On the role of the hippocampus in the acquisition, long-term retention and semanticisation of memory

Sarah M Gingell

Doctor of Philosophy

Institute for Adaptive and Neural Computation
School of Informatics
University of Edinburgh

2005

Abstract

A consensus on how to characterise the anterograde and retrograde memory processes that are lost or spared after hippocampal damage has not been reached. In this thesis, I critically re-examine the empirical literature and the assumptions behind current theories. I formulate a coherent view of what makes a task hippocampally-dependent at acquisition and how this relates to its long-term fate. Findings from a neural net simulation indicate the plausibility of my proposals.

My proposals both extend and constrain current views on the role of the hippocampus in the rapid acquisition of information and in learning complex associations. In general, tasks are most likely to require the hippocampus for acquisition if they involve rapid, associative learning about unfamiliar, complex, low salience stimuli. However, none of these factors alone is sufficient to obligatorily implicate the hippocampus in acquisition. With the exception of associations with supra-modal information that are always dependent on the hippocampus, it is the combination of factors that is important.

Detailed, complex information that is obligatorily hippocampally-dependent at acquisition remains so for its lifetime. However, all memories are semanticised as they age through the loss of detailed context-specific information and because generic cortically-represented information starts to dominate recall. Initially hippocampally-dependent memories may appear to become independent of the hippocampus over time, but recall changes qualitatively. Multi-stage, lifelong post-acquisition memory processes produce semanticised re-representations of memories of differing specificity and complexity, that can serve different purposes.

The model simulates hippocampal and cortical interactions in the acquisition and maintenance of episodic and semantic events, and behaves in accordance with my proposals. In particular, conceptualising episodic and semantic memory as representing points on a continuum of memory types appears viable. Support is also found for proposals on the relative importance of the hippocampus and cortex in the rapid acquisition of information and the acquisition of complex multi-model information; and the effect of existing knowledge on new learning. Furthermore, episodic and semantic events become differentially dependent on cortical and hippocampal components. Finally, as a memory ages, it is automatically semanticised and becomes cortically-dependent.

Acknowledgements

My supervisors were Professor David Willshaw and (initially) Dr. Richard Shillcock.

Many people play a role in the processes surrounding the production of a PhD thesis. Most obviously this includes supervisors, but when problems arise, an important role should also be played by Faculty staff, Graduate advisors, those involved in the University's complaint procedures, the Dean, the Head of Department *etc*. I would like to thank these people for their involvement as appropriate.

I would specifically like to thank David Willshaw for staying on board for the second submission of this thesis and especially for his help in the last few weeks, and Alexander Gutkin for help with *that* line of code. Thanks should also go to the various office-mates I have had over the years, and to other friends and family, for stopping asking...

I thank the EPSRC for providing funding for three years, The President's Trust for a small grant, and Mark Corti for a generous hand-out.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(*Sarah M Gingell*)

Contents

1 Introduction **1**

1.1 Motivation for the thesis 1

1.2 Aims of the thesis 2

1.3 Boundaries of thesis 2

1.4 Outline of the thesis 3

1.5 Introduction to the hippocampus 5

2 Theories of hippocampal function **9**

2.1 Introduction 9

2.2 Qualitative theories of hippocampal function 11

2.2.1 "Declarative" versus "procedural" memory 11

2.2.2 Spatial theories 15

2.2.3 Other hypotheses about the role of the hippocampus 18

2.3 The long-term role of the hippocampus 23

2.3.1 The standard view: temporary storage 23

2.3.2 Long-term involvement of the hippocampus 24

2.4 Computational theories of the hippocampus 25

2.4.1 Models of declarative function 25

2.4.2 Models of sequence learning 28

2.4.3 Models of spatial function 29

2.4.4 Other models 31

2.5 Summary 32

2.5.1 A note about terminology: 'episodic' and 'semantic' memory . . 32

3 The elusive role of the hippocampus **35**
3.1 Introduction . . . 35
3.1.1 Issues for hippocampal theories . . . 36
3.2 Task acquisition and hippocampal dependency . . . 41
3.2.1 Empirical considerations . . . 43
3.2.2 The acquisition of some tasks is unaffected by hippocampal damage . . . 45
3.2.3 The acquisition of some tasks is impeded by hippocampal damage 53
3.2.4 The acquisition of some tasks is totally prevented by hippocampal damage . . . 59
3.2.5 The acquisition of a few tasks is improved by hippocampal damage . . . 66
3.3 Summary and conclusions . . . 67
3.3.1 The nature of hippocampally-dependent tasks . . . 70

4 The nature of hippocampally-dependent learning **73**
4.1 Introduction . . . 74
4.2 Speed of learning . . . 77
4.3 Convergence of information . . . 81
4.4 Incidental and automatic learning . . . 85
4.5 Representational structure . . . 89
4.6 Conclusions . . . 91
4.6.1 Relationship to other proposals . . . 93
4.6.2 Episodic and semantic memory . . . 94

5 Long-term role of the hippocampus **99**
5.1 Introduction . . . 99
5.2 Retrograde amnesia after hippocampal damage . . . 100
5.2.1 Empirical issues . . . 100
5.2.2 Information whose acquisition is unaffected by hippocampal damage may be retained or lost after hippocampal damage . . . 103
5.2.3 Information whose acquisition is impeded by hippocampal damage shows graded retrograde amnesia after hippocampal damage . . . 105

5.2.4 Information whose acquisition is prevented by hippocampal damage shows flat retrograde amnesia after hippocampal damage 110
5.2.5 The retrograde effect of hippocampal damage on tasks whose acquisition is improved by hippocampal damage is currently not known . . . 116
5.2.6 Degenerative and non-discrete disorders . . . 117
5.3 Summary and conclusions . . . 119
5.3.1 The long-term role of the hippocampus . . . 121

6 Semanticisation and the role of the hippocampus 123
6.1 Introduction . . . 123
6.1.1 Traditional Consolidation Theory . . . 124
6.1.2 Multiple Trace Theory . . . 126
6.1.3 Semanticisation theory . . . 128
6.2 Graded retrograde amnesia results from the semanticisation of memories 130
6.2.1 Decay of detailed information . . . 130
6.2.2 Increased ability of semantic traces to mediate recall . . . 131
6.2.3 Change in recall strategy . . . 133
6.2.4 Episodic and semantic memory revisited . . . 135
6.3 Semanticisation is a multi-stage process . . . 137
6.3.1 Scaffolding and inter-regional information transfer . . . 139
6.4 The temporal extent of graded retrograde amnesia . . . 141
6.5 Regional redundancy and trace decay . . . 145
6.6 Interleaved learning and semanticisation . . . 147
6.7 Reconsolidation and semanticisation . . . 148
6.8 The long term role of the hippocampus . . . 152
6.8.1 Key features of Gingell's Semanticisation Theory . . . 152
6.9 Summary and conclusions . . . 158
6.9.1 Relationship to other proposals . . . 161

7 A simple model of episodic and semantic learning 165
7.1 Introduction . . . 166
7.2 Network architecture . . . 168
7.3 Implementation of the model . . . 169

7.3.1 Training phase . . . 170
7.3.2 Test phase . . . 170
7.3.3 Modulation of learning . . . 171
7.3.4 Offline learning . . . 172
7.3.5 Training and testing schedule . . . 172
7.3.6 Representation of events . . . 172
7.3.7 Parameter values . . . 175
7.3.8 Visualising the data . . . 175
7.4 Findings . . . 177
7.4.1 Control data . . . 177
7.4.2 Learning and decay rates . . . 180
7.4.3 The role of different sets of weights . . . 183
7.4.4 Learning different kinds of information . . . 188
7.4.5 Modulation of learning . . . 194
7.4.6 Semanticisation of memory . . . 198
7.5 Summary and conclusions . . . 203
7.5.1 Limitations of model and further work . . . 206

8 Conclusions **209**
8.1 Main conclusions of thesis . . . 210
8.2 Contributions of the thesis . . . 214
8.3 Predictions . . . 215
8.4 Final words . . . 216

Bibliography **217**

Chapter 1

Introduction

1.1 Motivation for the thesis

The hippocampus has been implicated in memory function since Scoville and Milner (1957) noted a link between hippocampal damage and amnesia in human patients. Since then there has been a burgeoning experimental literature on the hippocampus — several thousand such papers are now published each year — and there are many theories of hippocampal function. One might expect that an ever-expanding body of data should more tightly constrain speculation so that theories would increasingly converge. However, a consensus on the role of the hippocampus remains elusive, and many current proposals appear incompatible. Even when there is agreement as to what data is important (which often there is not), there are widely different interpretations of even the same data. Although most authors agree that the hippocampus has *some* role in *certain* sorts of memory, there is little agreement about the fundamental nature of hippocampally-dependent tasks, about the role of the hippocampus in the acquisition of tasks, or the long-term role of the hippocampus in the maintenance of information.

A re-examination of the empirical data shows that the acquisition of tasks that are typically considered to be "hippocampally-dependent" is not equally impaired by hippocampal lesions. Similarly, different types of information, and even different sub-components of information acquired in the same learning episode, appear to have different long-term fates in terms of hippocampal dependency. Since deficits that are different in degree or that refer to different material are likely to depend on quanti-

tatively and qualitatively different facets of hippocampal function, it is important to make such distinctions for the purposes of formulating hippocampal theories. Broadly speaking, the motivation for this thesis was the belief that by acknowledging such issues, I would be able to reach a better understanding of the role of the hippocampus.

1.2 Aims of the thesis

The research question addressed by this thesis is:

What is the nature of the relationship between the dependence of information on the hippocampus at acquisition and its long-term dependency on the hippocampus?

More specifically, this thesis addresses the question by attempting to:

- Strictly define the nature of tasks that are dependent on the hippocampus for their acquisition, taking into account differences in the degree of effect of hippocampal damage on acquisition.
- Identify common traits within groups of tasks whose acquisition is affected similarly by hippocampal damage, and explain how these traits relate to the role of the hippocampus in acquisition.
- Strictly define the nature of hippocampal long-term involvement in the maintenance of information.
- Determine how the long-term maintenance of information by the hippocampus relates to the nature of the information initially acquired.
- Formulate proposals about the fate of memories after initial acquisition by the hippocampus.
- Provide a 'proof of concept' for the key ideas proposed in this thesis through the design, implementation and testing of a computational model.

1.3 Boundaries of thesis

This thesis focuses on what the hippocampus itself does, and therefore does not deal in detail with the role of related memory structures.

Similarly, as the thesis is concerned with understanding the role of the hippocampus at a cognitive rather than mechanistic level, issues relating to the specific functional mechanisms of the hippocampus are not addressed.

The model makes no attempt at biological realism, instead providing a method for testing the plausibility of proposals put forward in this thesis.

1.4 Outline of the thesis

In this chapter I outline the motivation, aims and structure of this thesis. I briefly describe the key features of hippocampal anatomy in section 1.5.

In *Chapter 2*, I review theories and perspectives on the hippocampus, and introduce existing models of hippocampal function.

Despite years of research, there is no consistent view on how to characterise the memory processes that are lost or spared when the hippocampal formation is damaged. In *Chapter 3*, I explore several issues that may have contributed to this confusion, and re-examine the empirical literature on task acquisition and the hippocampus. Crucially, tasks that are traditionally considered to be hippocampally-dependent at acquisition are in fact affected to different extents and in different ways by hippocampal damage: some information (such as task-dependent allocentric spatial information) is obligatorily dependent on the hippocampus for acquisition, whereas other information (such as conditional motor learning) is merely acquired faster or more easily when the hippocampus is intact. In general, tasks are most likely to require the hippocampus for acquisition if they involve rapid, complex, arbitrary associative learning about novel, low salience stimuli. So-called 'procedural' factors, such as the number of trials available or the intensity of task stimuli can have as profound effect on task acquisition after hippocampal damage as the logical high-level nature of the task.

In *Chapter 4* I argue that there are both quantitative and qualitative differences between the functions of the hippocampus and other areas: it is the interaction between these features and the nature of the information to be acquired that determines the relative extent to which the hippocampus and other areas can acquire information under given task conditions. Many of the features of tasks that were identified in chapter 3 as being hippocampally-dependent, such as their speed of acquisition, complexity, salience or novelty, can be envisaged as continuous dimensions. I suggest that tasks that 'score' highly on one or more of these dimensions, or contain many such features

even if not at extreme values, will require the hippocampus if learning is to proceed in the normal fashion. I discuss the contribution of these different task features to information acquisition, and thereby both extend and constrain previous proposals relating to rapid learning and complex associative learning by the hippocampus. I also show that traditional definitions of 'episodic' and 'semantic' memory need refinement, and provide definitions for use in the rest of the thesis.

The long-term role of the hippocampus in supporting recall is logically separable from its role in trace acquisition. In *Chapter 5*, I examine the long-term role of the hippocampus in mediating the recall of tasks that are affected by it at acquisition. I conclude that memories that are obligatorily hippocampally-dependent at acquisition remain hippocampally-dependent for their lifetime (which may be less than the lifetime of the animal), whereas the recall of traces whose acquisition is merely facilitated by the hippocampus may gradually become independent of the hippocampus. Typically as memories age they become less detailed and more generic or 'semanticised': this is central to understanding the long-term role of the hippocampus in memory.

In *Chapter 6* I argue that the graded retrograde amnesia seen after hippocampal damage on tasks whose acquisition is merely facilitated by an intact hippocampus stems from the gradual semanticisation of memories over time. Semanticisation can result from the decay of the most detailed information, the enhancement of semanticised representations, or from a change over time in the recall strategy used to access stored information. Multi-stage, lifelong post-acquisition memory processes act to produce semanticised re-representations of memories of differing degrees of complexity, which can support performance on different tasks depending on their requirements. Information that is obligatorily hippocampally-dependent at acquisition is difficult to re-represent outside the hippocampus, and therefore stays dependent on the hippocampus for its lifetime. However, since all information tends to decay with time, a memory's lifetime in its original form will usually be less than the lifetime of the organism.

In *Chapter 7* I present findings from a computational model that simulates the acquisition and maintenance of episodic and semantic information by the hippocampus and cortex. The simulation findings provide clear 'proof of concept' for the ideas put forward in earlier chapters. The model consists of a quickly learning 'hippocampal' component that stores orthogonalised traces; and an input 'cortical' component

which maps information topographically, and has very slow-learning long-range connections and relatively faster local connections. Initially, the hippocampal component supports the best recall of the specific details of all memories. As memories age, the cortical components become relatively more important for their recall. Recall via the cortical components is dominated by semantic generic information and does not support good recall of specific random information. The decay of detailed information from the hippocampus and the shift in the neural basis of recall with memory age therefore contributes to the semanticisation of memories.

Finally, in *Chapter 8* I conclude by briefly summarising my proposals, and outlining the main contributions of the thesis. I also make some predictions that could be tested empirically.

In the next section, I introduce the key features of the hippocampus.

1.5 Introduction to the hippocampus

Comprehensive and detailed reviews of the anatomy of the hippocampal system can be found in papers such as Amaral and Witter (1989) or Arbib *et al.* (1998). General introductions can be found in neuroscience texts such as Shepherd (1994) and Kandel *et al.* (1995). In this section I provide a brief introduction to the key elements of hippocampal anatomy and physiology.

The hippocampus is situated in the medial temporal lobe (MTL). It consists of two elongated C-shaped structures (located symmetrically, one in each hemisphere) oriented perpendicularly (in rodents) to the corpus callosum. Physically the hippocampus is continuous with the fornix (a major efferent pathway) and there are numerous transverse fibres between the posterior columns of the fornix which allows information arising from the hippocampus in each hemisphere to be integrated. The functional "hippocampal system" is generally considered to be composed of the hippocampus-proper or *cornu ammonis* (CA), the *dentate gyrus* (DG), the *subicular complex* (SUB) and the *entorhinal cortex* (EC).

A slice orthogonal to the longitudinal axis of the hippocampus reveals the basic neural circuitry. The cell layers appear as two interlocking curves of cells, the CA region and the dentate fascia (see Figure 1.1). The hippocampus-proper (CA) consists of large *pyramidal cells* that are the major output cells, whilst the dentate fascia (DG) has smaller pyramidal output neurons called *granule cells*. This basic hippocampal struc-

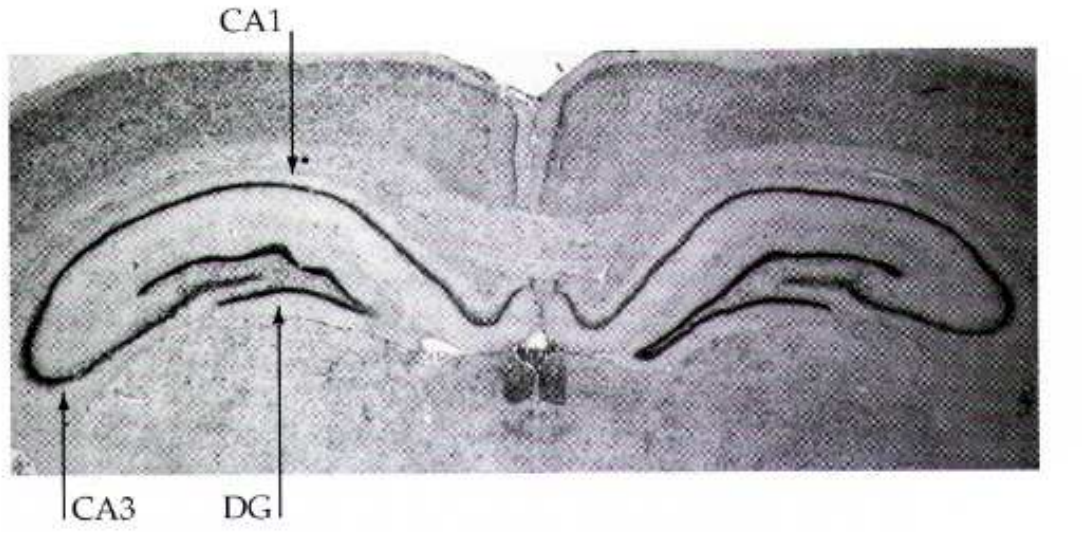

Figure 1.1: Cresyl violet stain of a coronal slice of rat brain, showing the two interlocking **C**s of the dentate gyrus (DG) and the hippocampus-proper (CA3 and CA1) on each side of the brain. Reproduced from Redish (1999), original picture courtesy of C. Barnes.

ture is well established, and dates back to early Golgi studies such as those of Ramón y Cajal (1893) and Lorente de Nó (1934). More recent anatomical methods have however revealed much additional detail. Lorente de Nó (1934) distinguished four CA subregions (CA1 – CA4) and this is the most commonly used labelling scheme. However CA4 is now not usually considered to be a separate region, and the boundaries of the small field CA2 are unclear (Arbib *et al.* (1998)), so the hippocampus is usually referred to as containing DG, CA1 and CA3.

Internally, the major excitatory hippocampal circuits follow a distinct pattern (see Figure 1.2): inputs from the entorhinal cortex terminate mainly in the dentate fascia. The dentate granule cells project to CA3 via the "mossy fibres". CA3 cells project to the septum through the fornix and send a collateral, the Schaffer collateral, to CA1. These receiving cells project to the subiculum, and also send outputs to the septum via the fornix. A more direct perforant path route from the entorhinal cortex to area CA3 bypasses the dentate gyrus. In addition, there are direct projections between most subfields in this "trisynaptic" circuit, especially recurrent CA3 connections. In summary, there is great inter-connectivity within the hippocampus, although there is probably some functional segregation: the details of these connections are still being

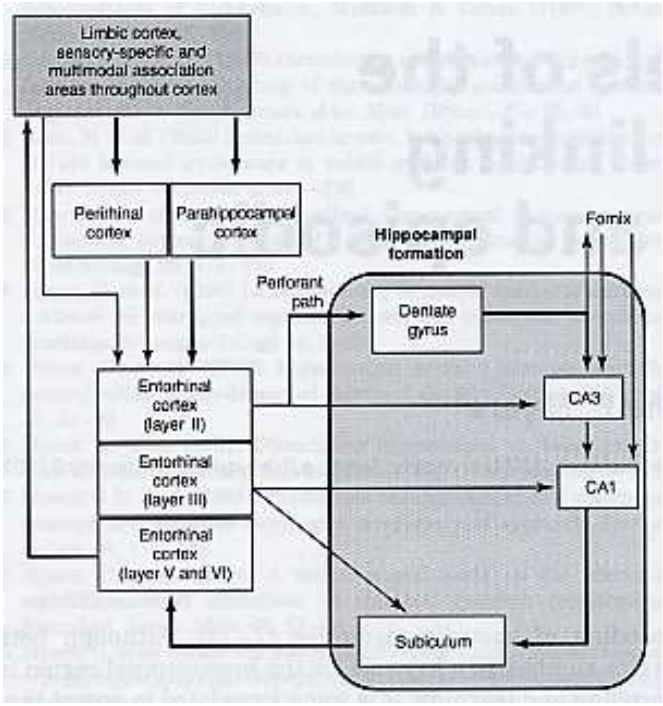

Figure 1.2: The neocortical-hippocampal processing hierarchy; and major connections between the entorhinal cortex and the hippocampus, and within the hippocampus.

worked out.

From the point of view of this thesis, what is perhaps more important is that the hippocampus can be conceived of as the 'top-node' in a neocortical-hippocampal processing hierarchy (see Figure 1.2). It receives massive inputs from tertiary and other association areas and from the supra-modal integrative areas, and heavy but relatively neglected inputs from sub-cortical areas. It also projects extremely widely on the output side.

Comparative anatomists have traditionally identified a homologue of the hippocampus in all vertebrates. However, without comparative behavioural experiments across species, it is difficult to separate functional and structural homologies, and to rule out the possibility that apparently similar neural structures compute different functions, or that physically different structures perform the same function(s) (O'Keefe and Nadel (1978)). However, in this thesis I consider data from many species in order to arrive at my conclusions as to what the hippocampus does (see section 3.1.1.2 for a discussion of this approach).

It is becoming clear that there are differences in the roles of left and right hippocampus, which are especially prominent in humans. For example, the right hippocampus has been implicated in spatial learning (e.g., Maguire *et al.* (1997, 1998a)), and the left in autobiographical memory and public events (e.g., Burgess *et al.* (2002)). There may also be a differential involvement of the left and right hippocampus in the long-term maintenance of information (Maguire and Frith, 2003). There is also evidence that different regions in the hippocampus are specialised for acquiring different information (Moser and Moser, 1998), or for acquiring information at different rates (de Hoz *et al.*, 2003). Whilst clearly important, such inter-regional differences are not considered in this thesis. Similarly, whilst there is no wish to imply that the hippocampus operates in isolation, the focus is on what the hippocampus itself does, rather than on the wider learning and memory systems of which it is part.

Chapter 2

Theories of hippocampal function

This chapter reviews existing theories of the hippocampus, drawing out the main themes. Existing models of hippocampal function are also introduced.

2.1 Introduction

Learning and memory have been topics of interest since ancient Greek times: the study of memory is integral to understanding the mind. Modern ideas about memory appear from about 1800 onward. For example de Biran (1804) distinguished several types of memory systems; Gall (1835) proposed that each mental faculty has a separate memory; James (1890) separated habits and memories, and "primary" and "secondary" memory (the latter was revised in the 1950's as short-term and long-term memory); and Ribot (1881) discussed the inverse relationship between the age of a memory and the likelihood of it being lost after brain damage. These ideas remain central to modern memory study. Hebb (1949)'s highly influential idea that memory properties must be based within individual cells forms the basis of most current memory models. Lashley (1950)'s conclusion that the "engram" for a particular memory is represented throughout an entire region paved the way for modern distributed memory theories.

Three years after Lashley gave up on his attempt to localise the engram, William Scoville operated on the patient "HM". HM had his temporal lobes removed to cure intractable epilepsy; the seizures were reduced, but he was left profoundly and unexpectedly amnesic. Scoville and Milner (1957) subsequently found evidence of amnesia

in another eight subjects who had received similar operations, and concluded that it was the removal specifically of the hippocampus and amygdala that had led to the severe amnesia. The hippocampus was firmly placed under the spotlight of memory research, and HM became the most famous and most studied patient in the field of cognitive science.

Initially it was proposed that the hippocampus was an all-encompassing learning structure, but it soon became clear that HM had intact motor learning abilities, although he could not explicitly remember the learning episodes (e.g., Corkin (1968)). Furthermore, animal studies showed that lesions restricted to the hippocampus did not produce the massive memory deficits seen in HM (e.g., Orbach *et al.* (1960); Kimble and Pribram (1963)). The current view is that the hippocampus is part of a larger learning system (including the entorhinal, perirhinal and parahippocampal cortices) and that some of HM's deficits are due to damage to the amygdala and overlying cortices.

It is now generally accepted that several anatomically and operationally different special purpose memory systems have evolved to acquire and store different types of information and solve different tasks. For example, response learning requires the caudate putamen (Kesner *et al.* (1993)), and conditioned emotional responses require the amygdaloid circuits (Phillips and Le Doux (1992)). In contrast, the role of the hippocampus remains vigorously debated. A consideration of neuropsychological data (such as that from HM-type amnesics) has led to the view that the hippocampus is involved primarily in remembering daily events (e.g. Squire (1992)), whereas data from other animals (primarily rodents) has led to the hypothesis that the hippocampus is involved specifically in spatial tasks (e.g. O'Keefe and Nadel (1978)). The hippocampus has also been implicated in innumerable other processes, such as novelty detection (Knight (1996)), emotional behaviour (Papez (1937)), working memory (Olton (1979)), consciousness (Clark and Squire (1998)), configural learning (Sutherland and Rudy (1989)), contextual learning (Phillips and Le Doux (1992)) and recognition memory (Reed and Squire (1998)), amongst others.

In this chapter I provide a review of theories and models of the hippocampus. In section 2.2 I review qualitative theories of the hippocampus – that is, those that attempt to define the essence of hippocampal function. In section 2.4 I review computational theories and models.

2.2 Qualitative theories of hippocampal function

The main aim of this section is to provide a conceptual overview of extant theories of the hippocampus. Full historical reviews are available elsewhere (e.g., Rosenzweig (1998)). Here I attempt to draw out the main themes.

Many early theories have been superseded, although some provide the conceptual basis for current views. On the basis that hippocampal lesions can disrupt spontaneous alternation and consistently produce hyperactivity, it was argued that the hippocampus served as a basis for Pavlovian inhibition (Douglas (1967); Kimble (1968)). Since there are alternative plausible explanations for these phenomena, and such theories are too under-specified to explain current data, it has been largely abandoned. Similarly, Gray (1982)'s suggestion that animals have a behavioural inhibition system, which causes feelings of anxiety when activated, and Rawlins (1985)'s proposal that the role of the hippocampus is to act as a temporary memory buffer to span delays have been largely forgotten. On the other hand, Hirsh (1974)'s early view, that many of the learning deficits associated with hippocampal damage can best be described as context effects, and O'Keefe and Nadel (1978)'s proposal that the hippocampus is the site of "cognitive map" formation, inform current spatial theories of the hippocampus.

Many of the functions that have at one time been ascribed to the hippocampus do in fact depend on different local substructures. For example, early lesion studies suggested that the hippocampus was involved primarily in mediating the emotional state of an animal, and that it, together with other parts of the heavily interconnected "limbic system", was essential for emotional behaviour (Papez (1937), Isaacson (1974)). In fact, structures adjacent to the hippocampus (such as the amygdala and anterior cingulate cortex) and other fibres of passage were probably inadvertently lesioned. Recognition memory was also thought to be hippocampally-dependent (e.g. Gaffan (1974)), but current studies suggest that the parahippocampal areas are responsible for recognition functions, although this is still hotly debated.

2.2.1 "Declarative" versus "procedural" memory

An early view in psychology was that memory was a dual-store system, consisting of a separate small capacity short-term memory and a larger capacity long-term memory. Amnesia was proposed to be due to an impairment of transfer of information into the

long-term store from an otherwise intact and functioning short-term store (e.g., Atkinson and Shiffrin (1968); Baddeley and Warrington (1970)). Slightly later process views posited that amnesia arose as a result of too-shallow encoding (e.g., Cermak and Butters (1972); Cermak *et al.* (1973)), or increased susceptibility to interference at retrieval (e.g., Warrington and Weiskrantz (1968, 1970)). Thus differential impairment of different tasks after brain damage cannot be accommodated by these hypotheses. However, hippocampal amnesics show normal learning on tasks such as mirror writing or learning a tactile maze (so-called *procedural* tasks), despite a complete lack of recollection of the learning experience (Corkin (1968); Milner (1962)). Later studies also showed that with appropriate implicit tests of memory ("...complete these word stems...") amnesics showed normal performance, whereas only with explicit tests ("...write down those words we just showed you...") were they impaired (Graf *et al.* (1984)). A new conception of memory was required.

It is now generally accepted that there are multiple memory systems (Squire (1992)). Most researchers today would accept a division of some kind between (1) hippocampally-independent memory that is typically only revealed by a change in the facility of task performance (such as some skills and dispositions, priming, habituation and sensitisation), and (2) memory for facts and/or events and/or experiences that depends (at least initially) on the integrity of the hippocampus and related structures.

There is converging agreement on the identity and anatomical basis of various aspects of the former systems. For example, classical conditioning requires the cerebellum (Thompson (1988)), fear conditioning requires the amygdala (Phillips and Le Doux (1992)), priming requires the systems engaged at perception, stimulus-response learning requires the striatum (Packard *et al.* (1989)), and response learning requires the caudate putamen (Kesner *et al.* (1993)). These types of learning are collectively known by various terms, including *non-declarative* (Squire (1983); Squire and Zola-Morgan (1988)), *procedural* (Cohen and Squire (1981), *habit* (Mishkin and Petri (1984)) or *implicit* (Schacter (1987)) memory. The detailed description and anatomical substrate of the purportedly hippocampal latter system(s) (for facts and/or events and/or experiences) is rather more controversial. Different authors have made different distinctions, and the system is known as *declarative* (Cohen and Squire (1981); Squire (1992), *explicit* (Schacter (1987)), *relational* (Eichenbaum *et al.* (1992)) and *config-*

ural (Sutherland and Rudy (1989)) memory (see figure 2.1).

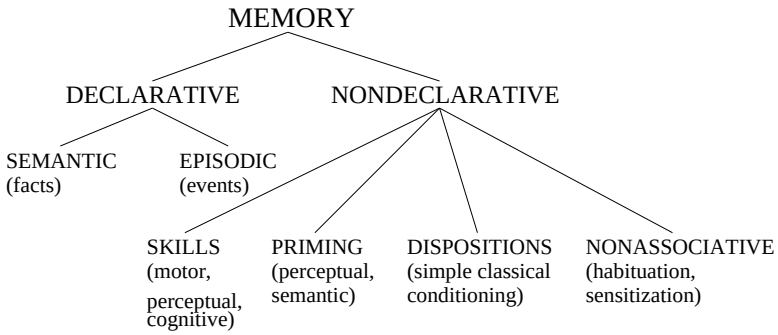

Figure 2.1: A memory classification (Redrawn from Squire (1992): It is generally accepted that *non-declarative* (procedural/implicit) memory is independent of the hippocampus. Which aspects of *declarative* (explicit) memory are hippocampally-dependent is somewhat controversial.

Cohen and Squire (1981) first made the distinction between *declarative* and *procedural* memory. The term "declarative" derives from work with human subjects and is often linked to the notion of conscious memory. Declarative memory was characterised as the record of everyday facts and events that can be brought to conscious recollection and typically is subject to verbal reflection. Procedural memory was characterised as the non-conscious acquisition of a bias or adaptation that typically is only revealed by implicit measures of performance.

This initial conception of declarative memory was based on work with humans and is difficult to apply to animals. We do not have means for monitoring conscious recollection in animals (if indeed it exists, see Eichenbaum *et al.* (1992) for a discussion). The definitions can be operationalised by making further distinctions between declarative and procedural memory that do not rely on verbal expression or subjective awareness. For example, several authors have proposed that a defining feature of the declarative code is that it is common across different processes or processing systems and allows information from different sources to be compared and contrasted; inferences can then be made in novel situations on the basis of what had happened before in another situation (e.g., Eichenbaum *et al.* (1992); Shapiro and Olton (1994)). As evidence accumulated that amnesics could perform normally on a variety of tasks given appropriate testing, "procedural" memory was renamed *non-declarative* memory (e.g., Squire (1983); Squire and Zola-Morgan (1988)).

Numerous other similar distinctions have been made. Tulving (1972) made a distinction between a hippocampally-dependent *episodic* memory and hippocampally-

independent *semantic memory*. Episodic memory was defined as a memory system for learning material presented in a particular place and time. Semantic memory was defined in terms of language-related declarative memory operations. A similar distinction was made between *explicit* memory (responsible for intentional or conscious recollection, and impaired in amnesia) and *implicit* memory (responsible for non-intentional recollection tasks) (Schacter (1987)). Olton (1983) argued for a hippocampally-dependent *working* memory (for the specific, personal or temporal context of a situation) and a hippocampally-independent *reference* memory (for rules and procedures that re-occur across specific situations, so-called "general knowledge"). Squire (1992) argued that the essential feature of the hippocampally dependent declarative memory system is the ability to rapidly establish *novel associations* in memory. In contrast, non-declarative memory is specialised for "incremental, cumulative change", so that new associations can be acquired but only after many repetitions. Placing a different emphasis on representational format, Eichenbaum *et al.* (1992) distinguished a hippocampally-dependent *declarative* memory, supporting the conscious processing of information and distinguished by the relational representation and representational flexibility and a non-hippocampal *procedural* memory based on the representations of single stimuli or configurations of stimuli. A distinction has also been made between hippocampally-based *remembering* (recollecting an experience) and *knowing* which permits a recognition judgement to be made on the basis of familiarity (Mandler (1980)). According to Tulving and Markowitsch (1997) episodic remembering always implies semantic knowing, whereas knowing does not imply remembering. More recently, Tulving and Markowitsch (1998) have suggested that declarative memory should be defined in terms of the overlap between semantic and episodic memory.

2.2.1.1 Current "declarative" views: semantic and episodic memory

Two main streams of thought have emerged with respect to the role of the hippocampus in the acquisition of 'declarative'[1] information. Some authors argue that both semantic (fact) and episodic (event) memory are dependent on the hippocampus for

[1]The term 'declarative' is still in widespread use, even though authors continue to use the term differently, and despite the potentially confusing overtones of conscious processing. For clarity I will use the terms *semantic* and *episodic* memory to refer respectively to memory for facts and commonalities across events, and to detailed memory for events that occurred once. In non-human animals, such memories are usually referred to as context-independent, and context-specific respectively.

acquisition (eg., Murre (1996)), whereas others argue that the hippocampus is obligatory only for episodic learning (e.g., Nadel and Moscovitch (1997)). For an instructive insight into these differences see the conflicting discussions of Vargha-Khadem *et al.* (1997)'s key finding of apparently intact semantic memory in the face of impaired episodic memory in young hippocampal amnesic patients (Eichenbaum (1997); Tulving and Markowitsch (1998); Squire and Zola (1998); Mishkin *et al.* (1998)).

Those involved in the study argue that selective hippocampal damage causes an impairment only in context-rich episodic memory but not context-free semantic memory (Vargha-Khadem *et al.* (1997), Mishkin *et al.* (1998) and Gadian *et al.* (2000)). This is in accord with the distinction made by Tulving (1972)) between semantic and hippocampally-dependent episodic memory, and is the view espoused by Tulving and Markowitsch (1998). Tulving's "episodic theory" views episodic memory as an extension of semantic memory. The *Serial encoding, Parallel storage and Independent retrieval* (SPI) model (Tulving (1995)) posits that encoding information into the episodic system critically depends on the semantic system, whereas semantic encoding does not require an intact episodic system. Retrieval from the episodic or semantic store can be independent (Tulving and Markowitsch (1998)).

The alternative view is that both episodic and semantic memory ("declarative" memory) are equally dependent on the hippocampus (e.g., Squire and Zola (1996, 1998); Cohen *et al.* (1999)). In this "unitary declarative memory" view, damage to both semantic and episodic memory is proportional to the degree of damage to the 'hippocampal system'. The apparently preserved semantic abilities of Vargha-Khadem's patients are claimed to be incomplete and to depend on similarly partially preserved episodic capacities (Squire and Zola (1998)). Thus in this view, episodic memory is the gateway to semantic memory (Squire and Zola (1998)).

2.2.2 Spatial theories

The declarative/non-declarative memory distinction, and the view that the hippocampus is a general cross-domain learning structure, derives largely from work with human amnesics. Experimental work with non-human animals has led to various different proposals; the most influential of which is that the hippocampus has a specifically spatial learning role. In the spatial approach, impairments in ostensibly non-spatial tasks after hippocampal lesions are explained in terms of supporting functions that

depend on spatial processing (e.g., episodic memory may depend on recall of context, Gaffan (1994b); language processing may have co-opted a spatial mechanism, O'Keefe (1996)). In more general memory hypotheses, spatial learning is seen as merely one of the functions of the hippocampus, albeit an important one.

In my opinion there is little evidence for an exclusively spatial role for the hippocampus; however, spatial theories have been very influential and make up the majority of implemented models.

2.2.2.1 Cognitive Mapping

Animals can show short-cut and detour behaviour through previously unexplored areas of an environment (Tolman (1948), although see Bennett (1996) for some scepticism about such abilities). Many authors have argued that the most parsimonious explanation for these abilities is that animals can create and store *cognitive maps*, that is, stored neural representations of environments that permits an animal to solve navigational problems using information about the structure or geometry of the environment (Muller and Stead (1996), p709). On the basis of evidence that hippocampal cells encode spatial information, O'Keefe and Nadel (1978) argued that such cognitive maps were created and stored in the hippocampus. This *Cognitive Mapping Theory* is perhaps the most influential theory of hippocampal function and has spawned innumerable variants.

O'Keefe and Nadel (1978) and Nadel (1994) proposed two distinct systems for processing spatial information: The hippocampal *locale* system encodes places in the environment into allocentric (world-centred) cognitive maps, whilst the hippocampally-independent *taxon* system codes motor responses in terms of specific orientations within a spatial environment. The systems differ in their susceptibility to interference (the locale system is more sensitive and acts to separate memory traces, whilst the taxon system combines memory traces based on overlapping features) and consolidation characteristics (learning in the locale system is all-or-nothing, whereas traces in the taxon system are built up incrementally). There are many models based on these ideas (e.g, Zipser (1985); Hetherington and Shapiro (1993); Muller and Stead (1996); Burgess and O'Keefe (1996)) and variants of the theory.

Cognitive Mapping theories propose that any ostensibly non-spatial phenomena that are sensitive to hippocampal damage emerge as secondary properties of spatial

mapping. For example, spatial context can be seen as integral to episodic memories, or as a vital cue for recall. The theory has been extended to account for the language deficits observed after hippocampal lesion (O'Keefe (1996)) and for sequence recall, by adding a temporal aspect. It is somewhat unclear when such theories cease to be "specifically spatial".

2.2.2.2 Scene and context memory

Scene or context memory refers to memory for the spatial arrangements of objects; it has also been described as "object in place" memory (Gaffan and Parker (1996)) and "memory for the location of objects" (Parkinson *et al.* (1988)). Several authors have proposed that one of the functions of the hippocampus is to store snapshot-like views of the environmental context (e.g., Gaffan (1994b,a); Gaffan and Parker (1996)). Many of the general associative memory theories of the hippocampus have similarly posited that an ability to store scene snapshots would naturally arise out of the hippocampus's role in binding together convergent inputs from various neural processors (e.g., Squire (1992); Cohen and Eichenbaum (1993); Eichenbaum *et al.* (1994).)

Scene information could be used in several ways; the information about the learning context in which a task is acquired could generally aid recall (e.g., Hirsh (1974)) or more specifically allow disambiguation of similar tasks learnt in different contexts (e.g., Gaffan (1994b)). In some ways, these theories are like the configural association theory in that they propose that the role of the hippocampus is to augment a basic stimulus-response learning ability.

2.2.2.3 Path integration

Path integration (PI) is the process of integrating information acquired in the process of self-movement to compute a current position with respect to a starting position. It provides the capacity to return directly to a starting point after following a circuitous out-bound path. Few experimental studies have linked PI to a specific brain structure, perhaps because PI most likely involves several processes including identifying an initial reference point, monitoring various idiothetic cue sources and computing current position (Taube (1999)). Several researchers have suggested that the hippocampus is the substrate for path integration (e.g., McNaughton *et al.* (1996); Samsonovich and McNaughton (1997); Whishaw *et al.* (1997)). Whilst it seems likely that path integration

information is represented in the hippocampus, the lesion evidence does not appear to support the strong claim that the hippocampus is the path-integrator proper.

2.2.3 Other hypotheses about the role of the hippocampus

Declarative and Spatial theories of hippocampal function dominate today's literature. However, there are many important issues and hypotheses that cut across these proposals:

2.2.3.1 The hippocampus as a convergence zone

The hippocampus receives massive convergent projections from both cortical and subcortical areas. Many authors have argued that this makes the hippocampus ideal for forming traces of events occurring throughout the brain.

Teyler and Discenna (1986) posited that the role of the hippocampus is to form and retain an *index* of neocortical areas activated by experiential events. Reactivation of the stored hippocampal trace of an event via a proposed one-to-one mapping between hippocampal loci and neocortical modules serves to re-instantiate the associated activity in the neocortex. The Convergence Zone theory (Damasio (1989a,b)) can be seen as an extension of the Indexing theory. In Damasio's theory, *fragments* of information stored in the cortex make up the building blocks of memory and *convergence zones* distributed throughout the brain (of which the hippocampus is one) act to trigger activity in lower convergence zones and fragments. Other authors have suggested that the hippocampus has a general binding function (e.g., Alvarez and Squire (1994); Murre (1996)). It is unclear how the binding theory differs from the index theory (Milner (1989)); the hippocampal component of a memory trace is not necessarily qualitatively different from that part stored in the cortex.

On the other hand some authors argue that memories stored in the hippocampus are not complete copies of cortical patterns of activation but "reduced descriptions that exploit redundancies in the cortical patterns" (McClelland and Goddard (1996), p655). In other words, the hippocampus is the (temporary) store for all the information acquired in a learning experience, and there is initially no cortical component. This view is explicitly espoused by Rolls (e.g., Rolls (1996); Rolls and Treves (1998)), and implicitly by others (e.g., Rawlins (1985); Marr (1971); McClelland *et al.* (1995)).

2.2.3.2 The hippocampus as a fast learner

A pervasive idea is that the hippocampus is a fast (but temporary) storage system for the patterns of activity that are induced in it from cortical areas (e.g, Murre (1996); Alvarez and Squire (1994); McClelland *et al.* (1995); Morris and Frey (1997); Treves and Rolls (1994)). This idea of a fast, small temporary hippocampal store and large incremental cortical store can be traced back to Marr (1970, 1971) and the idea forms the basis for many models of hippocampal function. In a generalised Hebb-Marr model the neocortex stores large complex event memories. To prevent interference between traces, Marr suggested that a separate processor was required — the hippocampus — that could rapidly store events, and then allow gradual transfer to the neocortex, which could reorganise and classify this information before incorporating it into the existing knowledge base. In a similar vein, but on the basis of considerations of learning in artificial connectionist systems, McClelland *et al.* (1995) argued that the role of the hippocampus is to store and play back new memories to the cortex to allow interleaved cortical learning.

On the other hand, authors such as Murre (1996) have argued that the hippocampus is required for fast on-line learning of episodic information not to prevent interference or because the cortex can only do slow learning, but because there is limited a priori long-range connectivity in the cortex. Recently Wise and Murray (1999, 2000) have proposed that the hippocampus is required (amongst other things) for the rapid acquisition of arbitrary antecedent-to-action mappings (e.g. see a yellow card — wave your right paw). The hippocampus is posited to be required for its speed of acquisition of these new associations.

A somewhat different set of ideas derives from the proposal of Morris and Frey (1997) that the hippocampus is involved in automatic fast encoding of attended experience, irrespective of the importance of an event at the time of its occurrence. They argue that this entails a mechanism such as "synaptic tagging" (Frey and Morris (1998)) that allows the long-term strength of a trace to be determined sometime after acquisition at which time the affective value of the information may have become apparent, rather than at storage.

2.2.3.3 Novelty detection and comparator functions

Various authors have proposed that the hippocampus is particularly concerned with the detection of novelty, and that it acts to guide behavioural responses on the basis of the nature of impinging stimuli and their familiarity (e.g., Gray (1982); Knight (1996); Tulving *et al.* (1994); Parkin (1997)). Gray (1982) proposed that the hippocampal formation acts as a comparator of incoming information with old stored information, to set up conditions for changes in behavioural inhibition, arousal and attention. Common current opinion is that if there is a behavioural inhibition system, it is likely to be elsewhere, but there are moves to rehabilitate the theory (e.g., Lemaire *et al.* (1999)).

A more recent proposal is that the hippocampus plays a critical role in associative mismatch processes, beyond that of simple stimulus novelty processing (e.g., Squire (1992); Bunsey and Eichenbaum (1996); Honey *et al.* (1998)). Various studies (e.g., imaging (Henke *et al.* (1997)), cell recording (Wood *et al.* (1999)) and *c-fos* (Wan *et al.* (1999))) show that the hippocampus is activated by novel or complex *arrangements* of stimuli, but not novel stimuli per se. Honey *et al.* (1998) claim that simple stimulus habituation can proceed without the hippocampus. It seems likely that the mechanisms that generate increased signals for novel individual stimuli are extra-hippocampal. Of course, independently from any specific "cognitive" role in the processing of novelty, differential activity in response to novel versus familiar stimuli might be expected in any putative memory system, since it is desirable to encode novel information.

2.2.3.4 Incidental learning

The stimuli that are available when a task is learnt may either be integral to task performance (e.g., a tone on a tone-conditioning task), or incidental and unimportant to the performance of the current task (e.g., room layout or door colour on a tone-conditioning task). Hippocampal animals frequently show deficits on tests of incidental learning; that is, they acquire less information than controls about stimuli which need not be processed for the task for which the animal is rewarded. This has led several authors to propose that the hippocampus has a critical role in the encoding of incidental information (e.g., Phillips and LeDoux (1994); Good *et al.* (1998)). One common view is that this incidental learning may provide retrieval cues. Another view is that it is important for an animal in the real-world to encode information even if it is apparently incidental and unimportant when it occurs, since it might turn out to be

highly salient (Morris and Frey (1997)). If the event was not initially stored, associative learning across time (e.g., that exposure to an apparently meaningless tone predicts a shock in 20-min) would not be possible.

Rudy and Sutherland (1995)'s suggestion that the hippocampus can act to enhance the activation or salience of certain representations encoded outside of the hippocampus may be of relevance to putative hippocampal incidental learning, in that it would provide a mechanism for modulating the encoding of low salience information. In Rudy and Sutherland's theory, the differential reinforcement provided by the hippocampus primarily works to aid extra-hippocampal configural learning mechanisms.

2.2.3.5 Relational representation

In the *Relational Representation* theory (Eichenbaum *et al.* (1992, 1994); Eichenbaum (1997)) the hippocampus is required for the creation and use of flexible *relational* representations that permit the inferential use of knowledge in novel situations, whilst the "para-hippocampal region" encodes isolated *individual* representations within those modules engaged at learning which are thus inflexible in that they will only be reactivated by a restricted range of events similar to those occurring at storage. Both relational and individual representations can support declarative memory, thus the distinction here is not along the episodic/semantic line.

2.2.3.6 Configural representation

An event in which two stimuli **A** and **B** are together paired with a stimulus **C**, could be stored as the two *elemental* associations **A-C** & **B-C** or as the single *configural* association [**AB**]-**C**. Various tasks (e.g., negative patterning, trans-switching) can only be solved using configural encodings, and many of these tasks were found to be impaired after hippocampal lesions. The *Configural Association* theory (Sutherland and Rudy (1989)) distinguished between a (1) hippocampally-independent *simple associative system* that stores experiences as changes in strengths of associations between single stimuli as presented to the organism and (2) a hippocampal *configural associative system* which can combine these representations of elementary stimuli to form new configural representations, which can then be associated with other configural or elemental representations. Wickelgren (1979)'s "chunking" proposals can be considered

a predecessor to this Configural Processing theory. Squire (1992) similarly proposed that the hippocampus is necessary for memory tasks that involve the development of configural as opposed to simple associations.

2.2.3.7 Sequence and temporal learning

Several authors have proposed that the essential role of the hippocampus is to store and replay sequences of events (e.g., Minai and Levy (1994); Levy (1996); Skaggs and McNaughton (1996); Wu *et al.* (1998)). The encoding of sequences of events could be seen as integral to episodic memory, if episodic memory is viewed as a sequence of static views of the world. Levy and colleagues show that the recurrent connections of CA3 are suited for the storage of sequences. They and others (e.g., Wallenstein and Hasselmo (1997)) argue that conceptualising the hippocampus as a sequence learner unifies various theoretic and paradigmatic perspectives. Interestingly, the most recent proposal from the Eichenbaum lab (that the hippocampus is an associator of discontiguous spatial or temporal events) is based on a sequence learning model (Wallenstein *et al.* (1998)).

Hippocampal lesions disrupt not only memory for temporal order, but also the frequency of occurrence of items in a list, duration estimates and relative recency judgements. This has led to proposals that the hippocampus is required for the processing of temporal information per se (e.g., Kesner and DiMattia (1987); Kesner (1998); Wallenstein *et al.* (1998)).

2.2.3.8 Awareness, attention and consciousness

Memory performance on hippocampally-dependent tasks in humans, such as recalling what you had for breakfast, or a list of words, is usually accompanied by conscious awareness (hence the term "declarative" memory). It has recently been proposed that it is this factor that governs whether an event requires the hippocampus (e.g., Clark and Squire (1998); Manns *et al.* (2000)). This conclusion was based on studies that showed that only those patients who became aware of certain stimulus contingencies showed good performance on trace-conditioning tasks. The argument is thus circular, and more recent studies have dissociated awareness from performance (e.g., Chun and Phelps (1999)). It is more likely that awareness is not a prerequisite for learning per se, but both tend to occur as a consequence of hippocampal processing. Moscovitch

proposed that conscious memory is only special because the conscious experience of the remembered event is stored with other event information. Similarly, Eichenbaum (1999) suggested that the hippocampus could be the gateway by which awareness enters memory. It is clear that the hippocampus could not be the sole "source" of consciousness, since hippocampal amnesics remain self-aware.

Some authors have attributed attentional control mechanisms to the hippocampus (e.g., Buhusi and Schmajuk (1996)). The direction of attentional control processes could be related to putative novelty detecting functions of the hippocampus.

2.3 The long-term role of the hippocampus

This section reviews theories that address the long-term role of the hippocampus in memory.

2.3.1 The standard view: temporary storage

Hippocampal damage commonly leads to both anterograde and retrograde amnesia, that is, to an inability to lay down new memories of a particular kind, and a loss of information acquired before damage. Retrograde amnesia is commonly graded such that more remote memories are disproportionately less affected than recently acquired memories. This observation underpins the "standard model" (so called by Nadel and Moscovitch (1997)). In this view, medial temporal lobe structures store traces initially, after which there is gradual reorganisation within long-term memory such that the hippocampus becomes less important for the recall of a given trace and a more permanent memory is set up elsewhere. Thus medial temporal lobe structures are merely temporary stores of memory traces (e.g., Scoville and Milner (1957); Teyler and Discenna (1986); Squire (1992); Murre (1996); Squire and Alvarez (1995); Kim and Fanselow (1992); Shen and McNaughton (1996)). In most standard views, the hippocampus is involved in some way in consolidating memories into their non-hippocampal long-term form.

2.3.1.1 Mechanisms of consolidation

Numerous writers have proposed that sleep (specifically REM sleep) is important for memory processing. Vertes and Eastman (2000) trace the proposal back to Jenkins and

Dallenbach (1924). Marr proposed that neocortical long-term memories are trained by the hippocampus during sleep whilst there is low interference from daily events (Marr, 1970, 1971). Interest in the idea of sleep as a consolidation mechanism has waxed and waned, but recent neurophysiological data have provided some new support for this idea, although it is somewhat circumstantial.

Many recent authors (e.g., McClelland *et al.* (1995); Chrobak and Buzsáki (1994)) have similarly suggested that sleep acts to replay events to the cortex so that traces can be slowly potentiated there. A few have proposed that consolidation may occur during engagement in "type II" behaviour in rats (such as awake grooming or eating, e.g., Vanderwolf *et al.* (1975)) or during conscious and unconscious rehearsal (e.g., Chrobak and Buzsáki (1994); Murre (1996)).

2.3.2 Long-term involvement of the hippocampus

Until recently the standard view has had few detractors. However, Nadel and Moscovitch (1997) have argued that the evidence cited in support of the standard view is equivocal. They propose instead that the "hippocampal system"[2] is the long-term store for traces that are obligatorily hippocampally-dependent at acquisition. In their *Multiple Trace Theory*, each time an event is recalled (whether internally or externally triggered) another slightly different trace of that event is stored. The *Headed Records* account of Morton *et al.* (1985) is a similar early view.

The influential Cognitive Map theory also proposed that the hippocampus is a long-term memory store, albeit only for maps of environments (O'Keefe and Nadel (1978)). Milner (1989)'s ideas are also of relevance, since they provide an intermediate view between the standard view and that of Nadel and Moscovitch. Milner proposed that easily modified but transient "soft" limbic system traces and weak "hard" neocortical traces are initially laid down in response to an event — after neocortical trace have been reactivated a sufficient number of times they may become strong enough to be of functional importance. Thus the only traces that will exist of non-reoccurring events (such as episodic traces) will be in the hippocampus.

[2]Moscovitch and Nadel (1998) define the "hippocampal system" as the entorhinal and perirhinal cortices, dentate gyrus, subiculum, parahippocampal cortices and the amygdala, as well as the hippocampal formation.

2.4 Computational theories of the hippocampus

The theories of hippocampal function reviewed so far are qualitative, consisting of a central concept or metaphor that attempts to characterise the essence of hippocampal-region function. In this section I briefly introduce extant high-level cognitive models of hippocampal function. I focus on declarative memory models, especially those that aim to model the development of graded retrograde amnesia, as they are closest in spirit to my model (Chapter 7). For a full exposition of implementation details and mathematical underpinnings, see the original journal articles.

Most theories of learning and memory (whether qualitative or computational) are predicated on *Hebbian type* learning mechanisms, in which the co-activation of connected neurons results in a progressive strengthening of the connections between them (Hebb (1949)). Furthermore, despite the diversity of current models of the hippocampus, with few exceptions (e.g., Gluck (1996)) they are influenced by Marr (1971)'s conceptualisation of the hippocampus as an auto-associative memory that performs pattern storage and retrieval. Marr established several precedents that remain influential today: the model was based on (the then-known) details of neuroanatomy, and incorporated a learning mechanism based on the idea of Hebbian synaptic plasticity; it emphasised the importance of interference (and how it can be reduced); and it recognised the importance of understanding the interaction between the hippocampus and the cortex in consolidation.

