B→food	L2→B	L2

Note: A and B denote a tone and a clicker; L1 and L2 denote left light and right light. B in grey colour indicates stimulus B with less intensity.

5.3.6. Assessment of the neural basis of short-term traces

The results of Experiment 8 show that the hippocampus has a selective involvement in maintaining the immediate trace of a stimulus, with the maintenance of the short-term trace having a different neural basis. Two issues are prompted by this pattern of results. First, according to the analysis of hippocampal function that was just described, the role of the hippocampus should not be restricted to the acquisition of the immediate trace

discrimination, but should also be evident if the hippocampus was inactivated (e.g., by muscimol) following acquisition of this discrimination. Although the association in long-term memory would include the fact that stimulus A was in the A1 state when paired with food, the retrieval of this association would be precluded to the extent that the hippocampus could no longer maintain A in the A1 state at test. Second, the results of Experiment 8 suggest that there is a different neural basis for the short-term trace of a stimulus. The obvious question that follows from this is: what is the neural locus of the short-term trace? One obvious candidate is the prefrontal cortex. A number of studies have been suggested that this structure is involved in working memory in human and non-human animals (e.g., Goldman-Rakic, 1987; Kesner, Hunt, Williams and Long, 1996). Thus, Kesner et al (1996) suggested that the prelimbic and infralimbic prefrontal cortex is involved in visual object memory; and Porter, Burk and Mair (2000) also suggested that medial prefrontal cortex is involved recurring-choice delayed non-matching-to-sample tasks when the retention interval increased. Thus, it would be worth investigating the role of prefrontal cortex in the immediate and short-term trace discrimination tasks described in Chapter 4. The obvious prediction is that the pattern of results will be the opposite of that observed following hippocampal lesions in Experiment 8: Lesions of the prefrontal cortex should disrupt the short-term trace discrimination, but not the immediate trace discrimination.

5.4. General summary

The findings reported in this thesis have begun to increase our understanding of the content of associative learning in both complex configural tasks and simpler Pavlovian conditioning procedures. The evidence that has been presented shows that our current theories of associative learning (e.g., Pearce, 1994; Pearce & Hall, 1980; Mackintosh, 1975; Wagner, 1981) are too restrictive in their analysis of the content of associative learning: The long-term representations that animals can acquire involve both things that are not physically present and details of the mnemonic conditions that obtained during acquisition. The next challenge is to develop a formal model that captures these new insights and to analyse the neural bases of these mnemonic processes in greater detail.

References

- Allman, M. J., & Honey, R. C. (2006). Transfer of configural learning between the components of a preexposed stimulus compound: Implications for elemental and configural models of learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 32, 307-313.
- Asratyan, E. A. (1961). The initiation and localization of cortical inhibition in the conditioned reflex arc. *Annals of the New York Academy of Sciences*, 92, 1141-1159.
- Asratyan, E. A. (1965). *Compensatory adaptations, reflex activity and the brain*. Oxford: Pergamon.
- Atkinson, R. C., & Estes, W. K. (1963). Stimulus sampling theory. In R. D. Luce, R. B. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology, Vol. 3* (pp. 121-268). New York: Wiley.
- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: a proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation: advances in theory and research, Vol.2* (pp. 89-195). New York: Academic Press.
- Aydin, A., & Pearce, J. M. (1995). Summation in autoshaping with short- and long-duration stimuli. *Quarterly Journal of Experimental Psychology*, 48B, 215-234.
- Barnet, R. C., Arnold, H. M., & Miller, R. R. (1991). Simultaneous conditioning demonstrated in second-order conditioning: Evidence for similar associative structure in forward and simultaneous conditioning. *Learning and Motivation*, 22, 253-268.
- Beylin, A. V., Gandhi, C. C., Wood, G. E., Talk, A. C., Matzel, L. D., & Shors, T. J. (2001). The role of the hippocampus in trace conditioning: Temporal discontinuity or task difficulty? *Neurobiology of Learning and Memory*, 76, 447-461.