2.4.1 Models of declarative function

If episodic storage is conceived of as the storage of patterns of co-occurring activity in different areas of the cortex, then with an appropriate conceptualisation of what the inputs represent, a minimal auto-associative network could be said to implement episodic storage and recall. This has been considered to be a weakness when compared to implementations of putative spatial functions, since for "general memory models" it is often unclear what the inputs represent, and how recall performance can be compared with real-world "remembering".

In most models, the hippocampus acts to bind or index patterns of activity occurring in cortical areas (e.g., Damasio (1989b); Teyler and Discenna (1986); Moll *et al.* (1994); McClelland *et al.* (1995); Squire and Alvarez (1995); Murre (1996)) since the

hippocampus is so small with respect to the cortex. Others argue that this small size makes the hippocampus ideal for the storage of compressed representations of cortical activity (e.g., McClelland and Goddard (1996); Rolls (1996)). The relatively simple sequence learning models of Levy can also do one-trial (episodic) learning (e.g., Levy (1996)). The models of Gluck and Myers (1997) and McClelland and Goddard (1996) are based on the properties of hidden layers in multi-layer nets. Other models aim for more detailed neurobiological plausibility (e.g., Buzsáki (1989); Crick and Mitchison (1983); Rolls and Treves (1998); Hasselmo *et al.* (1996)).

Four connectionist models have simulated the long-term role of the hippocampus in protecting information initially stored there from the effects of subsequent hippocampal damage (Alvarez and Squire (1994); the Complementary Learning Systems Framework, McClelland *et al.* (1995); Tracelink, Murre (1996); and Multiple Trace Theory, Nadel and Moscovitch (1997)). The first three models share the assumption that the hippocampus learns quickly on exposure to an event so is initially important for recall; whereas the cortex learns slowly through hippocampally-supported interleaved learning, gradually making the hippocampus redundant over time for the recall of that information.

Alvarez and Squire's (1994) simple network model consists of two 'cortical' regions (each of 4 units) reciprocally connected with a 'MTL' region (of four units, see figure 2.2a). Two non-overlapping patterns are stored in the net, using a learning rate an order of magnitude higher in the MTL-cortical connections than the cortico-cortical connections. Random activity in the MTL then drives rehearsal allowing the slow incrementation of cortico-cortico connections. Global forgetting occurs in all weights in proportion to their current strength.

In the Complementary Learning Systems (CLS) framework, McClelland, McNaughton, O'Reilly and others have explored several analytical issues and implemented several models that they consider to be one unified model (O'Reilly and Norman, 2002). Instead of acting as an 'index' for information stored in the cortex, the fast-learning hippocampus is considered to initially store all information about an event, that it then teaches to the cortex. The models consist of several layers (see Figure 2.2b), trained using the back-propagation algorithm, making them sensitive to catastrophic interference. Training and hippocampal replay to the cortex are interleaved to protect the cortex from interference - this slow learning allows the extraction of structure.

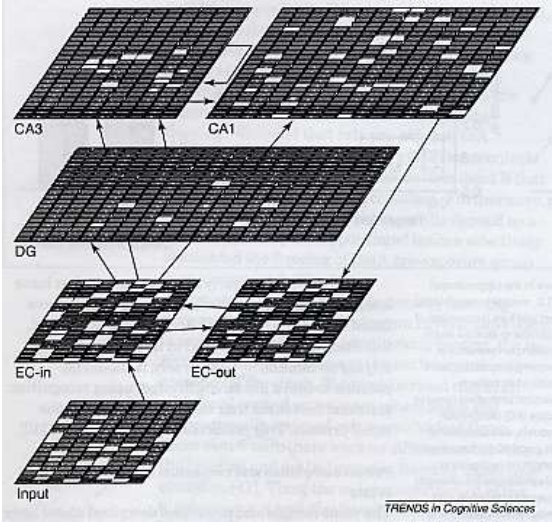

(a) McClelland, McNaughton & O'Reilly's Complementary Learning Systems model

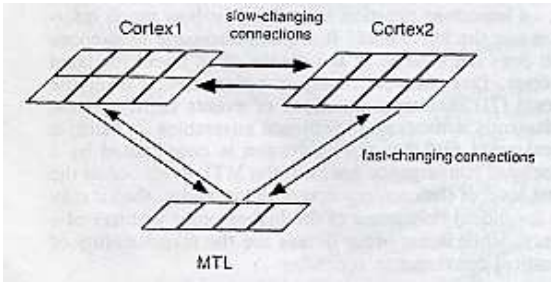

(b) Alvarez & Squire's (1994) model

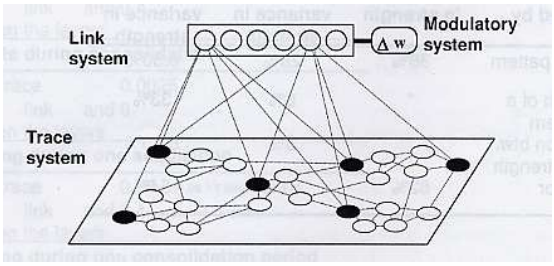

(c) Murre's (1996) Tracelink' model

Figure 2.2: Network models of consolidation share several characteristics. The 'hippocampal/MTL' component has relatively more convergent inputs, relatively more orthogonalisation of patterns, and a higher learning rate compared to the 'cortical' component.

Weight decay occurs in the 'hippocampal' component.

Tracelink (Murre, 1996) consists of three systems: a trace system (the neocortex), a smaller link system (the hippocampus), and a modulatory system (amygdala and other areas) that alters the rate at which memories are stored (see figure 2.2c). In a typical run, the link systems consists of 42 nodes, and the trace system 200 nodes. Link-link and trace-link connections are formed more rapidly than trace-trace connec-

tions, although all connections are active during learning. Random noise in the link system activates a random subset of nodes, and the net is allowed to settle into an attractor for the strengthening of trace-trace connections. There is no weight decay.

In contrast to the preceding models, the Multiple Trace model (Nadel *et al.*, 2000) allows multiple traces to build up over time in the hippocampus for a given memory, which protects memories from *partial* hippocampal damage as they age. The recall of episodic and spatial detail remains dependent on the hippocampus indefinitely. The neural network consists of a 'hippocampal' (HC) component and a 'cortical' (NC) component, each consisting of 1000 units; which are fully-connected. Training and replay are interleaved, with events chosen randomly for replay. Replay is initiated by activating a NC trace, which activates a HC pattern. Noise is then added to the HC pattern, and the pattern re-stored in NC-HC connections producing a new trace.

2.4.2 Models of sequence learning

The recurrent architectures of auto-associative networks are suited to sequence learning. Given a partial input of the present state, an auto-associative network can perform pattern completion and also then retrieve the predicted next state. Levy argues that a general sequence prediction paradigm can provide a computational unification of various putative hippocampally-dependent functions, such as one-trial episodic learning, short-cut behaviour and inference tasks (Levy (1996)). Skaggs and McNaughton (1996) also propose that the hippocampus can store and replay sequences to drive long-term consolidation processes in the cortex. Liaw and Berger (1996)'s model of "dynamic synapses" encodes spike-trains into spatio-temporal network patterns. Each synapse is sensitive only to action potentials occurring in a relatively small time window, thus "temporal chunking" occurs; this acts as a mechanism for sequential pattern recognition.

Whilst most sequence learning models focus on CA3, Granger *et al.* (1996) developed a model of CA1 which incorporates a (not strictly Hebbian) temporally-dependent LTP learning rule: the amount of potentiation depends upon the order of arrival of afferent activity to a target neuron. Together with putative asynchronous inputs from CA3, the CA1 model can store brief simulated temporal sequences of inputs.

2.4.3 Models of spatial function

Many connectionist models of hippocampal spatial processing have been based on an auto-associative model of the hippocampus. The basic idea is that the broad memory of a place could be evoked by aspects of the scene, such as different views or specific features (e.g., McNaughton and Morris (1987); Muller and Stead (1996); Recce and Harris (1996); Sharp (1991)). If it can be shown that apparently non-spatial functions of the hippocampus (e.g., general episodic memory or transitive inference) can be modelled by the same networks as spatial functions, this would suggest that spatial learning is not necessarily fundamentally different to other kinds of representative learning.

Spatial theories of hippocampal function are largely founded on the existence of so-called *place cells*. Most such models thus aim to capture at least some of the characteristics of place cell activity.

Several theories have proposed that place cells show place fields because they are sensitive to combinations of visual landmark cues in the environment (e.g., Zipser (1985); Sharp (1991); Hetherington and Shapiro (1993); Shapiro and Hetherington (1993)[3]). These place cell models are thus in line with O'Keefe and Nadel (1978)'s original hypothesis that the hippocampus was required for *locale navigation*, which depends on a combination of cues (rather than *taxon* or *praxic* navigation which depends on a single cue or motor strategy).

However, place cells can continue to show place fields in the dark (O'Keefe (1976)), so the hypothesis that place cells are solely driven by (visual) local view is inadequate. Many theories rely on purported associative memory properties to complete missing details on the basis of what is perceivable (e.g., McNaughton and Morris (1987); Rolls (1996); Recce and Harris (1996)). However, place cells can also show place fields in the dark even if an animal first enters an environment in the dark, although the place fields have a tendency to drift (O'Keefe (1976); Quirk *et al.* (1990); Knierim *et al.* (1995)). One likely possibility is that place cell activity in the dark can be driven by non-visual sensory inputs (e.g., O'Keefe and Nadel (1978)), although some authors argue that this would lead to large errors (e.g., Redish (1997)). Alternatively, internally generated information such as self-motion may be used to keep track of position (and predict the

[3]Unlike the other associative models, Shapiro and Hetherington (1993) identify place cells with the recurrent hidden layer of a three-layer (back-propagation of error) net, and train using supervised back-propagation of error methods.

next expected local view) in the absence of any external sensory cues (Wan *et al.* (1994); Recce and Harris (1996)).

Some models place heavy emphasis on inputs from path integration systems, and it has been proposed that place cell firing in the absence of sensory cues is the direct result of path integration processes (e.g., O'Keefe (1976); Muller *et al.* (1991b); McNaughton *et al.* (1996); Redish (1999)). McNaughton and collaborators propose that the hippocampus is an auto-associator that associates local views with movements to predict future local views, forming a sort of "transition table" (e.g., McNaughton (1989); McNaughton *et al.* (1991)). An ability to navigate in the dark thus arises by updating place representations with self-motion information. Before the model can show path integration abilities (and thereby the ability to navigate in the dark) in a given environment, the environment must be explored to set up the correct associations between self-motion and local-view information.

In the Cognitive Graph theory, Muller and colleagues argue that Hebb-type correlational learning along with random exploration of an environment will produce a synaptic weight function such that the weight between two place cells is inversely proportional to the overlap between their place fields (Muller *et al.* (1991a, 1996); Muller and Stead (1996)). The synaptic weights thus represent the distance between place field centres, and the connection matrix represents the topology of the space. They argue that a graph-search algorithm could plan paths using this structure. There is some evidence that similarly structured connections get set up after exploration (e.g. Wilson and McNaughton (1994)), but it is difficult to see how the search algorithms could be implemented neurally. Another proposal is that a combination of asymmetric LTP and phase precession[4] produces an asymmetric connection matrix in the recurrent connections of CA3 that can represent recently travelled routes (Blum and Abbott (1996); Skaggs *et al.* (1996)); these asymmetric connections could be used to guide navigation (Blum and Abbott (1996)). However, recent empirical data does not support these hypotheses. Levy's sequence learning model implements goal-directed navigation without a search algorithm, by assuming that a goal representation (e.g., of a source of water if an animal is thirsty) is partially activated at navigation and acts to pull path attractors towards it (Levy (1996)).

[4]Phase precession is a rapid LTP-dependent experience-dependent expansion of place fields in the direction that the animal enters a place field on route-following tasks (Mehta *et al.* (1997); Shen *et al.* (1997)).

An alternative proposal is that the hippocampus is the path integrator proper (McNaughton *et al.* (1996); Samsonovich and McNaughton (1997)). The model is based on the cognitive graph theory (Muller *et al.* (1991a)). A loop between the hippocampus and the subiculum is proposed to perform path integration. The cognitive graph needs to be pre-wired into the system before an animal explores an environment. This extension of the cognitive graph theory is known as the *multi-chart* model of the hippocampus (each map is a "chart").

2.4.4 Other models

Several apparently non-declarative functions are lost after hippocampal region damage, such as classical conditioning tasks that involve learning about unconditioned stimuli, configurations of stimuli, contextual information, or relationships spanning short-term delays (see chapter 3). Several models have addressed such non-declarative roles of the hippocampus (e.g. Hirsh (1974); Buhusi and Schmajuk (1996); Myers *et al.* (1996); Sutherland and Rudy (1989)). Most associative models of incremental learning assume that the hippocampus is required for some complicated forms of stimulus association (e.g. relational, configural or contextual), whereas the neocortex is sufficient for simpler stimulus-response associations, such as those underlying classical conditioning. A few of the models addressing non-episodic encoding are not associative models, but incorporate multi-layer networks (e.g. for classical conditioning (Gluck and Myers (1997)) or attentional and configurational mechanisms (Buhusi and Schmajuk (1996)). For example, in the Gluck and Myers model, the hippocampus is modelled by a 3-layer feed-forward network; the "hippocampal" hidden layer representation is used to teach the hidden layer of a "cortical" auto-encoder network. It is currently unclear whether such nets can be trained in a biologically plausible way. Notwithstanding, the higher proportion of non-associative models of classical conditioning compared to those for episodic learning suggests that these functions may be incompatible[5]. A recent paper acknowledges that the 'incremental' functions may in fact be based outside the hippocampus proper (Gluck *et al.*, 2003).

Other models have been concerned with predictive differentiation in the dentate gyrus or hippocampus proper (e.g., Levy (1985); Lynch and Granger (1992); Rolls and

[5]Given that classical conditioning, as modelled by Gluck and Myers, is not obligatorily hippocampally-dependent (see chapter 3) it would hardly be surprising if episodic memory and classical conditioning required different implementations.

Treves (1998)). This is consistent with several suggestions that the hippocampus is involved in predicting future events such as the arrival of an unconditioned stimulus, given current inputs. Other models have proposed that the hippocampus is responsible for forming configural associations (Schmajuk and DiCarlo (1992)) or relational representations (e.g., Buhusi and Schmajuk (1996); Hasselmo (1995); Hasselmo *et al.* (1996)). Whilst relational processing could support episodic memory, it could serve other functions too.

2.5 Summary

In this chapter I have outlined the important ideas and theories associated with the hippocampus. Most researchers agree that the hippocampus is important for autobiographical recall in humans and spatial learning in non-human animals; and that the importance of the hippocampus in the recall of at least some tasks, wanes over time. However, there is a lack of consensus on several key issues. There is little agreement about the specific types of learning and memory tasks that the hippocampus is crucial for, or about the long-term role of the hippocampus in memory recall. Few theories to date have attempted to directly relate the nature of hippocampally-dependent tasks to their long-term fate in the brain.

2.5.1 A note about terminology: 'episodic' and 'semantic' memory

Episodic and semantic memory have been defined in several different ways, initially in terms of the nature of the information being retrieved (e.g., episodic memory refers to memory for personally experienced events, whilst semantic memory corresponds with knowledge of the world, Tulving (1972)), and later on increasingly in terms of the phenomenal subjective experiences accompanying retrieval (episodic memory is accompanied by self-knowing 'autonoetic' consciousness that gives rise to feelings of mental time travel, whereas semantic knowledge is accompanied by knowing 'noetic' consciousness that gives awareness of familiarity with facts, Tulving (1983)). The terms 'episodic' and 'autobiographical' are often used interchangeably, although for some authors they have specific (and even contradictory) meanings. For example, Kopelman and Kapur (2001) state that the term 'autobiographical memory' is typically used to refer to a person's recollection of past incidents and events, whilst 'episodic'

is a broader term including performance on certain learning tasks such as the recall of word lists in addition to autobiographical memories; whereas Conway (2001) conceives of episodic memories as traces of sensory-perceptual details of recently experienced events that are of short duration (lasting less than 24-h), and autobiographical memory as a permanent system that represents knowledge of the self. The term 'semantic' memory is usually used more consistently by authors, but that is perhaps because it refers to such a broad range of types of information from grammar and word meanings, though categories and relationships between objects, to the knowledge about public events and famous personalities.

Clearly, episodic and semantic memory must be strictly operationally-defined if sense is to be made of empirical data, and especially if one is to adequately describe a theory of progressive semanticisation of memories with age (as I do in chapter 6). I therefore return to the issue of the relationship between episodic and semantic memory in several sections in this thesis.

Chapter 3

The elusive role of the hippocampus

In this chapter I examine empirical data on the role of the hippocampus in the acquisition of information. I conclude that task-dependent allocentric information and detailed episodic information cannot be acquired at all in the absence of the hippocampus. The acquisition of different types of semantic memory is impaired to different extents, depending on how much detail is required at recall and how much exposure there has been to information. The acquisition of associative information may be unaffected if the information to be acquired is relatively simple. However, as learning demands increase (such as when the information to be acquired is complex and cross-modal, of low salience or must be acquired quickly), associative learning is more likely to be affected by hippocampal damage. Generally, the acquisition of low-salience or incidental information is reduced by hippocampal damage. Procedural features of tasks such as the number of trials or stimulus-salience and the animals' learning history are as important in determining hippocampal-dependency as the high-level features that are typically used to describe tasks.

3.1 Introduction

This chapter focuses on elucidating the nature of tasks that require the hippocampus for their acquisition. I start from the observation that 'hippocampally-dependent' tasks can be divided into those that are *obligatorily* hippocampally-dependent and cannot be acquired at all after hippocampal damage; and those whose acquisition is *facilitated* by an intact hippocampus and probably mediated by the hippocampus under

normal conditions, but can be acquired to some extent after hippocampal damage. Since such deficits are likely to depend on qualitatively different facets of hippocampal function, it is important to make such a distinction for the purposes of formulating hippocampal theories.

In this chapter I use the term "hippocampally-dependent" to refer to tasks that require the hippocampus *at acquisition*, irrespective of whether the dependency is transient or permanent. Empirically, this refers to tasks whose acquisition is impaired by interference with the normal function of the hippocampus. The logically separate question of the long-term role of the hippocampus on such tasks is addressed in Chapters 5 and 6. Furthermore, whilst I focus on the hippocampal system there is no intention to imply that it alone supports any of the functions discussed: different "hippocampally-dependent" tasks are affected to varying extents by various extra-hippocampal lesions.

3.1.1 Issues for hippocampal theories

Most dominant current theories of hippocampal function are unitary and one-dimensional (Bannerman *et al.* (1999)), in that each hinges on a single central concept that attempts to encompass all hippocampally-dependent tasks. Of course, the "single concepts" evoked may be capable of capturing a wide range of tasks. For example, Levy (1996) argues that conceptualising the hippocampus as a sequence learner would unify various perspectives and predict that one-trial learning, finding short-cuts and transverse patterning would be hippocampally-dependent. I am not making an *a priori* assumption that unidimensional theories are *necessarily* inadequate. However, it is important to recognise that explanations *do* tend to be couched in terms of a single qualitative feature by which the hippocampus differs from other areas, even though we have no evidence that the functions of the hippocampus differ from those of other areas in a uni-dimensional qualitative functional manner.

The main motivation for positing that some single feature distinguishes the hippocampus from other areas appears to be parsimony. However, parsimony at the expense of truth is undesirable. Electrophysiological (e.g., Jung *et al.* (1994)), anatomical (e.g., Swanson *et al.* (1978); Amaral and Witter (1995)), *c-fos*[1] activation (e.g., Vann

[1]Expression of the *c-fos* gene is an indirect correlate of increased neural activity, and is induced under conditions of learning. Thus it can be used to detect differential brain activation, although it is not expressed in all areas.

et al. (2000)), lesion (e.g., Moser *et al.* (1995); Hock and Bunsey (1998); Richmond *et al.* (1999)) and inactivation studies (e.g., Moser and Moser (1998)) all provide support for the idea that there *may* be functional differences along the septotemporal axis of the hippocampus, and there is evidence that CA1, CA3 and the dentate gyrus can be differentially activated on certain tasks (Nitz *et al.* (1997); Wan *et al.* (1999)). These data imply the hippocampus *may* be involved in several independent computational functions and that it is unreasonable to attempt *a priori* to straitjacket all functions attributed to the hippocampus into a single explanatory concept. Of course, it may turn out that the hippocampus functions as a single functional unit, given the huge longitudinal interconnectivity in the hippocampus (Amaral and Witter (1995); Ishizuka *et al.* (1990)).

Characterising the nature of hippocampally-dependent tasks must be the starting point for any theory of hippocampal function. However, standard task descriptions often omit some of the key features that determine the extent to which a task will be hippocampally-dependent. Tasks are routinely described in terms of the type of stimuli thought to be involved and the logical nature of the task to be learnt. Thus a task might be described as an "allocentric spatial learning and navigation task in a watermaze with a hidden platform", or as an "eyeblink conditioning task using a tone as a conditioned stimuli and a puff of air to the cornea as an unconditioned stimulus": these are high-level descriptions of tasks in that they have abstracted away from more low-level detailed aspects of tasks, such as the salience of cues or the number of trials. There is accumulating evidence that lower-level features of tasks (such as the number of trials, and stimulus type and intensity) may in some cases be just as important as higher-level abstract task features in determining whether a task is dependent on the hippocampus. In addition, an animal's learning history may also affect the outcome of attempts to learn a task under given conditions. Greater consideration of these procedural level aspects of tasks is crucial to a proper understanding of what the hippocampus does.

The hippocampus is likely to be configured for particular kinds of information processing which will be required to varying degrees by various different tasks, rather than for the performance of particular logical type of task as defined by the experimenter. Therefore it seems probable that a description of hippocampal function at a low level of abstraction (say, that the hippocampus has a faster learning rate than

other areas or is most suited to associating disparate information) would be preferable to those theories that describe what the hippocampus does at a higher level (such as allocentric spatial learning, or storing episodic memories). A "low-level" characterisation of the hippocampus (say, that the hippocampus has a faster learning rate) may underpin hippocampal dependency in many tasks that appear unrelated when described at a higher level.

Moreover, since the high-level tasks that the hippocampus normally mediates in different species are likely to be different because the type of information that is important to an individual or encountered by an individual varies across species, if we are to capture any commonalities in the role of the hippocampus across species and across individuals with different life experiences, then we are likely to need descriptions and explanations in terms of lower-level functional features, rather than types of information processed.

Any hypothetical problem with the level of description cannot explain why different authors have reached such different conclusions about the role of the hippocampus. At least in part, one (understandable) reason is that researchers simply have different interests and have focused on different data. Authors have tended to address or emphasise different bodies of experimental literature, leading to many apparently incompatible theories of hippocampal function. In reality, these hypotheses may reflect different functions of the hippocampus. Although these theories may successfully account for the data set they address, it is difficult to see how this approach can lead to a broader inclusive understanding.

In sum, multiple types of constraints at several levels of description are likely to be needed to adequately characterise the types of tasks that are hippocampally-dependent. Procedural level features are at least as important as more abstract features. A wide range of literature should be surveyed.

3.1.1.1 Interpreting the data

A more insidious reason for why different authors have reached such different conclusions about the role of the hippocampus appears to stem from conflating tasks that show a differing severity of post-lesion impairment. For example, after hippocampal damage some tasks can be performed only at chance irrespective of the training period (e.g. the specifically allocentric aspects of spatial learning as evidenced by an-

nulus crossings (Pearce *et al.* (1998); McDonald and Hong (2000)), whereas on other tasks animals may show deficits when compared with controls after a fixed number of learning trials, but show asymptote performance as good as controls when allowed more training trials (e.g. arbitrary visuo-motor learning in primates Wise and Murray (1999); and some configural tasks, in rats, McDonald *et al.* (1997)). In both cases – that is, on tasks that show no improvement with training and those that do – performance may appear similar on many standard tests, such as immediate post-lesion performance, performance after initial training trials, and possibly on performance after a fixed number of trials (if there are insufficient trials to show learning in the second group). Perhaps because of this, several authors appear to treat such tasks as equivalently impaired after hippocampal lesion (see for example, Squire's treatment of semantic learning in humans and recognition memory in other animals).

In my opinion, these types of deficits must be interpreted differently for the purposes of formulating hippocampal theories, because they are likely to depend on qualitatively different facets of hippocampal function. Tasks that are performed at chance and cannot be relearnt after hippocampal damage are "truly" hippocampally-dependent, in that acquisition *can* only occur with an intact hippocampus. Those that can be learnt or relearnt to some extent after hippocampal damage (albeit slowly) are not obligatorily hippocampally-dependent: the acquisition of such tasks is facilitated by an intact hippocampus, and probably depends on it in an intact brain at least initially, but *can* be acquired by other brain regions under duress. Of course, we must be sure that such learning does not reflect learning by a partial hippocampus.

A priori it might be thought that hippocampally and non-hippocampally mediated behaviour will necessarily be performed by a different "mechanism". However in some cases the hippocampus might in effect act to improve non-hippocampally based learning (for example by providing a more quickly acquired scaffold for the recall of information represented in the cortical regions), whilst in other cases the 'solution' found by the hippocampus and by other regions might differ. Very careful investigations of the characteristics of learning and performance are needed to tease apart these possibilities.

This distinction between (1) what *can only be done* by the hippocampus, that cannot be compensated for, no matter how, or for how long a hippocampally-damaged animal is trained; and (2) what the hippocampus *normally does* in an intact brain is

crucial. Current theories appear to be agnostic as to whether they are addressing the former or the latter issue (at least, I can find no paper that explicitly addresses this issue). It seems plausible that the hippocampus has one or more computational features that are *qualitatively* different even from those of other parahippocampal areas (and wider brain structures) and which necessarily mediate performance on obligatorily hippocampally-dependent tasks; and other features that are only *quantitatively* different, such that the rate of learning of some tasks will change after hippocampal ablation.

3.1.1.2 Cross-species investigations of hippocampal function

There is converging evidence that the hippocampus or a homologous region plays a role in spatial learning in many species from reptiles through birds to mammals (Squire, 1992; Salas *et al.*, 2003; Day, 2003); as well as playing a role in non-spatial tasks such as reversal learning, extinction and context learning in species where this has been tested (Day, 2003). This implies that there are functional similarities across species, and that cross-species data could be potentially complementary. In my opinion, considering data from different species might help us unmask the underlying function of the hippocampus without getting seduced by very evident applications of hippocampal function in a given species. For example, theories of hippocampal function in humans have tended to focus on episodic memory, whilst in rodents they have focused on spatial abilities – this situation has undoubtedly arisen largely because of the relative ease with which rodents' spatial skills can be tested when compared with more specific "what/when/where memory", and the obviousness of general episodic memory deficits in amnesic patients compared to their spatial deficits. In addition, pooling data from several species might also allow us to partially compensate for empirical difficulties in different species (such as the paucity of well-defined hippocampal lesions in humans, and the difficulties in examining analogues of episodic memory in non-human animals). I therefore believe that even if we wished to understand the functions of the hippocampus in only a single species, then it would still be potentially informative to consider data from other species, as this data might lead us to a less blinkered and more rational interpretation of the data from the species under question.

The possible evolutionary effects of differences in species-specific ecology should

not however be ignored. Although there are similarities between the hippocampi and hippocampal homologues of different species in terms of neural architecture, connectivity and neurochemistry (indeed this is what makes them homologues), there can be quite large differences in internal cell field structure and connectivity with other regions in the brain (Clayton and Krebs (1995); Day (2003)), as well as large differences in the wider brain. In addition, there can be large, apparently heritable inter-individual differences in hippocampal anatomy and chemistry even within one species (e.g., Crusio (1996); Lemaire *et al.* (1999); Zilles *et al.* (2000)).

It is not currently known in any detail how specific hippocampal differences between species or between individuals affects function. However, accumulating evidence suggests it does. For example, the size of the hippocampus in food-storing birds correlates well with spatial abilities (e.g., Sherry *et al.* (1992); hippocampal differences in anatomy and chemistry between inbred strains of mice can have considerable consequences for an animals' behaviour (e.g., Crusio (1996); Lemaire *et al.* (1999); Zilles *et al.* (2000)), and cyclical changes in the hippocampus of female rats correlate with differences in spatial behaviour and memory (Desmond and Levy, 1997; Rudick and Woolley, 2000). Furthermore, there is some evidence that hippocampal damage might differentially affect even closely related species (Day (2003)). So, in considering data from different species we must be aware that a) the *tasks* normally carried out by the hippocampus of different species might not be the same (for example, supporting verbal recall of autobiographical experience versus migratory navigation), although this merely underlines the weakness of task-level descriptions of hippocampal function; and b) that there might be variations in how the hippocampus of different species carry out even the same functions.

Given our current state of knowledge and the lack of consensus about the function of the hippocampus despite the vast amount of data available, I believe that we can currently gain most from a rational consideration of the data available from all species. Therefore I draw on data from a range of species (though mainly rodents and primates, as the most detailed work has been done with these species).

3.2 Task acquisition and hippocampal dependency

In this section I investigate which tasks depend on the hippocampus for acquisition. The discussion is divided into sections focusing on:

1. Tasks whose acquisition is unaffected by hippocampal damage.

2. Tasks whose acquisition is facilitated by the hippocampus, so that acquisition is impeded but not prevented by hippocampal damage.

3. Tasks whose acquisition is obligatorily hippocampally-dependent, so that acquisition is completely prevented by hippocampal damage.

4. Tasks whose acquisition is inhibited by an intact hippocampus, so that hippocampal damage improves performance.

Subtly different variations on tasks that are usually defined as 'the same type' may belong in one or more of these categories depending on the fine details of training and testing procedures, such as whether quadrant occupancy or heading vectors are used to assess spatial learning, and whether implicit or explicit responses are used to assess recognition memory in humans.

The categories outlined may be useful for expositional purposes, but they mask some crucial issues, both empirical and conceptual. One problem is the necessarily tentative nature of the classification of "obligatory" and "facilitated" tasks when data is limited. Categorising a task as belonging to the "obligatory" category (with acquisition at chance irrespective of learning opportunities) assumes that documented training has been sufficiently extended to rule out slowed learning, which is often not the case. A recent study demonstrates this point well: de Hoz *et al.* (2003) replicated Moser and Moser (1998)'s study of spatial learning in rats with hippocampal lesions and found no evidence of learning after 4 trials, as reported by Moser and Moser (1998). However, de Hoz *et al.* found that with further training, rats could learn this task. I would be forced to categorise the task administered by de Hoz *et al.* (2003) as facilitated by the hippocampus and the task administered by Moser and Moser (1998) as obligatorily hippocampally-dependent. It should also be borne in mind that for some tasks whose acquisition is facilitated by an intact hippocampus, it is difficult to decide whether they are performed in a similar way in the presence and absence of an intact hippocampus. This has a bearing on whether the hippocampus is merely better at doing something that can also be done by other regions, or whether areas outside the hippocampus compensate for a damaged hippocampus on these tasks using a different "method".

One important observation is that procedural factors (such as the number of trials available, testing procedure, or the intensity of the stimuli on a task) can have a very profound effect on task acquisition after hippocampal damage. In some cases, whilst no learning may be apparent under one set of circumstances (leading to an 'obligatory" classification), slight changes in procedures allow learning to be demonstrated (leading to a classification of "facilitated" or even "unaffected"). The classifications of 'obligatory', 'facilitated' and 'no effect' in reality fall on a continuum, and often do not reflect categorical differences. Furthermore, within the 'facilitated' classification itself (of tasks whose acquisition is slowed or reduced by hippocampal damage) there is a continuum of effect.

3.2.1 Empirical considerations

The most relevant data on which to assess the dependency of specific task acquisition on the hippocampus is the traditional lesion/behavioural paradigm. Hippocampal activation on a task (as recorded by electroencephalograph, evoked response potential, imaging or cell recording studies, for example) cannot be used as an index of obligatory hippocampal involvement because activation is seen on many tasks which are not apparently affected by hippocampal lesions[2], and which presumably reflect functions that are not essential to performance on the experimenter-defined task.

There are several well-known problems with the traditional lesion approach, some of which are likely to be especially problematic for hippocampal studies. The most obvious of these is that lesion studies can only tell us what the brain can do in the absence of an area, not what that area does in an intact brain; and many of the tasks commonly used to probe hippocampal function can be solved using several different strategies dependent on different brain areas. Careful behavioural observation and tests of performance are therefore required to distinguish possible contributions to a learning task from different regions.

Another common problem with lesion studies is un-identified damage to structures other than those intended. Many tasks and functions that were once ascribed to the hippocampus are now known to be dependent on neighbouring structures. Thus

[2]For example, imaging studies show hippocampal activation on tasks such as sitting with closed eyes or visual fixation (Binder *et al.* (1999)), and recording studies find hippocampal complex-spike cell activity on tasks such as random foraging (Muller *et al.* (1987)) and delay conditioning (Berger *et al.* (1983)), none of which are hippocampally-dependent.

in the rest of this section, I will focus on studies that use modern stereotaxic neurotoxic lesions which generally reduce though do not eliminate unintended damage to surrounding areas and fibres of passage, and that use well-designed behavioural testing paradigms that allow various performance factors to be teased apart. The most pertinent data thus tend to be the more recent and better controlled studies, therefore many much-cited "historical" studies will be omitted (where other studies exist) because the lack of specificity in lesion technique and/or testing renders them difficult to interpret with respect to the role of the hippocampus.

3.2.1.1 Defining 'episodic' and 'semantic' memory

Generally speaking, a memory is usually designated to be 'episodic' if it refers to some details that are unique to one occasion, and 'semantic' if it refers to factual or generic information about the world. However, in reality most recall elicits both episodic and semantic components. Indeed, it is difficult to imagine how an episodic event (such as losing one's hat whilst feeding ducks in a park) could be recalled without recalling semantic information (such as the generic appearance of ducks and the local park). Thus I prefer to use the term episodic and semantic *aspects* of memory to make it clear that I am referring specifically to traces mediating the recall of particular components of a memory; even when the whole recalled memory may well contain both episodic and semantic components.

Episodic aspects of a memory encode information about unique occurrences of perceptual features or cognitive events, or unique co-occurrences of features or events that might not in themselves be individually unique. So, for example, a trace that could mediate the recall of a one particular unfamiliar goose with unique markings could be episodic, as would the recall of a single specific time that a swan and a goose were seen together on your local lake.

Semantic aspects of memory refer to traces that mediate the recall of information that is common to repeated occurrences of a similar event or experience. So, for example, a general memory of what a generic Canada goose or the local lake looks like would depend on the recall of semantic information. Similarly, recalling the appearance of a particular uniquely-identifiable goose that had been encountered *many* times might come to depend on semantic recall. Of course, in humans at least, semantic information can apparently be acquired in one exposure, for example, through being

told that "the capital of Mongolia is Ulaanbataar". However, I would argue that this is a learnt ability that depends on previous related learning experiences and possibly on the development of a symbol system, and thus is not really "one-shot" learning. I return to this is important issue in section 4.2.

One well-established way to visualise memory is as a hierarchy of linked representations (see figure 3.1). Assume that the activation of a node activates the nodes below it. Thus the activation of the "goose" node would reactivate the nodes representing features such as a goose's foot, wing and colour. Such a pattern of activation would therefore represent semantic recall. Activating the "specific visit to lake" node would activate nodes representing the co-occurrence of a goose, a swan and the lake. Thus the connections between that top node and the lake, swan and goose nodes could be said to represent the episodic aspects of the memory, although the semantic "content" of those nodes (the downward connections) must be activated to experience episodic recall.

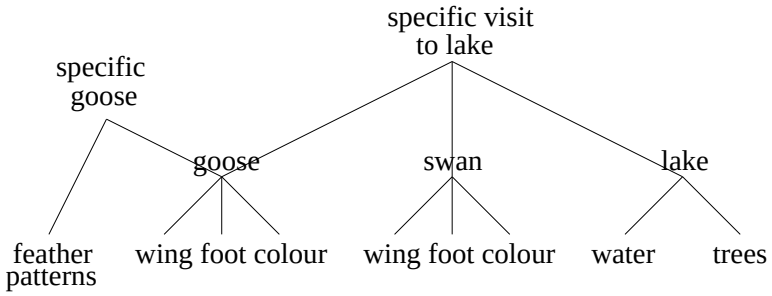

Figure 3.1: Memory can be visualised as a hierarchy of linked representations

The examples given to illustrate the terms 'episodic aspects' and 'semantic aspects' of memory are relatively clear-cut, which is useful for expositional purposes. However, I will show that there is a continuum between these two 'types' of memory, rather than a clear-cut biological category distinction. I return in detail to the relationship between episodic and semantic memory in section 4.6.2.

3.2.2 The acquisition of some tasks is unaffected by hippocampal damage

3.2.2.1 Skill learning and priming is unaffected

Initially, hippocampal amnesia was thought to affect all types of memory. However, it quickly became apparent that so-called "procedural" or skill learning tasks,

priming and some implicit tasks are preserved in amnesia. Thus patients with hippocampal damage can acquire perceptual/motor tasks such as mirror drawing (Milner (1962); Cohen and Squire (1981)), reading novel mirror-transformed words (Cohen and Squire, 1980), rotary pursuit (Corkin (1968)) and robot control (Shadmehr *et al.* (1998)) at relatively normal rates despite denying at the start of each new practise session that they have any familiarity with the task.

People generally have a lower perceptual identification threshold (and thus faster reaction times) for recently experienced stimuli: this effect is known as priming. Priming is normal in amnesics if they are tested appropriately. Perceptual priming is very long-lasting when compared to retention on standard explicit memory tasks, even in amnesics (e.g., Cave and Squire (1992); Tulving *et al.* (1991)) and is sensitive to the frequency of repetition even in people who cannot explicitly remember the stimuli (Wiggs *et al.* (1997)). The relationship between priming and recognition memory is controversial (see, for example, the discussion in Aggleton and Brown (1999) and associated commentaries), but it seems likely that priming can at least contribute to recognition/familiarity judgements.

3.2.2.2 Many implicit tasks are unaffected by hippocampal lesions

Implicit tests of memory for recently presented items can produce performance in amnesics indistinguishable to that of controls (e.g., Graf *et al.* (1984); Cohen and Squire (1981); Kitchener *et al.* (1998)) if the retention test provides partial cuing information (Squire, 1992). Amnesic patients can also learn novel single associations as normal in sensitive tasks that rely on perceptual identification, especially when the information is presented over multiple trials (Gabrieli *et al.* (1997); Musen and Squire (1993a,b)); and show intact implicit memory for newly formed verbal associations following single study trials (Goshen-Gottstein *et al.*, 2000). Amnesics can perform as well as normals on categorisation tasks such as classifying novel dot patterns or artificial grammar strings, despite having chance recognition of specific exemplars (e.g., Knowlton and Squire (1993); Squire and Knowlton (1995)).

Densely amnesic patients exhibit near normal performance on implicit sequence learning tasks such as serial reaction time tasks, (Nissen and Bullemer, 1987; Reber and Squire, 1994) and artificial grammar learning (Knowlton *et al.*, 1992; Knowlton and Squire, 1996). Several reviews have concluded that the areas that are damaged in

medial temporal lobe amnesics and Korsakoff's patients are not involved in sequence learning (Curran (1995); Clegg *et al.* (1998), but see section 3.2.4.5).

3.2.2.3 Delay conditioning is unaffected by hippocampal lesions

In the commonly used eyeblink conditioning paradigm, a tone (the conditioned stimulus or CS) is presented before an airpuff to the eye (the unconditioned stimuli or US) – the index of learning is blinking at an appropriate time after the tone. In *delay* conditioning, the CS and the US overlap; in *trace* conditioning, the CS terminates before the US begins. Hippocampal lesions do not impair acquisition of the basic delay paradigm (e.g., Solomon and Moore (1975); Moyer *et al.* (1990)), and humans with hippocampal damage are unimpaired on the delay eyeblink conditioned response task, despite being unable to describe it (e.g., Weiskrantz and Warrington (1979); Clark and Squire (1998)).

3.2.2.4 Recognition itself is unaffected by hippocampal damage

There is increasing acceptance that the perirhinal regions are crucial for recognition memory (for review see Murray *et al.* (2000), also Meunier *et al.* (1996); Murray (1996)). What role – if any – the hippocampus plays in recognition is less clear: the human neuropsychological data has been used both to support the view that the hippocampus is required for recognition memory (e.g., Squire (1992); Knowlton and Squire (1995); Manns and Squire (1999)) and that it is not (e.g. Eichenbaum (1994); Murray (1996); Aggleton and Brown (1999); Vargha-Khadem *et al.* (1997)).

Clinically, a profound recognition deficit is often reported on various tasks in patients with MTL damage (e.g., Stark and Squire (2003)). However, recognition abilities can be normal in patients with suspected *focal* hippocampal damage (for references see Mumby (2001)), implying that the recognition deficits reported in many amnesics may be due to extra-hippocampal damage. Clearly, given doubts about the extent of functional damage in human patients, it is difficult to draw conclusions about the role of the hippocampus from such studies. Interestingly, two groups of patients that apparently have very localised hippocampal damage (those with developmental amnesia such as Vargha-Khadem *et al.* (1997)'s well-known patients and others; and patients with selective damage to the fornix (see Easton and Parker, 2003 for a review) are unimpaired on many recognition memory tasks, including one-trial recognition for

lists of words and non-words, and familiar and unfamiliar faces. Several studies have also reported relatively preserved recognition memory in adult amnesic patients with hippocampal damage despite totally impaired explicit recall of the same information (Holdstock *et al.*, 2002; Bastin *et al.*, 2004).

This data suggests that the hippocampus is not needed for item recognition at short delays, although long-term storage might benefit from hippocampal processing. Several factors in addition to hidden pathology may have confounded research on hippocampal involvement in human recognition memory. Firstly, it is possible that recognition tasks can be mediated by recollection of the stimulus and its context ("remembering") which may depend on the hippocampus, and/or detecting stimulus familiarity ("knowing") which may depend on the perirhinal cortex (Mandler, 1980; Aggleton and Brown, 1999). Therefore, whether hippocampal damage leads to measurable "recognition"impairments would depend on whether "know" responses were sufficient to mediate a given task. Secondly, the common use of explicit verbal tests might disadvantage amnesic patients, although a capacity for explicit responses is not usually integral to a definition of "recognition". A third possible confound is that recognition memory in humans that takes place against a background of many already-acquired memories might depend on the disambiguation of episodes, which may depend on hippocampally-dependent scene memory (Gaffan, 1994b).

Studies of recognition memory in hippocampal monkeys are rather contradictory, even if we disregard early studies that are undoubtedly confounded by parahippocampal damage. Broadly speaking, studies show no impairments on recognition tasks such as delayed-non-match-to-sample (DNMS) or delayed-recognition-span (DRS) tasks with short delays (up to 1 or 2 minutes), with increasing impairments at longer delays (Zola-Morgan *et al.*, 1992; Alvarez *et al.*, 1995; Murray and Mishkin, 1998; Beason-Held *et al.*, 1999; Zola *et al.*, 2000). Various factors could explain the specific differences between the findings of different studies. These include 1) extra-hippocampal damage – some lesions are likely to transect perirhinal cortex efferent fibres such as those used in the studies of Alvarez *et al.* (1995) and Zola *et al.* (2000); 2) timing of training – pre-operatively in Murray and Mishkin (1998) and post-operatively in Alvarez *et al.* (1995); 3) the size of hippocampal lesions – larger hippocampal lesions result in less impairment at longer delays on DNMS tasks, (Murray and Mishkin, 1998; Baxter and Murray, 2001); 4) whether the animals were removed from the apparatus

before testing; and 5) the small number of subjects used in studies – for example, no significant differences were reported between different lesion groups in the study of Zola *et al.* (2000), but this is likely to result from a lack of power in the statistical tests[3], amongst other possibilities. The primate recognition literature is somewhat messy, but there is little compelling reason to believe that the hippocampus is required for object recognition per se, at least not with short delay periods. On balance, the fact that hippocampally lesioned animals can be unimpaired even when tested with a "list" of 40 objects for recognition is pretty impressive (Murray and Mishkin, 1998).

Studies of rodents with hippocampal damage find little or no impairment on object, odour and social recognition memory with possible signs of an impairment only at longer delays (e.g., Wood *et al.* (1993); Mumby *et al.* (1996); Dudchenko *et al.* (2000); Clark *et al.* (2001); Kogan *et al.* (2000), see also references in Hampson *et al.* (1999)). They also show an impressive intact memory for long lists of non-spatial objects (Dudchenko *et al.* (2000)). Mumby (2001)'s meta-review of studies using DMS and DNMS tasks to assess object recognition in rats with hippocampal formation damage concluded that there was no impairment in 17/18 studies, with a possible mild impairment at long delays in 2 studies. In addition, fornix lesions in both monkeys and rats have little effect on DNMS tasks (e.g., Rothblat and Kromer (1991); Zola-Morgan *et al.* (1989)). Since hippocampal and fornix lesions often (though not always) produce similar learning deficits, this is suggestive evidence that the hippocampus is not necessarily required for DNMS tasks.

Extra-hippocampal areas appear to be able to support recognition memory in the absence of the hippocampus. However, the hippocampus in an intact brain may normally be involved in the acquisition of the information that supports performance on recognition tasks. After all, when the information is experienced, a subject would not know how long information is to be retained. In accord with this, whilst rats with hippocampal damage show normal anterograde performance on novelty preference tasks, they show ungraded retrograde amnesia for novel objects experienced up to 7 weeks prior to lesion (Gaskin *et al.*, 2003). Lesioning the hippocampus has sometimes

[3] At acquisition, the number of trials required to reach criterion performance is reported as not being significantly different across the RF, IBO, ISCH and normal groups, although the initial trials-to-criterion range from a low group average of 19 to a high group average of 352, with the control group averaging 118. Furthermore, the relative pattern of acquisition data (with the RF groups most impaired, and the IBO group least) is repeated on the second administration of the task (4-9 months later) which would be unexpected if the differences between groups were not consistent.

been reported to increase the number of trials required to learn the DNMS task to criterion in primates (e.g., Alvarez *et al.* (1995); Beason-Held *et al.* (1999)), suggesting that the hippocampus might be involved in the normal acquisition of recognition tasks. However, given the possibility of the lesion in Alvarez *et al.* (1995)'s study transecting some perirhinal cortex fibres, and the partial nature of lesions in Beason-Held *et al.* (1999)'s study, then at most this result must be considered tentative. Rodent studies have found no significant difference in re-acquisition rates for a DNMS task comparing rats with hippocampal lesions and controls (e.g., Dudchenko *et al.* (2000)) or when initial acquisition is post-lesion (Clark *et al.* (2001)). However, it is possible that rodents and primates differ in this respect, with recognition memory being more dependent on the hippocampus in the primate.

In summary, I conclude that recognition tasks do not necessarily depend on the hippocampus. In most cases, damage to the hippocampus has no effect on the recognition component per se of recognition memory tasks, although it may play a role with increasing delay, suggesting that it may be essential for maintaining recognition information over longer periods. It is possible that the hippocampus may normally play a role in encoding information needed for recognition tasks in an intact brain.

3.2.2.5 Some spatial information can be acquired without the hippocampus

Tasks commonly used to assess spatial learning (such as finding food on a multiply-armed maze, escaping onto a platform in the Morris watermaze, or delayed matching-to-place) can be solved using several different strategies. On the Morris watermaze, for example, the platform can be found on the basis of *random* navigation (swimming blindly until an animal hits the platform – as must happen on the first trial of a hidden platform maze); *praxic* strategies (responses based on body movements); *taxon* methods (navigating on the basis on approaches to particular cues); *route* navigation (stringing praxic and taxon strategies together; *spatial*, *locale* or *allocentric* methods (using the geometric arrangements of constellations of cues) ; or combinations thereof, such as navigating randomly until a familiar view is recognised. Only the spatial/allocentric method of navigation is impaired in animals with proscribed hippocampal damage.

However, evidence is accumulating that the hippocampus is *not* needed for an animal to benefit from pretraining on standard spatial tasks such as the water-maze and radial mazes (in rats, Kimble *et al*, 1982; Parron *et al.* (2001); Poucet *et al.* (1991), see

also Moser and Moser (2000)). That is, it is not needed for the initial (usually untested) acquisition of spatial information about the environment that is then used to support subsequent navigation behaviour. In accord with this, Ramírez-Amaya *et al.* (1999) reported that synaptic reorganisation in the hippocampus was not seen in animals that were merely allowed to swim in a water-maze for a day (but only after 3 days of training on a maze task). This suggests that the hippocampus is probably *not* the site of a "cognitive map"[4]. Electrophysiological studies strongly support this interpretation, as hippocampal place cell responses can change, sometimes profoundly, without any change in the spatial lay-out of an environment (e.g., Bures *et al.* (1997); Jeffrey and Anderson (2003)).

Recent studies that have similarly divorced the *acquisition* of spatial information itself from the *demonstration* of such information through behaviour support the idea that learning about a spatial environment need not be hippocampally-dependent. Hippocampal lesions have no effect on the incidental acquisition of "pure spatial information" where learning is inferred from effects on subsequent reinforced learning: in rats, dorsal hippocampal lesions do not eliminate conditioned place preference retardation (White *et al.* (2003)); and hippocampal lesions (White and Wallet (2000), Gaffan *et al*, 2000), and fornix or entorhinal cortex lesions (Gaffan *et al.* (2003)) do not prevent normal enhancement of constant-negative learning by acquired allocentric information[5] (but see Good *et al.* (1998)). On the other hand, in humans, hippocampal damage has been reported to impair even implicit learning about a 2D spatial scene (Whitlow *et al.* (1995); Chun and Phelps (1999)). It remains to be determined whether this is due

[4]Definitions of cognitive maps (CMs) can be split into two main categories: (1) in the sense of Tolman (1948) and O'Keefe and Nadel (1978), in which the CM is a powerful representation allowing novel shortcut formation, amongst other things; and (2) in the sense of Gallistel (1990) in which a CM is simply any representation of space held by an animal (Bennett (1996)). Many researchers believe that the ability to swiftly find the position of a known platform from a new location in a maze after experience in that maze, say, depends on a cognitive map in the first sense. However, most studies of maze learning to date have allowed subjects to explore, see and travel through the areas from which they would subsequently be required to start on probe trials. Studies which have controlled the extent to which a subject has had access to the new region from which the platform must now be found, strongly suggest that accurate transfer performance depends on prior experience viewing distal cues from that region whilst navigating in that region (in rats, Sutherland *et al.* (1987), Alyan, 1994, and humans, Hamilton *et al.* (2002)): that is, short-cut behaviour is unlikely to depend on a CM 'calculation'. Interestingly, it is now widely accepted that despite their great navigational powers insects probably do not possess 'Tolman-like' cognitive maps either – instead they navigate using vectors, snapshots and landmark-based routes (Giurfa and Calpaldi, 1999). It may be that any true capacity for deductions based on 'cognitive map' use is limited to modern humans with a conception of 2-D maps, and is not a product of hippocampal function per se.

[5]When both egocentric and allocentric cues are available, acquisition of the allocentric incidental information is overshadowed in the lesioned animals of Gaffan *et al.* (2003).

to a non-equivalence of tasks, whether it results from extra-hippocampal damage in human amnesics, or whether it represents a difference between the species in terms of the role of the hippocampus.

Single-unit recording studies in the hippocampus have demonstrated that a large proportion of complex-spike cells have spatial correlates (e.g., see reviews in O'Keefe and Nadel (1978); Muller (1996); Eichenbaum *et al.* (1999)). However, to deduce from the prevalence of so-called "place cells" that the hippocampus has a specific role in spatial processing is a non sequitur, yet many authors do so. Instead, the predominance is likely to reflect the fact we observe and operate in a 3D world.

In summary, the data available suggests that the hippocampus is not involved in the acquisition of tasks that can solved using a non-allocentric strategy; and that it is not needed for the passive (non-rewarded) acquisition of information about the layout of an environment.

3.2.2.6 The acquisition of various associative tasks is relatively unaffected by hippocampal damage

Patients with selective hippocampal damage show little impairment on one-trial associative learning of word-word and face-face associative learning when tested implicitly (e.g., Vargha-Khadem *et al.* (1997); Gadian *et al.* (2000); Holdstock *et al.* (2002)). In monkeys, slowly acquired visual-visual paired associate learning is unimpaired by hippocampal lesions (Murray *et al.*, 1993), as is single-trial learning of object-reward associations (i.e. learning which of two objects should be approached for a reward, after fornix transections in monkeys, Gaffan *et al.* (1984); Spiegler and Mishkin (1996)). Rats with hippocampal or perirhinal/entorhinal lesions learn slowly-acquired individual visual-visual pairwise discriminations at a normal rates (Bunsey and Eichenbaum, 1996; Dusek and Eichenbaum, 1997, 1998), although human amnesics are impaired on a similar task (Reed and Squire, 1999). Rats with hippocampal lesions are also unimpaired at acquiring and retaining pair-wise olfactory discriminations (Jonasson *et al.*, 2004). Unimpaired learning on a context-object association task after fornix transections in monkeys has also been reported (Gaffan *et al.*, 1984), although Ridley and Baker (1997) reported deficits on a similar context-object tasks after fornix lesions. Perhaps, as has been found in other situations involving learning about a context, whether a deficit is observed depends on whether a single prominent stimulus in a

context can form an association with the object in the context-object association test. Together, these data suggest that the hippocampus might not be required for acquiring associative information about similar stimuli within one modality.

Some tasks that involve learning associations between stimuli of different modalities are also unaffected, for example, visual imprinting (e.g., in chicks, Horn (1998)), conditioned taste aversion acquisition (e.g., in rats, Yamamoto (1993) and chicks, Rose (1994)) and visuo-motor tracking (such as catching a ball, Lang and Bastian (2001)). These tasks seem somehow less "cognitive" than most of those typically discussed in the hippocampal literature, and may depend on pre-wired propensities that support the learning of species-specific useful associations.

In some cases, associative information can be acquired normally in the absence of the hippocampus only under very particular learning conditions. For example, odour-odour mappings can be acquired in one trial by rats with hippocampal damage only after extensive experience with similar mappings (Eichenbaum *et al.* (1986)); and fear conditioning can be achieved in one trial without the hippocampus only if the stimuli is sufficiently aversive (in rats, Izquierdo *et al.* (1999)). These data suggest that previous learning experiences and/or the salience of relevant stimuli determine whether the hippocampus is required for the acquisition of particular kinds of associative information.

3.2.3 The acquisition of some tasks is impeded by hippocampal damage

For some tasks, hippocampal lesions lead to *slower acquisition*, but normal levels of performance can be achieved with enough trials. The hippocampus could be said to facilitate the acquisition of these tasks, but is not essential for it. This contrasts with tasks in the next section (3.2.4) for which performance never improves beyond chance after hippocampal damage, despite any amount of training.

3.2.3.1 The acquisition of semantic memory is slowed after hippocampal damage

Semantic information is acquired gradually through life, making it difficult to assess the normal acquisition of such information after adult-onset hippocampal damage. Cases of childhood-onset amnesia are therefore particularly informative. A study that has recently received much attention is that of Vargha-Khadem *et al.* (1997), which presented findings from three young amnesic patients who had received bilateral

damage to the hippocampus at birth, age 4 or age 8. They were severely amnesic for everyday events and items on 90-min delayed recall (as would be predicted given the common view that the hippocampus is involved in episodic memory processes), but all attended mainstream school and had attained levels of speech and language competency, literacy and factual knowledge that were in the low-average to average range. Thus these patients have relatively intact general semantic learning in the face of severe episodic deficits. They were however impaired at semantic laboratory tests such as story recall. Similar data from other young subjects with apparently circumscribed hippocampal lesions has been presented by Gadian *et al.* (2000) and Isaacs *et al.* (2003), and data broadly in line with these findings obtained in young subjects with hippocampal and probable additional damage (Ostergaard, 1987; Broman *et al.*, 1997; Brizzolara *et al.*, 2003). One possible confound with these patients is a possible compensatory reorganisation of function, especially given the age of the patients. However, an unpublished study reported in Gadian *et al.* (2000) found that patients suffering anoxia in early or later childhood are indistinguishable on memory tests, and Isaacs *et al.* (2003) recently reported no differences in the relative preservation of semantic memory in the face of episodic impairments, suggesting that age-related compensation is unlikely to be the whole story. Admittedly, a recent imaging study suggests that the remaining hippocampus in at least one of Vargha-Khadem *et al.* (1997)'s subjects shows a similar pattern of activity to control patients (Maguire *et al.*, 2001b), suggesting that it may be functionally active to some extent. However, the point here is that relatively normal semantic information can be accumulated from the world during life, in the face of severe episodic memory deficits.

In apparent contrast, clinically it is widely accepted that amnesics with presumed hippocampal damage are impaired at acquiring general factual (semantic) knowledge. Studies that have explicitly trained amnesics in the laboratory (e.g., Glisky *et al.* (1986a,b); Shimamura and Squire (1987); Reed *et al.* (1997); Vargha-Khadem *et al.* (1997); Holdstock *et al.* (2002)) have generally found none or a little semantic learning (for example, of new vocabulary, facts or stories). Any acquisition is slow and arduous in these cases, and subjects usually show inflexibility in applying the knowledge learnt. It has also been claimed that amnesic patients cannot acquire new facts from 'real-world' exposure. It has been reported that HM has retained almost no new facts since his lesion (Gabrieli et al. (1988)); and amnesic patients with hippocampal

and extra-hippocampal damage are reported to show no (Verfaellie *et al.*, 1995) or extremely limited (Cipolotti *et al.*, 2001; Reed and Squire, 1998) learning of vocabulary that had entered public language post-lesion. Cipolotti *et al.* (2001)'s patient also showed no acquisition for post-morbidly experienced public knowledge.

However, 'real-world' learning has been reported in a few cases: Kitchener *et al.* (1998) reported post-morbid learning about public events and vocabulary in their severely amnesic patient. Similarly, a patient (YR) with selective bilateral hippocampal lesions showed normal discrimination of famous events and names from non-famous ones for the post-morbid period, with some ability to categorise people according to the nature of their fame, and a lesser ability for event categorisation (Holdstock *et al.*, 2002). Interestingly, O'Kane *et al.* (2004) has now reported clear evidence for semantic learning about famous personalities in the severely amnesic HM. Therefore it may simply be that clinical assessments have tended to refer to information to which there has been insufficient exposure to demonstrate acquisition in the adult. Interestingly, imaging studies support the idea that the hippocampus is selectively important for the recall of material that has only been infrequently encountered: whilst activity in the rhinal cortex is predictive of subsequent memory only for high frequency words, hippocampal activity was predictive of memory for both high and low frequency words (Fernandez *et al.*, 2002).

More generally, fewer semantic details on tests or public events and knowledge of personalities are recalled by amnesic patients (Nadel *et al.*, 2000; Holdstock *et al.*, 2002). Spatial semantic information can also be lacking in detail and higher-order complexity in amnesic patients. Rosenbaum *et al.* (2000) reported that their amnesic patient KC was progressively more impaired as more specific information was required: his performance at identifying oceans and continents on a map was normal, but he was impaired at identifying cities, and even more impaired at locating geographical features. Therefore, even when some semantic information can be acquired or retained, it may be deficient in detail.

3.2.3.2 The acquisition of trace conditioning tasks is very much reduced by hippocampal damage

Tasks which require timing across a temporal gap are impaired by hippocampal lesions. For example, acquisition of a trace fear conditioning task (in which a 15-s

tone precedes a shock by a 30-s trace interval) is impaired by damage that includes the hippocampus (McEchron *et al.*, 1998). Similarly, Huerta *et al.* (2000) found that knock-out mice that lacked NMDA receptors in CA1 were slow to acquire a trace fear-conditioning task in which the white noise (CS) and shock (US) were separated by 30-s, but unimpaired when the trace interval was removed. It is possible that in this task, the gap that needs to be bridged may simply exceed the short-term memory capacities remaining after hippocampal damage, as proposed by Nichelli for human amnesics.

However, the hippocampus is required for bridging very small temporal gaps too. If the trace interval exceeds ∼500-ms, hippocampal lesions severely impair the subsequent acquisition of eyeblink conditioned responses (e.g., Moyer *et al.* (1990); Solomon *et al.* (1986)). Human amnesics are also impaired at trace tasks when the trace period exceeds ∼500-ms (e.g., Clark and Squire (1998)). On well controlled studies, a limited acquisition of conditioned responses is evident after hippocampal lesions (Takehara *et al.*, 2003). Therefore, I categorise trace-conditioning as a task that is merely facilitated by an intact hippocampus, rather than obligatorily-dependent on the hippocampus for acquisition, although the acquisition of trace eyeblink conditioning is very severely impaired by hippocampal lesions. This perhaps underlines the fact that tasks that are affected by damage to the hippocampus fall on a continuum of effect, and that the tasks that I have separated out as 'obligatorily hippocampally-dependent' are merely at one end of that continuum.

Why trace eyeblink conditioning is impaired after hippocampal lesions is difficult to explain, since the time interval employed (0.5-sec) is clearly within remaining short-term memory capacities that can support recognition memory over many seconds. However, *eyeblink* conditioning is itself particularly difficult to learn and this may interact in some way with the hippocampus. Eyeblink responses even in intact animals are poorly learnt when the trace interval is longer than 1 or 2-secs (Solomon *et al.* (1986); Moyer *et al.* (1990), whereas fear heart-rate responses can be learnt with trace intervals of up to 9-secs and shock-induced freezing responses to a tone with 30-sec delays (McEchron *et al.* (1998)). Similarly, trace eyeblink responses typically take many more trials to learn than trace fear responses — 7-14 days for eyeblink responses (Moyer *et al.* (1990)) compared to 1-2 days for trace fear (McEchron *et al.* (1998)). Even delay eyeblink conditioning takes longer to acquire than tone-freeze con-

ditioning, which can be learnt in one trial if the shock is sufficiently aversive (Izquierdo *et al.* (1999)). It is well-known that some associations are particularly difficult for some species to learn (e.g., a chick cannot learn to associate a tone with subsequent sickness, but will avoid a bitter tasting bead that precedes sickness (Rose, 1994)). It is possible that tone-eyeblink conditioning is one such association.

On trace eyeblink conditioning tasks, the performance measure is whether the CR – the eyeblink – is emitted in a particular narrow time window after the US. There is some evidence that different lesions affect the *timing* of CRs, which may be independent of whether lesioned animals can learn to make the US-CR association. In addition, it is becoming increasingly obvious that the procedural aspects such as cue modality and intensity, and whether the period to be timed is filled or empty of ostensibly timing tasks affects the behaviour of the animal. Therefore results cannot be interpreted solely in terms of timing processes (Buhusi and Meck (2000)) or in terms of whether the animal can acquire the association. Hippocampally lesioned animals may be differentially affected by some of these parameters. In accord with this, in lesioned animals whether the CR occurs before or after the US appears to depend on the nature of the US (O'Keefe (1999)).

In summary, the hippocampus is required for some tasks that must span temporal gaps. Trace eyeblink tasks are the most well-documented, but it is possible that eyeblink conditioning is unusual in some way.

3.2.3.3 The acquisition of some 'cross-modal' associative tasks is slowed by hippocampal damage

Patients with hippocampal damage are impaired at learning cross-modal associations, such as face-voice, object-location, picture-sound, and word-position associations, although they can eventually show good performance with extended training (e.g., Vargha-Khadem *et al.* (1997); Mayes *et al.* (2001); Holdstock *et al.* (2002)).

As noted by Brasted *et al.* (2003), many studies have reported that neither hippocampal or fornix lesions have a discernable effect on the learning of conditional visuomotor associations in rats. However, monkeys with removal of the hippocampus and subjacent cortex (Murray and Wise (1996); Wise and Murray (1999, 2000)) are extremely slow at acquiring visuo-motor mappings in which they must learn to make an arbitrary movement in response to a visual stimuli, although they do eventually reach

the same level of performance as controls. Brasted *et al.* (2003) reported a similar effect (with fornix lesions) when there is no spatial component to the task and responses are distinguished by temporal features (e.g., tapping or long presses), when more than two visuo-motor response pairs must be acquired. Although it is difficult to precisely compare the speed of learning in primates and rodents because the amount and nature of pretraining and existing knowledge is unclear, rats typically require hundreds or thousands of trials to acquire two conditional associations, whereas monkeys can acquire such information in a few trials.

It is plausible that the representation of conjunctions of stimuli is stored in a different area to the representation of the individual elements themselves. Therefore configural tasks can be considered to be 'cross-modal' if the term is used to refer to tasks that require the association of information represented in significantly different geographical areas of the brain. The literature on the effects of hippocampal damage on non-linear learning are mixed. Early data suggested that the hippocampus was obligatory for learning tasks that required a configural solution, i.e. differential reinforcement of elements and compounds (such as negative patterning, AB−, A+, B+; or bi-conditional learning, AB+, AC−, DB+, DC−), but accumulating data suggests that areas outside the hippocampus mediate such learning (see Rudy and Sutherland (1995) for a review). However, there is suggestive evidence that hippocampal lesions slow the acquisition of a negative patterning task (AB−, A+, B+; in rats, (McDonald *et al.*, 1997)) and mildly slow the acquisition of a bi-conditional auditory/visual discrimination task (AB+, AC−, DB+, DC−; in rats, (McDonald *et al.*, 1997)).

Rats with either dorsal or ventral hippocampal lesions have sometimes been reported to show slowed acquisition of tasks that use hunger signals as a conditioned stimulus (e.g., Davidson and Jarrard (1993); Hock and Bunsey (1998)), although Deacon *et al*, 2001 reported no impairment on a conditional object discrimination cued by internal state.

Honey *et al.* (1998) reported that rats with neurotoxic hippocampal lesions fail to recognise (orient) when combinations of familiar cross-modal stimuli (tone or click and constant or flashing light sequences) are rearranged, although they orient as normal when properties of the visual targets changed. However, whether this reflects changes in normal orienting behaviour or a failure to learn the requisite association is unknown, as the animals were not trained further to see if they could be trained to

behave 'normally' on this task.

Taken together, these data suggest that hippocampal damage might lead to slowed acquisition of information about associations between stimuli of different kinds. In accord with this idea, functional neuroimaging studies find greater hippocampal and para-hippocampal activation on associative tasks than during the learning of single items (Henke *et al.*, 1997, 1999; Yonelinas *et al.*, 2001).

3.2.4 The acquisition of some tasks is totally prevented by hippocampal damage

The tasks examined in this section can be considered to be *obligatorily* hippocampally-dependent as they cannot be acquired at all after complete hippocampal damage.

3.2.4.1 Normal episodic memory depends on the hippocampus

Several studies have linked discrete human hippocampal formation lesions with anterograde memory impairments (e.g, Zola-Morgan *et al.* (1986); Kartsounis *et al.* (1995); Rempel-Clower *et al.* (1996); Reed and Squire (1998); Kapur and Brooks (1999)[6]). Anterograde amnesia (AA) typically manifests itself as a severe deficit on both experimental and day-to-day memory tasks, such as recalling a list of unrelated words, or what was eaten for breakfast, as well as remembering significant personally-experienced events. Hippocampal patients are also widely reported to fail to orientate themselves in space and time.

One test that is commonly administered to test anterograde memory in amnesic patients is list-learning of words or pictures. Patients are quizzed on the list content, as well as the 'context' of items within it, such as whether items were from the first or second list or presented on the right or left. In general, there is very little evidence of such learning after hippocampal lesions, whether patients are tested hours, days or weeks after initial acquisition. For example, one amnesic patient with hippocampal damage showed chance performance in remembering which list correctly recognised words were from and for the recall of the correct order of word lists; and despite performing normally on a forced-choice word pair recognition task, was at chance in recognising the original order of presentation of items in the pairs (Mayes *et al.*,

[6]Since many cases of non-surgical hippocampal damage arise from anoxia, and there is evidence that anoxia can result in "invisible", but functionally relevant extra-hippocampal damage (Mumby *et al.* (1996)) there remains the possibility that even patients with apparently circumscribed damage at post mortem have sustained extra-hippocampal damage.

2001; Holdstock *et al.*, 2002). Vargha-Khadem *et al.* (1997)'s well-known patients are similarly impaired.

The majority of formal memory studies of personally experienced events probe memories that have been acquired years or even decades before the test date. Therefore the hippocampus' possible role in the initial acquisition of information (the topic of this chapter) and its long-term maintenance (Chapters 5 & 6) is confounded in much of the available data. However, for memories of all ages, it is the most specific details of memories that appear particularly affected by hippocampal damage. For example, in Oxbury *et al.* (1997)'s study, the patient developed severe AA combined with 16-months RA for autobiographical information, with a "patchy RA for earlier years mainly for sequencing and detail of events" after a series of convulsions that apparently destroyed the remaining hippocampus[7]. Sequencing and detailed information is of course exactly what distinguishes event-specific episodic information from generic semantic information or semanticised 'episodic' information. Similarly, one of the subjects in Kapur and Brooks (1999)'s study apparently confused places that he had visited alone with those he had visited with his wife during the period of RA. Again, it is the more specific features ("who with?") that are missing, not the semantic information ("I visited place X"). Patient KS (Kitchener *et al.* (1998)) was aware that his son had been accidentally shot in the eye, but did not know how, when or where the event had happened, or even if he had been present (he had). Therefore whilst episodic-like information can be recalled after hippocampal damage, it is deficient in important details and has the flavour of semantic information.

List learning and the recall of item context is widely assumed to be analogous to the recall of personally experienced memories. However, these tests of information are widely different in several ways, such as the typical retention period (minutes, hours or possibly weeks on list-learning tasks; possibly months, but usually years and decades for autobiographical recall); complexity of the context information; and the significance of the information to the individual. As I discuss in the next chapter, all of these factors play a role in determining the importance of the hippocampus in the acquisition of information. Furthermore, Gilboa *et al.* (2004) has recently shown that the functional anatomy of list recall and autobiographical recall is significantly

[7]The patient had previously had a left temporal lobectomy, which led to a mild verbal memory impairment. A later set of convulsions apparently led to atrophy of the right hippocampus (with sparing of the right EC, right PHG and rest of right temporal lobe).

different. Care should therefore be exercised in extrapolating from list-learning events in the laboratory to real-world autobiographical events.

To date, most studies with non-human animals have been unable to demonstrate learning that is analogous to human episodic memory (Griffiths *et al.*, 1999). To do so they would have to demonstrate that performance depended on the recall of specific detailed experiences ("remembering"), rather than on, say, simple familiarity for particular items or responses ("knowing"). Several tasks that were once assumed to assess the same memory systems as those lost in hippocampal amnesia (such as DNMS tasks; object-in-place tasks, Gaffan (1994b); object-in-scene tasks, Gaffan and Parker (1996); Murray and Mishkin (1998)) are typically learnt over several trials. Therefore (amongst other confounding factors) the animals may be depending on trial-*non*-specific information (Griffiths *et al.*, 1999), and acquisition could not be considered analogous to episodic information as usually defined. Griffiths and colleagues propose that food-storing behaviour that requires memory for the location of the cache, the identity of the food in the cache and when it was cached, and whether the cache has been emptied (e.g., in scrub jays, Clayton and Dickinson (1998), see review in Griffiths *et al.* (1999)) *does* provide evidence for episodic-like memory in animal. This is very reminiscent of Tulving (1972)'s original definition of "episodic" memory as providing information about the *what*, *where* and *when* of an event. In accord with this, interference with the hippocampus in food-storing birds disrupts spatial memory for the location of food (Shiflett *et al.*, 2003).

In sum, the hippocampus appears to be particularly important for the storage and/or recall of typical detailed episodic or autobiographical information.

3.2.4.2 The *use* of allocentric spatial information is prevented by hippocampal damage

Morris *et al.* (1982) first reported that hippocampal lesions led to a pattern of impaired allocentric navigation and preserved cue- and response-based navigation. This finding has subsequently been replicated many times (e.g., in rats, Morris *et al.* (1990); Cho *et al.* (1999); Czurkó *et al.* (1997), and birds, Bingham et al. (1990); Ioalè *et al.* (2000)). Damage to the dorsal hippocampus is sufficient to produce similar deficits (e.g., Bannerman *et al.* (1999); Moser *et al.* (1995); Silva *et al.* (1998), although see de Hoz (2000)). Deficits after hippocampal lesions which can most easily be explained by the impaired use of allocentric information are also found on disparate tasks such as the lattice maze

(Okaichi, 1996); radial arm water maze (Diamond *et al.* (1999)); and dry radial arm mazes (Hunt *et al.* (1994)).

Closer investigation shows that hippocampal animals are more likely to be impaired when the reward site changes frequently so that new locations have to be continuously learnt, than if the reward site is held constant. Thus hippocampal animals are particularly impaired when the correct arm on a 8-arm radial maze task varies from trial to trial (e.g., Olton & Shapiro (1978); Hunt *et al.* (1994)); or when the platform position (e.g., Steele and Morris (1999)) or start position (Eichenbaum *et al.*, 1990; Compton *et al.*, 1997; McDonald and Hong, 2000) in a water-maze changes from one set of trials to the next. This makes sense only if the role of the hippocampus is to *use* spatial information or store ongoing information, rather than merely to form a 'map', as the spatial environment has not changed on these tasks. Similarly, although the map would be the same, damage to the hippocampus leads to deficits when the positions that must be distinguished are made more similar, for example, by being closer together (in rats, McDonald and White (1995); Gilbert *et al.* (1998); and mice, Etchamendy *et al.* (2003)). In accord with this, tasks that specifically require objects to be associated with particular locations are impaired after hippocampal damage (in rats, Sziklas and Petrides (2002)). Data from recording studies of hippocampal complex-spike cells also strongly supports the idea that these cells record associative information beyond spatial information: direction of travel (Markus *et al.*, 1994; Gothard *et al.*, 1996), task information (Markus *et al.*, 1995; Wiener *et al.*, 1989; Eichenbaum *et al.*, 1994; Wood *et al.*, 2000), and changing rewards (Breese *et al.*, 1989; Gothard *et al.*, 1996; Wood *et al.*, 1999), for example, have all been shown to affect CS cells' firing rates.

Hippocampal rats often significantly improve their performance with time on many apparently "allocentric" tasks, and merely take more trials than controls to reach criterion (e.g., Eichenbaum *et al.* (1990); Compton *et al.* (1997)). However, such performance is only observed when relatively insensitive measures of learning are used, and when tasks *can* be performed to some extent on the basis of non-allocentric strategies. If quadrant occupancy on a water-maze or arm choice on a multi-armed maze is used as the dependent variable then animals merely require increased trials to obtain criterion performance, because less accurate non-allocentric strategies can support performance on these measures. Probe trials suggest that any improvements

that occur are due to learning the procedural aspects of the task, cue-based strategies based on prominent olfactory, auditory or visual cues (e.g., Morris *et al.* (1982)), utilisation of odour trails (Means *et al.* (1992)), response strategies (e.g., Packard *et al.* (1989)) and/or increasingly accurate path integration strategies (e.g., Alyan and McNaughton (1999)). When more subtle tests of allocentric learning such as annulus crossings are used, chance asymptote performance is seen.

The data discussed in section 3.2.2.5 suggests that the hippocampus is not needed to form a representation of the spatial environment per se. One key difference between the tasks discussed in that section and the standard allocentric tasks discussed here, is that here the animals are required to perform a particular learnt response to demonstrate acquisition of spatial information. An alternative interpretation is that the hippocampus controls the motor/navigation processes themselves – Whishaw *et al.* (1995) suggested that the hippocampus might be responsible for "getting there as opposed to knowing where". However, this seems unlikely as animals are usually *not* impaired with respect to controls on the *first* trial of a spatial task suggesting that the deficit is in *acquiring* information, rather than in motor abilities. It seems more plausible that the hippocampus is needed to represent associations between what is present, happens or should be done at different places in an environment (which may be episodic or stable), with the representation of that environment.

Humans with hippocampal damage also show deficits on spatial tasks, particularly in memory for allocentric spatial information (Holdstock *et al.*, 2000; Burgess *et al.*, 2001; Kessels *et al.*, 2001). A similar pattern of impairments is seen in remembering neighbourhoods, both real and virtual. Amnesic patients may remember the broad gist of both real and virtual environments, but are typically unable to be specific about detailed items in the environment (Rosenbaum *et al.*, 2000).

As would be predicted if the hippocampus plays an important role in allocentric navigation, functional imaging studies find specific activation of the hippocampus on tests that tax topographical memory (e.g., recalling spatial routes, Maguire *et al.* (1997); recalling landmarks or locations, Maguire *et al.* (1997); Aguirre and D'Esposito (1997), mental navigation, Ghaem *et al.* (1997); performing an immersive virtual navigation task, Maguire *et al.* (1998a); or recalling environmental knowledge, Aguirre and D'Esposito (1997)). Interestingly, *c-fos* imaging studies have shown hippocampal activation proportional to the allocentric spatial demands of a radial maze task (Vann

et al. (2000)).

In conclusion, the hippocampus is implicated in associating allocentric information with what should be done, or what is present, at particular points in an environment.

3.2.4.3 "Normal behaviour" may depend on the hippocampus

Lesions to various components of the hippocampal formation and their interconnections can cause spontaneous hyperactivity in both familiar and novel environments in rodents (e.g., in rats, Douglas and Isaacson (1964); O'Keefe and Nadel (1978); Olton *et al.* (1979); Whishaw and Jarrard (1995); Cassel *et al.* (1998) and mice, Frankland *et al.* (1998)). Hippocampally-damaged animals also tend to show stereotypical motor behaviour in some circumstances (O'Keefe and Nadel (1978), p240); and can show changes in their behavioural repertoire such as reduced rearing and increases in running around the perimeter of enclosures (Harley and Martin, 1999); changes in exploratory behaviour (Whishaw and Jarrard, 1995; Save *et al.*, 1992a,b); changes in object marking (Harley and Martin, 1999); and changes in eating patterns (Clifton *et al.*, 1998). Exploratory differences between controls and hippocampals are generally exacerbated under novel conditions (O'Keefe and Nadel, 1978); and lesions enhance stress and startle responses to unexpected or aversive events (Anagnostaras *et al.*, 1999). The hippocampus may also be directly involved in regulating hormones, such as those affecting the hypothalamo-pituitary-adrenal axis (e.g., Lathe (2000); Lemaire *et al.* (1999)) which might affect general behaviour by changing baseline physiological stress hormones.

Without further information it is difficult to interpret such post-lesion changes in terms of a cognitive theory, although O'Keefe and Nadel (1978) proposed that all such changes reflect impaired spatial novelty processing. These changes in activity may confound typical tests of "pure" problem solving ability that involve, say, delayed responses or performance measurements that depend on an animals' rate or speed of responding. Therefore in this chapter emphasis is placed on studies that have attempted to control for these factors (for example, by measuring heading vector rather than speed to reach platform in a water-maze, or comparing response rates to the new baseline of activity). It has been proposed that hyperactivity underlies many hippocampus lesion-induced deficits. However this cannot be the whole story:

whilst fornix lesioned rats are initially (e.g., 15-d after lesion) significantly more active (and more spatially impaired) than hippocampally-lesioned rats (Whishaw and Jarrard (1995)), after a 4.5-month recovery period, both groups are equally active, but the hippocampal group is more impaired on the radial maze (whilst both are equally impaired on the Morris water-maze, Cassel *et al.* (1998)). Therefore hyperactivity can be divorced from task performance.

The visual paired-comparison (VPC) task assesses recognition abilities by assessing spontaneous novelty preferences. VPC studies consistently show that hippocampal damage leads to a decreased preference for looking at new stimuli after short delays (in humans, McKee and Squire (1993); and monkeys, Bachevalier *et al.* (1999); Zola *et al.* (2000), although Gaskin *et al.* (2003) has reported no impairment in rats on a novelty-preference paradigm which depends on rats' natural propensity to explore novel objects). Thus deficits on the VPC task do seem more profound than that of other recognition tasks after hippocampal lesions (see section 3.2.2.4). Interestingly, Baxter and Murray (2001)'s meta-analysis of the extant studies of recognition memory in monkeys strongly suggests that VPC and DNMS tasks dissociate in terms of hippocampal dependence, with greater hippocampal damage associated with increasing deficits on the VPC tasks, but with less impairment on the DNMS task. Given the conclusions of section 3.2.2.4, it seems likely that it is not the recognition component of the VPC task that is hippocampally-dependent, but perhaps the "natural" untrained motivation to explore. In DNMS tasks, animals are trained to make certain responses. It is currently unknown whether animals on the VPC task can be trained to follow a 'rule' to orientate and respond to the novel item.

3.2.4.4 Transverse patterning may depend on the hippocampus

Transverse patterning tasks (A+ vs B−, B+ vs C−) are sensitive to hippocampal damage. Close to chance performance is seen on the configural stage despite relatively normal learning of the non-configural components (in rats, Alvarado and Rudy (1995, 1993) and humans, Rickard and Grafman (1998)). This implies that information is represented differently in the absence of a hippocampus.

3.2.4.5 Tasks that depend on explicit sequence information are performed at chance after hippocampal lesions

Although amnesic patients can show relatively normal levels of implicit sequence learning (see section 3.2.2.2), more subtle testing on those tasks reveals that the patients do not acquire higher order information as well as controls (Curran (1997); Hopkins *et al.* (2004)). Many amnesic patients show chance knowledge of the correct order of word lists, even when recognition tests are used (Mayes *et al.*, 2001; Holdstock *et al.*, 2002).

Non-human animals may also be impaired on true sequence learning. For example, the performance of hippocampally-lesioned monkeys on a delayed-recognition span tasks is not improved by repeated sequences, in contrast to controls (Beason-Held *et al.* (1999) but see Murray and Mishkin (1998)). Like humans, rats have been shown to implicitly acquire sequence information, but cannot use this information explicitly (Kesner *et al.*, 2002; Hopkins *et al.*, 2004). Honey *et al.* (1998) also reported that rats with neurotoxic hippocampal lesions fail to orient when combinations of familiar cross-modal stimuli (tone or click and constant or flashing light sequences) are rearranged, although they orient as normal when properties of the visual targets changed. However, without further information it is unclear whether this results from deficits in sequence learning, factors such as the cross-modal nature of the information to be associated, or the fact that learning is indexed by an orienting response.

3.2.5 The acquisition of a few tasks is improved by hippocampal damage

3.2.5.1 Learning egocentric responses is enhanced by hippocampal damage

On some tasks, more than one learning/memory system in the brain can provide a solution. For example, many spatial tasks can be solved either by learning to make particular egocentric body turns (which would depend on the caudate putamen) or by allocentric navigation (which requires the hippocampus). Whether the caudate or hippocampal strategy dominates usually depends on the amount of training and the time elapsed since beginning training. However, inactivation of the hippocampus results in caudate-dependent response learning being seen at all stages in training (Packard and McGaugh (1996); Shroeder *et al.* (2002)), thereby speeding the development of a reponse strategy. The acquisition of several other tasks can be speeded by

hippocampal lesions. These include performance on a two-way active avoidance task that requires the animal to return to a place in which it has been shocked (O'Keefe and Nadel (1978)), and a win-stay radial maze task which requires the animal to revisit the area from which it has removed food in order to receive another food reward (McDonald and White (1993); Packard *et al.* (1989)). This suggests that the hippocampus and other memory/learning systems may sometimes interact competitively.

The fact that the hippocampus inhibits the acquisition or expression of information in other regions is very important. I examine the question of competition between the hippocampus and other memory systems in detail in section 6.5.

3.2.5.2 Some tasks may be performed faster after hippocampal damage

Fornix and ventral hippocampus lesioned rats have been shown to be superior on some measures of learning (such as the latency to find a platform, Bannerman *et al.* (1999)). This appears to be dependent on lesion-induced hyperactivity and increased swim speed under stress (for more information see section 3.2.4.3).

3.3 Summary and conclusions

The data discussed in this chapter clearly shows that the acquisition of different types of tasks, including those that are traditionally considered to be 'hippocampally dependent' are differentially affected by hippocampal damage. Several broad conclusions about the effect of hippocampal damage on the acquisition of information can be drawn:

Skill learning (e.g., mirror drawing, rotary pursuit) is generally unimpaired by hippocampal lesions. However, conditional visuo-motor learning in which subjects must learn to perform an *arbitrary* motor movement in response to a visual stimuli is significantly slowed after hippocampal lesions. This interesting finding implies that the acquisition of information that is not predictable or similar to what has gone before might be particularly dependent on the hippocampus.

Perceptual priming is normal after hippocampal damage. If anything, hippocampal amnesic patients are more affected by changes in the perceptual features of items than control subjects, suggesting that the hippocampus may usually allow stimuli to be processed in terms of higher order information. *Implicit learning* on tasks such as

word-stem completion, categorisation and sequence learning is generally unimpaired after hippocampal lesions, although higher-order information may not be acquired normally. Implicit tests of complex semantic information sometimes reveals relatively well preserved memory compared to explicit tests. This suggests that part of the problem after hippocampal damage may be in accessing at least partially intact memories, rather than due to a complete destruction of stored components of a memory.

Recognition abilities themselves are not usually affected by discrete and complete hippocampal damage, although partial hippocampal damage can paradoxically lead to greater deficits. It is possible that delayed recognition memory might be poorer after hippocampal damage, but the data is currently equivocal. The hippocampus may normally play a role in the acquisition of recognition memory, even if it is not essential.

Associative learning about stimuli in the same sensory modality is often unaffected by hippocampal damage (e.g., learning odour-odour or visual-visual paired associates). These tasks are often acquired slowly. Tasks that depend on associating information from different modalities show mixed effects after hippocampal damage. The acquisition of cross-modal tasks such as arbitrary visuo-motor learning, learning face-voice or picture-sound associations is often slowed. However, the acquisition of some less 'cognitive' species-specific associative learning tasks (e.g., conditioned taste aversion learning or visual imprinting) is unaffected by hippocampal damage. Hippocampal damage also leads to deficits in associating supra-modal spatial or temporal information with other information. The hippocampus therefore becomes more important for the acquisition of associative information when associations must be made between more diverse types of information.

The hippocampus does not seem to be required for the formation of a *spatial representation* of the environment per se; but navigation or other behaviour based on allocentric spatial representations, or on the arrangement of spatial stimuli in a scene, is totally prevented by hippocampal lesions. The hippocampus therefore appears to be necessary for learning what should be done at particular regions in an environment or scene.

Tasks that are best served by *ego-centric strategies* may be acquired faster after hippocampal lesions, which suggests that the hippocampus may initially suppress learning in that system.

The acquisition of tasks which depend on associating information across a *temporal* gap (e.g., trace fear responses, trace eyeblink conditioning) is completely or very severely impaired by hippocampal damage. Subjects with hippocampal damage also appear unable to remember sequence information, or acquire implicit high-order sequence information even when training is extended. Therefore, the hippocampus may play a major role in representing high-order sequence information, and/or associating supra-modal temporal and sequence information with behaviour.

Normal detailed anterograde *episodic memories* are rarely evident after hippocampal damage. Performance on memory tasks such as remembering a list of words or a story are also severely impaired in amnesic patients. With very extended day-to-day exposure over years or decades, new generic *semantic information* about the world (e.g., new vocabulary, world events) can be acquired in patients with hippocampal damage, although detailed information may be omitted. Explicit training on new semantic information in a laboratory situation can lead *very* slowly and arduously to the acquisition of a little new semantic information, but any information acquired in this way is often very inflexible and quickly lost. Such data suggests that the acquisition of declarative information after hippocampal damage is sensitive to the degree of repetition of information: one-off episodic information cannot be acquired, whereas semantic information can be acquired after very extensive exposure. There is some indication that detailed information is most affected by hippocampal damage.

There is some evidence that the *incidental acquisition* of information (such as that involved in episodic learning or acquiring contextual information) is more affected by hippocampal damage than the acquisition of knowledge about information that is the central focus of attention. Furthermore, in the absence of the hippocampus, information can sometimes only be acquired when stimuli are more prominent than would be required when the hippocampus is intact (for example, one-trial fear-conditioning to an object can be achieved without the hippocampus only if the shock is sufficiently strong). The hippocampus may therefore be more important than other areas for acquiring information that does not appear to be of great importance when it was experienced.

"Untrained behaviour" such as a tendency to orient to novel items or to make inferences about the relationships between stimulus pairs which have not been presented together, may be disrupted after hippocampal damage. There has been relatively lit-

tle exploration of this issue to date, and it is difficult to explore experimentally. One interpretation of these findings is that 'natural behavioural propensities' may be disrupted after hippocampal damage. Alternatively, in the absence of the hippocampus, information might be represented differently and in a way that impedes the discovery of higher order associations within acquired information.

In general, the hippocampi of *different species* appear to perform a largely similar functional role, in that damage to the hippocampus causes deficits on similar types of tasks in different species, when analogous tests have been devised. However, damage focused on the hippocampus has been reported to lead to deficits for primates but not rodents on a few apparently similar tasks (e.g., some recognition and conditional motor learning tasks). This may reflect the relative ease with which intra-regional connections can be made in the smaller-brained rodents compared to primates. Brain size increases as we ascend the phylogenetic scale, with the number of possible connections between neurons increasing exponentially as neuron number increases. Larger brains may be associated with increased cortical specialisation and reduced connectivity between regions (Murre and Sturdy, 1996; Schoenemann, 2001). Therefore if the role of the hippocampus is to allow the formation of associations between information that cannot be easily represented by cortico-cortical connections, we might expect it to be relatively more important for the fast acquisition of complex information in species with larger brains. Of course, many other factors must be considered, too, as different species may acquire apparently equivalent tasks by different mechanisms, and at massively different rates.

Table 3.3 summarises the effects of hippocampal damage on task acquisition.

3.3.1 The nature of hippocampally-dependent tasks

I conclude that hippocampal damage leads to characteristically different effects on the acquisition of different groups of tasks. Learning cannot be demonstrated at all on tasks that use spatial or temporal relational information, although some spatial and temporal information can be acquired. Memory for detailed information is so impoverished that in effect few real episodic memories can be demonstrated. Different types of semantic tasks are impaired in proportion to their dependence on detailed information, or the amount of exposure to information. The acquisition of semantic information can be relatively normal after extended lifetime exposure, but laboratory-based

Acquisition Unaffected	**Slowed or reduced acquisition**
• Learning responses to single cues	• Context learning based on single prominent cues
• Associating stimuli within same modality	• Associating stimuli of different modalities
• Some one-trial cross-modal associations (e.g conditioned taste aversion, visual imprinting)	
• Tasks that are sufficiently aversive (e.g. fear conditioning with strong shock)	• Low salience information
• Priming	• Acquisition of (non-sensori-percep-tual) event & fact details
• Well-cued implicit tasks	• Incidental learning
• Some slowly acquired tasks (e.g. categorisation and grammar-learning)	• Semantic information
• Basic timing of periods	• Behaving in response to temporal information (e.g. trace conditioning, explicit use of sequence info)
• Recognition memory (short delay)	
• Skill learning	
No acquisition	**Improved acquisition**
• Behaving in reponse to allocentric spatial layout	• Egocentric strategies
• Normal detailed episodic memory	• (Reaction speed often faster)
• Some material that is tested explicitly	
• Tasks that rely on 'natural behaviour'	
• Tasks that depend on specific representations	

Table 3.1: Summary of the effects of hippocampal damage on task acquisition.

semantic learning is very much impaired, even when trials are extended over several weeks. The acquisition of information about associations between stimuli from different modalities is often reduced or prevented, whilst associative learning within one modality apparently proceeds as normal after hippocampal damage. The acquisition of low-salience or incidental information is also affected by hippocampal damage.

Tasks that are often considered to be hippocampally-dependent differ fundamentally in their degree of dependence on the hippocampus. Under given conditions some tasks are obligatorily hippocampally-dependent and only the hippocampus *can* form a viable trace that can mediate recall; whereas others are merely facilitated by hippocampal activity as the hippocampus is merely *faster* or *better* at forming traces than other areas. On these latter tasks, changing the training type, length or intensity can affect the degree of deficit seen after hippocampal damage. It seems reasonable therefore to assume that there are both quantitative and qualitative differences in the learning abilities of the hippocampus and extra-hippocampal areas. In the next chapter, I attempt to characterise these important features of hippocampally-dependent learning.

Chapter 4

The nature of hippocampally-dependent learning

In this chapter I explore the features of a task that make it hippocampally-dependent. In general, tasks are most likely to require the hippocampus for acquisition if they rely on the rapid acquisition of information, or if the information to be acquired is highly complex, unfamiliar or of low salience. The presence of any one of these features alone is not sufficient to obligatorily implicate the hippocampus in acquisition. Instead, it is the combination of these factors, taking into account their individual notional positions on a scale of severity, that determines hippocampal dependency. All other things being equal, tasks that are learnt very fast, or involve the acquisition of low salience or high complexity information are likely to require the hippocampus, whereas tasks that score low are not. Tasks with intermediate scores can be at least partially acquired by regions outside the hippocampus under normal conditions, although the hippocampus is likely to mediate their performance in an intact brain. Whether the hippocampus is necessary for normal performance in such cases will depend on the specific conditions of learning and testing. The acquisition of task-dependent supramodal information is however always dependent on the hippocampus; it can be seen as representing the extreme point on a scale of complexity.

4.1 Introduction

The rapidity of hippocampal learning and the highly convergent inputs to the hippocampus are widely considered to be key factors underlying the learning capacities of the hippocampus. However, the data reviewed in the previous chapter shows that neither the speed of learning nor a requirement for complex associative learning alone determines the involvement of the hippocampus. This chapter therefore refines and extends these approaches to understanding hippocampal function.

I start this chapter by outlining what could be called the 'consensus' view of neurobiology relating to memory:

1. The brain consists of relatively independent systems that process different types of information. All (or nearly all) such systems contain plasticity mechanisms which modify operations by experience – i.e. all (most) regions store memories. Multiple such learning systems can be activated simultaneously and in parallel on different learning tasks, each performing different operations with a subset of 'active' data.

2. Perceptual ('lower') areas are arranged by modality and are relatively self-contained. Increasingly diverse types of information converge at progressively higher levels of the brain, through uni-modal, multi-modal, and associational cortices to the medial temporal areas.

3. Some kind of Hebbian learning scheme is generally assumed, whereby the strength of connections between neurones, and internal mechanisms within neurones, can change progressively as a function of their coincident activity. Such changes represent memory.

4. The type of memory stored in an area is likely to be the same as that which is processed there.

Interestingly, even these very simple and widely accepted ideas about memory systems often go against the grain of popular consolidation views about hippocampal function, in which the hippocampus is the 'gateway' for storage of information outwith the hippocampus, and information is initially stored only in the hippocampus.

In terms of the hippocampus, it is widely accepted that:

1. The hippocampus is one of the areas of the brain that receives information from the widest range of sources.

2. The hippocampus is capable of more rapid learning than other brain regions.

3. The structure of the hippocampus means that it is more able than other areas to represent conjunctions between different types of information: hippocampal complex-spike cells each receive inputs from a wide range of sources, and are massively inter-connected.

4. Pattern separation is likely to occur most strongly in the hippocampus, so that traces for similar events would be expected to overlap less there than in other areas which are arranged more topographically.

Taken together, if we assume that the hippocampus is not fundamentally different to other areas (although it may differ quantitatively in speed of learning or degree of convergence, say), this implies that a basic framework for understanding the relationship between memory in the hippocampus and wider brain areas might be as follows. Presentation of similar events causes largely overlapping activity in areas outwith the hippocampus, and less overlapped activity in the hippocampus. Fast and strong encoding in the hippocampus allows the robust storage of traces of individual events, which can support the recall of individual traces when cued appropriately, even after only one or a few encounters. In the cortex, weaker event encoding leads to less robust traces for each event, which are likely to be too weak to successfully mediate recall. However, the parts of a trace that overlap in cortical areas on re-experiencing similar events or items will slowly potentiate, so that in effect a robust representation of a generic item or event, shorn of its "episodic" (specific event) information might eventually become available for recall.

Such a scheme could, for example, explain why the hippocampus might be required for the rapid storage of complex episodic memories, and how the hippocampus could support "episodized" semantic recall after damage to semantic areas. Similarly, it would explain how semantic information could get extracted from repeated exposures to a learning event, why specific event recall tends to become semanticised over time, and how semantic learning can proceed in the absence of the hippocampus. However, this basic framework cannot adequately capture data the complex findings discussed in the previous chapter.

The hippocampus' capacity for rapid learning does not really tell us what tasks should be hippocampally-dependent, although it would imply that tasks that need to be learnt quickly are more likely to be hippocampally-dependent than those that are not. In fact, several tasks that can be acquired in one trial are *not* hippocampally-dependent, so a capacity for fast learning cannot be the only feature that distinguishes hippocampal learning from that of other areas. Conversely, some slowly acquired tasks (e.g., transverse patterning and eye-blink conditioning) generally cannot be acquired after hippocampal lesions.

Other theories have emphasised the role of the hippocampus as a convergence zone for information from many brain regions (e.g., Teyler and Discenna (1986); Damasio (1989b); Milner (1989); Alvarez and Squire (1994); Murre (1996)). In such theories, the hippocampus is usually required to form traces that bind incoming cross-modal information together, when this cannot be stored in a single neocortical module. These approaches have been criticised for being under-constrained and merely postponing a characterisation of the role of the hippocampus until the nature of stimulus processing in the cortical modules is better understood (O'Keefe, 1999). Certainly the hypothesis that the hippocampus is required for the formation of associations per se does not explain why some apparently cross-modal tasks (e.g. rotary pursuit, conditioned taste-aversion) are not impaired after hippocampal lesions, nor would it predict that, say, incidental (but not contingent) learning and trace (but not delay) conditioning would be impaired. From an anatomical point of view too it would be surprising if the hippocampus was the only place capable of cross-modal associative learning: basic sensory information from different modalities appears to be fully integrated by the tertiary association cortices, and certainly by the entorhinal cortex. Thus general associative theories of hippocampal function are also in need of refinement.

In the following sections I discuss the factors inherent to tasks that determine how important the hippocampus is likely to be for the acquisition of that task. I conclude that a coherent explanation of why particular tasks are hippocampally-dependent requires a consideration of the combined effect of several factors, including the speed of learning and the complexity of the information to be acquired; but also the salience and novelty of the information, and its dependence on the default representational characteristics of the hippocampus.

4.2 Speed of learning

There is much evidence to support the idea that the hippocampus is particularly important for the rapid acquisition of information. On many tasks the primary deficit of hippocampally-damaged animals is slowed acquisition, so that if animals are tested after the number of trials at which control subjects reach asymptote performance, they are impaired. Such tasks include arbitrary visuo-motor mappings (e.g., Murray and Wise (1996)), some configural tasks (e.g., Whishaw and Tomie (1991); McDonald *et al.* (1997)), learning face-voice associations (e.g., Vargha-Khadem *et al.* (1997)) and recognition memory tasks (e.g., Alvarez *et al.* (1995); Beason-Held *et al.* (1999)). Inactivation of the hippocampus also impairs learning on massed more than spaced trials (Poucet *et al.*, 1991). Episodic learning (which by definition involves fast one-trial acquisition) is also impaired after hippocampal lesions (e.g., Scoville and Milner (1957); Clayton and Dickinson (1998)). Lesioned animals perform normally on category learning (whilst being significantly impaired at remembering exemplars, e.g., Knowlton and Squire (1993)), and semantic information can be acquired slowly over a lifetime by humans amnesics with severely disrupted episodic learning (e.g., Kitchener *et al.* (1998); Vargha-Khadem *et al.* (1997); Holdstock *et al.* (2002)) – such learning must therefore depend on the repetition inherent to categorical and semantic information, that allows slow learning in the absence of the hippocampus.

However, not all rapid learning is impaired after hippocampal lesions, so the number of training trials available cannot be the sole determinant of hippocampal dependency. For example, conditioned taste aversion (in rats, Yamamoto (1993)) and visual imprinting (in chicks, Horn (1998)) do not require the hippocampus for acquisition even with only one or a few trials. These abilities are present from birth, which is good evidence that the hippocampus is not the only area that can perform one-trial learning, since the hippocampus may not be operational at birth (at least in rats, Waters *et al.* (1997)). However, these tasks could be said to be 'less cognitive' than most of the tasks examined in the hippocampal literature, and perhaps tap species-specific, relatively hard-wired abilities. It seems plausible that different species have evolved designated circuitry that bypasses the hippocampus to perform certain valuable learning tasks such as learning about poisonous foods. This suggests that the hippocampus need not be required for the rapid acquisition of information that an animal is hard-wired to acquire.

Clearly, one-trial acquisition of information can only occur when pre-existing neural circuitry is very similar to that that is needed to represent current information, so that existing neural connections can be enhanced without the need for time-consuming large-scale 're-wiring'. Therefore, the capacity for learning on any task must depend on the number of trials available *and* on how similar existing representational structure is to what is required to represent the current task. An innately-biased "hardwired" learning capacity would result from gross neural patterns that bring representations of certain information (e.g., memory for recently tasted food, and feelings of nausea) together in a way that enhances an ability to learn about associations between them.

When associations are to be made between information that is represented in geographically close areas of the cortex, it is more likely that relatively small changes in existing circuitry that can be effected rapidly will be able to capture the associations that need to be stored. Therefore the hippocampus will not necessarily be needed for the rapid storage of information that is represented within geographically local areas outside the hippocampus, or within areas that are already heavily interconnected (perhaps as a result of past learning experiences). The rapid acquisition of simple tasks that depend on learning associations between information from the same modality therefore need not depend on the hippocampus if the local circuitry is appropriate.

Performance on a given learning problem would be aided if previous experience with similar items and problems over the lifetime of the animal had led to the development of a set of representations (a 'learning set'[1]) outside the hippocampus that are similar to what is needed to represent or 'solve' the task at hand. This would allow rapid acquisition of new similar information. This implies that fast learning of some types of information should sometimes be possible in the absence of the hippocampus after extensive experience with similar tasks, even on tasks whose initial acquisition might normally depend on the hippocampus. In accord with this, hippocampectomized animals can perform one-trial odour-odour learning after extensive experience with similar tasks (e.g., Eichenbaum *et al.* (1986); Reid and Morris (1992)). Similarly, pre-training in the presence of an intact hippocampus can sometimes protect animals from the anterograde learning deficits that would be expected to occur following interference with the hippocampus (Bannerman *et al.*, 1995; Saucier

[1]This definition is more general than the sense used in Reid and Morris (1992), for example, in which a "learning set" refers to abstract strategies developed to solve certain tasks.

and Cain, 1995; Roesler *et al.*, 1998; Hoh *et al.*, 1999; Moser and Moser, 2000; Takehara *et al.*, 2003). However, such pre-training would be unable to counteract the effects of subsequent hippocampal damage on tasks which have a high ongoing demand for the fast acquisition of new complex information as it would impossible to pre-learn enough information and to adequately represent it in the absence of the hippocampus. Therefore, pre-training does not prevent deficits on spatial learning on a water-maze with delayed testing after 2-h, when the platform moves from day to day (Steele and Morris, 1999).

In humans it is particularly clear that we "learn how to learn". Throughout childhood and beyond we learn, for example, how the world is structured into hierarchical categories that share features, and the types of attributes that are likely to be constant across items (e.g., that hairy animals that bark usually have four legs): this underpins our adult learning abilities. It is only once we have learnt what kinds of semantic information are likely to generalise (e.g., if we are told that the capital of France is Paris today, that it is likely to be 'Paris' the next time we ask) that generic semantic information can be acquired in one exposure. Prior to the establishment of such learning sets, novel information must be repeatedly experienced in order to establish which *are* the generic semantic components, and which are the episodic idiosyncrasies of that particular presentation of an event.

Some types of tasks consistently involve similar mappings between certain types of information (e.g. standard visuo-motor tracking tasks involve "congruent" associations between visual inputs and motor outputs to allow tracking). On the grounds that the speed and efficacy of responses would be best served by the shortest connections between relevant perceptual and motor areas, it is to be expected that such information processing would bypass the hippocampus, as indeed seems the case, at least in the adult once basic tracking abilities are established. It is therefore of great interest to find that when *arbitrary* visuo-motor mappings must be acquired (e.g., see a yellow card – wave your right paw) the speed of learning is hugely impaired by hippocampal lesions, and performance is close to chance after a number of trials that results in good asymptote performance in controls (Murray and Wise (1996)). Since any mapping (whether it turns out to be commonly occurring or not) is arbitrary the first time it is encountered, this provides indirect support for the idea that the hippocampus is involved when learning sets elsewhere are inadequate and information

must be acquired fast. The absence of a learning set for a task is effectively equivalent to that task requiring "arbitrary" mappings with respect to a given animal's learning history.

The hippocampus' capacity for rapid acquisition of information derives not only from the fact that hippocampal synapses potentiate and depotentiate faster than those in other areas (Ivanco and Racine, 2000), but also from the way information converges onto individual complex-spike cells. Information from many sources converges onto each complex-spike cell so that novel associations can often be registered by single or already-interconnected neurons. In effect, the hippocampus is built with the potential to quickly learn about relations between unpredictably-related, and possibly complex, stimuli. Of course, not *every* possible association will be easily acquired by an existing hippocampal structure, and some tasks typically take many hundreds of trials to acquire even with an intact hippocampus – perhaps in order to extract probabilistic information or complicated mappings or rules, or to represent certain novel associations sufficiently strongly. Clearly, the relative facility of the hippocampus for rapid acquisition alone cannot adequately explain why *these* tasks are hippocampally-dependent.

The hypothesis that the hippocampus is required for fast learning when there is no established learning set implies that learning early on in an animal's life would be more likely to be hippocampally-dependent. This implication is unlikely be true in its baldest form, since the hippocampus is probably not operational in young animals (Waters *et al.*, 1997). However, since childhood amnesia for specific events is well documented, it could be argued that fast learning of the type that depends on the hippocampus in adulthood simply does not occur in the young animal. On the other hand, periods of high plasticity are well-documented in many regions of the young nervous system (Kirkwood *et al.* (1995); Berardi *et al.* (2000)) so that areas other than the hippocampus are capable of relatively fast learning in young animals (i.e. in the period before the hippocampus comes "online"). This would allow learning sets for commonly encountered task mappings and stimuli to be constructively set up in young animals. Perhaps part of the hippocampus's uniqueness is that it loses its capacity for plasticity with age less precipitously than other brain areas, and comes online only when 'pre-learning' has been completed in other regions.

In adulthood, neurogenesis has been reported only in the hippocampus (in rodents, Altman and Das (1965), and humans, Eriksson *et al.* (1998)) and olfactory cortex

(Bayer, 1986). There is speculation that new hippocampal memories might be stored in new neurons. Such a mechanism might underpin quantitative differences in the learning abilities of the hippocampus and other areas. For example, the high plasticity of new neurons might allow relatively fast acquisition of new information, and allow the associations between previously unrelated information to be represented easily. This may also lead to fundamental qualitative differences in the nature of the information that *can* be represented in different areas, after plausible opportunities for learning.

In summary, the blunt assertion that the hippocampus is required when information must be acquired quickly is untenable. Several qualifications are required. Firstly, the hippocampus is only necessary for the fast learning of certain kinds of quite complex information – simple uni-modal information and some kinds of low-order multi-modal information will generally be relatively easy to represent in regions other than the hippocampus. Secondly, whether the hippocampus is required for an animal to learn a task rapidly in a few trials depends on the animal's learning history, as representations for previously learnt information can facilitate the acquisition of new related information. In other words, the speed with which information can be acquired, the nature of the information to be acquired, and existing local convergence of information cannot really be separated.

4.3 Convergence of information

We have established that the hippocampus is likely to be involved in the fast storage of many types of information, especially as complexity increases. However, given unlimited training trials, the hippocampus remains obligatory for the acquisition of tasks that depend on associating supra-modal information with other information, such as navigating on the basis of allocentric information and using timing information to control behaviour. The acquisition of episodic and autobiographical information is also massively impaired by hippocampal damage. Of course since such information cannot by definition be repeated, we do know whether such information could be retained after repeated exposure. Typical episodic recall may depend on supra-modal information, although deficits are observed even when allocentric and temporal information are not essential to recall. More generally, episodic information requires the association of many pieces of unpredictably-related complex information.

It is possible that some sources of information do not converge prior to the hip-

pocampus in a way that allows the associations that are necessary to the performance of certain tasks to be discovered or represented. For example, it is widely believed that in the visual information pathway, there are two segregated streams from the primary visual cortex (as first suggested by Mishkin *et al.* (1983)): a dorsal "what" pathway concerned with object information and a ventral "where" pathway concerned with location. More extreme views have come increasingly under attack as evidence accumulates for intermixing between pathways (e.g., Goodale and Milner (1992); Merigan and Maunsell (1993)) and the theoretical "proofs" of the separation are brought into doubt (Goodhill *et al.* (1995)). However, object information from the ventral stream is largely sent to the EC via the PrC, which projects primarily to the anterior and lateral portions of the EC; whilst visuo-spatial information from the dorsal stream projects largely to the posterior portions of the EC, via the PHG (Suzuki and Amaral, 1994). Witter *et al.* (2000) has recently suggested that even within the hippocampus there may be largely parallel paths originating in the PrC and post-rhinal cortices, although communication between these streams seems likely. If the segregation is maintained outwith the hippocampus, then this would explain why hippocampally-lesioned animals fail to acquire various tasks that require the association of 'what' and 'where' with behaviours ('what to do'), even with unlimited training trials.