- Blanchard, R. J., & Blanchard, D. C. (1969). Crouching as an index of fear. *Journal of Comparative and Physiological Psychology*, *67*, 370-375.
- Brandon, S. E., & Wagner, A. R. (1998). Occasion Setting: Influences of conditioned emotion response and configural cues. In N. Schmajuk (Ed.), *Occasion setting: Associative learning and cognition in animals* (pp. 343-282). Washington, DC: American Psychological Association.
- Brandon, S. E., Vogel, E. H., & Wagner, A. R. (2000). A componential view of configural cues in generalization and discrimination in Pavlovian conditioning. *Behavioural Brain Research*, *110*, 67-72.
- Brandon, S. E., Vogel, E. H., & Wagner, A. R. (2003). Stimulus representation in SOP: I Theoretical rationalization and some implications. *Behavioural Processes*, *62*, 5-25.
- Brogden, W. J. (1939). Sensory pre-conditioning. *Journal of Experimental Psychology*, *25*, 323-332.
- Cole, R. P., Barnet, R. C., & Miller, R. R. (1995). Temporal encoding in trace conditioning. *Animal Learning & Behavior*, *23*, 144-153.
- Craik, F. I. M., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General*, *104*, 268-294.
- Desmond, J. E., & Moore, J. W. (1988). Adaptive timing in neural networks: The conditioned responses. *Biological Cybernetics*, *58*, 405-415.
- Desmond, J. E., & Moore, J. W. (1991). Altering the synchrony of stimulus trace processes: test of a neural-network model. *Biological Cybernetics*, *65*, 161-169.
- Dickinson, A., & Burke, K. (1996). Within-compound associations mediate the retrospective revaluation of causality judgements. *Quarterly Journal of Experimental Psychology*, *49B*, 60-80.

- Donegan, N. H. (1981). Priming-produced facilitation or diminution of responding to a Pavlovian unconditioned stimulus. *Journal of Experimental Psychology: Animal Behavior Processes*, 7, 295-312.
- Estes, W. K. (1959). The statistical approach to learning theory. In S. Koch (Eds). *Psychology: A study of a science, Vol. 2* (pp. 380-491). New York: McGraw-Hill.
- Gallistel, C. R. (1990). *The organisation of behaviour*. Cambridge, MA: Bradford Books/MIT Press.
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate and conditioning. *Psychological Review*, 107, 289-344.
- Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behaviour by representational memory. In V. B. Mountcastle & F. Plum (Eds.), *Handbook of physiology: The nervous system, higher functions of the brain, Vol. 5* (pp. 373-417). Bethesda, MD: American Physiological Society.
- Good, M., & Honey, R. C. (1991). Conditioning and contextual retrieval in hippocampal rats. *Behavioral Neuroscience*, 105, 499-509.
- Grand, C., Close, J., Hale, J., & Honey, R. C. (2007). The role of similarity in human associative learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 33, 64-71.
- Grand, C., & Honey, R. C. (2008). Solving XOR. *Journal of Experimental Psychology: Animal Behavior Processes*, 34, 486-493.
- Hall, G. (1996). Learning about associatively activated stimulus representations: Implications for acquired equivalence and perceptual learning. *Animal Learning & Behavior*, 24, 233-255.
- Hall, G., & Honey, R. C. (1989). Contextual effects in conditioning, latent inhibition, and habituation: Associative and retrieval functions of contextual cues. *Journal of Experimental Psychology: Animal Behavior Processes*, 15, 232-241.

- Holland, P. C. (1981). Acquisition of representation-mediated conditioned food aversions, *Learning and Motivation*, 12, 1-18.
- Holland, P. C. (1983). Representation-mediated overshadowing and potentiation of conditioned aversions. *Journal of Experimental Psychology: Animal Behavior Processes*, 9, 1-13.
- Honey, R. C., & Good, M. (2000a). Associative modulation of the orienting responses: distinct effects revealed by hippocampal lesions. *Journal of Experimental Psychology: Animal Behavior Processes*, 26, 3-14.
- Honey, R. C., & Good, M. (2000b). Associative components of recognition memory. *Current Opinion in Neurobiology*, 10, 200-204.
- Honey, R. C., & Hall, G. (1992). CS memory after trace conditioning. *Learning and Motivation*, 23, 145-155.
- Honey, R. C., Marshall, V. J., McGregor, A., Futter, J., & Good, M. (2007). Revisiting places passed: Sensitization of exploratory activity in rats with hippocampal lesions. *Quarterly Journal of Experimental Psychology*, 60B, 625-634.
- Honey, R. C., & Ward-Robinson, J. (2002). Acquired equivalence and distinctiveness of cues: I. Exploring a neural network approach. *Journal of Experimental Psychology: Animal Behavior Processes*, 28, 378-387.
- Honey, R. C., & Watt, A. (1998). Acquired relational equivalence: implications for the nature of associative structures. *Journal of Experimental Psychology: Animal Behavior Processes*, 24, 325-334.
- Honey, R. C., Watt, A., & Good, M. (1998). Hippocampal lesions disrupt an associative mismatch process. *The Journal of Neuroscience*, 18, 2226-2230.
- Hull, C. (1943). *Principles of Behavior*. New York : Appleton-Century-Crofts.