'What should be done' in specific spatial locations or at particular times obviously varies unpredictably from task to task and occasion to occasion. Therefore it would not be possible to build up a general learning set about how spatial or temporal information relates to action. If the hippocampus is required for learning when there is no established learning set, this would explain deficits on such tasks. Severe impairments at learning new vocabulary after hippocampal lesions could be similarly explained, because it is not possible to build up general representations that link phonemes with meaning. It is also possible that the difficulties faced by patients with hippocampal damage on explicit recall depend on severing the links between high-level representations of information (that would tie together the aspects of an episodic event, for example) and language systems (that may be represented in the left hippocampus). Abstracting and generalising from specific instances (which is also often impaired in hippocampal patients) may also depend on the convergence of linguistic and high-level representations of information.

The MTL is organised as a hierarchy of associational networks (Lavenex and Ama-

ral (2000)), with extensive reciprocal inter-connections between areas, and extensive intra-level connectivity. Neuroanatomically, the hippocampus is considered to be at the top of this hierarchy (as shown in figure 1.2), with information integration and complexity increasing up the hierarchy. In fact, the whole brain can be seen as a hierarchy of inter-connected 'convergence zones' (Damasio, 1989b), with progressively more and more diverse information converging onto regions as one progresses from the perceptual 'lower' areas through the uni-modal, multi-modal and associational cortices to the 'higher' medial temporal regions. Ivanco and Racine (2000) found the highest rates of LTP induction and decay in intra-hippocampal pathways and the lowest in the cortex, with intermediate rates in connections between the hippocampus and perirhinal and pre-frontal cortices. A learning hierarchy can be envisaged in the neocortical-hippocampal axis, with increasingly fast learning rates and increasingly convergent inputs as one progresses from the cortex through the para-hippocampal areas to the hippocampus. Thus the hippocampus would be expected to be faster in the acquisition of some associative information than other lower regions, if those regions have not already developed representations of similar material through previous learning experiences. This might explain the slowed acquisition of cross-modal associative tasks that do not depend on supra-modal information (e.g., learning face-voice pairs) after hippocampal damage. Similarly, since complex semantic information such as new vocabulary or memory for public events is unlikely to be sufficiently similar to any information currently stored, the hippocampus would be implicated in normal acquisition at the usual learning rate. However, because semantic information by its nature is usually repeated, albeit embedded in different one-off 'episodic' events, areas outside the hippocampus may be able to form traces of this generic information over many exposures.

As one proceeds up the hierarchy it becomes increasingly easy to bind complex, unique and diverse traces of details of events into one memory. This might explain why the hippocampus and MTL regions are particularly important for the recall of episodic and complex semantic information. However, whether information is considered to be a specific 'detail' or is considered to be generic and 'core' depends, of course, on the consistency with it occurs with particular other information and whether it is a necessary component. This in turn often correlates with the number of times that the information has been encountered. There can be no absolute definition of what is

'core' to an event or memory, as it must be defined with respect to instances of that event, and/or to knowledge and expectations of an individual[2]. If something truly novel is experienced only once, it is impossible to know which are the generic core 'semantic' features and which features were merely incidental to that occurrence of the event. Since a memory is usually operationally-defined as being episodic on the basis of the type and number of details that can be recalled, the recall of episodic memory will be most affected by the loss of detailed information. The data reviewed here suggests that the recall of such information is mediated disproportionately by the hippocampus. This is not the same as saying that the hippocampus is 'configured' for processing episodic information. The way episodic and semantic memory is defined cuts across several confounding factors such as the degree of detail needed at recall, and the amount of exposure to information (see section 4.6.2).

In this hierarchical view, hippocampally-based memories will be formed on the basis of any inputs received, even if there is damage to some of the lower regions; whilst the recall of existing hippocampally-dependent memories would be affected to the extent they depended on fragments in now-damaged regions. The dissociation of episodic and semantic function seen in semantic dementia (Graham *et al.*, 1999) and of anterograde and retrograde autobiographical memory in visual memory-deficit amnesia (Rubin and Greenberg, 1998) would therefore be predicted.

In summary, the hippocampus is likely to be involved in the acquisition of complex information, especially when learning involves task-dependent supra-modal spatial or temporal information. The acquisition of complex high-order associative information will usually involve the hippocampus in the intact brain, but complex information (apart from that depending on associations with supra-modal information) may sometimes be acquired in the absence of the hippocampus over many training trials. These trials may be spread over the lifetime of an animal, so that acquisition may occur on a given occasion after apparently only a few trials.

[2]For example, if someone is attacked by a person carrying a baseball bat and wearing a green shirt, the colour of the shirt is likely to be considered irrelevant and unimportant and may well not be reliably encoded and recalled over time, in contrast to memory for the baseball bat. However, if someone is attacked several times by different people carrying different implements, but all wearing green shirts, the core features of such as event are likely to be reconsidered and the shirt colour given more prominence. If however, the attacked person had noticed on the first encounter that the green shirt was actually a military uniform, then they might have tentatively 'upgraded' this feature to a core feature sooner.

4.4 Incidental and automatic learning

Hippocampectomized animals have sometimes been reported to show deficits in incidental learning such as knowledge about the context on a stimulus-response task. The acquisition of information about low salience, discrete stimuli is also affected by hippocampal damage. In many ways, day-to-day episodic memory, which is severely affected by hippocampal damage, can be considered a paradigmatic example of incidental learning. On a test of autobiographical memory, subjects may recall incidental details of an event such as the colour of the shirt someone was wearing, or who got on the bus first, but it is unlikely that there was a deliberate intention to memorise these incidental details. Semantic information is more mixed in this respect: some might be acquired incidentally, say whilst passively listening to the radio, whilst other semantic information may be deliberately acquired through observation or deliberate research.

More generally, hippocampal activation is seen on most if not all tasks, including those that are *not* hippocampally-dependent (such as the random foraging task, O'Keefe and Dostrovsky (1971); Ranck (1973); O'Keefe and Nadel (1978); delay eyeblink conditioning, Berger *et al.* (1983); McEchron and Disterhoft (1997); when watching a film (Maguire *et al.* (1996, 1998a)) or sitting at rest with eyes closed, Martin (1999); Binder *et al.* (1999)); and tasks whose performance *improves* after hippocampal lesions (Eichenbaum *et al.* (1987); McNaughton *et al.* (1989)). This suggests that the hippocampus may be automatically engaged on all tasks, and may in some cases be acquiring information that is not necessary to the experimenter-defined task at hand. One suggestion is that the hippocampus automatically encodes all attended experience as 'snapshots' of experience (Morris and Frey (1997)). However, Poldrack *et al.* (2001) have shown that the hippocampus must be active on a categorisation learning task at a time when hippocampal damage causes no obvious deficit if learning is to proceed normally in the later stages of task acquisition. This suggests that in some cases hippocampal activation might reflect the acquisition of higher-order, more 'processed' information about experienced events, rather than, or in addition to, merely accumulating snapshots of experienced events.

Putative continuous non-selective storage of ongoing information would raise questions about capacity. However, the value of information is often not known when that event is experienced, so it is desirable to encode unfolding information as it occurs in case later events show its importance. A plausible solution is to *retain* information

on the basis of its apparent value: typically, only the most salient or important information is remembered well.

If the hippocampus was particularly important for the automatic storage of information (such as autobiographical events), it would be expected to be particularly sensitive to the modulation of memory strength and maintenance – as indeed appears to be the case. The long-term maintenance of memories can be affected by processes acting at acquisition and/or in a 'critical period' after acquisition. These modulatory processes reflect the importance attached by the animal to the information it has experienced. In general, the degree of hippocampal engagement may be modulated by gross factors such as global brain state (e.g., theta versus gamma electrical activity, sleep versus waking) as well as arousal, attention and motivation to learn; these factors may also influence initial trace storage strength. Furthermore, the hippocampus has substantial amounts of stress-related glucocorticoid receptors and expresses both types of receptors unlike most other regions (de Kloet *et al.* (1999)); many of the neuromodulatory substances that are released after stress (e.g., β-endorphin, vasopressin, adrenocorticotrophic hormone, substance P and cholecystokinin, McEwen (1999)) are known to act on the hippocampus; and it receives direct projections from the amygdala and the medial septum which are implicated in processing anxiety and alertness. All of these systems (and others) are thought to affect memory storage. In fact, high levels of stress have been shown to selectively impair hippocampally-dependent learning (such as spatial versus cued learning, Ohl and Fuchs (1999), and spatial memory versus navigation to a visible goal on a radial arm water-maze, Diamond *et al.* (1999)).

Making incidental stimuli more salient would aid cortical learning through modulatory processes (and hippocampal learning if it were intact) and potentially ameliorate the learning deficits produced by hippocampal lesions on some tasks. Indeed, the presence of reinforcers has been shown to produce learning in systems other than the hippocampus (White and Wallet (2000)). Higher foot-shocks on hippocampally-dependent tasks such as inhibitory avoidance can protect against the usually disruptive effects of interference with synaptic activity in structures such as the amygdala, neostriatum and thalamus (see references in Cobos-Zapiain *et al.* (1996)), probably via the effect of norepinephrine released in response to stressful stimuli (Seidenbecher *et al.* (1997)). Interestingly, whilst fear conditioning to a sufficiently aversive stimuli

can be achieved after experience with only one or a few trials, fear conditioning to less aversive stimuli takes several trials (Izquierdo and Medina, 1997). Thus we would expect to see deficits on some incidental tasks after hippocampal damage, but not necessarily find deficits in acquisition when the same information was on the foreground (in the same way we might predict deficits on fast but not slow acquisition of information). I would also predict that after protracted exposure, the cortex would develop viable traces of incidental stimuli through slow incremental learning processes.

A qualitatively similar pattern of activation in brain areas including the hippocampus is found on tasks incorporating deliberate instructions to memorise stimuli for later recall or recognition (e.g., Maguire *et al.* (1998b); Tulving *et al.* (1994); Nyberg *et al.* (1995); Eustache *et al.* (1995)); and those in which any learning is incidental and may be evidenced by, for example, an unexpected recall test (e.g., Price *et al.* (1994); Bookheimer *et al.* (1995); Martin *et al.* (1996); Zelkowicz *et al.* (1998); Sergent *et al.* (1992)). However, intentional learning produces significantly stronger activation than incidental acquisition, which implies that activity for typical incidental learning would be low. Similarly, salient stimuli produce greater cortical and hippocampal activation than low salience stimuli (Wiggs and Martin (1998)). All other things being equal, lower activity levels would be expected to lead to weaker trace storage. Therefore, the fast acquisition and incidental acquisition of information are similar in the sense that both would be expected to produce relatively weak traces especially in the more slowly learning cortex. The more robust learning of the hippocampus would produce relatively strong trace even in response to low levels of activity.

Deficits have sometimes been reported on tasks that depend on learning associations between interoceptive information (e.g., hunger state) and external stimuli, although again, the data is inconsistent. It seems reasonable to assume that since external stimuli are more likely to be causative in a standard learning situation, that interoceptive cues are normally consigned to the 'attentional background' on most tasks. Indeed, learning about interoceptive cues takes longer than learning about external stimuli even in intact animals (Davidson and Jarrard, 1993; Hock and Bunsey, 1998). Therefore tasks that require attention to internal states might be supported by the hippocampus partly because it can compensate for low activity levels by relatively robust trace storage. As already discussed, the hippocampus provides a site where novel, unpredictable associations can be easily registered; this may be important in learning

about unusual relationships between endogenous stimuli and external events.

Animals with hippocampal lesions perform at chance on tasks such as delayed visual-paired comparison (VPC) that measures looking-preferences towards novel stimuli (e.g., Zola *et al.* (2000); McKee and Squire (1993)), and orienting to changes in spatial or non-spatial arrangements of stimuli (e.g. Honey *et al.* (1998)). Although incidental learning is usually defined with respect to a motivated task, these tasks can be seen as dependent on incidental learning in that the *responses* are not explicitly trained. In these tests the animal does not "know" what the experimenter wants from it, thus any behaviour reflects natural behavioural propensities. This is also true of the more conventional "incidental" studies already discussed. Therefore the hippocampus may be involved in mediating automatic motor responses to stimuli such as orienting to novelty. Behavioural alterations after hippocampal lesions on tasks that assess untrained responses to stimuli may reflect changes in the control of responses, rather than an inability to make the distinctions on which to base responses. Indirect support for this idea comes from data showing that hippocampal damage leads to changes in various non-cognitive behaviours (as outlined in section 3.2.4.3). To my knowledge there is no data on whether rewarding a hippocampal animal for "behaving normally" (e.g. for looking at the novel object on the VPC task) can ameliorate the so-called "deficit", so this issue remains unresolved. Animals with compromised hippocampi are also impaired at certain inferential tasks (Dusek and Eichenbaum (1997); Bunsey and Eichenbaum (1996)) which can similarly be construed as involving unrewarded incidental learning.

In summary, the hippocampus appears to be important for the automatic acquisition of information. However, all areas that are involved in information processing are also likely to be continually registering some trace of ongoing activity. The hippocampus's importance is especially evident for the acquisition of low salience information, which probably reflects the relative ease with which robust changes can be made in the hippocampus compared to other areas. Tasks such as conditioned taste-aversion and contextual conditioning are not demonstrably affected by hippocampal damage, although they depend on the initial unreinforced acquisition of information. Therefore factors other than the incidental nature of the information presented must also be important, such as the complexity of information required for performance, or testing protocols. The hippocampus is presumably necessary for the storage, incidental

or otherwise, of associations involving supra-modal information and complex rapidly acquired information, such as that underlying autobiographical information.

4.5 Representational structure

Much of the preceding discussion refers to issues that cannot really be separated from the nature of representation employed by the hippocampus. The hippocampus's crucial role in the acquisition of complex information results largely from the numerous, disparate, non-topographical inputs to each hippocampal complex-spike cell, their massive interconnectivity and high synaptic plasticity, and the fact that the hippocampus is the 'top-node' in a series of associative indexing regions – such factors inherently affect the representations that can be employed by the hippocampus.

If it is assumed that associative conditioning requires a temporal overlap of neural activities representing the US and the CS, then it is possible that the hippocampus is required for tasks involving the association of elements across delays because it provides a mechanism for bridging these temporal gaps. There is some evidence for this view although it is not conclusive, and several authors have recently proposed variations on this theme (e.g., Levy (1996); Wallenstein *et al.* (1998); Lisman (1999); Huerta *et al.* (2000)). The hippocampus might play a role in, say, maintaining reverberatory activity in the earlier of the to-be-associated traces. Alternatively, if the hippocampus represents aspects of an unchanging physical environment, then it could provide an indirect associative bridge between information represented at different times in the same environment.

The hippocampus also appears to play an important role in 'extracting' high-order information from repeated presentations of similar information, (for example, on sequence learning, Curran (1997) or probabilistic weather prediction tasks, Poldrack *et al.* (2001)). This may depend on the hippocampus's ability to 'oversee' learning in other areas, by receiving information about activity in several regions that do not communicate directly. Fast episodic-type encoding and a role for the hippocampus in bridging temporal gaps or learning about complex associations that are not immediately evident from observed data, may be incompatible at the cellular level. The latter tasks require that representations for similar re-presented data are not orthogonalised. However, there is a direct path from entorhinal cortex to CA1 that bypasses the orthogonalising dentate gyrus, and it is interesting to speculate that this may allow

a division of labour in the hippocampus.

In humans, the ability to re-represent incoming information in terms of a symbol system which has its own set of manipulation rules (i.e. language) might also allow us to go 'beyond' straightforward statistical relations between stimuli. In accord with this, the left hippocampus has been implicated in verbal processing (Papanicolaou *et al.*, 2002), and analogical reasoning activates the hippocampus (Luo *et al.*, 2003).

Recall of a memory depends on providing cues that can trigger recall of parts of a memory representation, that can then trigger recall of the complete attractor representing other aspects of that memory. Therefore the ability to recall specific information depends fundamentally on the cues employed, which interacts with the nature and robustness of the representations employed. One relatively consistent finding is that after hippocampal damage, implicitly tested knowledge tends to be better preserved than explicit on complex recall tasks. In many cases, implicit tests involve cues that are more similar to those that were present at acquisition. Such cues can therefore act relatively directly to co-activate many of the areas that were initially engaged in learning an event, and thereby support the reactivating of attractors. In effect, sufficiently complex cues that are similar to those originally experienced at learning can act like high-level indexing traces (such as those in the hippocampus) that are supposed to co-ordinate the co-activation of fragments in other lower-level areas. This would also explain why explicit tests that use more cues that are similar to those available at acquisition (such as photographs) produce better performance than those that depend on, say, verbal descriptions or requests for information. It is the nature of the cues per se, rather than whether a task is explicit or implicit, that is crucial.

Complex implicit cues would be particularly expected to aid recall in the following situations: 1) after hippocampal damage for traces whose recall would normally be mediated by the hippocampus - since the hippocampus is supposed to be the top-level indexing region; 2) on the recall of older memories in neurologically normal people where memories have started to fade – since traces appear to decay in a top-down fashion; and 3) in patients where semantic 'fragments', or the inter-connections between them, are breaking down – since such damage would reduce the ease with which existing attractor networks across the brain could be activated by activity in particular other semantic regions. These expectations are borne out by the data.

4.6 Conclusions

A unifying theme underlying this chapter is the *relative* amount of learning that can be expected in the hippocampus and the cortex on a given task. On some tasks under particular learning conditions, only the hippocampus is capable of mediating the acquisition of information that can support performance on a particular behavioural task; on others, extra-hippocampal regions may be able to mediate acquisition to some extent. In general, the hippocampus can learn faster and more robustly for a given event than the cortex, both in terms of the strength of changes to pre-existing connections, and the formation of new required associations. Hippocampal and extra-hippocampal areas are also different in terms of 'wiring' and default representational structure, and patterns of afferent information. Thus the hippocampus is usually more able than the cortex to learn tasks that need to be acquired quickly or depend on learning about low-salience information; that depend on associations between information that converges on the hippocampus; or that depends on using a representational scheme that is the hippocampal default (where this is different to that of the cortex). Learning is also more likely to be hippocampally-dependent when it is arbitrary with respect to the animal's learning history, since the lack of an established learning set will place emphasis on the hippocampal capacities for fast acquisition of new associations.

The hippocampus is *essential* for the acquisition of: 1) information that involves associations between high-level abstract supra-modal information and other information; 2) information that is learnt rapidly and is highly complex; 3) information that has been acquired automatically when the information is of low salience and/or highly complex; and 4) tasks that depend on the default nature of learning or representation in the hippocampus, or on natural behaviour mediated by the hippocampus. Items 2 and 3 refer to combinations of continuous factors such as salience, complexity and speed of learning. These factors clearly affect learnability, both by the hippocampus and by other areas. A particular instance of such a factor on a learning task could described by its position on a notional scale that describes how 'extreme' that instance is (e.g., from 'very high' to 'very low' salience or complexity). It is the *combination* of such factors, taking into account their positions on a continuum, that determines the extent to which the hippocampus is better than other areas for the acquisition of particular information, and if it is better at all. For example, the high speed acquisition of low complexity uni-modal associative information does not depend on the hippocam-

pus, whereas the high speed acquisition of highly complex cross-modal information does (item 2). Associating supra-modal allocentric information with other information (item 1) always acquires the hippocampus and could perhaps be considered to be at the extreme end of a 'complexity' scale. In addition, a few tasks may crucially depend on the inherent representational properties of the hippocampus, such as its ability to make associations between temporally separated information (item 4).

On some tasks such as learning cross-modal associations, the hippocampus and closely-related medial temporal regions and tertiary associational cortices may acquire information that is different only in degree — a stronger, more easily accessible, and more quickly acquired trace is likely to be set up in the hippocampus in response to a given event. Since weaker traces are likely to need stronger cuing for successful recall, this provides an explanation for the relative flexibility of recall in intact animals: Without a hippocampus, cues that are more similar to those experienced at acquisition (i.e., more implicit in nature) are likely to be required for successful recall since they provide stronger cuing. Implicit cues may also automatically allow for the disambiguation of associations and thereby negate a requirement for hippocampal-based incidental learning of context cues.

Tasks which do not obviously require processing of the types of information outlined above, but depend on it to differentiate pre-existing associations from those newly learnt will also obviously be affected.

In summary, the need to acquire information that is complex and multi-modal, that is incidental or of low salience, that depends on identifying high-order associations, or that must be acquired quickly increases the probability that the hippocampus is involved in its acquisition. When several of these factors are *combined*, or specifically when the complexity of the information to be acquired is very high, then it is very likely that the hippocampus is involved in the normal acquisition of such information. Areas outside the hippocampus may be able to acquire some such information with sufficient trials or with different training methods, especially if the information to be learnt is made more salient, is relatively simple or is similar to what has already been acquired. This approach therefore provides a way of unifying various aspects of existing theories that are often treated as competing theories. These proposals both extend and constrain the predictions of earlier associative memory formulations.

4.6.1 Relationship to other proposals

Many theories have proposed that the hippocampus is necessary for the acquisition of a particular 'type' of information, such as allocentric spatial and/or episodic or declarative information (see Chapter 2). However, the discussions of this chapter and of chapter 3 show that the hippocampus is not limited to processing a particular type of information. I have argued that these 'types' often merely represent the extreme end of a continuum of factors that determine whether the hippocampus acquires a trace that is significantly better in supporting the performance of a given task than other areas.

Existing general associative theories of the hippocampus would not predict the specific pattern of findings seen in the previous chapter. For example, recent reworkings of the associative/relational theory which borrow heavily from sequence learning ideas (e.g., Wallenstein *et al.* (1998); Kesner (1998); Lisman (1999)) do not specifically predict deficits after hippocampal lesions on quickly acquired or incidental tasks, or on the processing of interoceptive stimuli. Wallenstein *et al.* (1998) also acknowledges that it is difficult to accommodate the persistence of conditioned taste-aversion learning after hippocampal lesions in their proposals. The Arbitrary Visuomotor theory (Wise and Murray (1999, 2000)) can potentially capture much of the data showing slowed learning in animals with hippocampal lesions, but cannot account for deficits that remain in hippocampal animals after unlimited opportunities for training. However, Brasted *et al.* (2003) has gone some way towards addressing these concerns. Together with Holdstock *et al.* (2002)'s ideas, progress has recently been made in more tightly delineating what the hippocampus is 'for'.

Holdstock *et al.* (2002) have recently put forward the view that "it is the extent to which information is repeatedly experienced, rather than the kind of information that may be the crucial determiner of [...] new learning following selective hippocampal damage". Whilst I agree that the amount of exposure to material is extremely important, Holdstock *et al.* (2002)'s wholesale switch from a 'type' to 'speed' explanation is far too simplistic. In my opinion, not only do factors inherent to the *type* of information to be learned affect the *speed* with which it *can* be acquired; but the *speed* of learning intrinsically affects the *type* of information that *can* be acquired. In addition, some specific types of information appear to be dependent on the hippocampus for acquisition irrespective of the number of trials available.

4.6.2 Episodic and semantic memory

Generally speaking, the term episodic memory is used to refer to 'personally experienced events', whilst semantic memory refers to 'knowledge of the world' (see also section 3.2.1.1). It is deceptively easy to understand what is meant by these definitions. However, are such memories really categorically distinct?

Operationally, episodic memory is usually distinguished from semantic memory on the basis of 1) the recalled *content*, such as whether specific event details can be remembered – episodic recall typically elicits contextual information, whereas semantic recall for facts and vocabulary, for example, does not; and 2) whether the memory refers to information that was experienced over an extended *period of time* or refers to a discrete 'episode'. Clearly, both these factors are continuous. Even if a subject recalls a full-house of 'what', 'when' and 'where' (which is sometimes considered to be the hallmark of episodic recall), the actual duration of the event would remain crucial to a definition. Remembering playing cards with a long-lost aunt for five minutes might be classed as memory for an episode, whereas remembering playing cards with the long-lost aunt over a week's visit would not. This 'content & duration-based' distinction gets progressively more fuzzy for the most 'episodic-like' semantic memories such as the recall of public events, as recall may depend on similar contextual information to that produced on episodic recall (e.g., who was where? at what time? and what were they wearing?). It is therefore not surprising that amnesic patients with temporal lobe damage tend to show proportionate impairment of memory for public events and autobiographical incidents (Kopelman, 2000).

Of course it is possible to artificially 'digitise' the definitions. For example, it could be specified that if recall elicits four details or more about an event with a duration of less than half an hour, then such recall will be designated 'episodic'. Recently there has been a move towards measuring the *number* of details recalled, which seems a better measure of episodic/semantic-ness than a categorical judgement (see section 5.2.4.1), but this merely underlines the continuous nature of the memories. The level of cuing employed in a test of memory introduces yet another 'continuous' factor into memory recall, as it strongly affects whether a specific episodic memory can be accessed, and thus how much detail can be recalled about a specific event.

To my knowledge, the only definition of episodic and semantic memory that is truly categorical refers to subjective consciousness: either a subject decides they have

a feeling of mental time travel and is therefore defined as experiencing episodic recall, or does not. Unfortunately, this concept is problematic in several ways. Firstly, the concept is untestable in non-human animals. Secondly, even in humans, a subjective feeling of recollection does not guarantee that the memory with which it is associated is 'real' or actually happened, so subjective feelings cannot really be used as a index of 'real' episodic recall.

Leaving aside problems of definition, typical real memories appear to combine episodic and semantic components to different extents, and may also depend for their recall on each component to different extents. As already discussed (section 3.2.1.1), normal episodic memory usually depends on the recall of semantic memory components. Indeed, a failure to reactivate appropriate semantic areas can lead to a failure of 'episodic' recall (Rubin and Greenberg (1998)). Whilst the recall of semantic information does not *necessarily* depend on recalling episodic aspects of memory, episodic information can enhance the recall of semantic information in normal individuals (Westmacott *et al.*, 2003). For example, when recalling that the Princess of Wales was killed in a car-crash, memories for where one was when one heard the news might support further recall of the story details. Furthermore, memories of ostensibly the same 'type' can consist of specific episodic and generic semantic components to different extents. For example, recall of 'semantic' information in patients with semantic dementia refers more to the patients' recent experiences that do typical semantic memories (Graham *et al.*, 1999), whilst 'episodic' memories in amnesics with hippocampal damage tend to be more generic and semanticised than typical episodic memories (see next chapter). It therefore seems practically impossible to separate semantic and episodic recall in the real world.

It seems much more plausible that real-world memories fall on a *continuum* of episodic/semantic-ness. Memories vary continuously in terms of the number, type and specificity of details recalled, and the extent to which recall depends on information that archetypally characterises episodic and semantic recall. A typical episodic memory is merely a memory at one extreme – of high complexity, specificity and novelty, referring to an event of short duration that has only been experienced once.

Several confounding factors cut across the standard definitions of episodic and semantic memory. One obvious confound is the level of detailed information usually required to demonstrate recall of the different categories of memory. To demon-

strate semantic memory for a public event one would need to recall, say, that Marilyn Munro fell downstairs getting off a plane; but to recall a similar personal event, not only would one need to recall that your friend fell down the stairs getting off a plane, but one would additionally need to demonstrate memory for other aspects of the occasion – for example, that you were on your way to a friend's wedding, it was cold, what you thought your friend thought about falling down the stairs, etc. The amount of complex context-like detail integral to recall is even less for more generic types of semantic information such as fact or vocabulary recall. Since damage to the hippocampus appears to reduce the amount of detailed information that can be recalled, irrespective of whether the information is ostensibly spatial, semantic or episodic (see sections 3.2.3.1, 3.2.4.1 and 3.2.4.2), and since the recall of typical episodic information inherently requires more detail, hippocampal damage might be expected to disproportionately affect the recall of episodic information on this basis alone.

Another confound is the frequency of exposure to the information: episodic events are by definition one-off occurrences, whereas semantic information can be abstracted from many encounters with a similar event or set of features. Again, the evidence suggests that hippocampal damage particularly affects the rapid acquisition of information (section 3.2.3.1), therefore on this basis also, hippocampal damage would be expected to disproportionately affect episodic recall.

As discussed in section 4.4, on average, typical episodic information is less likely to have been acquired deliberately than typical semantic information. Therefore, another likely confound is the intention to learn at acquisition, which is known to affect the engagement of memory systems.

The key factors that set episodic and semantic information apart are, in many cases, continuous (see Figure 4.1). That is, episodic information is acquired more quickly than semantic information, is more detailed, depends on a wider range of information to demonstrate recall, is more likely to be acquired incidentally, and is more likely to incorporate supra-modal information such as sequence or spatial information. Standard definitions of episodic and semantic memory will in many cases overlap with definitions that relate to the amount of specific detail recalled, or the amount of exposure there has been to information. This arises because by definition higher order conjunctions of features occur less frequently (or at least no more frequently) than lower order conjunctions of features. For example, an event such as seeing a mallard

duck land on water in your local park occurs less often (or at least no more often) than seeing a bird land on water. However, a consideration of the amount of detail, speed of acquisition and nature of the learning task more accurately taps into the reasons why information is hippocampally-dependent, and therefore allows us explain some of the anomalies in the data that are not captured by assuming that the hippocampus is necessary for episodic learning defined by specific content. Of course, specific instances of episodic and semantic information might differ qualitatively and quantitatively in one or more of these aspects.

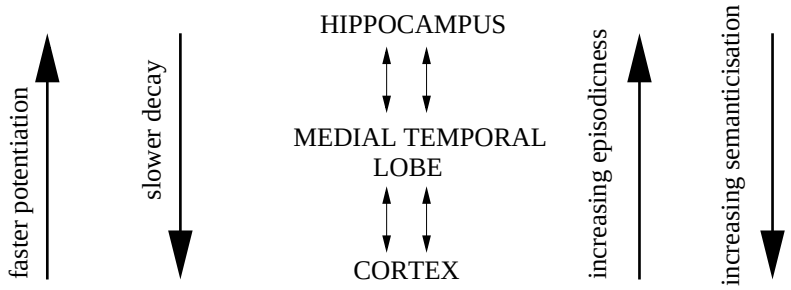

Figure 4.1: Simplified diagram of continuous features of the neocortical-hippocampal learning hierarchy. The learning rate is highest in the hippocampus, and the decay rate lowest (shown), and the degree of convergence of information similarly highest in the hippocampus and lowest in individual cortical regions. Recall mediated by the hippocampus is more likely to elicit the most detailed, specific information; whilst information that can be recalled via lower regions is progressively more generic and semanticised.

In the rest of this thesis I will use the term 'true episodic memory' to refer to a memory for an specific event that is similar to that initially stored, in that it contains a similar amount of specific detail and refers to a specific event that really happened. In next chapter I will show that as memories age, they become more semanticised. That is why it is crucial to strictly define what is meant by episodic information.

Chapter 5

Long-term role of the hippocampus

In this chapter, I examine empirical data on retrograde amnesia after hippocampal damage. I conclude that information that obligatorily requires the hippocampus at acquisition (such as the use of allocentric spatial information, or detailed true episodic information) depends on the hippocampus for the lifetime of the trace, although that may be less than the lifetime of the animal. The recall of information whose acquisition is merely facilitated by the presence of an intact hippocampus (such as semantic information or conditional motor learning) may become independent of the hippocampus over time. In general, those tasks whose acquisition is more impaired by pre-acquisition hippocampal damage show longer periods of retrograde amnesia after post-acquisition hippocampal damage. It is also clear that old memories tend to be more generic and less detailed than more recently acquired information.

5.1 Introduction

Retrograde amnesia (RA) refers to the loss of information acquired before the onset of amnesia. *Graded* RA is widely reported after hippocampal damage, that is, recently acquired memories are more affected than older ones (the so-called 'Ribot gradient', Ribot (1881)). This pattern of memory loss is traditionally thought to result from the progressive 'consolidation' or strengthening of traces outside the hippocampus, which can mediate the recall of information that was once hippocampally-dependent.

Two major kinds of consolidation processes should be distinguished: *local* consolidation processes that act to strengthen and/or prolong initial short-term connectiv-

ity changes in the areas initially involved in learning and mediating recall soon after acquisition; and *cross-regional* or *systems* consolidation processes that act to increase the strength of traces in areas outwith those areas that initially mediated acquisition of a trace. In the hippocampal literature, it is usually assumed that the hippocampus teaches or directs learning in the cortex, rather than it being under cortical control, making this type of consolidation an example of systems consolidation. At the biochemical level, similar mechanisms are assumed to be involved in both types of consolidation. Many current theoretical proposals lack clarity about the precise mechanisms of consolidation, and none provide constraints on a time scale. Most theories assume some version of the 'indexing theory', and implemented models of consolidation (e.g., Alvarez and Squire (1994); McClelland *et al.* (1995); Murre (1996)) tend to view the hippocampus as initially representing associative links between fragments of information represented elsewhere, until those regions develop the associative links themselves.

In this chapter I explore the relationship between the degree to which tasks are hippocampally-dependent at acquisition, and the extent of retrograde amnesia after post-acquisition hippocampal damage. Much of the data that I discuss has previously been presented as support for the standard view that all type of memories that depend on the hippocampus for acquisition are 1) stored in the hippocampus only temporarily; and 2) have a similar long-term fate that involves replication outwith the hippocampus. However, certain standard theoretical assumptions and interpretations of the data can be questioned; and studies using more subtle testing procedures have revealed important new findings that strongly suggest that a new theoretical perspective is needed. Whilst graded RA undoubtedly occurs on some tasks in amnesia with an etiology that includes damage to the hippocampus, there are alternative explanations to those given by the consolidation approach. Furthermore, some types of memory appear to depend indefinitely on the hippocampus.

5.2 Retrograde amnesia after hippocampal damage

5.2.1 Empirical issues

There are several key empirical concerns with studies of retrograde amnesia: some of these challenge the standard interpretation of graded RA.

Firstly, there are well-known problems in assessing the extent of brain damage, especially in humans; and a paucity of probable cases of "pure" hippocampal damage. The relevance and importance of this issue in attempting to correlate RA with damage to a particular brain area cannot be over-emphasised, especially in humans where amnesic subjects often present with imprecise non-surgical damage to the hippocampus with associated extra-hippocampal damage. In other animals, partial damage to the hippocampus has been shown to impair function in other intact regions (Baxter and Murray, 2001); and "invisible" extra-hippocampal damage has been shown to be functionally active in some cases (Mumby *et al.*, 1996). It is therefore to be expected that the human literature (especially the earlier studies) might be rather misleading at times. Of course, with careful testing and mapping of lesions, valuable information can be gleaned even from patients with widespread lesions (e.g., Rosenbaum *et al.* (2004)).

Most patients who have become amnesic have not had their memory assessed prior to the diagnosis of amnesia. Therefore, within-subjects methods can rarely be employed. Tests of human RA for personal events typically rely on asking patients to recall a few incidents from particular time periods in their lives. Because such events are necessarily personal, it is difficult to check the absolute truth of such recollections. Therefore, most studies compare patients' recollections with those of relatives or friends. However, in my opinion, the supposedly 'episodic' memories that can be recalled and verified under these conditions are likely to be quite different to 'true' episodic memories, which are detailed, refer to a particular event and have not been semanticised.

It is generally accepted in the wider field of memory research that memories undergo semanticisation with time. It is almost inescapable that the memories that can most easily be corroborated and thus used in a study are most likely to be semanticised and least likely to be "truly episodic". The existence of strong corroborative evidence such as testimony from family members, or even photographs or letters, implies that the recalled event is likely to be highly salient to a subject and/or his family – either at acquisition, or had become so since – and may have been rehearsed[1] many

[1]The psychological term "rehearsal" is often used to refer to repeating information (such as a phone number) to oneself to hold it in a memory store. This is more strictly termed "maintenance rehearsal", whilst "elaborative rehearsal" involves thinking about how the new information relates to previously stored information. The term is traditionally not used for processes in animals, however, I will use "rehearsal" to refer to any endogenously triggered reactivation of a memory trace that could change its subsequent strength.

times. Such rehearsal of the key points of an important event is likely to produce even more semanticisation than normal. In the most extreme cases, the content of memories that can be corroborated may well have been 'negotiated' since the event to produce an agreed 'family memory' that has little in common with the original event. Furthermore, as only a few memories are measured from each time period, and each time period may span several years or decades, then patients and controls are likely to choose only the most salient events. One patient with severe hippocampal damage has recalled the same few, poorly remembered 'episodic' events on different tasks and separate studies over the years (Rosenbaum *et al.*, 2004). The point here is that many tests that purportedly investigate *episodic* recall may well be accessing memories that are very far removed from episodic memories as usually defined, especially for older memories which are likely to be particularly semanticised. Therefore tests of recent and remote autobiographical memory may not be comparing like with like.

A related profound problem for retrospective memory studies is equating the difficulty of test information across time periods. In the majority of extant studies of human memory, an attempt is made to equate performance levels in the control group across the recent and remote periods tested (Brown, 2002). Since memories typically decay over time, this implies that the information evaluated from the earlier time periods was initially more strongly encoded, presumably because it had greater significance at the time. Therefore the remote and recent memories recalled on typical memory tasks might have been qualitatively even different from the outset.

Despite these concerns, studies of human amnesic patients are invaluable to our understanding of memory, because vastly more detailed and subtle information about deficits can be gleaned from the use of verbal and written tests than from mere observations of behaviour as in animal studies. Animal studies can however circumvent many of the empirical problems that arise in human studies, as the behaviour of an animal on a given memory test can be known in advance of a lesion, and the lesion can be relatively well-controlled and characterised. Specifically, the same known amount of training can be given at two different time points prior to a lesion and retrograde memory test, which simplifies the interpretation of any memory gradient observed. Of course, other issues arise, such as whether there might be an interaction between the speed of acquisition of information on earlier and later sets of learning trials and the possible role of the hippocampus in the fast acquisition of information (Murray

and Bussey, 2001) and the possibility that an animal's internal "context" is progressively more different to that at learning as time elapses between learning and testing, which is well-known by psychologists to affect memory performance. Furthermore, the validity of animal studies that claim to address the role of the hippocampus in memory have to date often been compromised by the nature of the task employed. As discussed in chapter 3, many of the tasks that have traditionally been used to investigate hippocampal function (such as delayed-non-match-to-sample or other recognition tasks), probably do not tap the same resources as those tasks that are most impaired by hippocampal lesions.

In chapter 3, I concluded that the different tasks that depend on the hippocampus for their normal acquisition depend on the hippocampus to differing extents. For expositional purposes, I divided these into tasks that are *obligatorily* hippocampally-dependent (i.e. tasks that are completely un-learnable without the hippocampus) and those whose acquisition is *facilitated* by an intact hippocampus (and are therefore acquired less easily without the hippocampus). In reality there is a continuum of 'hippocampal dependency' within the latter category, and across both categories: it is very important to remember that the division is an artificial one. In addition, a few tasks are actually acquired *more easily* after hippocampal lesions; and others are apparently *unaffected* by hippocampal damage at acquisition. I therefore divide the present discussion of the effects of hippocampal damage on the long-term maintenance of information into the same (continuous) 'categories' as used in chapter 3.

5.2.2 Information whose acquisition is unaffected by hippocampal damage may be retained or lost after hippocampal damage

Hippocampal damage has no discernible effect in a great many learning situations including category learning, priming and some implicit motor tasks. Consequently, for most such tasks there has been little interest in investigating whether hippocampal damage affects recall, but the assumption is that it does not.

More generally, whether retrograde amnesia can ever arise without anterograde amnesia is a contentious issue (see for example, the contrasting reviews of human literature by Kopelman (2000) and Kapur (2000)). In general there is a poor correlation between retrograde and anterograde amnesia (Kopelman, 2000), but it is rare that RA arises without *some* degree of AA (Poliakoff and Meudell, 2000). I can find no reports

of retrograde amnesia arising without anterograde deficits after damage limited to the hippocampus in humans[2]. There are a few reports of patients who show relatively normal acquisition and initial retention of new memories, but abnormally fast forgetting over the following days and weeks (Mayes *et al.*, 2003). However, most such patients show no evidence of hippocampal damage, but have damage or disruption in the temporal neocortex.

In non-human animals, one study has reported preserved retention of preoperative memories after hippocampal lesions on a forced-choice olfactory discrimination task in which post-lesion learning of new discriminations is normal (in rats, Jonasson *et al.* (2004)), in accord with the idea that retrograde memories remain intact if damage does not lead to anterograde deficits. However, retrograde amnesia without anterograde deficits has been reported after hippocampal damage. For example, Gaskin *et al.* (2003) reported that rats' retrograde memory was equally impaired by hippocampal lesions administered 5 weeks or 1 week after training on a novelty-preference task, but anterograde learning was normal. Intact AA with graded RA has also been reported after hippocampal damage on Pavlovian fear-conditioning tasks (Maren *et al.*, 1997; Sutherland *et al.*, 2001)). However, it is known that performance on this task can be based on several strategies, such as true contextual learning or cue-based learning, so it is possible that preserved anterograde learning reflects the acquisition of a different type of information to that lost in RA. This may also be the case in the previously mentioned studies. Therefore, RA without AA might reflect the retrograde loss of memories that were hippocampally-dependent at acquisition, combined with an unimpaired ability to acquire information using a non-hippocampally-dependent strategy. On the other hand, RA without AA might reflect the possibility that when the hippocampus is intact, it may sometimes mediate performance on a task even if it is not *essential* for acquiring the task, and may prevent concurrent learning in other areas. When the hippocampus is subsequently damaged, information that it had been instrumental in acquiring would be lost, and an intact alternative system for acquiring and storing that information 'revealed'. This would imply that learning may not proceed independently and in parallel in different memory systems on some

[2]The most widely accepted explanations for occurrences of disproportionate RA are a frontal lobe-related retrieval deficit (Levine *et al.*, 1998) or a visuo-spatial deficit (Ogden, 1993) that would affect the recall of visually-based memories, but would allow the storage of new memories without visual components. Alternatively, multi-focal damage to low level (semantic) features of memories or the pathways that connect them into an event might also result in disproportionate RA (Evans *et al.*, 2003).

tasks, but there may be inhibition and competition between regions. This may occur in other regions too, as RA without AA has been reported after damage to the rhinal cortex. Thornton *et al.* (1997) found that monkeys with entorhinal & perirhinal cortex lesions that impaired the recall of an already learnt two-object discrimination problem could learn a new object discrimination task apparently normally.

In summary, it is assumed that on most tasks whose acquisition is un-affected by hippocampal damage, there is no effect of hippocampal damage on similar retrograde memories. However, on a few tasks, hippocampal damage leads to a retrograde loss of information already acquired, but does not prevent new learning on a task. It is currently unclear whether the information lost is of an identical nature to that which can be acquired after hippocampal damage.

5.2.3 Information whose acquisition is impeded by hippocampal damage shows graded retrograde amnesia after hippocampal damage

In chapter 3, I concluded that semantic information, simple kinds of spatial information, delayed-non-match-to-sample tasks, paired-sample learning, and trace conditioning information could be acquired to varying extents after hippocampal damage. In this section, I conclude that after hippocampal damage, these types of memories depend on the hippocampus for their recall for characteristically different periods of time.

The type of information that is recalled on *typical* 'episodic memory' tests also shows graded RA. I will argue that this largely reflects the increasing semanticisation of memories with age, although other factors also play a role.

5.2.3.1 Semantic memory

Semantic information is defined as 'general knowledge' that can be made declarative, and includes vocabulary, grammar, object information, personal identity information, public information (such as knowledge of famous events or people) and facts about the world (such as knowledge about capital cities). Clearly, this is a very broad category, encompassing knowledge that is encountered in very different circumstances and with differing frequencies, and that is constituted by very different types of information of different levels of complexity. For example, grammar and most vocabulary is acquired relatively early in life, is experienced and used continually, and is usually

learnt through relatively implicit means. At the other extreme, famous events and personalities might only be fleetingly in the public eye, knowledge about them is constituted by detailed episodic-like information about 'who, where and what', and such knowledge may never be used again after its acquisition.

After hippocampal lesions, object identity information (such as that needed to perform match or non-match to sample tasks) can often be acquired (section 3.2.2.4); new vocabulary can sometimes be acquired if the exposure is extended though not if exposure is limited to a few laboratory-based sessions (section 3.2.3.1); and information about famous events and people can sometimes be acquired if the information is extremely salient and in the public eye for an extended period of time (section 3.2.3.1). These different kinds of semantic information that are associated with different patterns and contexts of exposure and require different amounts of detail to demonstrate knowledge would be expected to be differentially affected by hippocampus lesions.

In accord with this, most studies of retrograde memory in human amnesic patients have concluded that MTL damage leads to graded retrograde amnesia for semantic information such as public events, personalities and public faces, but leaves vocabulary, grammar and object identity knowledge intact (e.g., Squire (1992); Schnider et al. (1995); Verfaellie *et al.* (1995); Rempel-Clower *et al.* (1996); Nadel and Moscovitch (1997); Reed and Squire (1998); Kapur and Brooks (1999); Westmacott *et al.* (2001)), although Kitchener *et al.* (1998) reported apparently unaffected semantic memory for all time periods. That is, memories for different types of semantic information are differentially affected by hippocampal damage. More subtle distinctions may be seen too. Knowledge of personalities and public faces may show less extensive RA than memory for public events (Cipolotti *et al.* (2001)). The most complex or specific memories thus appear to be the most affected. Cipolotti *et al.* (2001)'s study is of further interest because their amnesic patient showed a more extensive and severe loss of memory for public events than that usually reported and the study specifically limited the public events tested to those that 'were no longer discussed in public life'. Therefore it is possible that these events were more like episodic memory in terms of exposure than tests in other studies.

Taking all these findings together, less well-learnt, less personally relevant, relatively complex semantic information to which there has been little exposure shows the greatest retrograde amnesia after hippocampal damage. The finding that vocabu-

lary or object identity information (which may be encountered innumerable times and in many different circumstances) tends to be preserved, whilst knowledge of events (especially those that were not in the public eye for an extended period of time) is often impaired, can be explained if semantic memory is affected in proportion to how many times the information has been encountered or how much relatively specific detail needs to be recalled. For example, if a patient has had little exposure to a "public event" (perhaps she only saw a politician once on the television), the memory would 'behave' more like an episodic memory. There therefore appears to be a correlation between the severity of the retrograde RA and the ease with which anterograde learning of such information proceeds after hippocampal damage.

Fewer semantic details about news events are recalled from all periods by amnesic patients when compared to normals (Nadel *et al.*, 2000). A similar amount of detail is recalled about self-selected news events from any period tested in both amnesics and controls (Nadel *et al.*, 2000), in contrast with the normal loss of detail that occurs as autobiographical memories age. This implies that the non-core incidental elements of semantic information (e.g., the colour of Kennedy's shirt when he was shot) may be lost very rapidly within the first few years after acquisition. This does not rule out the possibility that relevant episodic traces can aid the recall of recently acquired or currently relevant semantic information.

5.2.3.2 'Episodic' memory

It is well established that *graded* RA is seen on typical tests of 'episodic' memory after hippocampal damage, with older episodic memories relatively preserved compared to more recently acquired memories (e.g., Ribot (1881); Rempel-Clower *et al.* (1996); Reed and Squire (1998); Kapur and Brooks (1999); Schnider *et al.* (1994, 1995)). However, as argued by Nadel and Moscovitch (1997), many standard tests of ostensibly 'episodic' memory are not very sensitive. Until recently, most tests of episodic autobiographical memory have relied on a 3-point scoring technique introduced by Crovitz and Shiffman (1974) – points are obtained for specifying the time and location of an event, and supplying an 'appropriate' amount of detail. Therefore this procedure cannot distinguish between memory reports with varying amount of detail above the minimum subjective threshold (Nadel *et al.*, 2000). We cannot therefore be sure that memories recalled from recent and remote time periods are in fact equivalent, and should not

assume that the relative preservation of older memories reported on standard studies reflects the preservation of an identical type of memory to that lost from the recent period.

Two very fundamental facts about memory are widely accepted in the wider field of memory research: firstly, that autobiographical information is *very* unreliable; and secondly, that memories get semanticised over time. These ideas have not yet really had much impact in the field of hippocampal research. However, these facts provide little support for proposing a system that perfectly preserves episodic information (as seems implicit in Consolidation Theory). On the contrary, the 'failures' of memory may reflect an adaptive loss of information, particularly of contextual and source detail, that would otherwise overwhelm the memory system (Schacter and Dodson, 2001).

Undoubtedly, humans can store accurate detailed episodic information quickly and on an ongoing basis; and this ability depends at least in part on the hippocampus (section 3.2.4.1). However, equally undoubtedly, such memories are often poorly retained and are very inaccurate:

- Most episodic information is forgotten very quickly (days, weeks, months), unless it is important enough to be recalled before it has decayed.
- Events that are considered to be 'somewhat important' and are of sufficient personal interest to be discussed with others, can be completely forgotten, misremembered or confused with other events after only a few weeks (e.g., 44% of simple diary-recorded events are misremembered after only 1 - 5 weeks (Linton (1975); Odegard and Lampinen (2004)).
- Even events that are rated as important when they were experienced, or have high salience in the general population because they refer to culturally significant events, can be completely forgotten or mis-remembered (e.g., information self-rated as important in a diary study can be completely forgotten or confused with other events, and a memory for ostensibly well-remembered events such as hearing that Kennedy had died can be completely incorrect (Linton (1975)).
- False memories can be subjectively compelling (Odegard and Lampinen, 2004), and can occur even for so-called 'flash-bulb' memories (Neisser and Harsh, 1992). 'Re-experiencing' a memory is therefore no guarantee of truth.

- People regularly falsely import information that is suggested to them into their subsequent 're-experienced' memories, and construct appropriate narratives to incorporate such information (Loftus, 1989; Odegard and Lampinen (2004)). This presumably happens in normal everyday memory as well as in experimenter-induced situations.

Therefore any episodic memories that are actually retained relatively intact for long periods are likely to differ in important ways from an average recently-acquired memory whose fate may well be to be forgotten or mis-remembered.

There is a very small literature looking in detail at the maintenance of information over decades. With the exception of diary studies (which have received little attention in the hippocampal field), most investigations of episodic memory depend on the recall of a very few memories from one or more periods spanning a number of years, that can be corroborated. As I suggested in section 5.2.1 it is almost inevitable that such memories are semanticised versions of the original episodic memories. Little is known about old episodic memories in normal subjects (Poliakoff and Meudell (2000)), but it is possible that very few 'true' episodic memories are retained for longer than ∼2-years, or at least they might only be accessible with extremely specific cuing. That is, old 'episodic' memories even in normal subjects may consist largely of well-rehearsed semanticised fragments. If old retained 'episodic' memories tend to be semanticised versions of the original memories, and if the hippocampus is particularly important for the recall of detailed, complex, more truly episodic information, then graded RA for the 'episodic' memories typically recalled on episodic tests could result from the decreased dependency of semanticised memories on the hippocampus.

In summary, typical episodic recall shows graded RA after hippocampal damage. However, this likely reflects a normal shift to the recall of semanticised information as memories age. I discuss data on 'true' episodic recall in section 5.2.4.1.

5.2.3.3 Tasks that show graded RA in non-human animals

The majority of studies of retrograde amnesia in non-human animals have reported a temporally-graded retrograde loss of information after hippocampal lesions. These tasks include delayed non-match to sample (Zola-Morgan and Squire, 1990), contextual freezing (Kim and Fanselow, 1992; Anagnostaras *et al.*, 1999; Maren *et al.*, 1997), and socially transmitted food preferences (Winocur, 1990; Winocur *et al.*, 2001) (for

reviews see Squire (1992); Squire *et al.* (2001); table 2/p221, Nadel and Moscovitch (1997); table 1/p5, Murray and Bussey (2001)). Cue- and vector-based solutions on a watermaze also show graded RA (Ramos, 1998; Kubie *et al.*, 1999). Hippocampal lesions also lead to graded RA for trace-conditioning tasks (Kim *et al.* (1995); Moyer *et al.* (1990); Takehara *et al.* (2002)). Although the acquisition of trace eyeblink conditioning tasks is *massively* impaired by hippocampal lesions (unless there has been remote pretraining, Takehara *et al.* (2003)), a small amount of learning is possible after hippocampal damage (section 3.2.3.2).

The existing literature suggests that different tasks might show characteristically different lengths of RA, which may reflect the ease with which they can initially be acquired.For example, in rats, contextual-freezing tasks that can be acquired in one session of 7 CS-US pairings, depend on the hippocampus for only a few days after acquisition (Sachetti *et al.*, 1999); whilst the recall of trace eyeblink conditioning that requires 9-10 daily sessions of 90 CS-US pairings for acquisition, requires the hippocampus for up to 4 weeks (in rats, Takehara *et al.* (2003). Whilst this existing data is suggestive, it is currently too sparse to draw any strong conclusions.

5.2.3.4 Summary

Graded RA is seen after hippocampal damage for information that can be acquired to some extent in the absence of the hippocampus. In the human literature, this is documented on tests of semantic learning, non-allocentric spatial recall and semanticised 'episodic' recall. In non-human animals, graded RA has been reported for most tasks with the exception of allocentric spatial information. There may be a correlation between the severity and extent of retrograde memory deficits and the difficulty with which anterograde learning of such information proceeds after hippocampal damage.

5.2.4 Information whose acquisition is prevented by hippocampal damage shows flat retrograde amnesia after hippocampal damage

In chapter 3 I concluded that complex memories that are acquired quickly (such as episodic memory) and those that depend on associations between spatial or temporal information and other information (such as allocentric spatial navigation) can only be acquired with an intact hippocampus. In this section I discuss data that suggests that fully detailed and complex memories of these types remain dependent on the

hippocampus indefinitely: that is, this information shows a flat rather than temporally graded loss after hippocampal damage.

5.2.4.1 True episodic memory

Typical tests of retrograde memory do not distinguish between memories with different amounts of detail above a minimum threshold (section 5.2.3.2). In order to overcome this possible confound, Moscovitch and colleagues devised a new scoring system in which the *number* of details provided for each cued memory was counted (Moscovitch *et al*. 1998[3]). Using this scheme, Moscovitch *et al*. (1998) found an increased difference in the scores of control and amnesic groups, and that the small graded RA that is seen using the old scoring technique disappeared. Nadel *et al.* (2000) used the new more sensitive technique to assess the cued recall of memories from 5 periods from childhood to the recent period, and found that any recall of personal episodes by the amnesic patients was impoverished in detail at all time periods tested.

Other studies that have used a similar approach have similarly reported a loss of detail from all periods (Kitchener *et al.*, 1998; Cipolotti *et al.*, 2001; Westmacott *et al.*, 2001; Rosenbaum *et al.*, 2001, 2004). Subjects with hippocampal damage recall an equivalent (small) amount of detail about autobiographical memories from all time periods, whilst control subjects produce progressively fewer details as memories age (Nadel *et al.*, 2000; Rosenbaum *et al.*, 2004). Therefore for true episodic recall, patients show a Ribot-like pattern of recall when *compared* to control subjects. Clearly it would be wrong to interpret this as suggesting that the non-damaged areas in the amnesics were more able to mediate the recall of old rather than new memories of the same kind. As expected given the arguments of the previous chapter, the amount of detail recalled on semantic tests, such as memory for public events, has been reported to be similarly reduced at all time points (Nadel *et al.*, 2000).

Nadel and Moscovitch (1997) argued that flat losses extending up to 40-years for autobiographical information are seen after "hippocampal complex damage" (see table 1, p219, Nadel and Moscovitch (1997) for a summary of the relevant papers). However, with the exception of two of these studies, most patients have wide damage to

[3]Paper presented at the Brain, Behaviour and Cognitive Science meeting, Ottawa, Canada; discussed in Nadel *et al.* (2000).

MTL structures and 7/13 have additional damage beyond the MTL, so it is difficult to draw conclusions from these studies about the effect of hippocampal – or even MTL – lesions. However, a few studies of patients with relatively circumscribed hippocampal lesions do also report a virtually flat loss of autobiographical episodes (Kartsounis *et al.* (1995); Rempel-Clower *et al.* (1996); Cipolotti *et al.* (2001); Viskontas *et al.* (2000)[4]).

Thus the data suggests that it is *possible* that damage limited to the hippocampus itself and/or its closely related structures leads to the loss of detailed, complex, supra-modal associative information such as episodic knowledge from all time periods, although it is premature to implicate the hippocampus alone in the indefinite maintenance of detailed information. it seems likely that the hippocampus and other closely related structures play a similar role in the recall of many tasks.

The idea that the hippocampus might be required indefinitely for the recall of detailed information is given some support by the literature on imaging. Several imaging studies have reported equal activation centred on the hippocampus on the recall of recent and remote material (on verbal questioning about public events and autobiographical events, Maguire *et al.* (2001a); for mentally focusing on detailed aspects of events, Ryan *et al.* (2001); on verbal probing of autobiographical events, Nadel *et al.* (2000); on recall of detailed episodic information, Conway *et al.* (1999); and on strictly episodic tests when using SPM analysis, Piolino *et al.* (2004)). However, it is possible that the activation reflects new encoding. Furthermore, the studies of Maguire *et al.* (2001a) and Ryan *et al.* (2001) had previously re-activated the memories that were subsequently tested, which might confound a comparison of recent and 'remote' memories (see section 6.7).

In contrast, other studies have reported more activation centred on the hippocampus on the recall of recent rather than remote memories (for autobiographical memory retrieval, Piefke *et al.* (2003); for topographical details of episodic or semantic memories, Niki and Luo (2002); and for using cue words to recall episodes, Tsukiura *et al.* (2002); although Piolino *et al.* (2004) reported more activation of the hippocampus on remote versus recent memory on strictly episodic tests when using ROI analysis.) These findings might be interpreted as support the traditional view, or might equally plausibly result from memories becoming increasingly generic with time. Niki and

[4]Viskontas *et al.* (2000) reported extensive ungraded RA for episodic memories in patients with unilateral temporal lobe epilepsy. However, since unilateral hippocampal damage does not usually cause such profound deficits, it seems likely that the extensive RA reflects damage to regions in addition to the left hippocampus.

Luo (2002) claim that on their study at least, 'a similar amount of detail was recalled from all time periods', but unfortunately they did not control for episodic specificity. Even when attempts are made to control the similarity of the information recalled from recent and remote time periods, the quality of the mental image deteriorates with time (e.g., Piolino *et al.* (2004)), which suggests that the nature of the information recalled may be different.

Several other studies have reported differential effects for the left and right hippocampus, which may help to explain the contradictory findings. Several studies have found equivalent recent and remote activity but only centred on on the left hippocampus, with recent memories causing greater activation than remote in the right hippocampus (for the retrieval of autobiographical memory, Maguire and Frith (2003); and for episodic and semantic memories with a spatial component, Mayes *et al.* (2004)). Haist *et al.* (2001) reported greater recent than remote activation in the right hippocampus on the recall of knowledge about famous people in response to pictures. Different tasks are known to differentially engage the right or left hippocampus, so perhaps whether flat or graded RA is observed for whole-hippocampus measurements depends on the detailed nature of the tasks employed. Given that many studies to date have shown preferential left-sided activation in the recall of autobiographical memories irrespective of the age of the memory (Maguire and Mummery (1999); Ryan *et al.* (2001)), this suggests that the regions primarily involved in the recall of autobiographical memories may be involved indefinitely in recall.

To my knowledge there are no studies of retrograde amnesia in non-human animals on tests that could reasonably be said to be analogous to human episodic memory (such as Clayton and Dickinson (1998)'s "what, where, when" caching task).

In summary, true detailed autobiographical memory appears to remain indefinitely dependent on the hippocampus. However, the relative difference between the performance of controls and patients lessens as the memories age, suggesting that detailed memory information is normally lost as memories age. Detailed semantic information, especially that referring to relatively complex and infrequently encountered information such as public events and personalities, is also similarly dependent on the hippocampus.

5.2.4.2 Spatial information

Like episodic memory, spatial memory in amnesic patients has generally been reported to show Ribot-like graded RA, with remote memories intact (e.g., anecdotally in HM; Beatty et al. (1987); Teng and Squire (1999); Morris (1999)). However, also like episodic memory, recent evidence suggests that the spatial memory measured on standard tests may be abnormal. A recent extensive study of a severely amnesic patient reported good memory for general layout and major landmarks of the neighbourhood lived in since a child, but impairments at remembering details of landmarks and less salient features (Nadel *et al.*, 2000; Rosenbaum *et al.*, 2001). Thus after hippocampal damage, spatial memories might also be deficient in detail and complexity at all time periods.

In non-human animals, hippocampal damage has been reported to lead to ungraded (i.e. complete) memory loss on tasks such as the water-maze (e.g., Bolhuis *et al.* (1994); Ramos (1998); Sutherland *et al.* (2001), for reviews see Squire (1992); table 2, p221 Nadel and Moscovitch (1997); table 1, p5 Murray and Bussey (2001)); and Squire *et al.* (2001)). Two studies have reported *graded* RA after hippocampal lesions on an ostensibly spatial task in a water-maze (Ramos, 1998; Kubie *et al.*, 1999), but the nature of the residual performance of the task (and new learning) was qualitatively different — perhaps based on cues or vectors — so residual memory likely reflects the performance of areas outside the hippocampus. On the surface, these findings differ from those in humans, where some apparently allocentric spatial memory can be recalled from early time periods. However, the retrograde memory tests typically used in human and non-human animals differ in an important way: animals must perform a given task which they may or may not retain enough spatial information to perform, whereas humans are often free to verbally report any spatial information they retain. In other words, behavioural tests in non-human animals typically depend on the recall of more specific spatial information than verbal recall tests in humans. Therefore the complete failure of animals to perform specific allocentric tasks acquired in any pre-lesion period may be analogous to the failure of humans to retain detailed abstract spatial information; whilst the residual non-allocentric ability for animals to perform ego-centric or cue-based strategies that show graded RA may be more analogous to the general non-specific spatial memories recalled in amnesic humans on most tests to date.

Additionally, the nature of the exposure that humans and non-human animals have to spatial information differs, if, for example, the recall of neighbourhoods in humans is compared with an animal's memory for platform position in a water-maze. Rosenbaum *et al.* (2001) found that rats that were trained pre-operatively on a maze task that encouraged the formation of multiple representations were less impaired after hippocampal lesions than those trained on an invariant version (though they performed worse than controls which performed similarly after both types of training). This suggests that the 'multiple representations' training allows areas other than the hippocampus to gain some ability to mediate recall of spatial information. Such complex environmental exposure is probably more analogous to human spatial learning, which might also partially explain why humans show graded or generally reduced but not absent spatial memory after MTL damage; in contrast to the temporally-extended gradients seen in non-human animals.

Therefore the animal and human data may be consistent: flat RA is seen for detailed or truly allocentric spatial information that is crucially dependent on the hippocampus at acquisition, and graded RA for less detailed spatial information and for strategies that could be acquired by regions other than the hippocampus.

Knowlton and Fanselow (1998) have suggested that the tasks that Nadel and Moscovitch have identified as being stored in the hippocampus on the basis of their ungraded loss after lesion, are in fact dependent on the hippocampus because of online performance demands such as the use of working memory. This proposal is difficult to rule out experimentally as any inactivation of the hippocampus would be expected to simultaneously affect both the putative storage site and any 'online process'. Furthermore, the proposal might simply equate to suggesting a specific role for the hippocampus in the fast acquisition of information, such as when the information required to solve a task changes from trial to trial. However, ironically, anterograde deficits may be responsible for producing *graded* RA in some cases, which rather weakens the argument. In a replication of Bolhuis *et al.* (1994)'s study, Ramos (1998) found (as did Bolhuis and colleagues), that the retrograde amnesia reduced progressively as the time elapsed (1, 16, 32 or 64-days) between acquisition of a radial maze spatial task and dorsal hippocampal lesions, and no lesion-induced recall deficits were found after 64-days. This result is therefore in accord with traditional consolidation ideas. However, if the test data to be analysed was limited to an average of the first 5 tri-

als (compared to the original 18) — which helps to rule out possible contaminating re-learning over the test — there were no significant differences in retrograde amnesia amongst the experimental groups. That is, all lesioned groups (irrespective of the timing of the lesion with respect to training) were equally impaired. The "graded nature" of the deficit is apparent only with respect to the controls, since they become increasingly poorer at recall of more distant information, and thus increasingly close in performance to the subjects with no recall for any period. Furthermore, Knowlton and Fanselow (1998)'s 'anterograde deficit' explanation of flat RA is not very parsimonious, as an additional mechanism would then be needed to explain graded RA.

5.2.4.3 Summary

Memory for detailed information which is inherently likely to be relatively complex, unique, and/or of low salience appears to depend on the hippocampus indefinitely, irrespective of whether memories are ostensibly spatial, episodic or semantic. Because episodic memory is usually defined by reference to the number and/or specificity of the details that can be recalled, episodic memory appears disproportionately affected, with true episodic memory showing flat RA after hippocampal damage. At present there is no evidence that supramodal information such as temporal relational and allocentric spatial information becomes independent of the hippocampus over time. Instead, hippocampal damage appears to lead to an ungraded retrograde loss of usable allocentric information, with some sparing of less detailed and less specific spatial information.

5.2.5 The retrograde effect of hippocampal damage on tasks whose acquisition is improved by hippocampal damage is currently not known

In chapter 3 I identified a few tasks whose acquisition is speeded by hippocampal lesions. These were mainly tasks whose 'best' solution is an egocentric strategy mediated by the basal ganglia, such as caudate-dependent learning in a win-stay radial maze task (e.g., Packard *et al.* (1989); McDonald and White (1993)). On these tasks, the hippocampus appears to inhibit simultaneous strategy acquisition by other regions. Unfortunately, I can find no studies that report the retrograde effects on memory for such tasks after hippocampal lesions, so we can only speculate what those effects might be.

5.2.6 Degenerative and non-discrete disorders

So far I have focused on data that relates to patients with the most discrete hippocampal lesions, or where extra-hippocampal damage can be relatively easily quantified. In general, where there is evidence of damage to the hippocampus, the findings in neuro-degenerative disorders do not contradict the findings from studies with discrete hippocampal lesions. For example, in Alzheimer's disease (AD), lesions start in the hippocampal, entorhinal and trans-entorhinal regions, before expanding to neocortical regions. In accord with this, patients with AD generally present with anterograde episodic memory deficits and then progress to a breakdown of language, to perceptual and spatial function and possibly to semantic information. In a direct reflection of the 'discrete' studies already reviewed, patients with AD have commonly been reported to show a much greater loss of recent autobiographical memory than recent semantic memory, although there is better remote recall in both cases (Moss *et al.*, 2003); whilst truly episodic memories appear to be completely lost (Piolino *et al.*, 2003b). Many brain insults give rise to hippocampal damage, including Korsakoff's syndrome, Herpes encephalitis, hypoxia, vascular disorders, epilepsy etc. In all cases, there is severe amnesia as expected, but also extra-hippocampal damage (for a review see Kopelman (2002)).

Amnesia can also result from various insults to the brain that result in little hippocampal damage. Of these, semantic dementia (SD) has received the most attention from those interested in the hippocampus and its role in memory: it has been suggested that SD is in some ways a mirror-image of medial temporal lobe amnesia. The major insult in SD is in the infero-lateral temporal neocortex and temporal poles, and at least initially there is a relative preservation of the medial temporal lobe. Nestor *et al.* (2002) has however noted that similar levels of hippocampal atrophy have been reported in patients with SD and AD that show the typical memory profiles associated with these diseases. Patients with SD typically show a progressive impairment in semantic knowledge as the disease progresses, with a reverse temporal gradient so that more recently acquired semantic information is better retained. Patients initially remain well-oriented in time, and can learn new (anterograde) episodic events relatively normally for up to 2 years from the onset of disease (Graham and Hodges, 1997; Graham *et al.*, 1999).

It has been reported that *recent* autobiographical memories are better preserved

than remote memories (Graham and Hodges, 1997; Graham *et al.*, 1999; Nestor *et al.*, 2002; Piolino *et al.*, 2003a,b). Such data can be seen as support for Consolidation Theory in that the relatively preserved hippocampus might mediate recall of recently acquired, but not old, semantic and episodic memories. However, Moscovitch and Nadel (1999) have argued that this reverse-Ribot deficit largely reflects linguistic and semantic deficits, as current everyday experience provide progressively fewer relevant cues or vocabulary for older memories. In accord with this proposal, SD patients can show relatively well preserved memory for all time periods when cued with family photographs (Westmacott *et al.* (2001)[5], and anecdotally, Moss *et al.* (2003)) or with increasingly specific verbal cues (Moss *et al.*, 2003).

However, it is difficult to explain in terms of Moscovitch and Nadel's 'semantic deficit' hypothesis how increasingly detailed specific verbal cues could aid autobiographical recall in patients with SD (Moss *et al.*, 2003): it seems implausible that more specific subordinate cues would be more comprehensible to a patient with semantic deficits. Instead, tasks that provide very specific cues may simply be easier, which may disproportionately aid the recall of more impaired older memories, perhaps by directly activating parts of widespread attractors in the fragmenting semantic system. In accord with this, Moss *et al.* (2003)'s patient required significantly more cuing of older memories to achieve the same level of performance as with more recent memories: an effect which was absent in controls[6]. Interestingly, Moss *et al.* (2003) found a gentle reverse-Ribot graded RA for autobiographical information in IH for intermediate levels of cuing, like those normally reported in SD. They suggested that the questions asked in standard memory tests (e.g. AMI or modified Crovitz test) represent intermediate levels of cuing. This finding therefore provides a bridge between the two sets of findings.

The recall of autobiographical memories has been reported to be impaired compared to controls at all time periods (Moss *et al.*, 2003; Nestor *et al.*, 2002; Piolino *et al.*, 2003b). Deficits at all time points would be predicted if the recall of hippocampally-based traces depends on the integrity of the semantic fragments it indexes, as is usually assumed. Recently acquired memories would be relatively preserved not because any hippocampal trace alone can mediate recall, but because they will have been

[5]Unfortunately, Westmacott *et al.* (2001)'s study had no control subjects, and the criteria for 'episodic' recall was rather low.

[6]Unfortunately, in Moss *et al.* (2003)'s study, the control subjects were not cued in an equivalently personally-specific way as IH, so the lack of an effect on recall cannot be taken be taken at face value.

stored in conjunction with fewer subsequently damaged semantic fragments than older memories, as less time will have elapsed for the disease process to operate in the semantic areas. The possibility that older memories might generally depend more on semanticised memories might also lead to the relative preservation of 'episodic' memories. The relative preservation of the hippocampus would not protect the recall of truly semantic memories stored in the fragmenting semantic system, although it might be able to provide an alternative, more situation-specific method for recalling some semantic-like information. Given questions over the extent of hippocampal damage even early on in SD (Nestor *et al.*, 2002), autobiographical deficits may of course reflect hippocampal damage.

5.3 Summary and conclusions

In this chapter I have examined the effects of hippocampal damage on retrograde memory. Until recently, the established wisdom was that graded RA is seen after MTL damage for all information that depends on the hippocampus for its acquisition. Instead, the following points appear to more accurately sum up the effect of hippocampal damage on existing memories:

- Information that is obligatorily hippocampally-dependent at acquisition shows a complete or extremely extensive retrograde loss after hippocampal damage. This includes:
  - The unique non-generic *details* of complex information, whether ostensibly semantic, spatial or episodic. By definition, true episodic memories will be most affected by such a deficit, but all memories are affected to the extent that they depend on such information.
  - Information that involves associations between *supra-modal* information and other information (such as using allocentric or temporal information).
- Temporally-graded retrograde memory deficits are seen for information whose acquisition is facilitated by the hippocampus. This includes:
  - Most semantic information in humans.
  - Typical 'episodic' tasks – this probably reflects the increasing semanticisation of older memories.

  - Most tasks tested to date in non-human animals (excluding allocentric spatial tasks)

  There may be a correlation between the severity of anterograde learning deficits after hippocampal damage, and the extent of retrograde memory deficits for similar information.

- Memory for information whose acquisition is unaffected by hippocampal damage is probably largely unaffected by hippocampal damage. In a few cases, retrograde amnesia is evident after hippocampal damage, but it not clear whether the same type of information is being acquired before and after hippocampal damage.

- It is currently unknown how memory for information whose acquisition is impeded by an intact hippocampus is affected by hippocampal damage.

These findings are summarised in table 5.1.

Type of information	**Effect of hippocampal damage on:**	
	Anterograde learning	**Retrograde memory**
Detailed, complex, supramodal e.g., true episodic; complex semantic; allocentric navigation	Prevented	Flat retrograde amnesia
Generic, less complex e.g., semanticised episodic; public events; DNMS; taxon navigation	Impeded	Graded Ribot-like retrograde amnesia
Priming, basic skill learning, classical conditioning	Unaffected	No effect?
Ego-centric strategy learning	Improved	Unknown

Table 5.1: This table summarises the effect of hippocampal damage on anterograde learning and retrograde recall of information.

Crucially, hints of characteristically different patterns of recall between different categories of information have been seen even within the same patients. This helps to rule out explanations in terms of the different methodologies employed in different studies. For example, one patient showed a flat loss of detailed episodic information about events prompted by photographs, but graded retrograde deficits in naming

the people pictured in the photographs (Westmacott *et al.*, 2001); and flat RA for autobiographical memories, but graded RA for personal semantics (Rosenbaum *et al.*, 2004). Another patient showed flat RA for detailed episodic recall, and a trend towards graded RA on recall of famous faces (Cipolotti *et al.*, 2001).

5.3.1 The long-term role of the hippocampus

There is little evidence that memory components that are initially obligatorily hippocampally-dependent (traces that mediate the recall of spatial, semantic or episodic supra-ordinate detail) become independent of the hippocampus, at least not for normal memories. Perhaps this is not surprising. In chapter 3, I strictly defined obligatorily hippocampally-dependent tasks as those that could not be acquired to any extent after hippocampal damage irrespective of the type or extent of post-lesion training. I suggested that on these tasks, the hippocampus was the *only* area that could develop representations of the requisite information under the given task conditions due to its ability to represent complex associative, novel or supramodal information or to learn quickly. It is possible that the 'endogenous training' or rehearsal that is posited to underlie the relative strengthening of traces outside the hippocampus after acquisition is similarly incapable of supporting the development of traces for such information outside the hippocampus; or at least that it remains a very arduous proposition. Such information would therefore show a flat loss after hippocampal damage irrespective of the time elapsed since acquisition, as the hippocampus would be the only area that could acquire a representation of that information given the extant learning conditions.

Information that can be acquired to some extent in the absence of the hippocampus might be able to benefit from endogenous 'training trials'. Graded RA would therefore be seen after hippocampal damage on these tasks, if traces that could mediate recall were gradually built up outside the hippocampus. Tasks that could be acquired relatively easily after hippocampal damage, might therefore also be expected to show relatively short temporal extents of graded RA after hippocampal damage, as the necessary traces might also be relatively easily built up outside the hippocampus after acquisition.

However, on the basis of the evidence reviewed in this chapter, information that is initially represented in the hippocampus does not appear to be copied wholesale to

the extra-hippocampal regions. Instead, over time, regions lower in the neocortical-hippocampal hierarchy may develop representations that are less complete than those stored at higher levels such as the hippocampus, but that can nevertheless mediate performance on some tasks. I explore this and related issues in the next chapter. Crucially, the switch from a dependence on specific to more generic memories (and therefore the appearance of graded RA) can *only* occur on tasks that *can* be mediated by semanticised 'reduced' traces, and not on those in which all/some of the details initially stored remain important to performance.

In chapter 4 I concluded that different types of tasks are hippocampally-dependent at acquisition for different reasons: some rely on the convergence properties of the hippocampus, others on its ability for rapid learning, and yet others on prolonged delay activity that allows difficult associations to be formed. It seems that tasks also have different long-term fates in terms of hippocampal dependency. In the next chapter, I discuss how the semanticisation of memories, and its interaction with the specific nature of memory tasks, can inform an overall understanding of the role of the hippocampus.

Chapter 6

Semanticisation and the role of the hippocampus

In this chapter I argue that information that is obligatorily hippocampally-dependent at acquisition is indefinitely dependent on the hippocampus for recall. As all memories tend to decay with time, detailed information may not be maintained for the lifetime of the animal, although the core 'semantic' features may be retained outside of the hippocampus. The graded retrograde amnesia seen after hippocampal damage on tasks that can be acquired to some extent in the absence of the hippocampus reflects a qualitative change in the nature of memories recalled from recent and remote periods. That is, it reflects the semanticisation of memories. Semanticisation of a memory at recall results from a number of factors including the loss of detailed information, an increase in the strength of semantic components, or differences in the recall strategies underpinning the recall of new and old memories. Post-acquisition memory processing is multi-stage and lifelong, and acts to produce semanticised re-representations of information that can co-exist with more detailed traces. Different information and even different components of information acquired together may have different fates in terms of their long-term dependence on the hippocampus.

6.1 Introduction

It is instructive to begin this chapter by considering how the main extant theories of the long-term role of the hippocampus – that is, Consolidation Theory and Multiple

Trace Theory – cope with explaining the existing data. In the rest of the chapter I show that a theory of progressive semanticisation can better accommodate existing findings, focusing on the data that is most problematic for the aforementioned theories.

6.1.1 Traditional Consolidation Theory

Traditional Consolidation Theory (CT) approaches are motivated by the numerous reports of graded RA after hippocampal damage, which is assumed to reflect gradual hippocampally-driven consolidation in extra-hippocampal regions. It also rests on the assumption that the hippocampus is a small limited-capacity store, and that there is need for hippocampally-supported interleaved learning in the cortex to avoid catastrophic interference. Further support comes from studies that show that the hippocampus is progressively less active for recall as memories get older, whilst extra-hippocampal areas become more active (e.g. Bontempi *et al.* (1999); Piefke *et al.* (2003); Niki and Luo (2002)), although not all studies have shown this pattern. Studies showing an off-line replay of activity that was present during learning (e.g., Skaggs and McNaughton (1996); Qin *et al.* (1997)), and the neurodynamics of the hippocampus in different phases of sleep (e.g., Chrobak and Buzsáki (1994); Kudrimoti *et al.* (1999)) have also been interpreted as circumstantial evidence that the hippocampus might be involved in replaying information to the rest of the brain.

However, evidence that cannot be easily accommodated by this view is accumulating and several *a priori* assumptions can be questioned. For example:

- Finding extensive ungraded retrograde amnesia after hippocampal damage (section 5.2.4) implies that the hippocampus may be a permanent storage site for some information.

- The information recalled in normal individuals from different time periods is qualitatively different (section 5.2.3.2) – which suggests that memories stored initially in the hippocampus, and recalled subsequently from the cortex, may differ.

- Graded RA can also occur after damage to the structures other than the hippocampus, both in the MTL and elsewhere, which implies that the hippocampus/MTL may not be unique in its long-term memory role. (This issue is discussed further in section 6.3.)

- The state of activation of a trace rather than memory age per se may determine whether it is vulnerable to amnestic agents or hippocampal damage. This suggests that the idea that memories become less fragile with time is too simplistic. (This issue is discussed further in section 6.7).

- Lifelong neurogenesis in the hippocampus makes it questionable that the hippocampus has a limited storage capacity and can therefore only function as a temporary memory store, especially since capacity is inextricably linked to questions about the nature of information being stored, which is still open to debate.

- Graded RA has been reported to extend for decades in humans with large medial temporal lobe lesions. Therefore events from much of our ancestors' lifetime can be apparently maintained without the completion of 'consolidation' which is supposed to protect memories.

- Hippocampally-driven interleaved learning in the cortex is thought to be necessary to prevent catastrophic interference on memory storage in the cortex. However, this argument depends partly on the assumption that the cortex stores detailed episodic memories, which can be questioned.

Moreover, 'strong' interpretations of consolidation theory that supposes that initial memory storage takes place only in the hippocampus are refuted by large amounts of evidence showing learning-related changes in the cortex on the initial acquisition of hippocampally-dependent tasks (Greenough and Bailey, 1988), and by data suggesting that areas other than the hippocampus can initially hold information on learning tasks when the hippocampus is inactivated (Floresco *et al.*, 1996).

Furthermore, the concept of consolidation is divorced from the nature of the information that is putatively 'consolidated'. Therefore, it can provide no insight into the anterograde deficits that arise after hippocampal damage, or the different durations of RA reported for different tasks. Whilst this cannot be considered a criticism of the theory's ability to explain the existence of RA per se, one would expect AA and RA to be related and therefore a unified parsimonious explanation of acquisition and retention would be preferable. Similarly, since the consolidation view explains the Ribot gradient by proposing the need for a process called 'consolidation' that is supposed to make memories less vulnerable to interference, it clearly cannot provide a *reason* for

a process that makes memories vulnerable every time they are retrieved as is implied by the reconsolidation literature (see section 6.7).

In summary, consolidation remains a hypothetical construct, for which conclusive evidence is lacking. Flat retrograde amnesia can occur after hippocampal damage, as well as graded RA; and the information recalled from recent and remote periods may be qualitatively different. This strongly suggests that consolidation needs re-thinking.

6.1.2 Multiple Trace Theory

Nadel and Moscovitch (1997)'s Multiple Trace Theory (MTT) explains the Ribot gradient by proposing a gradual proliferation of traces in the 'hippocampal complex' (HC)[1] as memories are reactivated and rehearsed, that makes *partial* damage to this system less likely to lead to memory loss over time. The recall of all autobiographical memories, whether recent or remote is dependent on the HC, with flat RA resulting from complete HC lesions. This reflects the permanent role of the hippocampus in storing spatial contextual components of episodic memory. In later versions of MTT it was made stated that hippocampal-neocortical connections could contribute to the extraction of semantic information to be stored in the neocortex, although how episodic trace proliferation leads to graded RA for semantic information is never spelt out. Presumably this would provide a second consolidation-like mechanism that could generate graded RA in the model without reference to multiple traces, although this is not made explicit. Thus, MTT can accommodate findings of flat and graded RA after complete and partial HC lesions respectively. Furthermore, imaging and activation studies showing equal MTL activation on the recall of remote and recent autobiographical memories (e.g. Maguire *et al.* (2001a); Conway *et al.* (1999); Piolino *et al.* (2004)) can easily be accommodated.

However, other data is problematic for MTT.

- Apparently complete lesions of the hippocampus have been reported to produce graded RA (for spatial information in rats, Winocur *et al.* (2001); Squire *et al.* (2001) and for episodic, semantic, and spatial information in humans, Teng and Squire (1999)), whereas MTT would predict graded RA only for partial lesions.

[1] In Nadel and Moscovitch's terminology the 'hippocampal complex' effectively refers to the whole of the medial temporal lobe, and includes the hippocampal formation, the entorhinal cortex, the perirhinal cortex and the parahippocampal gyrus, and sometimes the amygdala too. Confusingly, the implemented computational model of MTT refers only to the hippocampus (Nadel *et al.*, 2000).

- Partial hippocampal lesions are sometimes more detrimental to memory function than complete lesions (Baxter and Murray, 2001); whereas in MTT, partial lesions should allow better recall performance than complete lesions.

- Flat and graded RA have been seen in the *same* patient depending on the task. Since the extent of RA in MTT is supposed to depend on the degree of damage to the 'hippocampal complex' then this result directly contradicts MTT.

- If trace replication is supposed to increasingly protect memories from partial hippocampal complex damage, then replicated traces in the hippocampus must be largely non-overlapping[2]. Therefore the reactivation of existing traces and the putative laying down of a new one would not be expected to interfere, although this is what the reconsolidation literature appears to imply (section 6.7).

- One of the few consistent findings in the re-consolidation literature is that the processes occurring at initial 'consolidation' are different and more laborious than those occurring at 're-consolidation'. This is very difficult to explain in terms of MTT, as the processes occurring at each reactivation – that is, new learning of non-overlapping traces – should be very similar.

- Data suggesting a change in the locus of recall from the hippocampus to elsewhere over time (e.g., Izquierdo and Medina (1997); Frankland *et al.* (2001); Bontempi *et al.* (1999)) is problematic for original versions of MTT, although it is unclear whether later versions of the theory allow for this if the information is not truly autobiographical.

- Similarly, a qualitative change in the nature of information recalled over time (see previous chapter) is not explicitly predicted by MTT, although may possibly be accommodated by later theoretical stances.

- Graded RA extending for decades cannot be easily accommodated by MTT, especially since Nadel and colleagues (2000) interpret their simulation results as implying that traces should best be replicated only a few times.

As with CT, there is little attempt to relate MTT proposals to the hippocampus's anterograde role, although MTT does make a distinction between the fate of autobi-

[2]MTT qualitative theory and the implemented model do however appear to differ in the degree to which replicated traces overlap, see Nadel *et al.* (2000).

ographical and semantic memories as regards the hippocampal complex. It is also not clear that such massive redundancy (each reactivation of a trace resulting in the storage of a new unrelated trace) is either plausible or beneficial. From a theoretical point of view, the idea of using truly random orthogonal traces to represent the same information is a little problematic, too, as it is not clear what they would 'mean'.

In summary, the idea that graded RA reflects trace proliferation in the 'hippocampal complex' is difficult to defend. Not only is there empirical evidence to the contrary, but the idea that the development of multiple physically-separate traces underlies graded RA remains hypothetical, and conclusive evidence is lacking.

6.1.3 Semanticisation theory

I believe it is time to seek an alternative to both Consolidation theory and Multiple Trace Theory. Like Nadel and Moscovitch I believe that some memories *do* remain dependent on the hippocampus for the lifetime of that information although not necessarily for the lifetime of the animal. In accord with Consolidation Theory, I agree that the neural substrate that is necessary for recall of *some* tasks *can* change over time. However, I believe that the existing data points to a new understanding which we might call 'Semanticisation Theory', which has not hitherto been worked out in much detail. In order to distinguish my proposals from the assumptions of related ideas embedded in other work (e.g.Rosenbaum *et al.* (2001)) I dub my views 'Gingell's Semanticisation Theory' or GST. In this view, different types of information, and even different components of information that were acquired at the same time, can have different fates in terms of maintenance by the hippocampus, and on whether and how other regions take over recall of this information.

GST shares some tenets with CT, but they differ fundamentally in several ways. Firstly, in GST, memories whose recall can be mediated by non-hippocampal regions are qualitatively different to those whose recall is necessarily mediated by the hippocampus. Graded RA after hippocampal damage results not from a hippocampally-driven consolidation process which gradually teaches the cortex the same information that was initially held by the hippocampus, but largely from a change in the nature of memories that are recalled from earlier periods. Secondly, in GST, the recall of non-semanticised memories can depend on the hippocampus or MTL for as long as those traces are retained. GST is also significantly different to MTT. Most obviously, trace

proliferation in the hippocampus plays no part in GST's explanation of graded RA after hippocampal damage. GST also explicitly connects the long-term fate of a trace in the brain to the nature of the task, a feature lacking in both CT and MTT.

The nature of the information recalled in normal individuals from recent and remote time periods is qualitatively different. In my opinion, this is fundamental to understanding the long-term role of the hippocampus and leads to a coherent explanation of the extent and nature of RA after hippocampal damage. In normal individuals, recent memories of all types tend to be more detailed, complex and 'situated' than remote memories, and there is a tendency for remote memories to be semanticised or 'tend toward the norm'. Patients with hippocampal damage tend to produce memories from all time periods that are more generic, less specific and less detailed than normals, for both typical semantic memory (e.g., knowledge of geography, Nadel *et al.* (2000)) and for unique events Nadel *et al.* (2000); Rosenbaum *et al.* (2004)). Flat RA for true (detailed) episodic information therefore arises straightforwardly because the hippocampus is needed indefinitely for the recall of detailed, complex, supra-modal information. Graded RA on memory tests after hippocampal damage can arise for several reasons, including 'process-related' reasons such an increasing dependence on semanticised memories for older information; as well 'methodological reasons' such as matching control performance on recent and remote memories, which effectively guarantees that recent and remote memories were qualitatively different from the outset.

Semanticisation of a given memory can result from the operation of several mechanisms including the normal decay of the details of memories, the relative or absolute enhancement of semanticised representations and the use of different memory recall strategies. I discuss these in section 6.2. In section 6.3 I argue that instead of a unitary hippocampal or MTL long-term memory processing system, there may be a series of regions that represent progressively more semanticised information. Semanticisation can also provide an explanation for the very long gradients of RA seen in human amnesic patients (section 6.4). Whilst it is implausible that a consolidation process should last for an organism's lifetime, it seems reasonable that memory could be continually re-organised and re-represented in order to make it more generalisable and to preserve the key important features even as detailed information decays. At a given time, multiple representations of different degrees of complexity and specificity may be avail-

able to mediate performance depending on the exact demands of a particular task (section 6.5). The realisation that the details of episodic information tend to remain dependent on the hippocampus for as long as they are retained, rather than being transferred to extra-hippocampal areas, also renders moot arguments about the necessity for a hippocampal mechanism that supports interleaved learning (section 6.6). Similarly, the idea that memories are continually being re-represented and progressively more robust semanticised traces being formed, can provide an understanding of why memories appear to undergo similar processes when they are recalled as they do when they are initially stored (section 6.7), an explanation which eludes existing theories.

6.2 Graded retrograde amnesia results from the semanticisation of memories

There is little evidence to support the idea that memories are retained for long periods of time with an undiminished ability to mediate recall. Instead, memories and components of memory are often forgotton, with the most specific detailed information usually being lost first. Specific memories therefore become more like average generic memories over time: this is termed semanticisation. Semanticisation of a memory at recall results from the relative decrease in the trace strength of detailed information compared to more generic information. This might occur because of 1) a decay of the detailed information; and/or 2) a relative or absolute increase in the strength of the semantic traces; and/or 3) from a change over time in the recall strategy employed to access memories as they age. I examine these possibilities in turn:

6.2.1 Decay of detailed information

On free recall tests, normal people produce many more details about recently acquired autobiographical memories than remote memories (Nadel *et al.*, 2000; Rosenbaum *et al.*, 2004). This normal decay of detailed information would in itself contribute to the semanticisation on average of autobiographical memories. Since hippocampal amnesic patients, show poor (but equivalent) recall of details about personal episodes from all periods back to childhood (Nadel *et al.*, 2000; Rosenbaum *et al.*, 2004), amnesic patients are more impaired with respect to controls in the recent period on the

recall of true detailed episodic information. This relative effect could contribute to the appearance of graded RA, depending on the test procedures used in a given study.

Existing data is limited, but suggests that even at the most remote time period, the performance of amnesics and control subjects on the recall of episodic details does not completely converge (Nadel *et al.*, 2000; Rosenbaum *et al.*, 2004). Therefore, some detailed information may be retained by the hippocampus/MTL throughout life in normal people.

6.2.2 Increased ability of semantic traces to mediate recall

Learning-related processes in the cortex are well-documented and there is a growing recognition that areas outside the hippocampus can learn independently, albeit more slowly (see section 3.2.3.1). Therefore, one possible way that semantic information might build up in the lower levels is from continued re-exposure to real-world events containing repeated features. Over time, areas outside the hippocampus would become more able to mediate recall of re-experienced generic information. Therefore older memories that are more dependent on generic information at recall would be less affected than recent memories by hippocamal damage.

By definition, each real-world episodic event (which may however incorporate repeated semantic fragments) will not be re-experienced, and therefore true episodic representations could not be strengthened by this mechanism. It has been argued that if episodes could be re-experienced (for example, by replaying a video, Holdstock *et al.* (2002)) then cortical regions could build up a trace of such episodic information. Whilst repeated exposures to an event might increase the strength of some aspects of the memory representation, I would stress that this does not necessarily mean that the non-hippocampal regions are acquiring the *same information* as that stored in response to a true episode. The aspects of an episodic event that can be captured on film are only a subset of what constitutes an episodic event – each time the film is viewed, another 'true' episodic event is occurring that incorporates the viewers current thoughts and feelings, expectations, recent experiences, attention paid to parts of the film, etc. These more transient truly episodic details will never be captured by the slower-learning cortex. There *may* be no fundamental logical reason why such information could not eventually end up being represented in the cortex, if events could be *accurately* replayed to the cortex over a *very* long period of time, but such a learning

scenario is practically impossible.

This does not mean that the hippocampus and other higher levels never help the lower levels to acquire information that they would not realistically be able to acquire by themselves. For example, extra-hippocampal areas alone can initially mediate normal learning and performance on a probablistic weather-prediction task, but the absence of the hippocampus impairs later performance (Poldrack *et al.*, 2001). The hippocampus may act as a scaffold to direct learning that occurs elsewhere on incremental tasks that are learnt over multiple trials. However, it is important to note that with the hippocampus' help, the extra-hippocampal areas develop representations of information that have been extracted from several trials, rather supporting learning about event-specific information.

Another possibility for increasing the strength of semantic traces is that the reactivation of traces stored at higher levels in the neocortical-hippocampal hierarchy could support learning at the lower levels, by orchestrating the offline co-activation of 'semantic areas' that might then strengthen shared connections. This is the mechanism usually posited to underlie consolidation. However, I do not believe that such rehearsal can support a complete replication of information from one level to the next. Different versions of an event would be recalled using traces stored at different levels in the hierarchy: memory recall mediated by traces stored at lower levels will be progressively more semanticised than those stored at the higher levels. There is no wholesale transfer of detailed episodic information to the cortex.

Replay-aided learning or rehearsal by definition occurs on the basis of the recalled memory, rather than in response to a real event. Overt recall that can be measured is notoriously unreliable (see section 5.2.3.2), and there is no a priori reason to believe that 'endogenous rehearsal' should be any more reliable. The neural environment (in terms of recent activity, intended action, circulating hormones, co-activation of other information *etc.*) will be different each time a trace is reactivated, which will also affect what is recalled, as well as how any reactivated information is re-stored. Therefore it is likely that only the core elements of a memory will be reliably reactivated at each replay trial, and therefore only those elements will be enhanced/retained over time.

As already noted, there can be no absolute definition of what is 'core' to an event or memory, as it must be defined with respect to instances of that event, and/or to the knowledge and expectations of an individual (section 4.3). Broadly speaking, core

features can be seen as the key semantic features of a memory that must be present if recall of that information can be said to have occurred. For the recall of events, it is most likely to be these (if any) features that are reliably recalled each time, whilst incidental features (such as shirt colour, who got on the train first) are inconsistently recalled. Therefore any rehearsal-driven memory storage would be expected to disproportionately benefit the most important semantic elements, even of specific autobiographical episodes. Furthermore, since information represented in the higher levels of the hierarchy decays at a relatively faster rate than information elsewhere, details may well be lost before they can be replicated. In addition, given 'wiring constraints' it is likely to be relatively more difficult to develop traces for novel complex detailed information in regions outside the hippocampus, which also implies that on average, the details of events are likely to be less reliably replicated than the generic aspects.

Figure 6.1 sums up these ideas. In Figure 6.1a, a memory is laid down initially most robustly in the hippocampus, but also in the cortex. In the lowest cortical regions, long-range connections between cortical regions do not reliably register new associations. Immediate recall via hippocampal traces results in the recall of the most detailed information (Figure 6.1b). Immediate cortically-mediated recall may elicit some information, especially if the recently-acquired information is similar to information already stored, because such related information will be easier to register in the slow-learning cortex. However, such information may be fragmentary and only locally robust. For all memories that are retained and do not decay immediately, fewer details are available as a memory ages. Recall via the hippocampus will however always elicit the most detailed information even though parts of the hippocampal indexing trace or the cortical fragments representing details of exemplars may deteriorate (Figure 6.1d). Cortical traces for the core features may become more robust over time through re-exposure to similar information embedded in different events, or attractor-based endogenous replay (Figure 6.1e), although in many cases the improvement in cortical recall may only be relative to hippocampally-mediated recall.

6.2.3 Change in recall strategy

Another possible source of semanticisation is that different recall strategies might be used to elicit recall for recently and remotely acquired memories. If asked to recall a specific event from the last week, information 'springs to mind' relatively easily be-

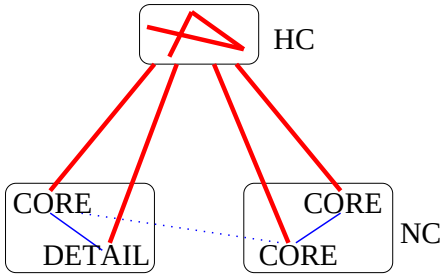

(a) Initial memory storage

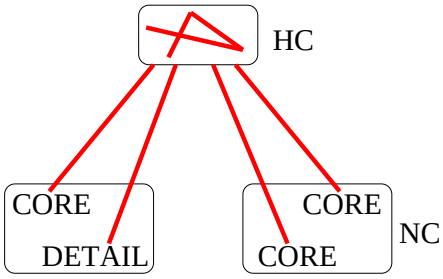

(b) Initial recall via HC

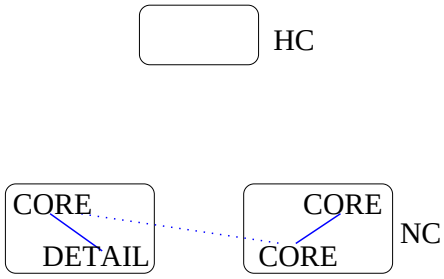

(c) Initial recall via NC

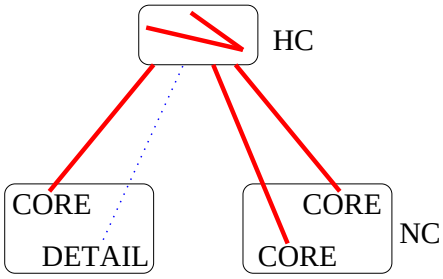

(d) Later recall via HC: semanticisation due to decay of HC trace

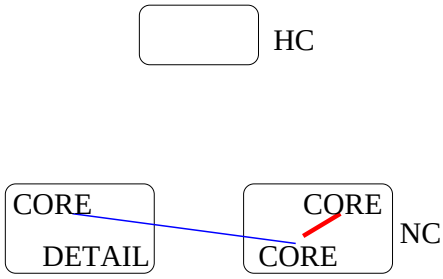

(e) Later recall via NC: semanticisation through cortical learning

Figure 6.1: Diagram of memory acquisition and long-term development. HC = hippocampal component, NC = neocortical component.

cause recall can be cued by currently relevant features of one's life. However, if asked to recall something from twenty years ago, the search must be focused using personal semantic knowledge (what age would I have then? Was I at university? Which university? Where was I living? *etc.*) until the recall of a salient incident is triggered.

Increasing evidence shows that memories are often created (not merely *re*-created) at recall, and that an estimate of the plausibility of a possible personal event may be based on personal semantic information (Odegard and Lampinen, 2004). This implies that the recall of older or currently less salient autobiographical events might be based on 'guesses' about what might reasonably have happened. Such a strategy would therefore contribute to the progressive semantication of older recalled memories, and tend to shift older memories towards a norm. If more detailed and specific cues are provided as recall cues (e.g., a photograph of the event in question) then recall would depend less on personal semantic memory. Therefore, the search strategies used for recent and remote memories might be more similar, and the relative degree of semanticisation of the older memories recalled might be reduced.

A related possibility is that there is a change in the balance of 'competitive power' between neural regions over time, so that the recall of older memories tends on average to be mediated by regions further down the neocortical-hippocampal hierarchy. This is perhaps merely another way of saying that the memories at higher levels might decay faster than those at lower levels. However, I return to the issue of competition between regions and the co-existence of multiple traces in section 6.5.

6.2.4 Episodic and semantic memory revisited

In section 4.6.2 I argued that episodic memory was best seen as an extreme point on a continuum of memory types. In the previous chapter, we saw that memories tend to become more semanticised over time, in that the most specific details are lost, and progressively fewer details are recalled when a particular memory is recollected. That is, archetypal 'episodic' memories may tend to be recalled in a semanticised form as time progresses. Above I discussed in detail the possible mechanisms underlying the semanticisation of memories. Loosely speaking, the repeated presentation of similar events allows the gradual build-up of a trace in the cortex of the repeated core (semantic) parts of the event, which will be largely shorn of the non-repeated episodic detail. Semantic traces will thus tend to build up over time as the rate of decay is relatively low in the cortical areas. The hippocampus will retain the most detailed traces for their lifetime. Semantic information ultimately derives from exposure to episodes since all experience of the world unfolds in episodes, but does not necessarily depend on a functioning episodic learning system that allows the demonstration of the recall

of each such individual episode.

The 'uniqueness' and specificity of memories reflects the amount of detail known about the entities involved (i.e. information about specific exemplars), and/or information about the combinations of entities involved, which may themselves be prototypical or specific. Semanticisation of information could therefore result from a loss of the higher traces that co-ordinate the recall of increasingly unique conjunctions of fragments, or from a change in the 'content' of the fragments from specific to prototypical. In terms of the neocortical-hippocampal representational hierarchy proposed, information about the combinations of fragments to be co-recalled will tend to depend on the higher regions, whilst the representations of exemplars may come to depend on the lower regions. Exemplars, such as the next-door neighbour's cat 'Cynthia', can be repeatedly experienced, therefore a representation of that specific cat can be built up through real-world experience. Such learning throughout life leads to the development of a range of information about entities from relatively specific to more generic, for example, representations of Cynthia... of Persian cats... of cats generally... of mammals etc.

On the other hand, specific events of which exemplars are part cannot be re-experienced in the real-world, and therefore representations of features of specific episodes cannot be built up in the same way. However, different types of events can be re-experienced allowing the development of 'scripts' ranging from the specific to the more generic, for example, representations of what is involved in walking a dog... or, more specifically, taking the next-door neighbours' dog for a walk... or what is likely to ensue when the next-door neighbours dog sees Cynthia *etc*. The whole system might therefore be considered to form a continuum from generic semantic information to progressively more more specific episodic representations (as suggested by Figure 4.1).

All other things being equal, traces stored in the hippocampus will decay very rapidly after initial acquisition, whilst traces stored in the cortex will be relatively stable. Episodic traces stored in the hippocampus will initially be relatively very strong, but in normal circumstances will decay very fast: most inconsequential episodic information appears to be largely lost within days of the experience. Any retained detailed information is likely to remain dependent on the hippocampus for its lifetime. A weak trace of some of the novel components of an episodic event might be retained in the

cortex, especially when this information can be easily represented in the existing architecture (usually because it is relatively similar to what has gone before). Any such episodic information will be very weak and probably fragmentary, and be unlikely to be accessible without very specific cuing; but relatively stable over time. Semantic information stored in the lower regions is relatively stable.

6.3 Semanticisation is a multi-stage process

Proponents of traditional Consolidation Theory originally proposed that all structures of the medial temporal lobe (MTL) namely the entorhinal and perirhinal cortices, hippocampal formation and parahippocampal cortex, plus the fornix, worked together as a single functional unit in the consolidation of memories (e.g., Alvarez and Squire (1994); Squire and Alvarez (1995)). However, recent evidence undermines the idea that the MTL plays a unitary role in long-term memory.

It is now largely accepted that different MTL structures make specialised contributions to *anterograde* learning and memory. For example, the perirhinal cortex (PrC) has a larger role in recognition memory (as indexed by DNMS tasks) than other areas (Suzuki, 1996; Murray *et al.*, 2000), and does not contribute to spatial memory (Aggleton *et al.*, 2000); whilst the entorhinal cortex (EC) contributes to object recognition memory (but less than the PrC) and to some spatial tasks (although less that the hippocampus, Suzuki *et al.* (1997); Aggleton *et al.* (2000)), and in some cases perform these functions apparently normally in the absence of other MTL structures. One would therefore expect different MTL structures to make distinct and perhaps independent contributions to the long-term maintenance of information, too.

Indeed, the steepness and extent of the graded RA seen after damage to MTL structures depends both on the region, and on the information being tested. For example, in rodents, recall of a one-trial step-down inhibitory avoidance task depends on the entorhinal cortex for more than 30-d (Quillfeldt *et al.*, 1996) and the parietal cortex for more than 60-d, with intra-parietal CNQX (an AMPA-receptor antagonist) only partially effective at causing amnesia at 90-d (for reviews see Izquierdo and Medina (1997); Ambrogi-Lorenzini *et al.* (1999)); recall of context- or tone-freezing in rats requires the hippocampus for more than 1.5-h, the basolateral amygdala for up to 2 days and the perirhinal cortex for 8 days (Sachetti *et al.*, 1999); and recall of trace eyeblink conditioning requires the hippocampus for up to 4 weeks (in rats, Takehara *et al.*

(2003) and mice, Takehara *et al.* (2002)). Other species show a similar pattern, for example, in Rhesus monkeys, recall of a two-object discrimination problem depends on the hippocampus for less than 4 weeks, but on the rhinal cortex for at least 16 weeks (Thornton *et al.*, 1997). This strongly suggests that different MTL structures have at least partly independent long-term roles, and that they become redundant in a fixed order.

Graded RA has occasionally been reported after damage to non-MTL structures such as the mamillary bodies, thalamic nuclei, substantia nigra and prefrontal cortex (Winocur, 1990; Aggleton and Brown, 1999). Therefore, we must also conclude that whatever mechanism underlies the production of graded RA after a lesion, it is not unique to the MTL. Furthermore, different patient groups with different brain damage, show characteristic differences in the nature and extent of graded RA, which also implies that graded RA is not caused by a single mechanism.

The 'unitary MTL' view of consolidation has clearly been superseded. Instead, the data outlined above shows that for tasks that show graded RA after hippocampal damage, individual regions in a series of regions (including structures in the MTL) are required for recall for different periods of time after acquisition, and that they become redundant in a predictable sequence. The hippocampus appears to become redundant first, followed by the EC, and then other areas such as the PrC and the parietal cortex. For each type of task, the relative importance of different regions appears to change, as does the length of time that regions are required.

It is commonly suggested that the hippocampus acts as a 'scaffold' or 'indexing system' (Teyler and Discenna, 1986) that can co-activate traces stored in cortical regions. However, the evidence outlined above suggests that this proposal should be extended at least to other areas in the MTL. This suggestion is close in spirit to Damasio (1989b)'s concept of 'fragments' and 'convergence zones' throughout the brain. In section 4.3 I introduced the idea of a 'neocortical-hippocampal hierarchy'. The aforementioned regions do indeed become redundant in the order predicted by this hierarchy, with flat retrograde memory loss usually reported for non-MTL structures, especially those lower in the hierarchy. In this view, the hippocampus is unique in its scaffolding role mainly because of its unique position at the top of the neocortical-hippocampal hierarchy, rather than because of qualitative differences in function; although its high plasticity and information convergence makes it distinct from other MTL structures in

other ways too.

Of course, different structures in the MTL specialise in processing different types of information and/or performing different functions, and are not distinguished solely on the basis of the complexity of the information they represent and the speed with which they acquire it. The perirhinal cortex is implicated in learning about conjunctions of visual features and memory for objects (Murray *et al.*, 2000); the parahippocampal gyrus in memory for landmarks and contexts (Owen *et al.*, 1996); and the hippocampus in memory for spatial locations (O'Keefe and Nadel, 1978). Damage to each of them leads to a complete loss of the type of information that they process. However, these regions could also be seen as representing increasingly specific information, from identity-free landmarks, through objects, to objects in particular places.

6.3.1 Scaffolding and inter-regional information transfer

There is substantial evidence for changes in neocortical structure as a result of learning occurring over weeks or months after training. For example, exposure to an enriched environment or to certain tasks leads to increases in dendritic length and synapse number in the relevant cortical areas (for a review see Greenough and Bailey (1988)); and somato-sensory and sensory cortical maps can reorganise and contract or expand after experience with certain sensory tasks or after lesions or sensory deprivation (see any neuroscience text e.g., Shepherd (1994); Kandel *et al.* (1995)). However, there is only scant direct evidence that the hippocampus or other areas in the hippocampal-neocortical hierarchy are *necessary* for the post-training changes in cortical traces that governs their time-dependent ability to mediate recall (although such data should be crucial to a defence of Consolidation Theory).

After the initial acquisition of information, memory-related processes do commonly occur sequentially in different regions, as might be expected if changes in a lower region depended on input from a higher region. However, the most well-documented of these are are relatively short-term memory-related processes (extending for minutes or hours[3]) which are therefore of a different order to the extensive

[3]For example, the timing of the onset of the initial NMDA-dependent phase of memory in the hippocampus, EC, and parietal cortex is sequential, starting immediately and extending for a few minutes in the hippocampus, starting 30-min after training in the EC, and after 60-min in the parietal cortex (Izquierdo and Medina (1997)). The sequence and time courses for impairments are similar (but not completely coincident in all areas) to those caused by the PKA inhibitor KT 5720 (Bernabeu *et al.* (1997); Izquierdo and Medina (1997); Ardenghi *et al.* (1997)).

periods of retrograde amnesia previously reported (days, weeks and months in rodents and non-human primates, and years in humans), and could not explain that phenomenon. Additionally, some studies have reported no obvious sequence[4].