- Iordanova, M. D., Burnett, D. J., Aggleton, J. P., Good, M., & Honey, R. C. (2009). The role of the hippocampus in mnemonic integration and retrieval: complementary evidence from lesion and inactivation studies. *European Journal of Neuroscience*, *30*, 2177-2189.
- James, W. (1890). *The Principles of Psychology*. New York: Holt.
- Jenkins, H. M., & Moore, B. R. (1973). The form of the autoshaped responses with food and water reinforcers. *Journal of the Experimental Analysis of Behavior*, *20*, 163-181.
- Kamin, L. J. (1965). Temporal and intensity characteristics of the conditioned stimulus. In W. F. Prokasy (Ed.), *Classical conditioning: A symposium* (pp. 118-147). New York: Appleton-Century-Crofts.
- Kehoe, E. J. (1988). A layered network model of associative learning: Learning to learn and configuration. *Psychological Review*, *95*, 411-433.
- Kehoe, E. J., & Gormezano, I. (1980). Configuration and combination laws in conditioning with compound stimuli. *Psychological Bulletin*, *87*, 351-378.
- Kesner, R. P., Hunt, M. E., Williams, J. M., & Long, J. M. (1996). Prefrontal cortex and working memory for spatial response, spatial location, and visual object information in the rat. *Cerebral Cortex*, *6*, 311-318.
- Kyd, R. J., Pearce, J. M., Haselgrove, M., Amin, E., & Aggleton, J. P. (2007). The effects of hippocampal system lesions on a novel temporal discrimination task for rats. *Behavioral Brain Research*, *187*, 159-171.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*, 276-298.
- Marshall, V. J., McGregor, A., Good, M., & Honey, R. C. (2004). Hippocampal lesions modulate both associative and nonassociative priming. *Behavioral Neuroscience*, *118*, 377-382.

- McDonald, R. J., Murphy, R. A., Guarraci, F. A., Gortler, J. R., White, N. M., & Baker, A.G. (1997). Systematic comparison of the effects of hippocampal and fornix-fimbria lesions on acquisition of three configural discriminations. *Hippocampus*, 7, 371-388.
- McLaren, I.P.L., & Mackintosh, N.J. (2000). An elemental model of associative learning: I. Latent inhibition and perceptual learning. *Animal Learning & Behavior*, 28, 211-246.
- McLaren, I.P.L., & Mackintosh, N.J. (2002). Associative learning and elemental representation: II. Generalization and discrimination. *Animal Learning & Behavior*, 30, 177-200.
- Meck, W. H. (2005). Neuropsychology of timing and time perception. *Brain and Cognition*, 58, 1-8.
- Miller, R. R., & Barnet, R. C. (1993). The role of time in elementary associations. *Current Directions in Psychological Science*, 2, 106-111.
- Olton, D.S., Becker, J.T., & Handelmann, G.E. (1979) Hippocampus, space and memory. *Behavioural Brain Sciences*, 2, 313-365.
- Oswald, C. J. P., Yee, B. K., Bannerman, D. B., Rawlins, J. N. P., Good, M., & Honey, R. C. (2002). The influence of selective lesions to components of the hippocampal system on the orienting response, habituation and latent inhibition. *European Journal of Neuroscience*, 15, 1983-1990.
- Pavlov, I. P. (1927). *Conditioned reflexes*. London: Oxford University Press.
- Paxinos, G., & Watson, C. (1998). *The rat brain in stereotaxic coordinates*, 4th Ed., Sydney: Academic Press.
- Pearce, J. M. (1987). A model of stimulus generalization for Pavlovian conditioning *Psychological Review*, 94, 61-73.