In terms of evidence that regions higher in the hierarchy might support learning at lower levels, one preliminary study found that the changes in total dendritic length seen after enriched experiences in the occipital lobe were reduced by damaging the hippocampal formation before the environmental manipulation, although this task did not in fact require the hippocampus for acquisition (Sutherland *et al.*, 1993). Miyashita and colleagues found that ibotenate lesions of EC/PrC in monkeys[5] disrupted the development of the code for paired associates in IT neurons – that is, the paired associates fail to elicit significantly correlated responses in the IT cells of the lesioned animals – without impairing the non-learnt response to each visual stimulus (Miyashita *et al*, 1994; Miyashita *et al.*, 1998). Thus it is possible that cells in the PrC/EC acts as a coincidence detector for representations of the paired associates, and then act a scaffold for the IT neurons to develop these responses. This sparse evidence therefore suggests that higher areas in the neocortical-hippocampal axis *might* support learning in lower levels, although the evidence remains far from conclusive.

As I have already indicated (section 6.2.2), information that is initially represented at one level is not generally *completely* replicated at a lower level. When a higher area becomes redundant for the performance of a given task, it is because a 'sufficient' amount of information to perform the task can be recalled via the lower regions, not because all the information that was initially represented in the higher regions has necessarily been transferred wholesale to the lower regions. Given the well-documented learning capacities of the cortex and other areas, it is likely that much of the information represented there is determined by those regions, rather than imposed on them by the hippocampus. Since the neocortical cells that originate projections to the PrC and PH cortices are not necessarily the same cells in particular cortical regions that receive the feedback connections (unpublished observations, (Lavenex and Amaral,

[4]For example, there is no obvious sequence for the induction of LTP in dorsal hippocampus, entorhinal and parietal cortex and amygdala, Walz *et al.* (2000)). There is also evidence for repeated waves of consolidation in areas that may correspond to gene expression and new protein synthesis (for review see Abel and Lattal (2001)), suggesting that these processes do not necessarily depend on a simple sequence of regional engagement.

[5]Bilateral lesions of the PrC/EC cortices would impair the monkey's behaviour on the pair-association task (Murray *et al.* (1993)). So the PrC/EC is lesioned unilaterally, and the anterior commisure is cut at the beginning of the experiment.

2000)), it is unclear how the hippocampus could re-impose particular cortical patterns on the cortex without aid from the cortex through, say, local already-stored attractors or synaptic tagging mechanisms. Since networks of neurons bear information in activity patterns by virtue of their connectivity and history of activation, it is also difficult to see what a pattern imposed by the hippocampus would "mean". Furthermore, there is unequal reciprocity between regions in the neocortical-hippocampal hierarchy, so it seem unlikely *a priori* that information could be 'transferred' in a similar fashion between all regions. For example, there are far fewer efferent projections from the PrC to frontal cortex than vice versa; more connections from the the superior temporal sulcus to the PrC than vice versa; and relatively few connections from visual areas V4 and TEO to the PrC, although the PrC sends widespread projections to these areas (Lavenex and Amaral, 2000).

The reconsolidation literature suggests that regions that have become redundant for the recall of a given memory become necessary again for recall for a period after that memory has been recalled. This implies that a one-way sequential redundancy of regions in the putative hierarchy might only occur if a memory is not triggered by real-world events again after acquisition. I discuss this further in section 6.7.

6.4 The temporal extent of graded retrograde amnesia

Graded retrograde amnesia extending for days, weeks, months or years has been found after specific brain damage in all species tested to date. The temporal extent of graded RA is traditionally assumed to reflect the time needed to build up or 'consolidate' representations outwith the hippocampus. However, there is no a priori reason for assuming that molecular changes underlying memory representations cannot be made more quickly than the typical periods of RA observed.

The length of graded RA usually seen after hippocampal damage increases from rodents to primates. RA has been reported to extend for days or weeks in mice, rats and rabbits; weeks or months in monkeys; and years, decades and whole lifetimes in humans. One explanation for an increase in the length of RA across species is that it reflects the increase in brain size across species: in larger brains, the structural changes outwith the hippocampus that would be necessary to re-represent the information initially stored in the hippocampus may be more onerous and hypothetically more time-consuming (see e.g. Murre and Sturdy (1996); Dash *et al.* (2004)).

Graded RA in humans has been reported to extend for decades in some cases, which is a contentious issue for traditional Consolidation Theory and variants such as MTT. As pointed out by Nadel and Moscovitch (1997), it seems unreasonable that a consolidation process should extend for nearly the whole lifetime of our ancestors, given that consolidation is supposed to overcome the twin constraints of a limited capacity hippocampus, and the supposed necessity of interleaved learning to safely store long-term memories. Unfortunately, Nadel and Moscovitch (1997) do not explain why trace replication in the hippocampus should plausibly take a lifetime either, and their simulation results imply that limited replication would anyway be more plausible (Nadel *et al.*, 2000). It is *theoretically* possible that the very long RA in humans *may* be related to brain size as suggested above, as human brains are disproportionately larger (2 - 3 times) than one would expect for a primate of human body weight (Gilissen and Simmons, 2001), which might create a step-like difference in the time taken for putative 'consolidation' processes. Specifically, larger brains may be associated with increased cortical specialisation and reduced connectivity between regions (Schoenemann, 2001), and it also been suggested that human brains might necessarily employ more energy efficient non-synaptic transmission mechanisms (Bach-y-Rita and Aiello, 2001) which might be associated with some as-yet-unknown slow 'consolidation' processes (although it might equally well be associated with a reduction in the need for time-consuming long distance synaptic re-wiring relays). However, whatever the explanation, the *value* of such an extended memory reorganisation scheme would remain unclear.

The length of RA in man is apparently of an entirely different order to that of other primates (i.e. years and decades rather than weeks or months), which may suggest additional *qualitative* differences between the processes underlying graded RA in humans and other animals. One possibility is that the recall of fundamentally different types of information is being compared. If, for example, recent and remote memories in humans depend to different extents on truly episodic and semanticised traces, then the measured memory gradient in humans would reflect the combined recall gradients of more than one memory system. The information typically required to demonstrate good performance on the behavioural tests employed in animal studies is often quite specific (e.g., the performance of a specific action at a particular place in response to a particular cue) and may necessarily depend on non-semanticised non-generic in-

formation, whereas a dim recollection of some pieces of information about where one lived might be counted as recall on human retrospective verbal tests. Semanticisation (which in animals canbe thought of as a tendency for memories to become progressively more generic and less context-specific over time), may be just as common in non-human animals as humans, but the point here is that the memory tests typically employed in different species may differ in the sensitivity to this factor.

In the semanticisation view, as time progresses after acquisition, memories that are important enough to be remembered or rehearsed may be reactivated, allowing parts of the memory trace to be strengthened, possibly in a process akin to that envisaged for 'consolidation' in the traditional view. However, as reactivated memories are never reactivated 'perfectly', only the core elements will be strengthened. Therefore, the long periods of time involved in graded RA (on those tasks for which it occurs) do not reflect a process that is necessary for the preservation of memories per se (because important complex memories will be retained indefinitely by the hippocampus), but reflects the building up of different, more generic representations that can independently mediate recall on some tasks, and that may have applicability to a wider range of situations than the original very specific event memory. New relationships between semantic information can be discovered over time. Building up such a store of reliable, generic information over a lifetime seems adaptive, especially given that the higher level highly convergent binding zones that co-ordinate the recall of higher order information are progressively more localised, and therefore progressively more prone to discrete brain damage. Indeed, the hippocampus itself is notoriously prone to many kinds of brain damage which likely relates to its high capacity for plasticity.

Although existing studies of humans are difficult to interpret and compare (Spiers *et al.*, 2001), a few general statements can be made about the length or completeness of RA: episodic recall is more affected than semantic recall; recall is more affected than recognition; more detailed episodic and semantic information is more affected than less detailed information; and more specific semantic information is more affected than information with more general applicability to which there would have been greater lifetime exposure. Similarly, in other animals, the length of RA after hippocampal damage depends on the task. In rats, for example, graded RA after hippocampal/MTL damage extends up to 2-d (Winocur *et al.*, 2001) or 5-d (Winocur, 1990) for socially transmitted food preference in rats, up to 4-wks for contextual fear

conditioning (Kim and Fanselow, 1992), and several months or indefinitely for place navigation (Ramos, 1998; Kubie *et al.*, 1999; Sutherland *et al.*, 2001). This suggests that any putative hippocampal 'consolidation' process must interact with the nature of the task. In the semanticisation approach, differences in the length of graded RA between tasks reflect tasks' differential sensitivity to the degree of episodic/semantic-ness of the information required at recall (with tasks depending on more specific information showing longer gradients after hippocampal damage), or differences in the ease with which certain kinds of information can be represented by regions other than the hippocampus.

Squire and colleagues have suggested that larger MTL lesions lead to more profound amnesia (Rempel-Clower *et al.*, 1996). However, very selective lesions sometimes apparently produce severe amnesia, leading others to conclude that the location and the completeness of a lesion are more important than its overall size (Spiers *et al.*, 2001). The current proposal can accommodate both these ideas. Recall depends on locally convergent 'index traces' stored at various levels in the neocortical-hippocampal hierarchy and throughout the brain. Larger MTL lesions will destroy progressively more such indexing traces for a given memory leaving fewer areas within the MTL that are able to recall any semblance of the original information. Very selective lesions, if they target regions which contain large concentrations of these bindings traces will have very profound effects on memory, especially if whole traces are destroyed leaving no partial traces to aid the recall attempts mediated by the remaining brain. On average, more discrete lesions will have a more profound and generalised effect on memory when they target regions higher up in the neocortical-hippocampal hierarchy, especially if the information to be recalled is relatively complex, because such regions represent increasingly specific information.

Morris (1999) has noted that the severity of RA damage appears to be related to the extent of cortical damage, too. Since recall mediated at a higher level in the hierarchy cannot proceed without activating appropriate components at lower levels, the extent of cortical damage should correlate with memory impairments irrespective of the age of a memory, if the cortical areas damaged contain fragments of a given memory. Discrete lesions at progressively lower levels of the hierarchy will have a less global effect on general memory, but may prevent the recall of specific memory components such as visual information.

6.5 Regional redundancy and trace decay

The data reviewed in the previous section suggests that over time, the areas that are *necessary* for the recall of a particular trace may change. However, this does not necessarily imply that the relevant traces in the areas that were previously required for recall have been lost, and indeed I am not aware of any studies that demonstrate the loss of hippocampal traces. In this section I explore the possibility that traces at each level in the hierarchy are retained at least for some time after that area has become redundant for a particular task (as indexed by lesion studies), and that these traces may continue to be involved in normal recall.

A series of studies by Packard and colleagues suggest that on at least some tasks, hippocampal traces are retained even when the mediation of recall has switched to another area. A cross-maze task, which requires a rat to learn the position of a consistent reward at the end of one arm, can be acquired either by a place (say, "turn to the west") or response ("turn to the right") learning mechanism. After seven days of training, saline-treated rats show predominantly place learning, and after another seven days training, they show predominantly response learning (Packard and McGaugh (1996); Packard (1999)). Lidocaine injections into the dorso-lateral caudate putamen (dl-CP) leave place learning intact at day 8, but block the expression of response learning at day 16 so that animals continue to express a 'place response'. In other words, neural inactivation of the dl-CP has uncovered the redundant, but intact, hippocampal trace. Clearly, such training differs from putative endogenously-driven consolidation or semanticisation[6]), but the findings underscore the possibility of multiple co-existing viable traces. To my knowledge, there is no direct evidence that hippocampal traces are lost over time. Admittedly, given the proposed asymmetrical nature of recall in higher and lower portions of the neocortical-hippocampal axis for 'consolidated' material (i.e. recall mediated via the higher areas additionally depends on reactivating traces in the lower regions, but not vice versa) this would be difficult to investigate

[6]Two important questions about the cross-maze task that would help us understand how these findings fit with the retrograde memory literature remain unanswered. Firstly, there are no studies of retrograde amnesia in the plus-maze task. Secondly, once a place response is initially established, it is not known whether animals will switch eventually to a response strategy merely with the passage of time. Such data could shed light on the consolidation/semanticisation issue, specifically on the question of whether what is initially learnt by the hippocampus is 'transferred' wholesale over time to other regions, or whether extra offline 'training trials' due to endogenous rehearsal allow other regions to learn their own region-specific solution to a task.

empirically. Packard's studies also clearly demonstrate that whilst there may be parallel learning of the same task by both hippocampus and CP, the 'solutions' that these different systems find are different. The fact that the CP-based reponse strategy is used when control has switched from the hippocampus implies that the information initially acquired by the hippocampus is not transferred or taught to other areas, at least in this case.

Interestingly, there is evidence to suggest that learning in these competing memory systems can be controlled with intra-regional glutamate injections (Packard, 1999; Packard and Teather, 1999), which implies that there may be biological mechanisms for managing the development of and access to multiple traces. The amygdala can modulate the separate types of memory mediated by the hippocampus and caudate nucleus (Packard and Cahill, 2001), both directly via efferent pathways, and indirectly by mediating the effects of drugs and hormones. Emotional state can determine the relative use of memory systems – blocking or damaging the amygdala prevents the typical stress-induced impairment of hippocampal learning. Additionally, the removal of cholinergic input to the hippocampus enhances the selection of place strategy in a water-maze task (Bizon *et al.*, 2003). Furthermore, the instructions given on a probablistic weather-prediction task (whether to focus on learning pairs of stimuli, or to give classification judgements on each presentation for which feedback would be given) lead to changes in the relative activity of the MTL and basal ganglia in humans (Poldrack *et al.*, 2001). Several other factors favour the develop of place- over response-strategies including the use of correction methods and massed rather than spaced training. Therefore it seems plausible that for a given task, multiple traces that could mediate recall might co-exist, and that the most currently useful trace could be selected or enhanced as appropriate.

We currently do not know what determines which area dominates recall when more than one viable trace exists. The simplest scenario would be one in which there is competition between regions for control over effectors, with the 'strongest trace' – in some as-yet-undefined way – winning. It would then be easy to explain the findings I have just outlined: glutamate injections into, say, the caudate putamen might aid the development of traces that ultimately becomes stronger than the individual episodic traces stored in the hippocampus. For tasks on which the hippocampus becomes redundant over time without any new training, the relative strength of traces

outwith the hippocampus might be increasing because the hippocampal trace decays faster and/or because extra-hippocampal trace strength is increasing in absolute terms, and/or due to changing use of recall strategies (as discussed in section 6.2). Active processes may also play a part: the level of metabolic activity in the hippocampus of control animals is higher than that of animals that had been exposed to spatial training 25-d earlier (Bontempi *et al.*, 1999; Maviel *et al.*, 2004), suggesting that the reduction of hippocampal activation might depend on active inhibitory processes from the areas that have come to mediate recall, rather than simply a reduction in recruitment.

6.6 Interleaved learning and semanticisation

It is widely accepted that the cortex may need the hippocampus to teach it in order to avoid catastrophic interference. However, this potential problem would be ameliorated if detailed episodic traces are not transferred to the cortex from the hippocampus.

Catastrophic interference arises in neural networks because weights that have been trained on earlier unrelated inputs get overwritten by later information. The cortex is divided into anatomical/functional regions that are relatively isolated from each other (especially in the lower areas) and information is mapped topographically, so that similar representations will map to a similar set of cells thereby reducing the overwriting of old weights by significantly different information. Because similar traces will be represented by similar patterns of activity, at recall one would expect some "confusion" between closely related traces. However, if the cortex is specialised for storing semantic information — that is, information that is common to all occurrences of the same type of entity or event — and is not primarily a long-term store for unique episodic memories then, as suggested in chapter 3, recalling the "overlapping" elements across similar traces (rather than exemplars per se) is exactly what is needed. In other words, if the cortex represents semanticised traces, and not truly episodic events, then the potential problem of catastrophic interference will automatically be limited. There would be no reason why the cortex could not store information without a hippocampal teacher.

The information stored in the hippocampus is relatively more specific and less overlapping than that stored in the cortex, and is acquired more quickly; which would be expected to lead to high levels of interference. However, the more fully-distributed

sparse representations found in the hippocampus intrinsically reduce interference, as even apparently similar inputs to the hippocampus will be orthogonalised. The 'place cell' recording literature suggests that similar events may be represented very differently in the hippocampus, depending on the animal's intentions, recent experiences, and exact behaviour at the time of acquisition (Markus *et al.*, 1995; McEchron and Disterhoft, 1999). It is usually assumed that catastrophic interference is avoided by the hippocampus because it is a temporary memory store. In the semanticisation approach, although detailed memories are dependent on the hippocampus for their lifetime, most detailed information is assumed to decay relatively quickly, so in this scheme also, the hippocampus could be considered to be a temporary store for most of the information it initially stores.

As already noted, neurogenesis continues into adulthood in the hippocampus. Furthermore, new neurons are stabilised when hippocampally-dependent tasks are learnt (Gould *et al.*, 1999). Therefore, it is possible that there is a unique memory storage mechanism available in the hippocampus (and olfactory system) – namely, that new memories are stored in new neurons – which would also reduce interference between memories. It may be that memories that are salient enough to be stored for the long-term in the hippocampus have sequestered and stabilised highly plastic new neurones for their storage. (On the other hand, adult neurogenesis may simply be a repair mechanism for a highly plastic brain region.)

6.7 Reconsolidation and semanticisation

Memories are typically sensitive to disruption by various agents such as electroconvulsive shock or protein synthesis inhibitors for a limited period of time after initial acquisition. An accumulating body of evidence suggests that if *old* memories (whose continued existence is no longer affected by such agents) are reactivated they are then vulnerable to amnesic agents again, as if they require new consolidation or 'reconsolidation' (e.g., Misanin *et al.* (1968); Mactutus *et al.* (1979); Robbins and Meyer (1970); Judge and Quartermain (1982); Nader *et al.* (2000); Przybyslawski and Sara (1997); Przybyslawski *et al.* (1999)). That is, it may be the *state*, rather than *age*, of a memory that determines its susceptibility to disruption. The biochemical processes initiated after the retrieval of old memories appear to be rather similar to those occurring after initial acquisition, although re-consolidation seems less laborious than consolidation

in several ways[7].

Recent work on the hippocampus extends these findings. Land *et al.* (2000) found that 30-d after training on a signalled avoidance task whose acquisition is facilitated by an intact hippocampus, when standard lesions to the dorsal hippocampus no longer affect memory, memory is impaired if animals are exposed to the experimental context before the lesion. Similarly, Debiec *et al.* (2002) reported that rats that had received hippocampal infusions of anisomycin (a protein synthesis inhibitor) or hippocampal lesions, after exposure to the context in which they had learnt an obligatorily hippocampally-dependent context-fear association, showed a large drop in fear memory. This loss of memory was seen at all acquisition-reactivation intervals tested, even when the test and infusion occurred 45-d from initial acquisition, at which time contextual fear memories are apparently unaffected by hippocampal lesions.

The implication is that 1) even when memory traces are 'consolidated' and have completed initial protein synthesis-dependent processes, protein synthesis is necessary for their continued maintenance after reactivation; and 2) even when memories that were initially hippocampally-dependent have been established outside the hippocampus, reactivated memories depend on an intact hippocampus for their continued expression. That is, re-consolidation appears to occur both in local circuits and cross-regional connections.

Milekic and Alberini (2002) have however reported a gradient of memory vulnerability with systemic anisomycin causing memory deficits when administered after a reactivation trial on an inhibitory avoidance task at days 2 and 7, but not days 14 and 28. However, it is interesting to note that Milekic and Alberini (2002) used systemic administration of anisomycin, whereas the studies reporting a flat gradient used hippocampal lesions (Land *et al.*, 2000) or high doses of intra-hippocampal ani-

[7] In both cases, the amnesia induced by interference is time-dependent (Mactutus *et al.* (1982); Przybyslawski and Sara (1997); Przybyslawski *et al.* (1999)); and both processes require the integrity of NMDA receptors, β-adrenergic receptors and protein synthesis mechanisms, and are affected by ECS. However, there appear to be several differences. For example, the time window for interfering with re-consolidation may be shorter than that for initial consolidation (Mactutus *et al.* (1982); Judge and Quartermain (1982); Przybyslawski and Sara (1997)); reactivated (but not new) memories may recover spontaneously (Mactutus *et al.* (1982); Judge and Quartermain (1982), but see **?**); the onset of RA for disrupted new and reactivated memories may be different — Mactutus *et al.* (1982) reported 24-h and 4-h, respectively; and re-consolidation and consolidation may be differentially vulnerable to different amnesic agents (Mactutus *et al.*, 1979; Przybyslawski *et al.*, 1999). Retraining is also relatively easy after interference with reconsolidating memories: Przybyslawski and Sara (1997) found that rats regained previous performance with a minimal number of retraining trials, even when training was massed in one session, and on the difficult version of the task.

somycin (Debiec *et al.*, 2002). It is therefore possible that an incomplete blocking of hippocampal protein synthesis (which would be more likely with systemic administration) could underlie the apparently graded vulnerability to protein-synthesis inhibition, even if protein synthesis demands were present, though decreasing, at all intervals.

In the semanticisation view, the Ribot gradient may reflect the semanticisation of memories, and reconsolidation may be the mechanism by which cortical memories become semanticised. Each experience that triggers the recall of related memory traces may allow the 'sub-traces' that represent the semantic aspects that are common to these traces, and the associations between them, to be enhanced. As material becomes more familiar, fewer large-scale learning-related neuronal changes will be needed at each reactivation or re-exposure, as the repeated information will be largely similar to that already represented. It therefore seems reasonable that older memories and stronger memories (such as those resulting from exposure to three rather than to one footshock) must be exposed to longer reminder trials than more recent or weaker memories if reconsolidation is to be triggered (as reported by Suzuki *et al.* (2004)), because it would be adaptive to retain well-established or important memories. (The use of the same short reminder trial for all ages of memory, might also explain the observed lack of reconsolidation for older memories seen in Milekic and Alberini (2002)'s study.)

Mechanisms acting at reconsolidation appear to allow the wholesale up- or down-regulation of trace strength on the basis of the co-occurring degree of arousal (Sara, 2000). This would allow the strength of traces of already acquired information to be significantly changed in one trial even if it had taken many trials to acquire; perhaps in response to new information about already acquired knowledge. When a memory is recalled, reconsolidation may generally act to enhance trace strength. This would selectively arrest the natural decay of memories that have recently been recalled, and that are therefore known to be currently relevant. Reconsolidation might also make a memory trace more retrievable by virtue of it becoming associated with an increasing number of learning situations. Of course, explaining reconsolidation solely in terms of adding new second order information would not predict the loss of old information (Nader *et al.*, 2000), although it is possible that the whole trace needs to be put into a labile state for modifications to be made. The semanticisation approach thus motivates

a *reason* for re-consolidation: each reactivation of a trace allows the content of given traces, and the associations between them to be built-up, lost or modified, allowing the development of semantic stores.

Hippocampal traces for similar events necessarily index overlapping traces in the lower cortical regions that represent generic features of the events, in addition to linking in traces for the unique features of a given specific occurrence of the event. When the hippocampus is intact, a trial that depends on the recall of even well-established generic/semantic information probably additionally (1) triggers recall of many other related traces in higher parts of the neocortical-hippocampal axis that reflect more specific instances of the occurrance of that generic information – even though such information would be redundant for performing the current task; and (2) triggers the storage of new information by the hippocampus about that reactivation event, which again indexes those established generic/semantic representations. Assuming some form of attractor dynamics, it is plausible that together the nodes in these activated representations could form a reverberating attractor, and that interference with any component of that attractor (such as that caused by hippocampal lesions) might interfere with the re-establishment of components elsewhere in that attractor. Hippocampal damage might thereby disrupt the successful re-storage of traces elsewhere in an attractor of which it is again part.

Another possible explanation for the effect of hippocampal damage on tasks that had become independent of the hippocampus is that a damaged or abnormal hippocampus might output information that interferes with or overwrites existing useful cortical traces. In accord with this, partial hippocampal lesions have sometimes been reported to be more damaging than complete ones. Empirically, we do not currently know whether interference with protein-synthesis outside of the hippocampus after a reactivation trial in the absence of the hippocampus leads to a loss of hippocampally-independent memory: this is clearly a key question.

In conclusion, both the age and the state of a memory determines its susceptibility to interference after reactivation. Because memories tend to be semanticised over time, an old semantic memory will be less likely to undergo large-scale alterations on reactivation because on average there will be less 'left to learn' about the relevant semantic components of the new experience. Therefore post-reactivation learning will be relatively less affected by the inhibition of protein synthesis. A failure to register

any new information would also be less evident. On the other hand, the more recently a memory has been reactivated or stored, the more vulnerable it will be to interference, as it will be put in a labile state to enable modifications to be made. Age and state interact with factors such as the nature of the information being stored and how similar it is to what is already stored, as well as its salience, the amount of times it has already been recalled and current arousal levels, which can all affect the nature and extent of consolidation and reconsolidation processes.

Whether post-reactivation interference with the hippocampus affects memory also depends both on age and state of the memories evoked. A typical old semanticised memory that has not recently been re-experienced will be unaffected by hippocampal damage, however any remaining trace of details of that event will be hippocampally-dependent and therefore affected by hippocampal loss irrespective of recent memory reactivation. Clearly, recently acquired episodic memories will be dependent on the hippocampus, and therefore susceptible to hippocampal damage. Semanticised memories that have recently been recalled in the presence of an intact hippocampus will be affected by subsequent hippocampal damage or interference, perhaps because control of recall switches to the now impaired hippocampus, or because a damaged hippocampal trace interferes in some way with traces stored elsewhere in the brain to which it is connected. This 'age-and-state' view can be contrasted with the traditional view, which sees consolidation as a one-way process that can be completed for a given a memory.

6.8 The long term role of the hippocampus

In this chapter I have explored data and arguments illuminating the role of the hippocampus in the long-term maintenance of memories. In my opinion, neither Consolidation Theory nor Multiple Trace Theory can adequately explain this data. Instead, a view I have dubbed 'Gingell's Semanticisation Theory' (or GST) provides a more satisfying account. In this section I sum up the key points of this proposal.

6.8.1 Key features of Gingell's Semanticisation Theory

Any learning event causes activation throughout the brain from the sensory tranducers to higher regions. In GST when a learning event occurs, traces are laid down

throughout the neocortical-hippocampal axis. These traces form an attractor with both horizontal (within regions, and between connected regions at the same level) and vertical components (between different levels in the neocortical-hippocampal axis). Neurons at the hippocampal end of the hierarchy are potentiated most per event, receive massively convergent inputs, are heavily inter-connected, and are not arranged topographically; so they are therefore most able to represent novel associations. Effectively, after one (episodic) presentation of information, only the hippocampus will contain a robust trace of the unique, detailed and complex associations that typify episodic memories. Episodic events might leave a weak trace in areas of the cortex that were involved in experiencing them, but because there is incomplete horizontal connectivity between cortical areas, the episodic trace would be fragmentary at the cortical level and unlikely to be able itself to mediate recall. Extremely precise cues that effectively act to co-activate different but largely unconnected areas of the cortical trace might in a few cases be able to elicit some cortical recall of parts of such an episodic trace.

At the cortical 'lower' end of the hierarchy, strong traces gradually build up for the generic 'exemplar' fragments that are common to many experiences or rehearsals. A complete loss of information from certain modalities would result from damage to individual lowest regions. Intermediate levels in the neocortical-hippocampal axis register information on a continuum from generic to specific. Of course, memories differ in ways other than their generic/specific-ness, and different regions in the hierarchy 'specialise' in the processing and representation of different types of information at the different levels of complexity.

Each presentation of an event, even if it is very similar to an event previously experienced, will cause a new event trace to be laid down in the hippocampus. Whilst the cortical regions indexed by a hippocampal trace for similar events will be largely overlapping, the actual hippocampal representation of each index may be rather different, due to the more profound effect of subtle differences in the real-world presentation of even relatively similar events on the pattern of activity in the hippocampus, and internal orthogonalising influences. Most such hippocampal traces will decay very quickly.

Each level in the neocortical-hippocampal hierarchy can act as a binding zone for the co-recollection of traces stored at a lower level. Recall mediated by the highest regions therefore has the potential to elicit the most detailed and complex trace reacti-

vations, by indexing the most numerous and unique associations. Whether a reduced representation stored at a particular lower level is sufficient to mediate recall on a specific task (i.e. whether deficits will be seen in task performance after damage to a region higher in the axis, such as the hippocampus) depends on the complexity and nature of the information needed to perform the task.

6.8.1.1 Decay and modulation

The detailed, complex, supra-modal traces that are stored in the hippocampus remain hippocampally-dependent for their lifetime, which may be much less than the lifetime of the animal. Without further intervention, traces of recently acquired information will begin to decay soon after acquisition, albeit at widely different rates in different regions: decay is fastest at the top of the neocortical-hippocampal hierarchy. Most day-to-day episodic information (what did that women on the bus this morning look like? exactly how much did lunch cost on Tuesday?) is lost automatically and quite quickly if it is not salient when experienced, or made so by subsequent events soon after its acquisition. Even semantic information to which we have extended exposure (e.g., school learning) will be forgotton within 3 - 5 years if it is not learnt well, although if adequately encoded it may end in 'permastore' and be retained relatively undiminished for 50 years (Bahrick, 1984).

There are, however, many factors that modulate the strength of initial storage, and the subsequent fate of a trace. As noted in section 4.4, the hippocampus may be particularly sensitive to the modulation of the maintenance and trace strength of information automatically laid down there. In exceptional learning situations that are associated with very high activity in systems that mediate the effects of stress (e.g., the hypothalamic-pituitary-adrenal axis, or the sympathetic nervous system), traces of an event appear to be laid down more robustly than usual throughout the whole neocortical-hippocampal hierarchy, with an abnormally unprocessed sensory nature in the cortex. Extreme cases may produce the so-called 'flash-bulb' memories of post-traumatic stress disorder.

6.8.1.2 Graded retrograde amnesia depends on the semanticisation of memories

The hippocampus and other regions necessarily store different kinds of information on exposure to the same event, even when a non-hippocampal trace is built up through

endogenously-driven rehearsal rather than re-exposure to similar events. The learning rate and information convergence in an area fundamentally affects the kind of information that *can* be stored under given conditions. For repeated information, acquisition can proceed as outlined in the 'basic model' in which overlapping generic aspects of similar events or entities that are mapped topographically in the lower regions are strengthened.

Tasks that can be performed on the basis of relatively semanticised or generic information may become progressively less dependent on the hippocampus. Several possible factors contribute to the appearance of graded RA after hippocampal damage, but the most important mechanistically is that over time, recall of a memory elicits more semanticised information which does not depend on the hippocampus for recall. Several factors (that were discussed in detail in section 6.2) contribute to semanticisation:

- *Loss of detailed information*
- *Development of semantic traces*
  - May result from ongoing exposure to real-world events containing repeated features.
  - Offline learning or attractor replay may enhance the semantic components of memory. Only the core elements will be reliably reactivated across reactivations, and therefore only those elements will be enhanced and/or retained over time.
- *Change in recall strategy* The recall of older memories depends more than recently acquired memories on generic past memories. In humans this means a search through personal semantic information, and on knowing what one might plausibly have done.

Semanticisation should not be understood solely as a process by which memories are 'diminished' through a loss of information. An ongoing process throughout life that allows extraneous detail to decay from the memory stores; but also supports the development of representations of useful, relevant and more generalisable semantic information; identifies associations between such information; and enhances trace storage of such information in lower regions, would clearly be of great value. Furthermore, as already noted, the hippocampus and other higher regions can fundamentally

change what is represented in other areas: the hippocampus's ability to bridge temporal gaps and to oversee learning in several unconnected areas allows it to to make associations between information that other areas could not make alone in a typical learning situation.

6.8.1.3 When does semanticisation occur?

In the view of most Consolidation Theorists, the spontaneous activation of hippocampal traces of events during sleep or other offline periods is the main mechanism orchestrating memory consolidation in neocortical circuits (Marr, 1971; Buzsáki, 1989; Squire and Alvarez, 1995; McClelland *et al.*, 1995)), perhaps by providing an increased number of 'learning trials' to support slowly developing synaptic reorganisation (McClelland *et al.* (1995)). A similar mechanism could theoretically underlie the partial replication of information from higher to lower cortical levels as suggested by GST.

However, the literature focusing on the effects of sleep on memory performance is both vast and contradictory. The majority of studies have focused on rapid eye movement (REM) sleep, perhaps because it is electrically most like the conscious state. However, a recent review concluded that REM deprivation (REMD) studies are approximately equally divided between those that do and those that do not show a disruptive effect on learning/memory (Vertes and Eastman, 2000), and that the time windows for REMD effects where they exist are extremely variable (Smith, 1996).

In fact, slow-wave sleep (SWS), rather than REM sleep, is probably a better candidate for replay-aided memory processing, as the strongest coherent hippocampal reactivations (or sharp waves) occur predominantly in SWS (Penttonen *et al.*, 1997; Kudrimoti *et al.*, 1999). During a sharp wave, the dynamics of layers II & III of the EC may allow large ensembles of hippocampal neurons to alter the synaptic connectivity of neocortical circuits (Buzsáki (1998); Chrobak *et al.* (2000)); plasticity in the hippocampus is low (Leonard *et al.*, 1987), returning to its normal waking level during REM sleep (Bramham and Srebro (1989); Leonard *et al.* (1987)); and most growth hormone (a promotor of protein synthesis) is released during SWS (Hobson and Steriade (1986)). Taken together, this might create good conditions for information transmission from the hippocampus to the neocortex (Chrobak and Buzsáki (1994)).

Several studies have shown that cells in CA1 that were co-active when acquiring a task also tend to fire together in SWS sleep immediately after task acquisition (Wilson

and McNaughton, 1994; Skaggs and McNaughton, 1996; Qin *et al.*, 1995, 1997). Furthermore, traces of two separate experiences can appear together in the same recording session after trace acquisition (Kudrimoti *et al.*, 1999), which suggests that reinstatement is not solely dependent on the persistence of certain memory traces. The 'coherent' replay of pairs of cells in the hippocampus and other parts of the neocortical-hippocampal axis has also been reported (Kudrimoti *et al.*, 1999; Qin *et al.*, 1997; Shen *et al.*, 1998), in accord with the idea that the hippocampus co-ordinates activity in other regions. However, although seductive, the proposal that neocortical consolidation necessarily involves hippocampal reactivation and recoding has not been verified (Sutherland and McNaughton, 2000). On the basis of current evidence it is plausible that there is no single (hippocampal) site of origin for "replay". Cortical replay might therefore arise (at least in part) due to local attractor properties in the cortex, and not exclusively through being driven by the hippocampus; and the reactivation of hippocampal (and other MTL) patterns might in turn (at least partially) serve processes internal to those regions. Of course, when sufficient activation occurs in any one area, coherent activation throughout the whole of an attractor spanning the neocortical-hippocampal axis might be expected, if the neural-transmission environment allowed it.

In at least some cases, the learning improvement that occurs during REM sleep is only that which would have occurred if the animal had stayed awake (Karni *et al.*, 1994). Therefore it is possible that any recoding that occurs in the waking state also occurs during sleep. Similarly, the sharp wave oscillations characteristic of SWS can occasionally occur during quiet wakefulness or *type II* behaviour (Vanderwolf (1969)) which includes eating, drinking and grooming, but not active exploration. Kudrimoti *et al.* (1999) found that there is a comparable dependency of replay patterns on prior experiences both in an alert motionless state and SWS, thus sleep itself is not necessary for memory trace reactivation. The strong hypothesis that there is a learning process exclusive to sleep has not yet been satisfactorily demonstrated. The difference may merely be quantitative – during sleep there are fewer competing inputs, and therefore on average perhaps more opportunities for reactivating extant attractors and/or making learning-related changes. To date, reconsolidation has only been demonstrated after real-world reminder triggers, so the evidence for post-acquisition memory modification processes specific to the awake state is in fact more robust.

In GST, the semanticisation and recoding of memories could plausibly occur as a result of any process that results in trace reactivation, whether due to recall triggered by specific real-world events, conscious 'free-wheeling reminiscence', or unconscious reactivation during different brain states. It seems reasonable that there would be quantitatively and qualitatively different re-structuring processes occurring during these different reactivation processes, both in terms of the extent and location of trace modification, and perhaps also of the nature of the relationships uncovered and modified between known information. Exposure to real-world information should be the dominant force in any biological memory system. It is possible, say, that the reactivation of memory attractors results in changes in trace strength only outside the hippocampus in certain stages of sleep (perhaps to different extents for trace components at different levels in the neocortical-hippocampal hierarchy), whilst conscious recall or re-experiencing might additionally lead to trace strength changes in the hippocampus, and the laying down of new episodic information. Such proposals remain speculative until definitive data is available. However, at present in GST, memory-modification processes occurring on re-experiencing an event, and after reactivation of existing traces, are both similarly assumed to result in the strengthening of overlapping parts of existing traces.

6.9 Summary and conclusions

This chapter has focused on understanding the role of the hippocampus in the long-term recall of information. Broadly speaking, tasks that are obligatorily-dependent on the hippocampus for their acquisition remain indefinitely dependent on the hippocampus for recall. There is no evidence that detailed context-specific episodic memories are transferred to the cortex, or that task-dependent allocentric information can become independent of the hippocampus. Tasks whose acquisition is merely facilitated by the hippocampus, on the other hand, may become independent of the hippocampus over time. However, the information recalled as the neural basis for recall changes is qualitatively different.

The change in nature of the information recalled from different periods results from a semanticisation of memories. Semanticisation reflects the increasing involvement in recall over time of areas representing more generic information. This may result from the loss of hippocampally-dependent detailed traces, an increase in the

ability of semantic regions to mediate recall, or a difference between the strategies used to elicit old and new memories. Figure 6.2 schematically summarises the contribution over time of the hippocampus and cortex to the recall of a typical episodic memory. Since the information recalled by the hippocampus and the cortex is always qualitatively different, the red hippocampal lines represent the recall of detailed information, whereas the blue semantic lines represents the recall of semantic information. Initially the hippocampus dominates overall recall and many details can be recalled; whereas later on, the neocortex dominates and information elicited is mainly semantic.

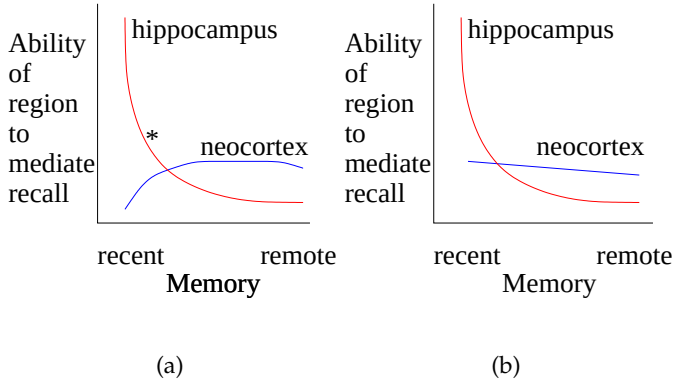

Figure 6.2: The cortex takes over the recall of memories as they age. The cortical ability to mediate recall may improve (a) absolutely, or (b) only relative to the hippocampus.

Most information stored in the hippocampus is lost very quickly. This initial information is rarely measured in neuropsychological studies requiring the recall of specific important events. Information that is maintained for longer periods by the hippocampus still decays, albeit at a slower rate. Let us suppose that the information that is retained and that may be recalled on typical tests starts at the values represented by the asterisk (subfigure 6.2a). Then the loss of detailed information over time in normal subjects would be represented by the portion of the hippocampal line to the right of the asterisk. In contrast, typical episodic memory tests that do not control for semanticisation effectively measure combined hippocampal and cortical output. Studies measuring the recall of true episodic information find flat retrograde amnesia after hippocampal damage, as this information is always dependent on the hippocampus.

In subfigure 6.2a there is an absolute increase over time in the cortex's ability to

mediate recall. Subfigure 6.2b shows a decreasing cortical recall ability over time. Cortical recall could also be equivalent at all time points. In all cases, the net effect is a relative increase in the ability of the cortical areas to mediate recall over time compared to the hippocampus. Current empirical data does not tell us which of these scenarios is correct.

Different regions throughout the neocortical-hippocampal hierarchy acquire information in parallel on a given learning task. Traces at each level support the acquisition and recall of information that is different in type, and that may be able to support recall on a given task to different extents. After the initial acquisition of some tasks, these regions become redundant in sequence from highest to lowest. The proposed semanticisation and freeing of specifically hippocampal dependency may progress down through the hierarchy. If a given region has become redundant for recall, it does not necessarily imply that the trace stored at that level has decayed. Indeed, multiple traces of different levels of complexity and specificity may be retained for different purposes; and may tend to be accessed under different cuing conditions. Different tasks have different representational requirements in terms of the type and complexity of information stored and recalled. This not only governs which regions are necessary at storage, but also whether regions can then become redundant and how long this is likely to take.

Semanticisation may generally act to identify and preserve the most important features of memory, and perhaps re-store them in a form that makes them more accessible. Most context-specific information is unimportant and represents noise in the identification of important world regularities. Reconsolidation may reflect the processes of semanticisation. When information that is related to ongoing real-world events is recalled, semanticisation processes may act to arrest the normal decay of these still relevant traces, enhance components that are common to the reactivated traces and traces for the new information, and make links between other aspects of active traces. The temporal extent of retrograde amnesia after hippocampal damage reflects semanticisation, rather than a process by which information initially stored in one area is transferred to another.

In GST, semantisation could theoretically occur in any situation involving the reactivation of memory traces. This would include the exposure to semantic and non-context-specific information that underlies normal learning in the awake state; and

the presumed reactivation of related information thought to occur on exposure to known cues in reconsolidation studies. The importance of offline learning perhaps during sleep remains to be conclusively demonstrated; as does a central role for the hippocampus in any such replay.

In a recent paper, Meeter and Murre (2004) identified several key findings relating to the long-term role of the hippocampus, and provided a table listing explantions offered by CT, MTT and a semanticisation approach (p22). Table 6.1 similarly summarises the explanations offered by GST for these important findings.

6.9.1 Relationship to other proposals

GST shares some features with consolidation theories and with the Multiple Trace theory. The points of overlap and disagreement have been spelt out in some detail in section 6.1. In brief, like CT, GST proposes that some memories become independent of the hippocampus over time. However, in GST, the memories recalled via the hippocampus and via other areas are qualitatively different. Like MTT, GST proposes that some memories remain indefinitely hippocampally-dependent; although trace replication plays no role in GST explanations.

My proposals share a central assumption with the Complementary Learning Systems approach (O'Reilly and Norman, 2002), namely that there is not a strict division of labour between brain systems. Differences in the basic functional features of the hippocampus and cortex, such as the degree of topographical mapping or orthogonalisation and the speed of learning, lead naturally to differences in their relative ability to perform particular tasks. However, they perform overlapping functions.

Milner (1989)'s views are also highly relevant to GST. Milner drew a distinction between quickly potentiating 'soft' synapses in the hippocampus; and 'hard' cortical synapses that are little affected by single bursts of activity, but hold changes almost indefinitely once they have been made. He suggested that occasional reactivations throughout the whole assembly would refresh soft synapses and increment hard ones. Later related views posit an 'indexing' role for the hippocampus (Teyler and Discenna, 1986), implying that the role of the hippocampus in a memory assembly is distinct from that of other components. Whilst it can be useful to think in terms of a hippocampal indexing function (and indeed I have used the term throughout the thesis), it obscures the similarity in the role of all the connections in the neocortical-

Finding	Explanations offered by GST
Graded RA after H damage	Older memories progressively more semanticised
Ungraded RA after H damage	'Obligatory' memories tested
RA for semantic information	i) Recent episodic memories may support semantic recall. ii) Episodic and semantic memory not categorically distinct
Semantic dementia	Recent episodic memories depend on fewer old semantic fragments
More H activity for recent than remote memory	'Facilitated' recent memories are more dependent on H than remote
H activity same for recent and remote memory	i) 'Obligatory' memories tested. ii) Storage of new H trace
Sequential regional redundancy, Izquierdo *et al.* (1997)	Progressive semanticisation of information
Reversible blocking of H impairs retention, Reidel *et al.* (1997)	Interference with maintenance of trace in H
Deficient LTP in cortex produces faster forgetting, Frankland *et al.* (2001)	Information cannot be stored in cortex
Immediate blocking of NMDA receptors in CA1 impairs later memory, Shimuzu *et al.* (2001)	Interference with maintenance of traces in H

Table 6.1: Summary of explanations offered by GST for some of the key findings relating to the hippocampus and retrograde amnesia as outlined by Meeter and Murre (2004). H = hippocampus. 'Facilitated' memories are those whose acquisition is aided by the presence of an intact hippocampus. 'Obligatory' memories are those whose acquisition is obligatorily-dependent on the hippocampus.

hippocampal assembly. This may have contributed to the idea that the hippocampus is necessarily the prime initiator of offline learning.

Clearly, GST also shares tenets with other semanticisation approaches. Cermak and O'Connor (1983) first noted that episodic memories become more fact-like and less event-like over time, but this idea has received surprisingly little attention. Nadel and Bohbot (2001) have recently focused on the idea that remote and recent memories are qualitatively different, and that recall mediated by the hippocampus and by other areas is qualitatively different. However, they have not yet incorporated these ideas into MTT.

The most detailed approach to semanticisation to date is that of Rosenbaum *et al.* (2001). Rosenbaum *et al.* (2001) argue that spatial detail and autobiographical information (because of its dependence of spatial context information) is indefinitely dependent on the hippocampus. In their view, semanticisation is a loss of context dependency in a relational sense. This contrasts with GST in which semanticisation results from a loss of the least robustly stored detailed specific information, or the least reliably reactivated information; which includes but is not limited to spatial contextual information. Rosenbaum states that 'as [semantic memories] become increasingly independent of context, extra-hippocampal regions are required for their recall' (p190, Rosenbaum *et al.* (2001)). This implies that recall is initially mediated by the hippocampus alone; and that information is then transferred to the cortex. This constrasts with GST which posits the involvement of the cortex in recall mediated by the hippocampus, and independent learning and semanticisation over time in the cortex. Rosenbaum and colleagues indicate that the initial memories 'break down' in some way, but do not suggest a mechanism by which semanticised representations might be built up. Thus GST extends these ideas by showing how semanticised traces might come to dominate recall through enhanced learning of repeated generic information and the automatic loss of detailed information from the hippocampus.

Chapter 7

A simple model of episodic and semantic learning

In this chapter I report on simulations with a neural net model that provide a 'proof of concept' for the ideas presented in other chapters. The focus of investigation is on how the different hippocampal and cortical components of the model interact with the acquisition of random 'episodic' events, and with 'semantic' events that share features with other events.

The findings provide clear support for the idea that episodic and semantic information can be usefully construed as representing positions on a continuum of memory types distinguished by speed of learning and the amount of specific, multi-modal detail needed to demonstrate recall. Support is also found for proposals on the relative importance of the hippocampus and cortex in the rapid acquisition of information and the acquisition of complex multi-modal information; and the effect of existing knowledge on new learning. Another key finding is that whilst recall for all events is initially dominated by the hippocampus, the cortex takes over as a memory ages. However, recall via the cortex is more robust for generic semantic components of events, and is progressively poorer for the 'episodic' details. Thus older memories become semanticised. Replay-aided learning also primarily benefits semantic sub-components at the expense of non-repeated elements.

7.1 Introduction

Many of the distinct functional features of the hippocampus and cortex result from trade-offs between systems designed for learning specifics and for extracting generalisations. The hippocampus is optimised for the rapid incidental encoding of detailed, novel, high-order information. Similar incoming information is represented in orthogonalised hippocampal traces which preserves specific information and reduces the brain's over-riding tendency to be biased by already acquired information. There is a wide consensus that the sparse dentate gyral inputs to CA3 may orthogonalise inputs to produce minimally overlapping representations for different episodes; and the recurrent connections of CA3 may provide a powerful auto-associative memory.

The cortex on the other hand learns slowly and incrementally on exposure to information, and information is mapped topographically, so that repeated exposure to information leads gradually to the strengthening of traces for repeated generic elements. The slow learning rate protects important well-established information from being disturbed without sufficient 'evidence' that such information has really been superseded. At progressively lower regions of the cortex, direct long-range connections or chains of connections are increasingly uncommon and difficult to form anew. Local cortical areas can however acquire new information about locally represented associations relatively quickly.

Several widely accepted ideas about the nature of information processing in the hippocampus and cortex inform my model. Many of these assumptions are shared with standard consolidation theories, and with MTT (see also Nadel and Moscovitch (1997)'s list, p223). Other key features derive naturally from GST (bracketed items indicate theories sharing the given view; CT = consolidation theories, MTT = Multiple Trace Theory, GST = Gingell's Semanticisation Theory):

1. The hippocampus rapidly and automatically encodes all experienced/attended information. (CT, MTT, GST)

2. The cortex automatically encodes all experienced/attended information, though less robustly than the hippocampus. (some CT, MTT, GST)

3. Information is encoded sparsely and non-topographically in the hippocampus; and topographically in the cortex. (CT, MTT, GST)

4. A hippocampal trace acts as an index to neocortical neurones representing information acquired in the same episode. (CT, MTT, GST)

5. The entire hippocampal-neocortical ensemble constitutes the memory trace for an episode. (some CT, MTT, GST)

6. The detail of memories is lost very quickly unless a memory is deemed important and receives robust encoding (GST).

7. Memory reactivation (however it is triggered) reliably activates only the core components of memory, and leads to the relative strengthening in cortex of core memory components, i.e. semanticisation (GST).

8. Repeated experiences with similar real-world information similarly benefits cortical traces representing overlapping semantic sub-components (GST).

9. Semanticisation and the extraction and re-representation of progressively more generalisable information is an ongoing process not tied to the acquisition of particular new memories (GST).

The neural network model presented in this chapter incorporates these features, and shows that certain characteristics of memory and learning arise from these features. In brief, the neural network model consists of (1) an input 'cortical' component which maps information topographically and learns at a low rate, and forms associations between locally represented information more easily that long-range associations; and (2) a smaller, more quickly learning, 'hippocampal' component that stores orthogonalised traces with equal facility for all information. The processing units in the model, or 'nodes', represent highly abstracted neurones or groups of neurones. At training, a set of patterns are presented to the cortical layer and learning takes place in the several different weight sets throughout the net. After each training trial, partial patterns are applied to the cortical layer, and the performance of the net and its sub-components at pattern completion is tested.

The net is fully connected, although sparse connectivity is more biologically plausible. In similar nets, the level of connectivity has been found to have little effect on the pattern of results observed (e.g., Murre (1999)). It does however affect capacity which is not the focus of these simulations. Patterns in the hippocampal component are sparser than those in the cortical component in accord with known biological facts,

but this feature has little effect on the pattern of behaviour seen. A k-winners-take-all strategy models the effect of inhibition between units in a layer.

My model thus belongs to a well-established class of models of hippocampal-cortical interaction. The main difference between my simulations and others is in the nature of the training and testing employed, and in the way the data is analysed and interpreted. Random episodic events and different types of semantic events that share features are applied to the net. Although there is nothing inherent to a single instance of an event that can distinguish it from events of other types, the net comes to distinguish different sub-components of information over training. Recall for episodic and semantic information, and for recently acquired and older information, becomes differentially dependent on the hippocampal and cortical components.

The simple model provided makes minimal assumptions, and cannot provide quantitative fits to empirical data. However, the findings clearly support the plausibility of proposals put forward throughout this thesis.

7.2 Network architecture

The model is implemented within a connectionist framework, based on a simple associative net. Input layer L1 – 'the cortical component' – consists of several sections, each with the same number of units. L2 – the 'hippocampal component' – consists of one section of units, whose size is less than the total size of L1 (see Figure 7.1). The model is highly abstracted from any "real" biological system, but sits squarely within the Cognitive Modelling tradition, in which a few simple input/output units arranged in layers can be said to represent cortex and hippocampus.

Recurrent connections within sections of L1 represent local connections within associative cortex, whilst recurrent connections in L2 represent connections within the hippocampus. Connections between sections in L1 represent long-range connections between cortical regions. The interconnections between units in L1 and L2 represent connections between cortical regions and the hippocampus, such as those through the parahippocampal regions.

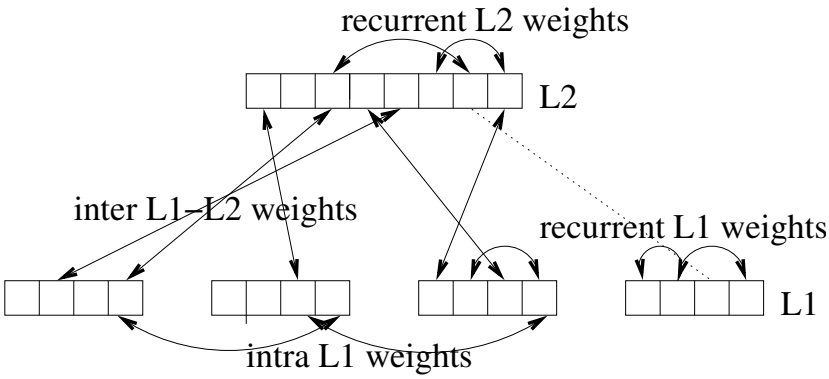

Figure 7.1: Basic 2-layer net. The net is fully connected, with four separate sets of weights as shown. For clarity, only a few connections are shown in the figure; and the weights are asymmetrical despite their representation.

7.3 Implementation of the model

Activation rule

To compute the binary valued activation (Y_i) in a node i, firstly calculate the raw activation y_i that increases with its net input:

$$y_i = \sum_{j=1}^{n} X_j w_{ij} \tag{7.1}$$

where X_j is the binary activation value of node j, w_{ij} is the connection weight from node j to node i, and n the number of nodes sending inputs to y_i via w_{ij}.

The binary activation Y_i is then:

1 if y_i is one of the k highest valued members of $[y_1, y_2, y_3 ... y_N]$

0 otherwise

where N is the number of units in a section.

There is therefore no fixed threshold for activation, and inhibition is simulated using a k-winners-take-all arrangement.

The values of k_1 (for sections in L1) and k_2 (for the single section in L2) are fixed at the beginning of training.

7.3.1 Training phase

All weights are initialised to a small random positive value between 0 and 0.1.

There are 4 sets of weight matrices: W_1 (for connections in sections in L1), W_{11} (for intra-sectional connections in L1), W_{12} (for connections between L1 and L2), and W_2 (for recurrent connections within L2). Each of the weight matrices have their own learning (λ) and decay (δ) rates. These are fixed at the beginning of training, and conform to the following order in the main simulations, unless otherwise specified.

$$\lambda_2 \geq \lambda_{12} > \lambda_1 > \lambda_{11} \tag{7.2}$$

$$\delta_2 \geq \delta_{12} > \delta_1 > \delta_{11} \tag{7.3}$$

Each set of weights is trained independently. Training of the weights in W_1 and W_{11} proceeds on the basis of the training pattern applied to L1.

For each training event applied to L1, a random pattern is generated in L2 with k_2 active units. W_{12} is trained on the basis of the associated activity in L1 and L2. W_2 is trained auto-associatively on the pattern in L2.

Learning rule

The learning rule is a simple Hebbian rule, with global weight decay.

Weights connecting co-active units are incremented. Forgetting is simulated by reducing all non-co-active connection strengths on each time step. The learning-associated change on each time step in a weight connecting units i and j (Δw_{ij}) is equal to:

$$\Delta w_{ij} = \lambda(1 - w_{ij})Y_iY_j - \delta w_{ij}(1 - Y_iY_j) \tag{7.4}$$

Equation 7.4 ensures that weights are kept within the interval [0-1].

7.3.2 Test phase

In the test phase, a partial pattern is applied to L1. A partial pattern is generated by switching d randomly chosen nodes that were active in the stored pattern from 1 to 0. The degree to which this partial pattern is completed by passing the pattern through the various sets of weights is then assessed. Units that are active in a test pattern are clamped at 1, and are always selected. Error is measured by computing the Hamming distance between the recalled pattern and the desired pattern.

Recall from W_1: Raw activations in L1 are calculated on the basis of W_1, and the k_1 most active units per section are set to 1 with other units set to 0.

Recall from W_{11}: Raw activations in L1 are calculated on the basis of W_{11}, and the k_1 most active units per section are set to 1 with other units set to 0.

Recall mediated by W_1 and W_{11} combined: is found by summing the raw activation arrays produced by each set of weights, before digitising them by setting the k_1 most active units in each section to 1, and the others set to 0.

Recall mediated by W_{12}: depends on propagating activity from L1 to L2 via W_{12}; setting the k_2 most active units in L2 to 1 and the others to 0, and propagating activity back to L1 via W_{12}. The k_1 most active units per section are set to 1, the others set to 0.

Recall mediated by W_2 and W_{12} combined: W_2 weights cannot be tested independently of W_{12} for their ability to support recall in L1. The combined recall elicited from W_{12} and W_2 initially proceeds as for W_{12}, resulting in a raw activity pattern in L2. This pattern is auto-associated using W_2, to create a new pattern of raw activation values in L2. The k_2 most active units are chosen in L2, and recall proceeds as above for W_{12}.

Whole net: The raw activations in L1 produced from the combined performance of W_1 & W_{11}, and from W_2 & W_{12} (as detailed above) are summed; the k_1 most active units per section are set to 1, the others set to 0.

7.3.3 Modulation of learning

Memory modulation in the model notionally models the effect of systems that increase arousal concommitant with experiencing an event or that act before an event trace has decayed to increase the initial strength with which an event trace is encoded. Memory modulation is engaged in the model by increasing the value of λ in one or more weight set for a sub-set of training event types. For ease of comparison with other presented data, memory modulation applies to either all episodic or all semantic information, rather than to a sub-set of such events.

7.3.4 Offline learning

When off-line learning is engaged in the model, it occurs after each training trial. The number of events to be trained offline after each training trial (p) and the number of the most recent patterns over which these events are to be selected (r) is fixed. A training event is randomly selected from the r most recently presented patterns. The L2 pattern associated with this training event is activated, and activity allowed to propagate to L1 via W_{12}. The k_1 most active units in the raw activation array in L1 are set to 1, and the others set to 0. W_{11} weights are then trained on this pattern as for a normal training trial. (Note that this pattern may not be one of those originally trained.) This off-line training process is repeated p times.

7.3.5 Training and testing schedule

Figure 7.2 summarises the model's schedule of operation.

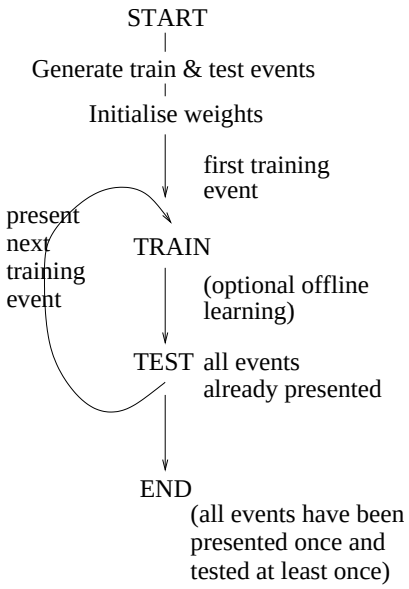

Figure 7.2: Training and testing schedule

7.3.6 Representation of events

A key feature of this model is the distinction that is made between episodic and semantic training events, as experience with events accrues.

Training items

All training items are represented by binary strings of size nm, with $k_1 m$ active units, where n is the number of units in a section in L1, m is the number of sections in L1, and k_1 is the number of active units in a section in L1.

There are two types of training item:

- *Episodic items* are represented by randomly generated strings. Items 1a and 1b in Figure 7.3 are examples of possible episodic events.

- Each *semantic item* belongs to a category of semantic events. A fixed and substantial proportion of a pattern in a category is identical with all other instances of that category. In the majority of simulations presented in this chapter, 2/3rds of the pattern is shared. The other 'variable' units are randomly generated for each instance of that category. In figure 7.3, items 2a and 2b are instances of one semantic category (and overlap on the first 3 sections) whilst items 3a and 3b are instances of another semantic category (and overlap in sections 1, 2 and 4).

  Two parameters are needed to specify the population of semantic events – the number of different categories of semantic event, and the number of instances of each presented to the net.

The 'episodic' item 1b in figure 7.3 could in fact also be identified as an event from the same semantic category as items 3a and 3b. This illustrates an important point: in this model there are no inherent differences between episodic and semantic events that are evident in a single pattern. All semantic information arrives in an 'episode'. What defines an event as episodic or semantic is whether there are many other similar events with substantially overlapping portions of activity. Thus if items from the same semantic category as that of items 3a and 3b were presented to the net along with item 1b, then item 1b could be considered to be a semantic item. If only events from the semantic category of items 2a and 2b were presented to the net, item 1b would be considered to be episodic.

Using a different conceptualisation, each presentation of an instance of a semantic event from the same category is effectively the re-presentation of the same event, but with some random noise. The pattern of activity in each episodic event, on the other hand, is not related in any predictable way to that of any other episodic event.

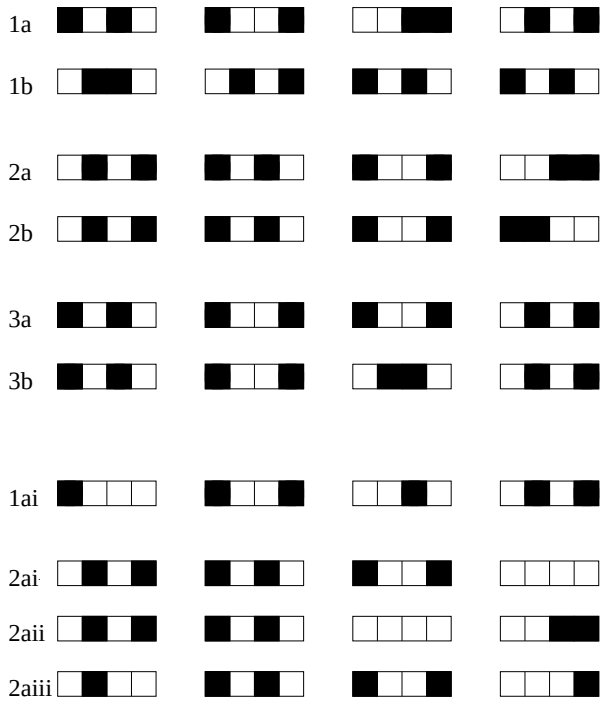

Figure 7.3: Examples of training and test items to be presented to L1. For easy visualisation the ones and zeros of the binary strings are represented here as black or white squares respectively. **1a** and **1b** show different possible episodic training events. **2a** and **2b** show possible semantic training events in the same category of semantic event as each other, as do **3a** and **3b**. Test events are formed by deleting active units from a given training event. **1ai** shows a possible test item associated with the training item in **1a**. Items **2ai-iii** are possible test events associated with the semantic event shown in **2a**. Test event **2ai** shows a test event where all the units omitted are from the "episodic" (non-overlapping) part of the semantic training event, **2aii** shows a test event where all the units omitted are from the "semantic" part of the semantic training event; and **2aiii** shows a test event with a mixture of semantic and episodic units omitted.

For simplicity, in the simulations presented the variable 'episodic' part of a semantic item is confined to a single section.

Test items

Test items are created by deleting d active elements from a given training pattern, so there are $k_1 m - d$ active units in a test pattern.

Either all d elements are deleted from one section, or $d/2$ elements are deleted from

each of two sections. There are therefore five possible types of test item, which are labelled to show (i) whether the event is *Episodic* or *Semantic* and whether (ii) deletions occur in a *variable (var)* or repeating *semantic (sem)* section. Note that episodic events have variable sections only. This labelling convention is used in the figure legends.

- The deleted units in episodic test events can come from either:
  1. The same section – *Episodic/var*
  2. From two different sections (e.g. item 1ai) – *Episodic/var var*
- The deleted units in semantic test events can come:
  1. From the same variable section (e.g. item 2ai) – *Semantic/var*
  2. From the same semantic section (e.g. item 2aii) – *Semantic/sem*
  3. From one variable and one semantic section (e.g. item 2aiii) – *Semantic/sem var*

7.3.7 Parameter values

After initial explorations, size- and activity-related parameter values were fixed (see Table 7.1). These 'standard values' are used in simulations except where otherwise stated. Approximately 80 training events are presented to the net on each of the main simulations[1]. The basic pattern of findings is similar across a wide range of parameter values. These values allow the phenomena of interest to be observed in a reasonable run-time.

7.3.8 Visualising the data

This code was written in the Matlab Programming Language, and simulated using Matlab Version 6.5.1.199709, Release 13 (The Mathworks, Inc).

In the graphs presented in this chapter, the x-axis plots the number of training trials intervening (ITT) between an event's initial presentation to the net for training,

[1]Test items come in multiples of twos (episodic items) and threes (semantic items) to ensure that equal numbers of the different types of test items are presented in the test phase. As one new test item is presented each time a new training event is presented to the net, there must be an identical number of training and test items. Therefore, depending on precise combinations of episodic and semantic events, and semantic types and instances of events, the total number of training events may not be precisely 80, although it is always at least 80.

Parameter	Value
Number of units in L1 (nm)	40 units x 3 sections = **120**
Number of units in L2	30 units x 1 section = **30**
Number of units active in pattern in L1 (k_1)	10 x 3 sections = **30** (25% of units)
Number of units active in pattern in L2 (k_2)	4 x 1 section = **4** (13% of units)
Number of nodes omitted to create test items (d)	**4** (either 4 in one section, or 2 each in 2 sections)
Learning rates [$\lambda_{11}\lambda_1\lambda_{12}\lambda_2$]	[**0.01 0.1 0.6 0.9**]
Decay rates [δ_{11} δ_1 δ_{12} δ_2]	[**0.001 0.05 0.3 0.8**]
Number of training events	– if episodic only: **80** – if semantic only: sc=**9**, si=**9** – if mixed: episodic: **44**; sc=**5**, si=**9**

Table 7.1: Parameters used in the main simulations presented in this chapter, unless otherwise stated. sc = number of categories of semantic event; si = number of instances of each semantic category.

and being tested. Therefore, low ITT represents 'recent memory' and high ITT refers to 'remote memory'.

The y-axis plots the performance of components of the net in completing partial L1 patterns. Performance is calculated as the number of nodes d omitted to form a test pattern minus the error, which is half the Hamming distance between test and target pattern. Percentage performance is calculated by dividing performance by d, and multiplying by 100. Some of the axes were mistakenly labelled in terms of absolute performance and some in terms of the percentage performance. However this does not affect the shape of the lines plotted.

The performance of different sets of weights can be tested separately (with the exception of W_2, as noted above), or in combination. Sub-plot captions indicate the weight set(s) producing the plotted performance. I refer to the combined performance of W_{11} & W_1 as the 'cortical component'; and of W_{12} & W_2 as the 'hippocampal component'. Hippocampal plots can be considered to represent the ability of the hippocampus to mediate recall without help from the cortex, and the cortical plots represents

the cortex's ability to mediate recall in the absence of the hippocampus. A comparison of such plots at a particular ITT value thus reflects the relative contribution that each component is making to the recall of particular types of information at that time-point after acquisition.

The graph legends refer to the five different types of testing events enumerated in section 7.3.6. I shall refer to the pattern of units within a single section of a training event as a sub-pattern. Each *Episodic/var* test item tests recall of a random sub-pattern that was acquired in association with other random sub-patterns. Each *Episodic/var var* item tests recall of two random sub-patterns that were acquired in association with random sub-patterns in other sections. *Semantic/var* items test recall of a random sub-pattern that was paired at acquisition with sub-patterns that repeat across different training patterns. *Semantic/sem* items test recall of a sub-pattern that repeats across different training items, and that was acquired in association with another repeating sub-pattern and a random sub-pattern. *Semantic/sem var* items tests the recall of a random sub-pattern and a sub-pattern that repeats in different patterns, that have been acquired in association with another repeating sub-pattern.

7.4 Findings

7.4.1 Control data

I first examine the standard deviations of error in the model, and explain the rationale for truncating the x-axis for the data plots.

7.4.1.1 Standard deviation of error

Events are presented sequentially, and recall for all events is tested after the presentation of each event. Each training event is associated with a single test item of a given type (one of 2 types for episodic training items, and one of 3 types for semantic items). Therefore, the maximum number of times that test items of a given type can be tested is the total number of training items minus the position of presentation of the first training item associated with a test item from that event type. If, for example, the first presentation of an episodic event associated with an 'Episodic/var' test item was on training trial 6 and there were 80 training trials, these test items could contribute data points only up to ITT = 74. Only the first training item can be tested the maximum

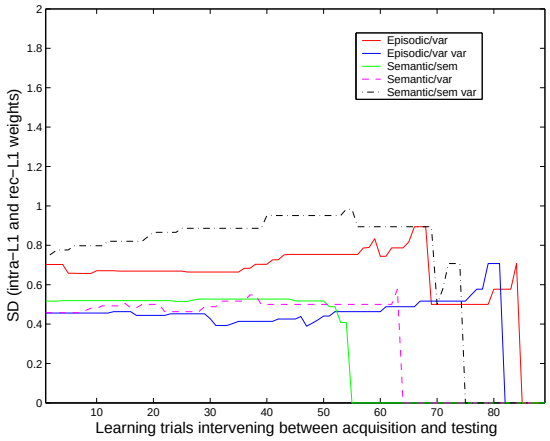

(a) Cortical recall

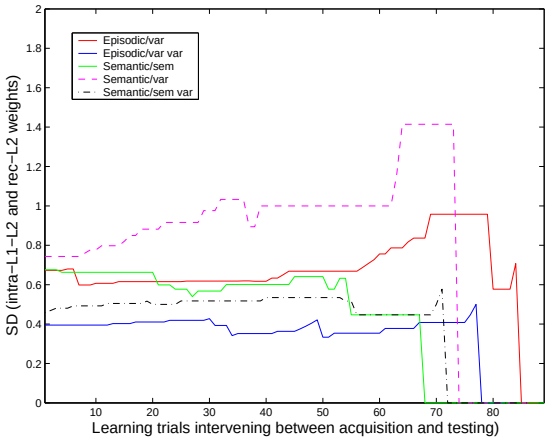

(b) Hippocampal recall

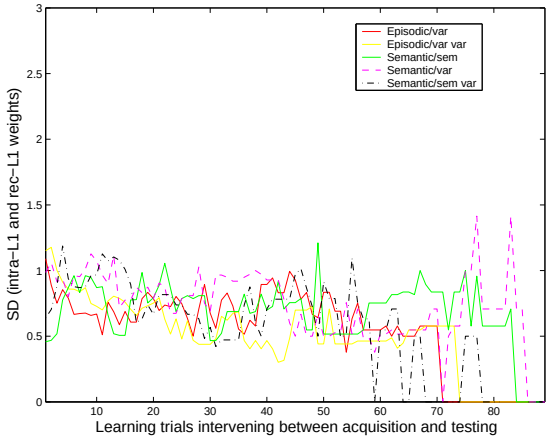

(c) Cortical recall

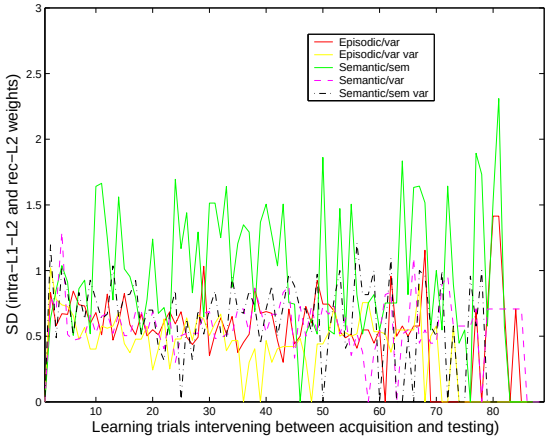

(d) Hippocampal recall

Figure 7.4: a & b: Standard deviation of error when learning and decay rates are zero. c & d: Standard deviation of error on a typical run.

number of times, and therefore only this type of event will contribute a data point for the maximum ITT value. Therefore, the standard deviation of error rises with increasing ITT, as the plotted values are averaged over progressively fewer data points (see figure 7.4, all sub-figures); and drops to zero at different ITT values for different event types (Figure 7.4a & b show this most clearly).

Standard deviation of error will be low on a simulation involving learning if

weights are performing well, as there will be no error (e.g., Figure 7.4d, ITT = 0). For this reason, there is an interaction between standard deviation and the different types of weights and test items when learning is allowed (compare Figure 7.4c & d; and compare the different test types within Figure 7.4d). In general, standard deviation is highest for 'Semantic/sem' items when acquired and tested by the W_{12} weights. This arises because similar overlapping sub-patterns in L1 are associated with random patterns in L2; and the learning and decay rates are relatively high in these weights.

In the main data plots reported in this chapter, the x-axis is truncated at half the maximum ITT because the increasing variance in the data at high ITT makes the means of these data unreliable.

7.4.1.2 Performance with no learning

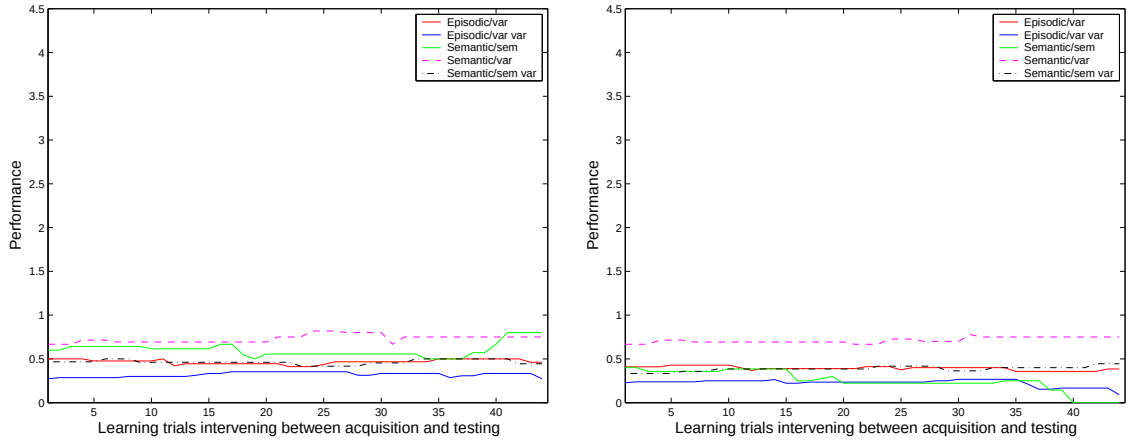

(a) Cortical recall

(b) Hippocampal recall

Figure 7.5: Recall performance with no learning. For this set of parameters, expected chance performance is approximately 0.66 for items in which units are deleted from only one section.

As expected when there is no learning in the model, that is, when λ and δ are zero in all weights, performance on test items is close to chance for all weight sets (Figure 7.5).

The probability of correct unit choice is higher for test items in which units are deleted from only one section. Therefore 'Episodic/var var' and 'Semantic/sem var' items show lower chance performance on average than other test items (Figure 7.5).

7.4.2 Learning and decay rates

These simulations investigate the effect of globally changing λ and δ values in all weights.

7.4.2.1 Changing learning and decay rates

The absolute and relative values of λ and δ in a given weight set affects immediate recall performance on a pattern when ITT = 0, how quickly performance deteriorates on a trace after initial acquisition, and the slope of performance deterioration as ITT rises. Figure 7.6 shows the effect of globally changing learning and decay rates.

When both λ and δ are high, events can be stored very robustly, and remembered accurately for a short period of time; by both the hippocampal and cortical components. That is, there is good performance at low ITT (subfigures 7.6a & 7.6b). Performance drops rapidly as ITT increases, because δ is high and stored information is quickly lost. Performance on the different types of test items is similar in all sets of weights.

With intermediate λ and δ, initial performance at low ITT may not be perfect because new traces are not made sufficiently distinct from existing traces (subfigures 7.6c & 7.6d). The ability to recall recently acquired information interacts with the nature of the weight matrices underpinning recall (compare different performance values at ITT = 0 in subfigures 7.6c & 7.6d; and 7.6e & 7.6f). Information is lost at a slower rate (that is, the deterioration in performance follows a shallower gradient, compare subfigures 7.6a& c; and b & d at low ITT), with overwriting playing a relatively greater role in memory loss than when δ is very high.

When sparser patterns are stored in the net at intermediate λ and δ rates, performance drops off less quickly as ITT increases (compare subfigures 7.6c & 7.6d, with subfigures 7.7a & 7.7b) because overwriting is reduced.

At very low λ and δ, recently stored events are not much more likely to be remembered than older events, as initial performance is close to chance, and decay is very low (subfigures 7.6e & 7.6f). Close to chance performance is seen for event recall mediated by the 'hippocampus'; and cortical recall is good only for semantic sub-patterns.

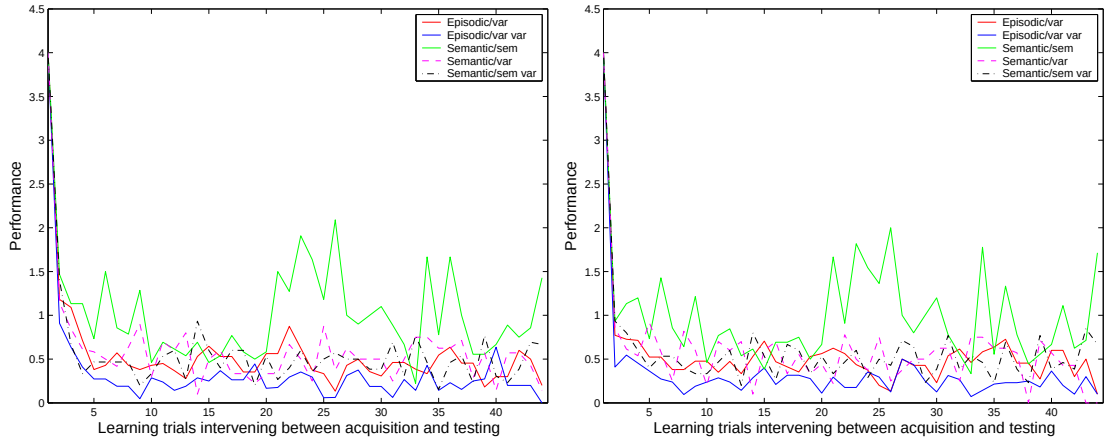

(a) Cortical recall

(b) Hippocampal recall

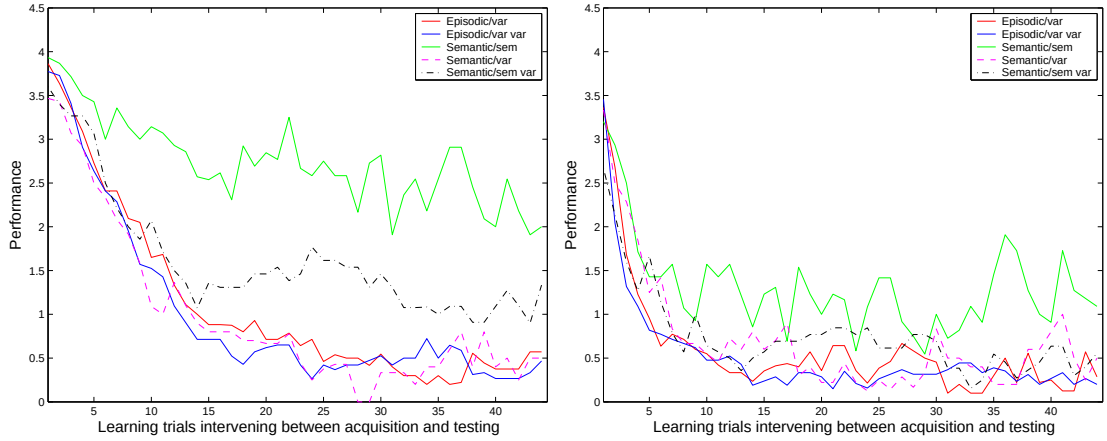

(c) Cortical recall

(d) Hippocampal recall

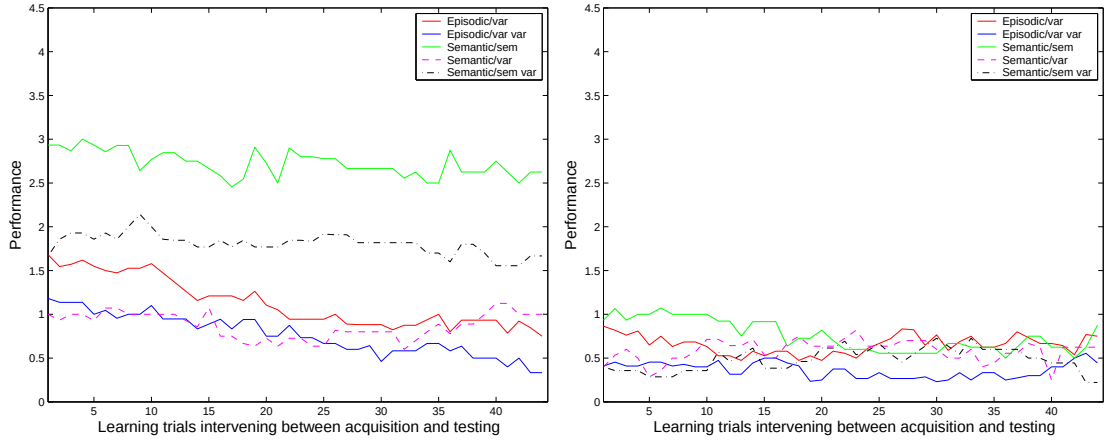

(e) Cortical recall

(f) Hippocampal recall

Figure 7.6: Learning and decay rates equal in all sets of weights. In (a) & (b) all $\lambda = 0.9$ & all $\delta = 0.9$; in (c) & (d) all $\lambda = 0.1$ & all $\delta = 0.1$; and in (e) & (f) all $\lambda = 0.01$ & all $\delta = 0.01$.

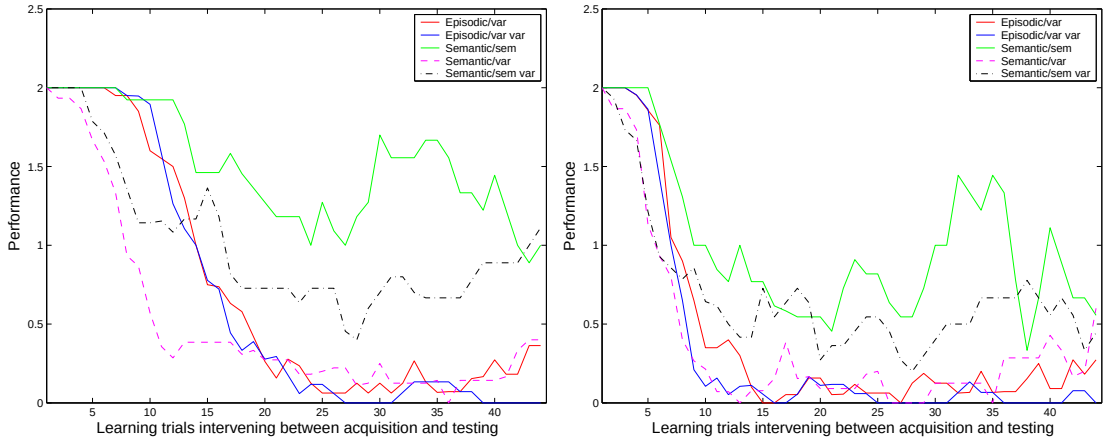

(a) Cortical recall

(b) Hippocampal recall

Figure 7.7: Performance drops off less quickly when sparser patterns are stored in the net. All $\lambda = 0.1$; all $\delta = 0.1$. $k_1 = 4$; $k_2 = 2$; $d = 2$.

7.4.2.2 Memory performance in hippocampal and cortical components

Figure 7.6 shows very clearly that even when λ and δ take the same values in all areas, hippocampal and cortical components behave very differently.

The most obvious feature is that the performance for different types of test item is more separated in the cortical component than the hippocampal component (compare subfigures 7.6c & d, and e & f). This arises because in the hippocampal component, similar L1 patterns are represented by orthogonalised L2 patterns, whereas information is mapped topographically in the cortical component. Overwriting by later patterns in the hippocampal component is always detrimental to performance, as there is no consistent relationship between elements of patterns representing different events. However, overwriting in the topographically organised cortical component can lead to better performance on the repeated semantic sub-patterns of events, as more robust representations build up for these features.

The initial performance at ITT = 0 reflects the combined effect of the recent learning trial (i.e. the most recent event for which ITT = 0), together with previous accumulated learning. Therefore this value is also affected by the degree to which overwriting is destructive or constructive.

The different weight sets also contain different numbers of connections, which par-

tially determines their susceptibility to overwriting.

7.4.2.3 Learning and decay rates: conclusions

The values of λ and δ interact with the nature of weight sets in hippocampal and cortical components to produce characteristically different patterns of recall behaviour. In general, high λ and δ support good immediate recall peformance, but poor long-term retention; whereas low λ and δ support relatively poor initial performance, but better long-term retention.

Memory studies have established that the hippocampus can acquire complex episodic information rapidly, but that such information decays rapidly; whilst the cortex cannot support robust one-trial learning of episodic information, but can gradually build up information about regularities of the world with time. In order to capture such behaviour, the cortical component of our model should produce a pattern of behaviour similar to that of subfigure 7.6e – that is, it should employ low λ and δ; whilst the hippocampal component should produce behaviour similar to subfigure 7.6b or d – and employ high λ and δ.

Preliminary investigations show that the general patterns of behaviour observed in the model are robust under a wide set of parameters, if λ and δ are higher for weights in the hippocampal component than in the cortical component. As the purpose of this model is to provide a 'proof of concept' rather than to exhaustively explore the range of possible λ and δ values, I have simply selected a set of λ and δ values that follow the pattern laid out in equation 7.3. The values indicated in table 7.1 are used in all simulations, unless otherwise stated. I in no way wish to imply that these values are the 'best' or the most realistic.

7.4.3 The role of different sets of weights

Broadly speaking, the weight matrices associated with the hippocampal component (W_{12} &W_2), and with the cortical component (W_1 & W_{11}) behave in a similar fashion. This is unsurprising because the cortical matrices both have low λ and δ, whilst the hippocampal matrices have much higher λ and δ; and whilst learning in cortical weights reflects associations within an L1 pattern, hippocampal learning reflects associations between the L1 pattern and a random L2 pattern. Many of the subfigures presented in this chapter plot the combined performance of the hippocampal component

and of the cortical component, rather than individual performance of the associated weights.

However, there are subtle differences in the performance of different weight matrices within the hippocampal and cortical components and in the role they play in supporting the acquisition and retention of different kinds of information. Figure 7.8 shows the contributions to performance mediated by each weight set on a typical run. This figure provides reference data for plots later in this chapter.

7.4.3.1 Recurrent connections within the hippocampus (W_2)

The ability of W_2 weights to support pattern completion in L1 depends on the integrity of weights in W_{12}. Therefore, increasing trace storage strengths or reducing forgetting or overwriting in W_2 alone has little effect on L1 pattern completion performance. (These factors do affect pattern completion within L2, but L2 pattern completion is not directly explored here.)

W_2 acts to 'clean up' the raw activity pattern created by W_{12} in response to the application of a partial test pattern to L1. Because W_2 weights (and the associated W_{12} weights) show a rapid loss of information, this 'clean up' adds a very sharp recency effect to the hippocampus's performance, allowing perfect immediate recall even when the bi-directional W_{12} connections alone produce only 50% correct performance (compare subfigures 7.8d & e).

7.4.3.2 Connections between the hippocampal and cortical component (W_{12})

W_{12} weights bi-directionally link the cortical and hippocampal component. W_{12} can therefore act indirectly to support associations between information stored in separate sections of the cortex. However, because λ and δ are high in W_{12} this associative information is available only transiently.

Because W_{12} learns associations between L1 training patterns and random L2 patterns, W_{12} representations benefit from repeating semantic sub-patterns only through chance associations that co-opt weights that have already been used in the representation of that sub-pattern.

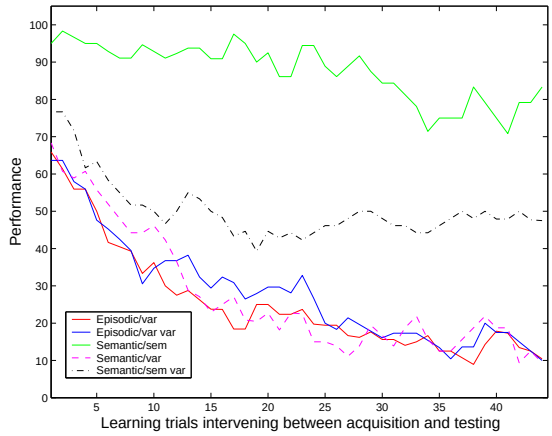

(a) Recurrent-L1 weights

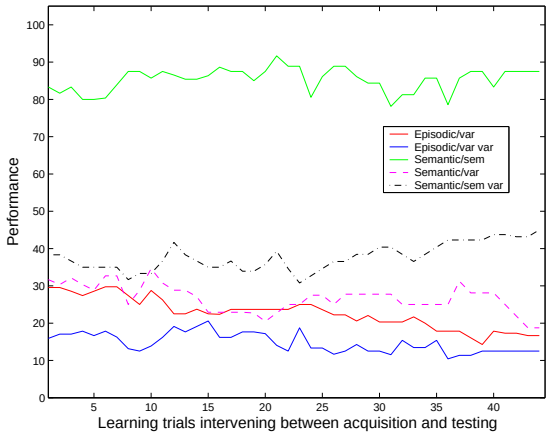

(b) Intra-L1 weights

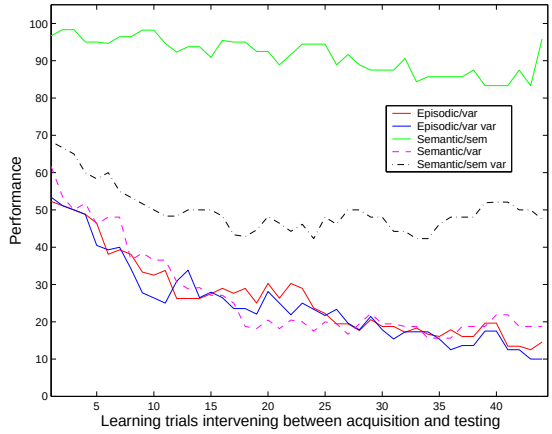

(c) Intra-L1 and recurrent-L1 weights

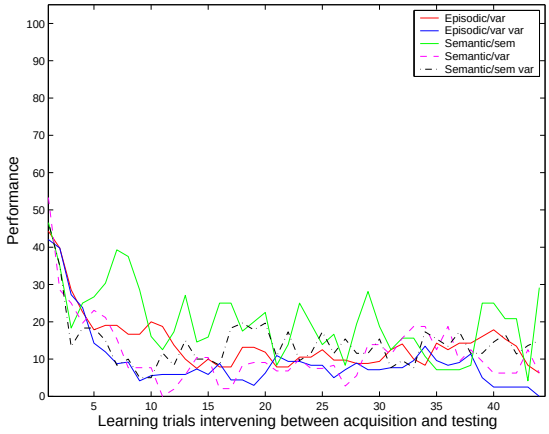

(d) Inter-L1-L2 weights

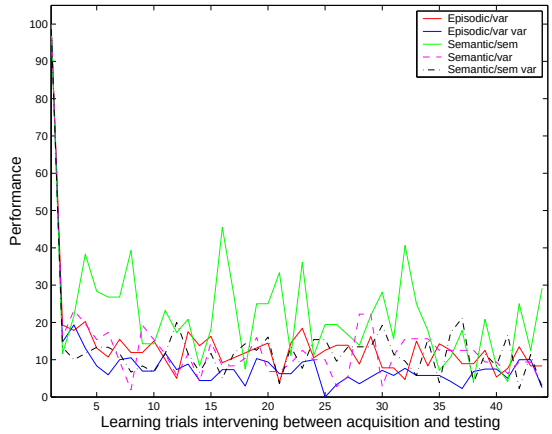

(e) Inter-L1-L2 and recurrent-L2 weights

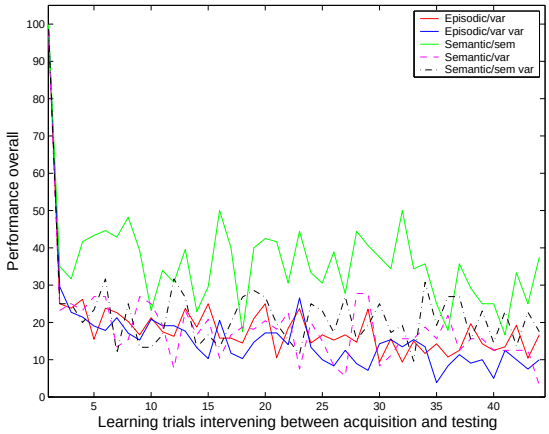

(f) Whole net

Figure 7.8: Performance on a typical simulation using standard parameters

7.4.3.3 Connections between different cortical sections (W_{11})

Performance mediated by W_{11} weights shows the lowest recency effect (that is, performance at low ITT is not much better than that at high ITT), and the greatest separation of performance for different types of test items (subfigure 7.8b). Both features result from low λ and δ, as weights cannot robustly store one-trial information but can gradually increment over time on exposure to repeated sub-patterns.

The 'constructive interference' due to repeated exposure to semantic sub-patterns interacts with δ and the conditions of testing to determine whether an increase, decrease or flat maintenance of performance for given test types is seen over the course of a simulation for W_{11}-mediated recall.

7.4.3.4 Recurrent connections within sections of the cortex (W_1)

W_1 weights associate units within individual cortical sections. Because sub-patterns are auto-associated without noise in W_1, semantic learning of local sub-patterns is very good. However, W_1 performance also shows a pronounced recency effect due to overwriting and decay (subfigure 7.8a). The W_1 weight matrix contains relatively few connections compared to the other matrices, increasing its vulnerability to destructive overwriting.

W_1 weights cannot register associations between sub-patterns occurring in different cortical sections. Clearly, if the partial pattern presented at testing entirely lacks active units in one of the sections, recall mediated via W_1 will produce chance performance for that section. Subfigure 7.9a shows such behaviour: W_1-mediated performance is at chance for test items in which $d = k_1$ for test items in which units are deleted from only one section. Subfigure 7.9c & d shows normal patterns of recall via W_{11} weights and the hippocampus when $d = 10$. Subfigure 7.9b shows typical recall when $d < k_1$ for reference.

7.4.3.5 Whole net

Recall performance for the net overall is dominated on average by weights that contribute most strongly to raw activations in L1. Generally speaking, for recently acquired information, the hippocampus contributes most to whole net output. As a memory ages, cortical weights take over recall (subfigure 7.8f).

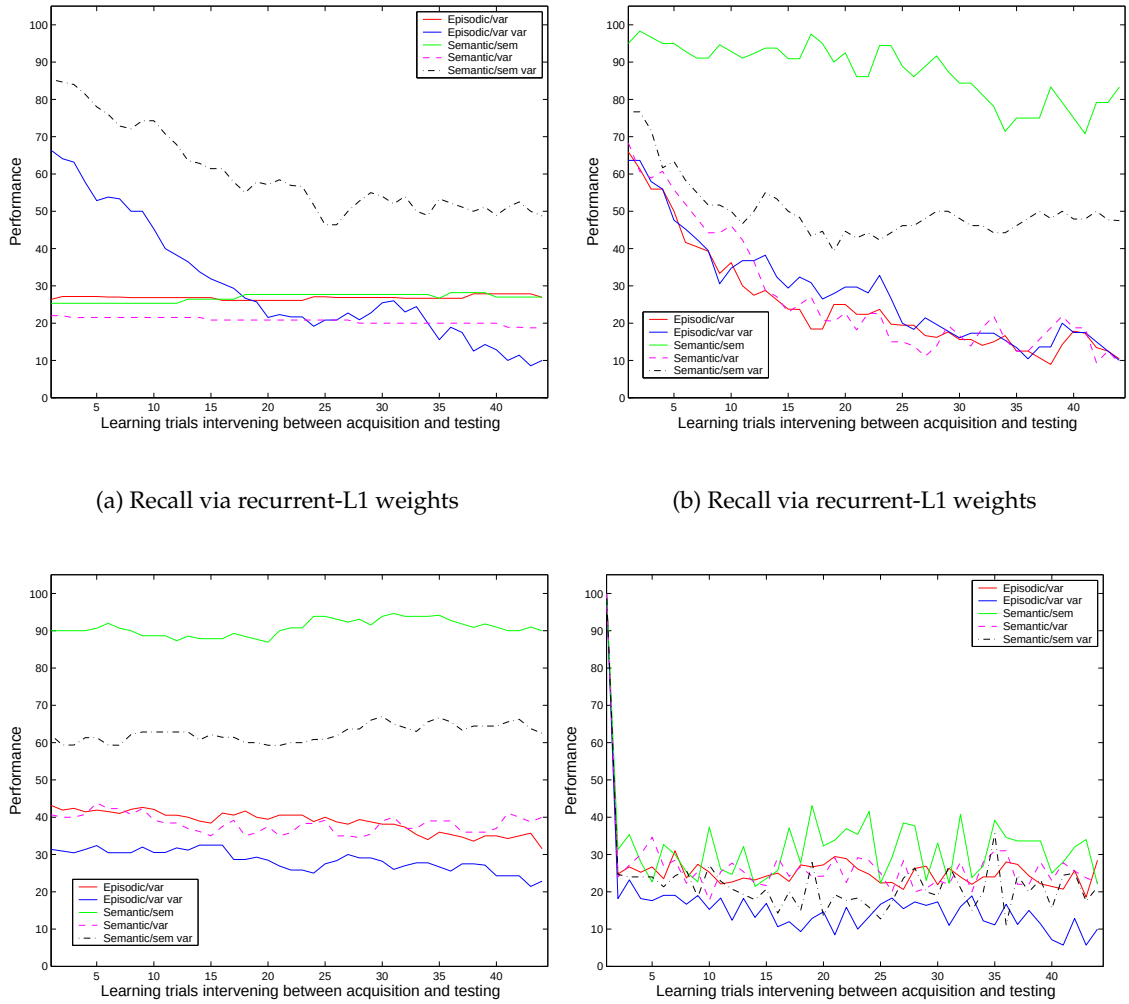

(a) Recall via recurrent-L1 weights

(b) Recall via recurrent-L1 weights

(c) Recall via intra-L1 weights

(d) Hippocampal recall

Figure 7.9: Local cortical connections cannot support associative recall of information represented in distant cortical sections. (a), (c) & (d): No partial pattern in one of the sections of cortex at testing (i.e. all local units deleted at test), d = 10; (b): Partial pattern in all sections at testing.

7.4.3.6 The role of different sets of weights: conclusions

Many differences in the behaviour of the hippocampal and cortical components with respect to episodic and semantic events can be traced back to a component's relative tendency to separate patterns or to topographically map similar information. Recall

performance for different test items is relatively more similar in the hippocampal component, as repeating semantic sub-patterns are represented in the hippocampus by randomly chosen patterns. In contrast, incremental learning leads to improving recall performance for semantic sub-patterns in the cortex.

Additionally, cortical learning distinguishes between local and long-range cortical associations, whereas the hippocampus is equally able to represent associations between any two nodes in L1. The rapid acquisition of associations between information represented in different areas of the cortex is therefore hippocampally-dependent, although the cortical component can acquire such information slowly. Local connections can acquire associations between locally represented information relatively quickly.

The choice of λ and δ enhances the a priori learning characteristics of hippocampally- and cortically-associated weights.

7.4.4 Learning different kinds of information

In these simulations, I examine recall performance for episodic and semantic events separately. Interleaving episodic events with semantic events effectively adds noise to the learning of semantic events, although this has little effect on the recall of episodic events (not shown). I also show how the model captures proposals on the relative importance of the hippocampus and cortex in the fast acquisition of information (discussed in section 4.2) and the aquisition of locally-convergent information (discussed in section 4.3).

7.4.4.1 Acquisition and maintenance of episodic information

Figure 7.10 shows recall performance for episodic items alone. The hippocampal region shows very good one-trial learning of this novel random information, with high performance when ITT = 0; but rapid decay of such information to chance.

In contrast, the cortical component shows poorer performance in initial recall of such information; but the rate of decay is lower so that at high ITT, a little of this information can still be recalled via the cortex. That is, as a memory ages, the cortex takes over the recall of 'episodic' information, although progressively poorer pattern completion (i.e. fewer details) are recalled over time.

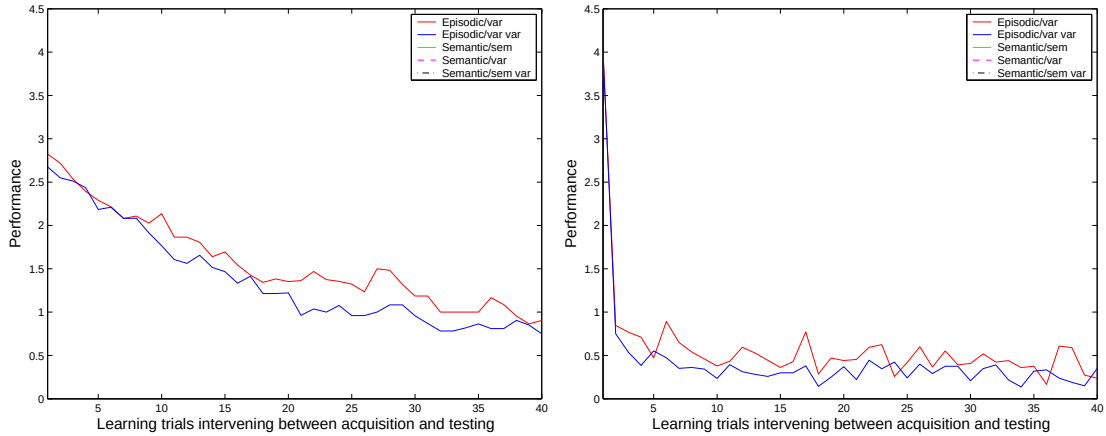

(a) Cortical recall

(b) Hippocampal recall

Figure 7.10: Recall performance on episodic information.

7.4.4.2 Recall performance on semantic information

A population of semantic training events is specified by the number of different categories of events and the number of instances of each. Figure 7.11 shows the effect on performance of presenting different populations of semantic training items to the net.

When a few types of semantic events with many instances of each are presented to the net, cortically mediated recall of semantic sub-patterns is very good (subfigure 7.11a). Since forgetting is low, robust representations for the overlapping semantic sub-patterns can be developed over time in cortical regions on exposure to different instances of the same category of event. 'Semantic/sem' items benefit the most as pattern completion is tested only within a semantic sub-pattern; 'Semantic/sem var' items benefit approximately half as much as they test pattern completion for a semantic sub-pattern and a random sub-pattern; and the random sub-patterns tested in 'Semantic/var' do not benefit at all. Since earlier test trials cannot benefit from the *subsequent* repetition of overlapping information, the first test of a repeated semantic sub-pattern is not necessarily going to produce the highest performance. This explains the increase in average performance seen for Semantic/sem items as ITT initially increases in subfigure 7.11a.

As the number of categories of semantic event increases and there are fewer instances of each category (figures 7.11 a & b; c & d; e & f), interference (overwriting)

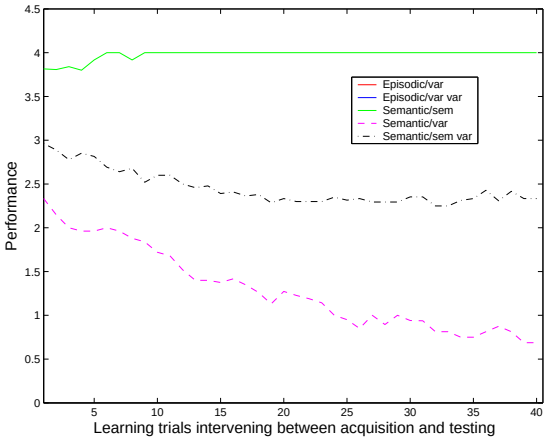

(a) Cortical recall

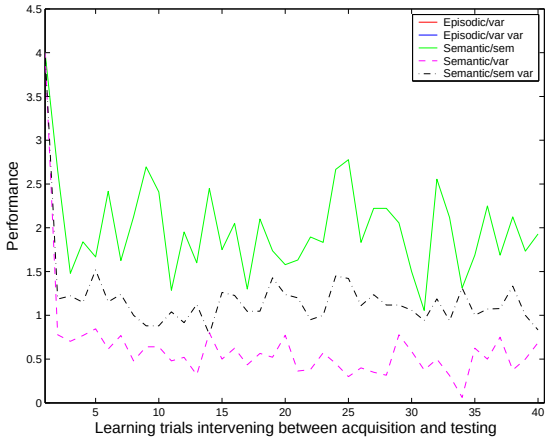

(b) Hippocampal recall

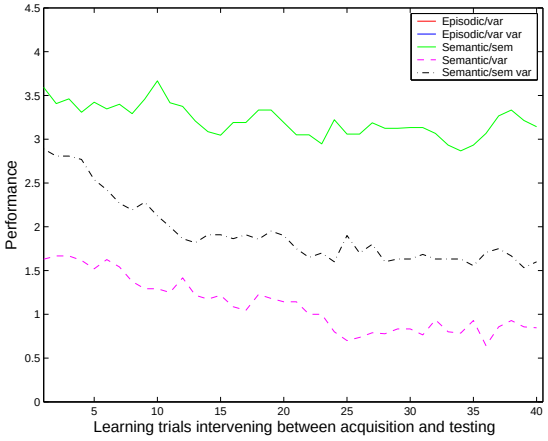

(c) Cortical recall

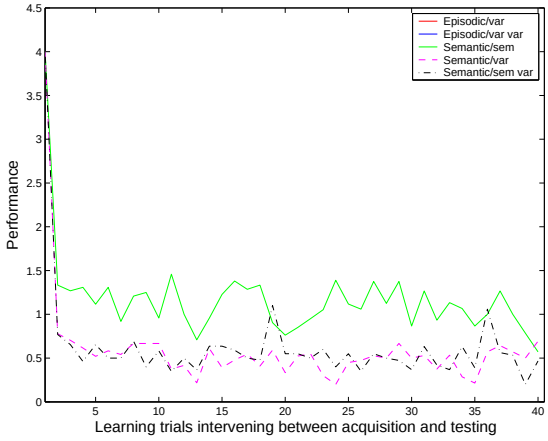

(d) Hippocampal recall

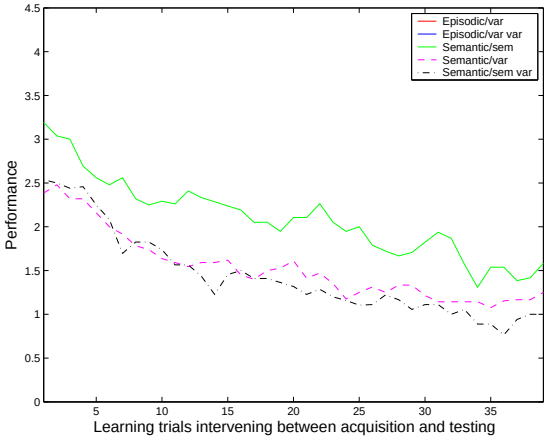

(e) Cortical recall

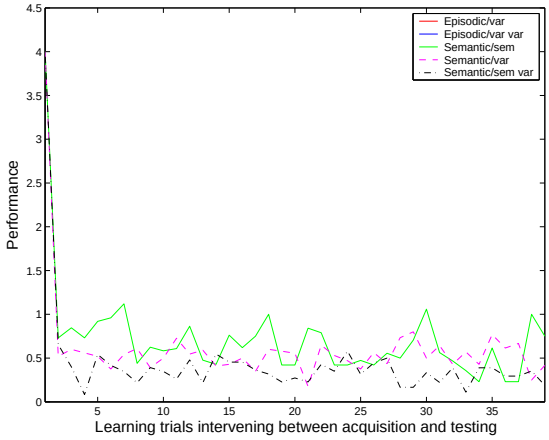

(f) Hippocampal recall

Figure 7.11: Recall performance for semantic information. (a) & (b) Semantic categories: 3; instances: 27 (c) & (d) Semantic categories: 9; instances: 9, and (e) &(f) Semantic categories: 26; instances 3.

between weights representing one semantic sub-pattern and those representing another is greater. In addition, on average there will be more trials intervening between re-exposure to repeating sub-patterns, so forgetting will be greater. These factors offset the gains from re-exposure to the same sub-patterns, and leads to a slight loss of semantic information in the cortical component as ITT increases (subfigure 7.11c). However, the very pronounced performance benefit for semantic over random information is retained, due to low forgetting and slow incrementation of weights across learning trials. The benefit of repeated learning is also indicated by the reduced gradient of performance decay as ITT rises for items with semantic sub-patterns.