- Pearce, J. M. (1994). Similarity and discrimination: A selective review and a connectionist model. *Psychological Review*, *101*, 587-607.
- Pearce, J. M. (2002). Evaluation and extension of a connectionist theory of configural learning. *Animal Learning & Behavior*, *30*, 73-95.
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not unconditioned stimuli. *Psychological Review*, *87*, 532-552.
- Porter, M. C., Burk, J. A., Mair, R.G. (2000). Distinctive effects of hippocampal and cortical lesions on delayed non-match-to-position: A comparison of tasks trained with retractable and in automated radial arm mazes. *Behavioural Brain Research*, *109*, 69-81.
- Rawlins, J. N. P. (1985). Associations across time: The hippocampus as a temporary memory store. *Behavioural Brain Sciences*, *8*, 479-496.
- Redhead, E. S., & Pearce, J. M. (1995). Similarity and discrimination learning. *Quarterly Journal of Experimental Psychology*, *48B*, 46-66.
- Rescorla, R. A. (1972). "Configural" conditioning in discrete-trial bar pressing. *Journal of Comparative and Physiological Psychology*, *79*, 307-317.
- Rescorla, R. A., & Cunningham, C. L. (1978). Within-compound flavour associations. *Journal of Experimental Psychology: Animal Behavior Processes*, *4*, 267-275.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.
- Revusky, S. H. (1968). Aversion to sucrose produced by contingent X-irradiation: Temporal and dosage parameters. *Journal of Comparative and Physiological Psychology*, *65*, 17-22.

- Rudy, J. W., & Sutherland, R. J. (1995). Configural association theory and the hippocampus formation: an appraisal and reconfiguration. *Hippocampus*, 5, 375-389.
- Saavedra, M. A. (1975). Pavlovian compound conditioning in the rabbit. *Learning and Motivation*, 6, 314-326.
- Schreurs, B. G., & Westbrook, R. F. (1982). The effect of changes in the CS-US interval during compound conditioning upon an otherwise blocked element. *Quarterly Journal of Experimental Psychology*, 34B, 19-30.
- Shores, T. J. (2004). Memory traces of trace memories: Neurogenesis, synaptogenesis and awareness. *Trends in Neurosciences*, 27, 250-256.
- Solomon, P. R., van der Schaaf, E. R., Weisz, D., & Thompson, R. F. (1986). Hippocampus and trace conditioning of the rabbit's classically conditioned nictitating membrane responses. *Neuroscience*, 100, 729-744.
- Sutherland, R. J., & Rudy, J. W. (1989). Configural association theory: The role of the hippocampal formation in learning, memory, and amnesia. *Psychobiology*, 17, 129-144.
- Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behaviour. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5-48). Hillsdale, NJ: Erlbaum.
- Wagner, A. R. (2003). Context-sensitive elemental theory. *Quarterly Journal of Experimental Psychology*, 56B, 7-29.
- Wagner, A. R. & Brandon, S. E. (2001). A componential theory of Pavlovian conditioning. In R. R. Mower & S. B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 23-64). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

- Wagner, A. R., & Rescorla, R. A. (1972). Inhibition in Pavlovian conditioning: Application of a theory. In R. A. Boakes & M. S. Halliday (Eds.), *Inhibition and learning* (pp. 301-336). New York: Academic Press.
- Ward-Robinson, J., & Hall, G. (1996). Backward sensory preconditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 22, 395-404.
- Ward-Robinson, J., & Honey, R. C. (2000). A novel contextual dimension for use with an operant chamber: From simple to hierarchical forms of learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 26, 358-363.
- Wilson, P.N., & Pearce, J.M. (1990). Selective transfer of responding in conditional discriminations. *Quarterly Journal of Experimental Psychology*, 42B, 41-58.
- Woodbury, C. B. (1943). The learning of stimulus patterns by dogs. *Journal of Comparative Psychology*, 35, 29-40.