When semantic events are derived from many categories and contain few instances of each, performance in the cortical component becomes progressively more similar for each type of test item (subfigure 7.11e & f), and progressively more like the recall of episodic information (compare subfigure 7.11e and subfigure 7.10a). If there was only one instance of each category, events would effectively *be* random episodic events, as they would not share sub-patterns with other traces except by chance. In accord with this, performance on the 'Semantic/var' items deteriorates in all cases at a similar rate as that seen for episodic information, as only recall for the random sub-pattern is being tested.

Recall performance mediated by the hippocampal component is relatively similar irrespective of whether there are many repeated instances of a given category, or only a few. Recall is similarly initially very high for all test items, followed by a very rapid loss of all information as ITT increases (subfigures 7.11b, d & f). The variance on 'Semantic/sem' items is highest in the hippocampal component, especially when there are few categories of events. A relative improvement for 'Semantic/sem', and to a lesser extent 'Semantic/sem var' events is seen as there are more instances of categories, but this is less evident than in the cortex. This results from the chance re-using of weights already used to represent information about a given semantic sub-patterns.

7.4.4.3 Convergence of information and associative learning

In section 4.3, I concluded that associative learning within one modality was not hippocampally-dependent, cross-modal learning between different types of information was facilitated by the hippocampus, and learning associations between supramodal information and other information was obligatorily hippocampally-dependent.

I also suggested that there was a hierarchy of learning areas from the neocortex, through the MTL areas, to the hippocampus that can learn and represent information of different complexity and acquire it at different rates.

Although the 'hierarchy' implemented in my model consists of only two regions, these distinctions can be seen in the behaviour exhibited by it. Let us posit that each section in the model's cortical component represents a cortical module in the brain that processes information about a particular modality. Associative learning within one modality can be performed by W_1 weights in the absence of the hippocampal component. For example, in subfigure 7.9b, W_1-mediated pattern completion within a section (i.e. the association of locally represented information) is good. Associations between information in two modalities cannot be performed by locally projecting W_1 weights (as we have seen, subfigure 7.9a), and W_{11} cannot mediate recall of such information without extensive experience (W_{11}-mediated pattern completion is good only for sub-patterns that are repeated across different training events, subfigure 7.8b). The hippocampal component is thus responsible for supporting acquisition of cross-modal (cross-sectional) information at the normal speed of acquisition (good recall at low ITT for 'Semantic/sem var' and Episodic/var var' associations, subfigure 7.8e): acquisition would be slower in the absence of the hippocampal component.

In accord with the data discussed in chapters 3 and 4, in the absence of the hippocampus the acquisition of associations between information that is not represented in local cortical regions requires more learning trials (compare 7.8c & e). Sufficiently complex cues that effectively provide partial patterns in poorly connected sections of the cortex (compare subfigures 7.9a & b) lead to better recall, as they reduce the dependence on long-range cortical connections. In addition, previous learning affects the ability of the cortical connections to acquire information in a given number of learning trials, although it has little effect on hippocampal learning (compare subfigure 7.8b & e at low ITT).

7.4.4.4 Speed of acquisition of information

In section 4.2, I argued that the fast acquisition of complex, novel information was generally hippocampally-dependent. The acquisition of random episodic information in the model is necessarily one-trial, whereas semantic sub-patterns may be acquired over several trials. Only the hippocampal component is capable of supporting

good recall of one-trial random novel information, even for immediate recall (compare episodic recall in subfigure 7.8b & e). The hippocampus supports good one-trial learning for the 'episodic event' of which semantic repeated information is a part. Such learning might support episodised 'semantic' recall in conditions such as semantic dementia in which long-range cortical connections are breaking down. Local cortical connections can support some recall of one-trial random information (subfigure 7.8a), although as we have seen this is only for locally represented information – perhaps for visual features of a specific object.

Recall performance for a recent experienced event reflects the combined effect of the recent learning trial together with previous accumulated learning. For weights with lower δ, such as those associated with the cortical component, performance progressively reflects information about all previously encountered information. Because all training events are treated similarly by the net at acquisition, any difference in the recall performance of different types of test item at ITT = 0, must reflect the extent to which previously acquired information can aid recall performance of the most recently presented event (subfigure 7.8a & b). Therefore, a particular level of performance in W_{11}, such as the good recall of semantic information at ITT = 1, in fact results from slow incremental learning over several trials.

That what is already known affects recall performance for subsequently presented events is trivially true in the model, because the sub-patterns to be recalled are identical across a sub-set of test items. However, it is clear that the model would also support faster acquisition in W_{11} of new associations that were similar though not identical to information already established (not shown).

7.4.4.5 Learning different kinds of information: conclusions

These simulations provide clear support for the idea that episodic and semantic information can be seen as points on a continuum of memory types. Episodic memory refers to the most quickly acquired, most complex novel information, whereas items in the broad category 'semantic memory' fall on a continuum from episodic-like, to very generic and simple information. Infrequently presented semantic information is treated progressively more like episodic information by the model; progressively less detailed episodic-like information is treated progressively more like semantic information.

More generally, learning associations between multi-modal information represented in different sections of the cortex requires the hippocampus, unless there are many learning trials. Less complex information can be acquired by local cortical regions. The hipppocampus is also generally more important for the rapid acquisition of information, especially when information is truly novel and cortical learning cannot take advantage of related pre-existing information represented in cortical weights.

The model initially treats each unfolding event the same – laying down a strong event representation in the hippocampus, and weaker one in the cortex. The hippocampus always supports the best recall of the complete, detailed, original traces; although a small amount of information for any memory is retained in the cortex beyond its complete decay from the hippocampus. However, recall for different types of information becomes differentiated with exposure to repeating sub-patterns. With experience, the cortex builds traces for generic semantic information, which tends to dominate recall for old information. The robustness of representations of semantic information in the cortex depends on the extent to which there has been re-exposure to the repeating elements. Thus, the 'neural basis' for recall of memories tends to segregate along the lines of the amount of detail required at recall and the amount of exposure to information, rather than on a 'category-distinction' between episodic and semantic information. Repeated exposure to information has little effect on hippocampal performance.

7.4.5 Modulation of learning

The standard λ and δ values employed in the model lead to very fast decay of information from the hippocampal component. Similarly, most episodic information probably decays rapidly in real life. If, notionally, the ITT values are considered to represent years since the acquisition of information, the data would imply that most detailed episodic information that is recalled via the hippocampus is lost in approximately 1-2 years, a figure supported by neuro-psychological studies. However, clearly, some episodic events are remembered longer than others, and in my opinion any such detailed episodic information is retained in the hippocampus.

Preliminary investigations showed that modulating storage in W_2 alone did not much affect hippocampal memory performance, since recall from these weights depends on W_{12} weights. Therefore modulation was applied to both W_2 and W_{12} in sim-

ulations investigating the effect of memory modulation in the hippocampal component. In order to more clearly demonstrate the effect of increasing learning rates in these weights, λ_2 & λ_{12} and δ_2 & δ_{12} were reduced from their high standard values (although they remain much higher than in the cortically-associated weights). New reference graphs are therefore provided (subfigures 7.12c, and 7.13c).

Increasing λ_2 & λ_{12} slightly improves the initial recall of the information whose storage is modulated, and extends the period for which this information is retained in the hippocampus (for episodic information, compare subfigures 7.12a & c; for semantic information, compare subfigures 7.12b & c). However, this improvement is at the expense of very much poorer initial performance for the non-modulated information stored in the hippocampus (see same graphs). The degree to which the performance of non-modulated information is impaired reflects the extent of overwriting of the non-modulated weights by the large modulated weights. When the net is trained on sparser patterns, the difference between performance at ITT = 0 for episodic and semantic information is reduced, and the gradient of performance decay as ITT rises is reduced (data not shown). The large oscillations seen in the recall of semantic information when semantic learning is modulated in the hippocampus, arises from an exaggeration of the effect of a chance re-using of weights that have already been used to represent information about a semantic sub-pattern.

In the brain, modulation of hippocampal storage probably occurs to a greater extent than memory modulation in other areas. However, arousal levels and the various putative memory modulation mechanisms may also affect storage in cortical regions, too. When memory modulation is applied to W_{11} weights, recall performance for the modulated memories is again enhanced at the expense of non-modulated information (for episodic information, compare subfigures 7.13a & c; for semantic information, compare subfigures 7.13b & c). However, strongly storing episodic information in the cortical component has a particularly detrimental effect on the recall of incrementally acquired semantic information (subfigure 7.13a).

7.4.5.1 Modulation of learning: conclusions

Enhancing the initial storage of a sub-set of traces in either the hippocampal or cortical components impairs the retention of all other information. The robust storage of episodic items in the cortex particularly disrupts the extraction of generic semantic

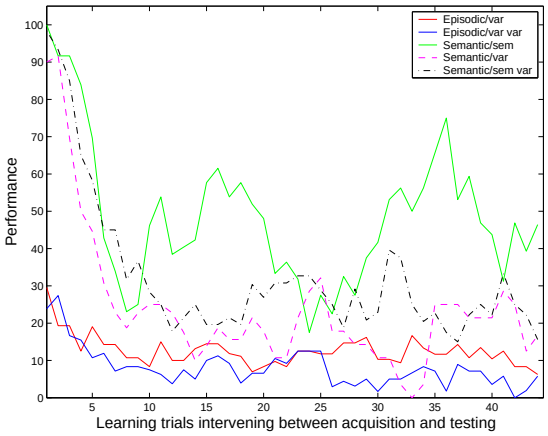

(a) Hippocampal recall

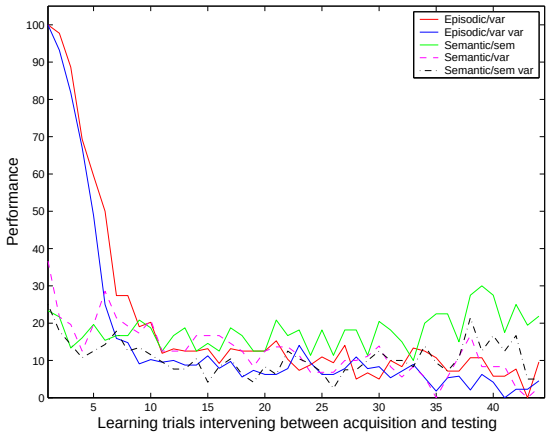

(b) Hippocampal recall

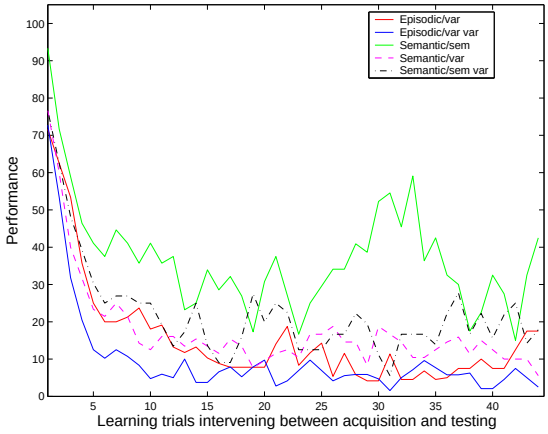

(c) Hippocampal recall

Figure 7.12: Memory modulation in the hippocampal component. (a) For the storage of episodic information, modulated λ_2 and λ_{12} = 0.5. (b) For the storage of semantic information, modulated λ_2 and λ_{12} = 0.5. (c) Reference data for non-standard parameters: λ = [0.01 0.1 0.2 0.2]; δ = [0.001 0.05 0.1 0.1].

information. This implies that storing episodic and semantic information in the same nodes may be incompatible.

However, caution should be applied in extrapolating from this model. The brain consists of many layers in addition to the two layers modelled here. I have suggested that there may be a hierarchy of memory with information of increasing complexity

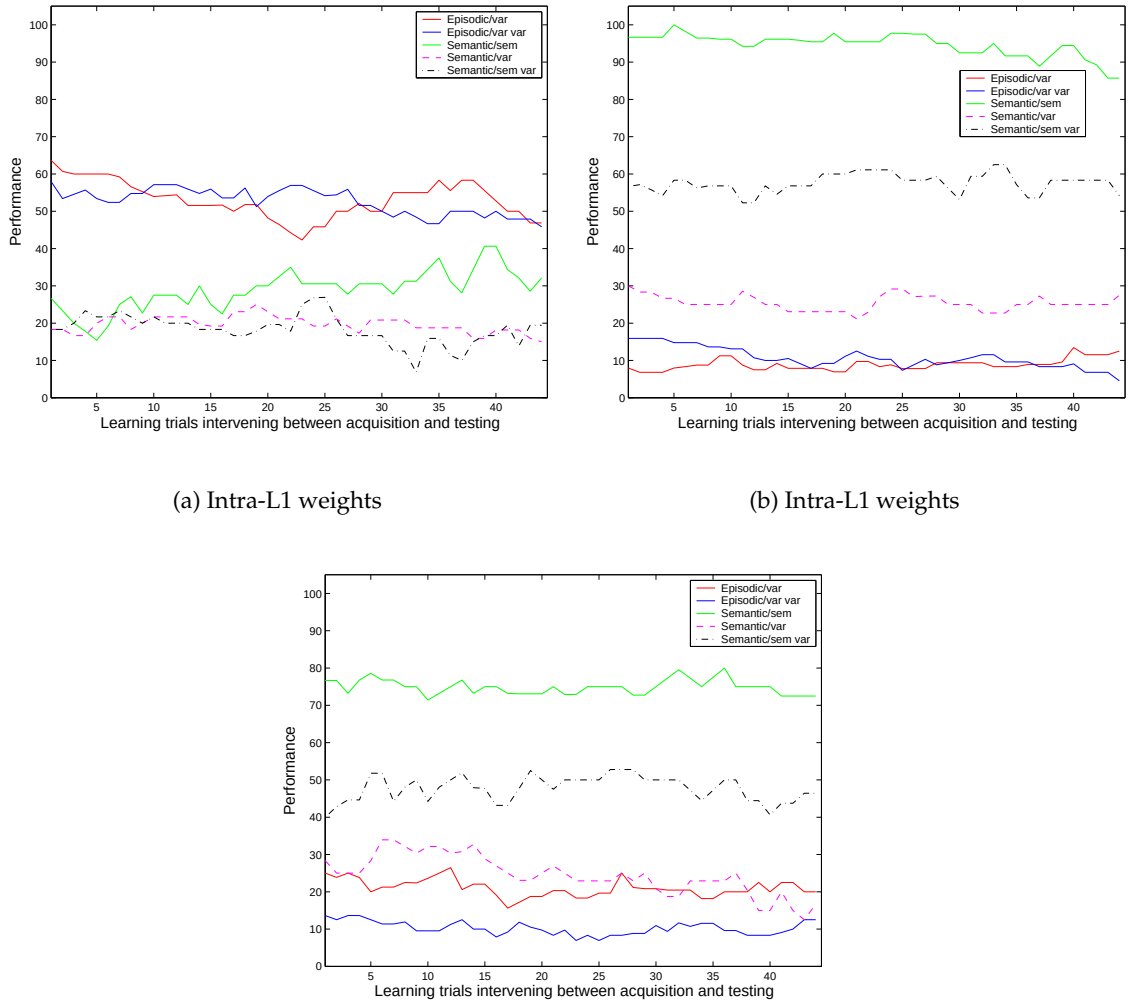

(a) Intra-L1 weights

(b) Intra-L1 weights

(c) Intra-L1 weights

Figure 7.13: Memory modulation in W_{11}. (a) For the storage of episodic information, modulated λ_{11} = 0.1. (b) For the storage of semantic information, modulated λ_{11} = 0.1. (c) Reference data for non-standard parameters: λ = [0.01 0.1 0.2 0.2]; δ = [0.001 0.05 0.1 0.1].

and specificity stored in the higher layers. In this scheme any detailed information that was represented outside of the hippocampus would be represented in a higher region to that for the most generic information. Furthermore, the enhanced storage of some traces would be less detrimental in a more realistic, less overloaded net.

The hippocampus stores all information relatively robustly on the first encounter, without the need for enhanced storage. The cortex, on the other hand, develops recall-

viable traces only gradually. Therefore, low salience, incidental information will normally be dependent on the hippocampus, although under high levels of arousal that trigger the modulation of storage the cortex may be able to acquire some information. Thus, as suggested in section 4.4, the acquisition and recall of incidental or low salience information, is similar to the rapid acquisition of information in the sense that both depend on the ability of the hippocampus to automatically lay down robust traces.

7.4.6 Semanticisation of memory

In the previous chapter, I discussed the semanticisation of memories that occurs with age. I suggested that semanticisation could arise through the loss of detailed information with age, or a relative increase in the ability of semanticised traces to mediate recall[2]. In this section, I explore the semanticisation that occurs in this model.

7.4.6.1 Semanticisation of memory through decay of details

In the model, the most accurate recall of all pattern elements (details) of an event is mediated by the hippocampus. The rapid decay of information from the hippocampal component means that as a memory ages, progressively fewer of the unique details can be recalled, although any sub-components that by chance (for episodic events) or by design (semantic events) overlaps with repeating sub-components will be retained. The cortex can initially support the recall of some random details, but this capacity also drops with time. Thus the recall of older events from the net will be less detailed than for more recent information, and tend to depend more on generic information.

The net's performance on recent memories is dominated by the contribution of the hippocampal component, whereas for old memories only semanticised information (i.e. information that is generic to several events) can be recalled above chance (e.g., subfigures 7.8f or 7.14d). The initial rapid drop in performance on 'Semantic/sem var' items (which contain a mixture of episodic detail and generic information) in subfigure 7.14d reflects the decay in recall of random components, whilst the long-term maintenance of performance reflects the recall of generic semanticised information that is supported by exposure to repeating sub-patterns.

[2] A change in the recall strategy used for the recall of recent and remote memories might also play a part. However, I do not explore this issue here.

7.4.6.2 Semanticisation of memory through 'real-world' exposure

In the model, exposure to training events containing repeating sub-patterns simulates the effect of real-world exposure to generic semantic information. As sub-patterns are re-experienced, progressively more robust traces are established in the slow-learning cortical component (subfigure 7.14c, and many others). As a memory ages and the fully-detailed hippocampal trace is lost, the cortical trace becomes relatively more important for recall, and the information elicited is progressively more semanticised.

7.4.6.3 Semanticisation of memory through offlne replay

I have argued that there may only be quantitative differences in the strengthening of connections that occurs on re-exposure to real-world information, and that which may occur with support from the hippocampal component or through local cortically-driven attractor reactivation. Therefore, the data presented so far could equally well be considered to reflect 'real-world' or 'replay-aided' enhancement of semantic components of traces. However, the model does also implement specifically 'off-line learning'.

Preliminary investigations showed that under standard conditions, information represented in W_{12} decays too quickly to allow a reasonable time-window for off-line learning. Therefore, the same lower-than-standard λ_2, λ_{12}, δ_2 & δ_{12} values that were used in the modulation simulations are also used here. Overwriting in W_{12} was still found to be a problem, so offline-learning simulations also used sparser-than-standard patterns ($k_1 = 4$, $k_2 = 2$). A similar pattern of behaviour was however found for both standard and sparse patterns.

Figure 7.14a & b shows the effect of allowing off-line learning in W_{11} for 3 out of 5 of the most recently experienced training items after each training trial. The initial rise in performance at low ITT reflects the effect of offline learning trials in the period shortly after acquisition. All test items benefit from offline learning, but this is especially evident for the episodic events, as they do not benefit in the same way as semantic information from incremental learning on re-exposure to sub-patterns. Episodic information is lost more slowly from the net as a whole (compare subfigures 7.14b & d), and semantic information is preserved almost perfectly.

When local W_1 weights in the cortical component are allowed to contribute to pattern completion in L1 (in a manner similar to the 'clean-up' performed by W_2 weights

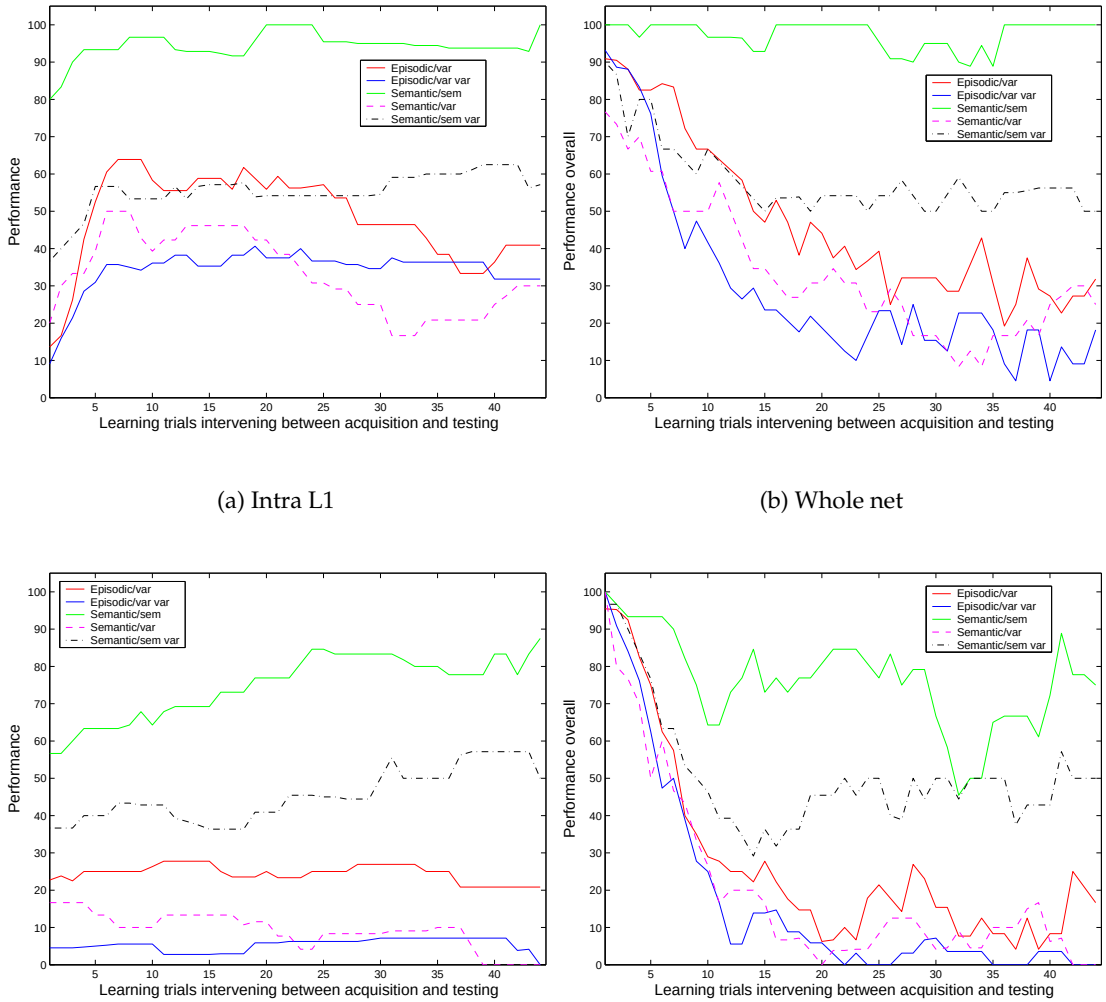

(a) Intra L1

(b) Whole net

(c) Intra L1

(d) Whole net

Figure 7.14: Offline learning enhances retention. (a) & (b): Offline learning in W_{11} using standard λ_{11}. 3 out of 5 of the most recently presented events undergo offline learning after the presentation of each training trial. (c) & (d): Reference data for non-standard parameters: λ = [0.01 0.1 0.2 0.2]; δ = [0.001 0.05 0.1 0.1]; K_1 = 4; k_2 = 2; $d = 2$.

in L2), then off-line learning in W_{12} very clearly benefits the recall of semantic subpatterns at the expense of performance on patterns with random components (compare subfigure 7.14a and 7.15a). Overall net performance also shows an enhanced advantage for semantic information, with lower initial performance for episodic information when ITT = 0, and a steeper decline in performance as ITT rises (compare

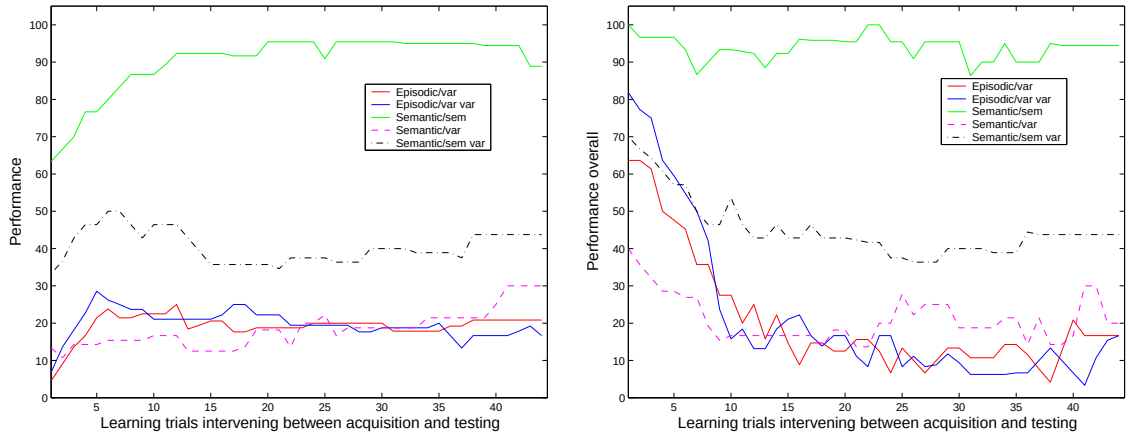

(a) Intra L1 (b) Whole net

Figure 7.15: Offline learning incorporating cortical 'clean-up' benefits semantic information. 3 out of 5 of the most recently presented events undergo offline learning after the presentation of each training trial. λ = [0.01 0.1 0.2 0.2]; δ = [0.001 0.05 0.1 0.1]; K_1 = 4; k_2 = 2; $d = 2$.

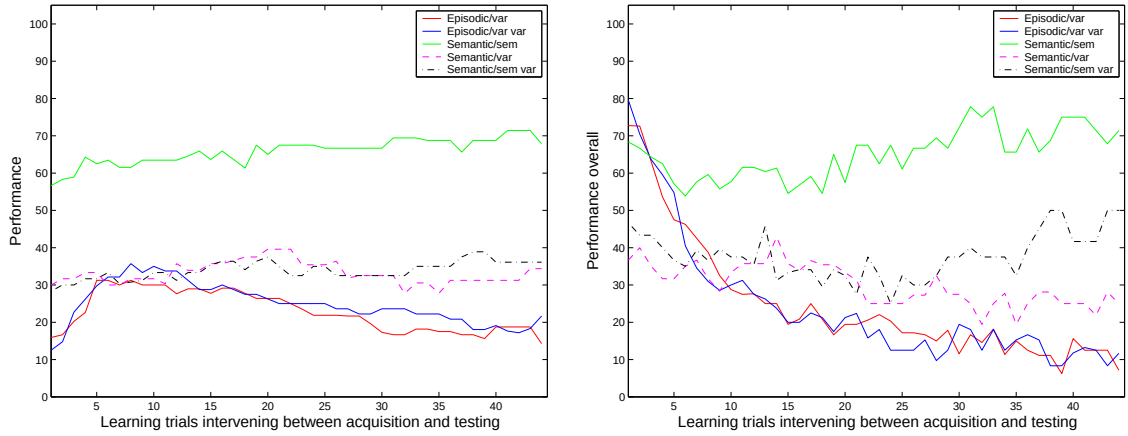

(a) Intra L1 (b) Whole net

Figure 7.16: Modulation of episodic storage in hippocampal component with off-line learning for all events. For the storage of episodic information, modulated λ_2 and λ_{12} = 0.5. 3 out of 5 of the most recently presented events undergo offline learning after the presentation of each training trial. λ = [0.01 0.1 0.2 0.2]; δ = [0.001 0.05 0.1 0.1]; K_1 = 4; k_2 = 2; $d = 2$.

subfigures 7.14b and 7.15b).

In the model, the decay of information in W_{12} determines the time window after acquisition in which offline learning driven by hippocampal reactivations is useful. At $r = 5$, enough information is retained in W_{12} for at least some parts of the original training patterns with which the L2 patterns are paired to be recreated (see subfigure 7.14d). However, as r rises and the patterns recreated in L1 become less similar to the original training patterns, W_{11} is effectively learning noise. This reduces performance on all patterns (not shown).

Extending the lifetime of traces in the hippocampal component extends the time window for offline learning. Combining memory modulation in W_2 and W_{12} for episodic information, say, with offline learning, leads to relatively more constructive offline learning for the modulated episodic events than others (Figure 7.16), as the patterns activated in L1 and learnt are more accurate for episodic than semantic events. (As noted earlier, memory modulation impairs performance for all non-modulated items.)

7.4.6.4 Semanticisation of memory: conclusions

Semanticisation of memories occurs in the model through a loss of hippocampally-dependent episodic detail, and the relative strengthening of cortically-represented core components of events that come to dominate recall as a memory ages. No specific assumptions are needed to generate this behaviour: it arises automatically from the storage characteristics and learning and forgetting rates in the different components of the model.

As noted earlier, evidence that the hippocampus is necessary for any post-acquisition semantic memory processing is currently equivocal. However, it seems reasonable that if the hippocampus does initiate activity in lower regions in support of offline learning, local attractor representations in the cortex would influence the triggered patterns of activation. In the model it is clear that allowing cortical involvement in pattern selection enhances the semanticisation effect that hippocampally-driven replay already has. Obviously the effect would be even greater if long-range cortical connections were allowed to contribute, as they represent the most generic information. This implies that 'consolidation-like' learning would lead to the increasing semanticisation of memory.

'Real-world' (online) exposure to events has a more profound effect on the maintenance of information than offline learning. This seems sensible. However, in a more realistic situation, online and offline learning would be less distinct. The biggest source of recall cues for memory reactivation is likely to be external cues, so offline learning for reactivated information that partially overlaps with new information would be expected to occur in a similar time period to the acquisition of the new information. Such behaviour might contribute to the uncovering of relationships between semantic information.

It is interesting to speculate that traces that were initially stored most robustly in the hippocampus (or elsewhere) could contribute longer to a post-acquisition semanticisation process, thereby having a greater effect on the ongoing re-organisation of memory.

7.5 Summary and conclusions

In this chapter I have presented findings from a simple neural network that simulates the acquisition and maintenance of episodic and semantic information by the hippocampus and cortex. The model makes few assumptions beyond the well-established view that the hippocampal component should be smaller than the cortical component, should store orthogonalised rather than topographically organised representations, and should be capable of more rapid acquisition of information than the cortex. However, the interaction between these features and aspects of the information to be learnt creates interesting and plausible recall performance. In particular, the findings presented in this chapter provide clear support for the proposals put forward in earlier chapters.

Almost all authors share the view that the hippocampus can rapidly learn conjunctive information, whereas many have argued that the cortex cannot. The model supports the view that relatively fast associative learning can occur without the hippocampus when the information to be associated is represented locally in the cortex, or there is a learning-set in the cortex that can support the acquisition of new information without the need for large-scale change. Information that depends on associations between information in more distant regions of the cortex can be learnt in the absence of the hippocampus, albeit at a slower rate. With sufficiently complex cues that act to partially activate weakly connected areas of the cortex, local cortical connections can

support local pattern completion allowing good overall recall.

This model implements only two layers and cannot therefore capture more subtle divisions of labour between regions in the neocortical-hippocampal learning hierarchy: associative learning between information of different types, and at different orders of complexity, would be accomplished at least initially in different layers in this hierarchy. However, although the model refers to a 'hippocampal' and a 'cortical component', at the level of abstraction employed, in some respects the model could be considered to represent any higher and lower region in the neocortical-hippocampal hierarchy. The specific features that distinguish the hippocampus, say, from other regions (such as particular cell fields, internal connectivity) are not modelled. Having said that, the use of random patterns in the 'hippocampal' component clearly represents the higher end of a spectrum from topographically organised to orthogonalised representations.

The hippocampus has been implicated more generally in the rapid acquisition of information. In the model, all one-off event traces are stored most robustly in the hippocampus, so the hippocampus will tend to be relatively more important for the recall of information to which there has been litle exposure. The cortex can acquire information independently from the hippocampus only with sufficient exposure to information. Therefore the storage of truly novel complex information is practically impossible in one trial in the cortex. Information that is already stored can increase the likelihood of the cortex developing a robust representation of related information in a given number of trials. Thus the development of a cortical learning set reduces dependency on the hippocampus's ability for rapid robust storage, and would allow subsequent related learning to be accomplished more quickly by the cortex.

The simulations also provide support for the plausibility of the idea that episodic and semantic memory represent points on a continuum of memory types. Typical episodic memory represents one extreme of one-off learning of very complex novel information; whereas different kinds of semantic memory vary in their dependency on generic and specific information, and in the exposure there has been to information. Although each individual episodic or semantic trace is initially treated in the same way by the net, as experience accrues, recall performance for different types of events via the hippocampal and cortical components becomes differentiated. The 'neural basis' for the recall of memories in the model tends to segregate on the basis

of the amount of random or generic information integral to a memory, and on how frequently the information has been encountered.

All information is initially most robustly recalled via the hippocampal route. The cortex can initially support the recall of some detailed information, but this too decays, albeit at a slower rate than in the hippocampus. Whilst the most detailed episodic information is always best recalled by the hippocampus, over time the cortex tends to take over the recall of all events. Semanticisation of older memories thus results naturally from the rapid loss of detailed episodic information, and from the concomitant tendency for the recall of old information to depend more on the cortex and thus elicit more generic information.

Traces for generic cortically-represented information may build up over time, or at least be lost at a slower rate, as a result of re-exposure to repeated patterns of 'real-world' information. Thus recall for a given specific event via the cortex tends to be dominated by traces representing well-established information. Offline learning in the model disproportionately benefits semantic learning at the expense of already weaker representations for random events, when cortical weights are allowed to contribute to pattern selection as seems reasonable. Thus replay-aided learning similar to that implemented in consolidation models (e.g., Alvarez and Squire (1994); Murre (1996)) leads to the semanticisation of memories.

In section 6.9 I suggested that the existing empirical data does not tell us whether there is an absolute increase in the ability of the cortex to mediate recall over time; or merely a relative increase in recall with respect to the hippocampus, in which case absolute cortical recall could increase, decrease or be constant over time. In the model, cortically-mediated performance interacts with forgetting rates and the amount of 'real-world' re-exposure to information, producing any of these patterns of performance. Offline learning as implemented in the model leads to an initial rise in the ability of the cortex to mediate recall, but this too interacts with the absolute value of a decay rate. In any case, the recall of recent events is dominated by inputs from the hippocampus, so the effect of the initial increase in the cortical recall is masked. Therefore the model does not take an unequivocal position on this issue.

In summary, this model captures earlier proposals about the relative importance of the hippocampal and non-hippocampal areas in the acquisition of certain types of information; and the long-term fate of such information with respect to the hip-

pocampus. Truly random information that requires associations across distant cortical regions is best recalled by the hippocampus for the lifetime of the memory; whereas information that could be acquired by the cortical regions albeit more slowly, and semanticised versions of all information, become relatively more dependent on the cortex as a memory ages.

7.5.1 Limitations of model and further work

The model is intended to provide a proof of concept for ideas presented earlier in the thesis, rather than aim for biological realism. To that end, the very small scale of the model, the extremely simple learning rule and the full connectivity should not be considered to be problematic. On the other hand, scaling the model up (which would require an optimisation of the code), and employing sparser connectivity would allow new features to be explored.

Obvious areas for further work within the current framework of the model include:

Hippocampal representation The hippocampal complex-spike recording literature implies that the representations used in the hippocampus for encoding similar information may be relatively stable over time. Thus selecting a random hippocampal pattern to represent each incoming event may be unreasonable. Employing sparse connectivity between the hippocampal and cortical components would allow the net to create hippocampal representations, although this kind of learning needs to be carefully controlled. Adding another set of L1-type nodes that send inputs to L2 that are randomly activated when an training event is applied to L1 might provide the right balance of input-driven, but partially orthogonalised, representations. Combining stable sparse point-to-point connectivity with weight changes within connections should be investigated.

Developing long-range connectivity The slow learning rate in the intra-cortical connections represents the difficulty of developing connections between distant nodes, as well as slow learning between already connected nodes. The current implementation of the model implies that all nodes in the cortex could potentially be usefully connected, which is unlikely. An improvement would be to allow intra-cortical weights to learn (as now), but to exclude weights below a certain threshold from contributing to recall. The value of the weight below the

threshold value would indicate the 'progress' made in the formation of a new viable connection.

Attractor dynamics The model simulates the effect of the interaction between connections in a very simple way, by summing raw activities produced by different weight sets, or allowing raw activations one pass through a weight set. A more dynamic system would support more realistic pattern selection for offline learning.

Comparisons with neuropsychological data Performance analyses in the model are averaged over all events of a particular type, so it is currently difficult to investigate the effect of 'lesions' at a particular point in training. An alternative method of data visualisation is required. However, conditions such as semantic dementia could be relatively easily modelled in the static state.

Learning hierarchy I argue for a hierarchy of learning areas in chapter 6. Implementing more than the current two levels would be informative.

Bridging gaps and higher-order learning This model simulates only fast episodic-type encoding in the hippocampus, and does not attempt to incorporate the hippocampus' probable role in bridging temporal gaps, or learning about complex associations that are not immediately evident from observed data. Such roles seem incompatible, as the latter tasks require that representations for similar represented data are not orthogonalised. It would be very interesting to explore the effect of inter-mixing directly activated and orthogonalised inputs in the hippocampus.

Chapter 8

Conclusions

This thesis has addressed the question of the relationship between the dependence of information on the hippocampus at acquisition and the hippocampus's role in its long-term maintenance. The approach taken was to critically re-examine existing empirical data and the claims that have been made on its basis, in order to formulate coherent proposals on the role of the hippocampus in the acquisition and long-term recall of information. I then tested the plausibility of these proposals with a computational model.

In short, I conclude that information whose acquisition is most severely impaired by hippocampal damage is dependent on the hippocampus for longer when it is acquired in the presense of the hippocampus. Tasks that are obligatorily hippocampally-dependent at acquisition, such as task-dependent use of allocentric information or detailed context-specific information, remain dependent on the hippocampus for recall for the memory's lifetime. Tasks that can be acquired to some extent in the absence of the hippocampus, such as conditional motor learning or semantic information, become independent of the hippocampus over time. In these cases, the type of information that is recalled from regions outwith the hippocampus is similar in nature to the kind of information that would have been acquired in the absence of the hippocampus.

8.1 Main conclusions of thesis

The starting point for this thesis was the observation that tasks whose acquisition is affected by hippocampal damage are not affected in an all-or-nothing manner. There are large differences in the extent to which the acquisition of different kinds of information is impaired by hippocampal damage. Under a given set of learning conditions some tasks are obligatorily hippocampally-dependent and only the hippocampus can form a viable trace that can mediate recall; whereas others are merely facilitated by hippocampal activity as the hippocampus is merely faster or better at forming traces than other areas. This division may reflect a distinction between what only the hippcampus can do, and what the hippocampus normally does in an intact brain. In the latter cases, in the intact brain, the automatically-acquired hippocampal trace improves learning because it is more quickly acquired under given circumstances, although the information acquired may not be strictly necessary for the task at hand.

Although a distinction has been made between 'obligatory' and 'facilitated' tasks for expositional purposes, all such information falls on a continuum of dependency. Furthermore, the difference between the ability of the hippocampus and other to areas to support learning and recall on a given task is not all or nothing. All areas activated on a particular task record information in parallel, but store different kinds of information that can support performance on a given task to different extents.

Some tasks are of course unaffected by damage to the hippocampus – the acquisition of such information is not addressed by this thesis. The acquisition of other information, such as ego-centric strategies, is actually improved by damage to the hippocampus. This suggests that the hippocampus and other areas may compete for control of recall.

Several common traits are shared by tasks whose acquisition is impaired by damage to the hippocampus. Tasks that depend on the task-dependent use of supramodal information are obligatorily hippocampally-dependent; this information cannot be acquired to any extent in the absence of the hippocampus under the conditions tested to date. Similarly, the rapid acquisition of detailed context-specific information, such as that underpinning episodic information, cannot be achieved without the hippocampus. In general, tasks are most likely to require the hippocampus for acquisition if they depend on rapid associative learning about unfamiliar, complex, low-salience information. Each of these 'qualifying' factors are continuous; for example, there can

be greater or fewer trials, or more or less complexity. None of these factors alone is sufficient to obligatorily implicate the hippocampus in acquisition. Instead, it is the combination of factors, weighted by their positions on a notional scale of severity that is important. Under a given set of learning conditions the hippocampus may be significantly better than others areas at acquiring such information, but other areas can learn when the conditions of training are altered. This general picture holds across data from many species.

Previous views of the hippocampus have emphasised the role of the hippocampus in the rapid acquisition of information. However, several tasks that can be acquired in one trial are not hippocampally-dependent, and some slowly acquired tasks cannot be acquired after hippocampal damage. Other views have emphasised the role of the hippocampus in complex associative learning. However, these approaches would not predict distinctions in performance between, say, implicit and explicit learning, or that incidental (but not contingent) or trace (but not delay) learning would be impaired by hippocampal damage. In this thesis, I make proposals that both constrain and extend existing views on the role of the hippocampus in rapid, associative and automatic learning.

I concluded that learning about relationships between stimuli of the same modality does not generally require the hippocampus. Higher-order cross-modal learning is generally only hippocampally-dependent when the information to be acquired a) must be stored quickly, but not if there are sufficient training trials; b) is truly novel and there is no established learning-set outwith the hippocampus; or c) is of low salience or to be acquired incidentally. The hippocampus is obligatory for learning about associations between supra-modal and task-related information that does not converge in regions outside the hippocampus, irrespective of learning conditions. The hippocampus is also always obligatory for the acquisition of true episodic memories, given the the inherent conditions of acquisition; that is, the fast incidental acquisition of low salience, complex information. It is also important for learning when the 'solution' to a task can be most easily discovered using the default high-order representational structure of the hippocampus.

Generally speaking, the term 'episodic memory' has been used to refer to recall for context-specific events, whereas 'semantic memory' refers to any of a wide range of information representing knowledge about the world. These standard definitions cut

across several confounding factors such as the amount and type of detail needed to demonstrate recall, and the amount of exposure there has been to information. I argue instead that real-world memories fall on a *continuum* of episodic/semantic-ness. Memories vary continuously in terms of the number, type and specificity of details recalled, and the extent to which recall depends on information that archetypally characterises episodic and semantic recall. A typical episodic memory is merely a memory at one extreme of the continuum of memory types. Episodic memories are by definition of high complexity, specificity and novelty, refering to an event of short duration that has only been experienced once. Semantic information, on the other hand, varies from episodic-like (such as information about a public event which has only been encountered once or a few times) to very generic (such as knowledge of common word meanings). Since the hippocampus is implicated in the rapid acquisition of complex incidental information, episodic information will be disproportionately affected by hippocampal damage. Semantic information will be affected to the extent that it depends on such information.

In the second half of the thesis, I examined the long-term role of the hippocampus in the recall of information that was hippocampally-dependent at acquisition. Information that obligatorily requires the hippocampus at acquisition depends on the hippocampus for the lifetime of the trace, although that may be less than the lifetime of the animal. Any episodic-like or task-dependent allocentric-like information recalled after hippocampal damage is qualitatively different to that occurring via the hippocampus. The type of information recalled under these conditions is similar to that which could have been acquired – albeit to a limited extent – in the absence of the hippocampus. In other words, detailed context-specific information and task-dependent allocentric spatial information does not appear to be replicated outside the hippocampus after initial acquisition by the hippocampus.

In contrast, the recall of information whose acquisition is merely facilitated by the presence of an intact hippocampus may become independent of the hippocampus over time. Structures in the medial temporal lobe, and perhaps also at lower levels of the neocortical-hippocampal hierarchy become redundant over time in an orderly sequence. This suggests that they may play a similar role to the hippocampus in this respect: each region may act as an indexing zone for information stored at lower levels. Redundancy does not arise from the wholesale transfer of information from the

hippocampus or other higher regions to other lower regions, but from a change in the nature of the information underlying recall. Regions higher in the hierarchy always support the recall of the most specific information. Clearly, any such change in the nature of the information underlying recall can only occur on tasks that *can* be performed by semanticised rather than very specific traces. The existing evidence suggests that tasks whose acquisition is most impaired by hippocampal damage show longer periods of retrograde amnesia after hippocampal damage when the task is acquired in the presence of the hippocampus.

Specific memories tend to become more like generic memories over time: this is termed semanticisation. Semanticisation results from the relative or absolute increase in the tendency or ability of semantic traces to mediate recall. There are several main contributing factors. Firstly, traces representing the most context-specific details of information decay faster than those for more generic information. This can lead to a relative increase in the ability of non-hippocampal areas to mediate recall over time, as the gap between recall abilities narrows. On the other hand, there may be an absolute increase over time in the ability of the non-hippocampal areas to mediate semanticised recall. Two main mechanisms are relevant here: 1) re-exposure to real-world events containing repeated sub-components could lead to a more robust traces for generic information that develop over time, and 2) offline learning through the reactivation of attractors. Unlike others, I argue that offline learning contributes to the development of representations of the core semantic components of memory, and does not support an equivalent consolidation of contextual and other details. Whilst there are reasons for believing that offline learning may be initiated by the hippocampus, this is currently unproven. Instead, or in addition, offline learning learning may be initiated by local processes in the cortex or throughout the neocortical-hippocampal hierarchy. The third main source of semanticisation is that the recall of old and recent memories may be supported by different strategies.

Semanticisation should not be understood as a negative process in which information is lost, but instead as a process that progressively strengthens the storage of the most valuable and widely-applicable information, and uncovers relationships between such information. The processes uncovered by reconsolidation studies may reflect such a semanticisation process. Post-acquisition memory processing is multi-stage and lifelong, and acts to produce semanticised re-representations of information

that can co-exist with remaining more detailed traces, and that serve a different purpose.

Finally, I implemented a neural model to test the plausibility of these proposals. In common with many other models, the model consists of a quickly learning 'hippocampal' component that stores orthogonalised traces; and an input 'cortical' component which maps information topographically and incorporates very slow-learning long-range connections and relatively faster local connections. The model also shows forgetting, with the fastest loss of information occurring from the hippocampus, and the slowest in the long-range cortical connections. The novelty of this model is in the nature of the training events and the way the results are analysed. When the model was trained on random 'episodic' events and 'semantic' events that share overlapping patterns, plausible behaviour developed with regard to the nature of the information recalled at particular periods after initial acquisition, and of the 'neural basis' for the recall of different types of information. The simulation findings provide clear 'proof of concept' for the proposals put forward in this thesis.

8.2 Contributions of the thesis

I have met all of the objectives outlined in section 1.2.

The main contributions of this thesis are to have:

- Drawn several coherent themes out of a vast body of empirical data.
- Explicitly assessed current theories against existing data.
- Redefined existing ideas on the role of the hippocampus in the acquisition of information.
- Established a new conception of episodic and semantic memory.
- Proposed a new conception of the role of the hippocampus in the long-term recall of information.
- Investigated the relationship between the role of the hippocampus in task acquisition and long-term recall.
- Developed a theory of semanticisation.

- Created a modelling tool.
- Demonstrated the plausibility of my proposals with the model.

8.3 Predictions

The development of any theory depends on an interaction between hypotheses and empirical investigation. This thesis make several predictions that could be empirically tested. A few examples are given here:

1. If the hippocampus is more important for acquiring and maintaining arbitrary information than predictable information, then hippocampal damage should lead to relatively greater impairment in associative memory for recently-encountered pairs of semantically unrelated words (teapot hang-glider) than of related words (teapot cup), say.
2. If memories become more semanticised and less context-specific over time, the effects of context would be expected to diminish at a faster rate than the ability to perform a generic task learnt in that context.
3. If the semantic regions can over time develop representations that allows them to support faster acquisition of related information, then the acquisition of some kinds of information may be more dependent on the hippocampus in younger or less experienced people than in older or more experienced people.
4. If the hippocampus becomes relatively less important for the recall of information to which there has been extensive exposure, the recall of low frequency words should be more affected by hippocampal damage than high frequency words. (There is some preliminary evidence for this position.)
5. If post-acquisition semantic memory processing is driven more by real-world triggers than by an autonomous endogenous process, then leaving cues related to a recently learnt task in a living environment may enhance the development of semantic representations over time.
6. Any generic model of 'consolidation' that involves cortical weights in the selection of patterns for offline replay, will naturally semanticise older memories when trained on patterns that contain overlapping sub-patterns.

8.4 Final words

The facts of nature and of life are more apt to be complex than simple.
Simplistic theories are generally one-sided and partial.
James Freeman Clarke (19th century)

Many hundreds of thousands of words have been written on the topic of the hippocampus, and no doubt the debate on the role of the hippocampus in the acquisition and long-term retention of memory will rage on. This thesis contributes to that debate, but will clearly not be the last word on the topic.

Indeed, there may be no 'last word'. Not only do different view-points determine the facts as we see them; but our desire for simple easily-stated explanations may be misplaced. The brain is vastly complex and its functions are unlikely to be organised along the discrete divisions favoured by researchers.

Bibliography

Abel, T. and Lattal, K. M. (2001). Molecular mechanisms of memory acquisition, consolidation and retrieval. *Current Opinion in Neurobiology*, **11**, 180–187.

Aggleton, J. and Brown, M. (1999). Episodic memory, amnesia and the hippocampal-anterior thalamic axis. *Behavioral and Brain Sciences*, **22**, 425–489.

Aggleton, J., Vann, S., Oswald, C., and Good, M. (2000). Identifying cortical inputs to the rat hippocampus that subserve allocentric spatial processes: A simple problem with a complex answer. *Hippocampus*, **10**(4), 466–474.

Aguirre, G. and D'Esposito, M. (1997). Environmental knowledge is subserved by separable dorsal/ventral neural areas. *Journal of Neuroscience*, **17**, 2512–2518.

Altman, J. and Das, G. D. (1965). Autoradiographic and histologic evidence of postnatal neurogenesis in rats. *Journal of Comparative Neurology*, **124**, 319–335.

Alvarado, M. C. and Rudy, J. W. (1993). Configural theory and the hippocampus: is lesion type critical? evidence from four tasks. *Society for Neuroscience Abstracts*, **19**, 363.

Alvarado, M. C. and Rudy, J. W. (1995). Rats with damage to the hippocampal-formation are impaired on the transverse patterning problem but not on elemental discriminations. *Behavioral Neuroscience*, **109**(2), 204–211.

Alvarez, P. and Squire, L. (1994). Memory consolidation and the medial temporal lobe: A simple network model. *Proceedings of the National Academy of Science USA*, **91**, 7041–7045.

Alvarez, P., Zola-Morgan, S., and Squire, L. R. (1995). Damage limited to the hippocampal region produces long-lasting memory impairment in monkeys. *Journal of Neuroscience*, **15**, 3796–3807.

Alyan, S. H. and McNaughton, B. L. (1999). Hippocampectomized rats are capable of homing by path integration. *Behavioral Neuroscience*, **113**(1), 19–31.

Amaral, D. and Witter, M. (1989). The three-dimensional organisation of the hippocampal formation: a review of anatomical data. *Neuroscience*, **3**, 571–591.

Amaral, D. G. and Witter, M. (1995). The hippocampal formation. In P. G. London, editor, *The rat nervous system*.

Ambrogi-Lorenzini, C. G., Baldi, E., Bucherelli, C., Sacchetti, B., and Tassoni, G. (1999). Neural topography and chronology of memory consolidation: A review of functional inactivation findings. *Neurobiology of Learning and Memory*, **71**, 1–18.

Anagnostaras, S. G., Maren, S., and Fanselow, M. S. (1999). Temporally graded retrograde amnesia of contextual fear after hippocampal damage in rats: Within-subjects examination. *Journal of Neuroscience*, **19**(3), 1106–1114.

Arbib, M., Érdi, P., and Szentágothai, J. (1998). *Neural Organization: Structure, function and dynamics*. MIT Press, Cambridge, MA; London, UK.

Ardenghi, P., Barros, D., Izquierdo, L. A., Bevilaqua, L., Schroder, N., Quevedo, J., Rodrigues, C., Madruga, M., Medina, J. H., and Izquierdo, I. (1997). Late and prolonged post-training memory modulation in entorhinal and parietal cortex by drugs acting on the camp/protein kinase a signalling pathway. *Behavioural Pharmacology*, **8**(8), 745–751.

Atkinson, R. and Shiffrin, R. (1968). Human memory: A proposed system and its control processes. In K. Spence and J. Spence, editors, *The psychology of learning and motivation: Vol. 2. Advances in research and theory*, pages 89–195. Academic Press, New York.

Bach-y-Rita, P. and Aiello, G. L. (2001). Brain energetics and evolution. *Behavioural and Brain Sciences*, **24**, 280 – 281.

Bachevalier, J., Beauregard, M., and Alvarado, M. (1999). Long-term effects of neonatal damage to the hippocampal formation and amygdaloid complex on object discrimination and object recognition in rhesus monkeys (macaca mulatta). *Behavioural Neuroscience*, **113**(6), 1127–1151.

Baddeley, A. and Warrington, E. (1970). Amnesia and the distinction between long- and short-term memory. *Journal of verbal learning and verbal behavior*, **9**, 176–189.

Bahrick, H. P. (1984). Semantic memory content in permastore – 50 years of memory for Spanish learned in school. *Journal of Experimental Psychology – General*, **113**(1), 1–29.

Bannerman, D., Good, M., Butcher, S., Ramsay, M., and Morris, R. (1995). Distinct components of spatial learning revealed by prior training and NMDA receptor blockade. *Nature*, **378**, 182–186.

Bannerman, D. M., Yee, B. K., Good, M. A., Heupel, M. J., Iversen, S. D., and Rawlins, J. N. P. (1999). Double dissociation of function within the hippocampus: a comparison of dorsal, ventral, and complete hippocampal cytotoxic lesions. *Behavioral Neuroscience*, **113**(6), 1170–1188.

Bastin, C., Van der Linden, M., Charnallet, A., Denby, C., Montaldi, D., Roberts, N., and Mayes, A. (2004). Dissociation between recall and recognition memory performance in an amnesic patient with hippocampal damage following carbon monoxide poisoning. *Neurocase*, **10**(4), 330–344.

Baxter, M. G. and Murray, E. A. (2001). Opposite relationship of hippocampal and rhinal cortex damage to delayed nonmatching-to-sample deficits in monkeys. *Hippocampus*, **11**(1), 61–71.

Bayer, S. A. (1986). Neurogenesis in the rat primary olfactory cortex. *International Journal of Developmental Neuroscience*, **4**(3), 251–271.

Beason-Held, L. L., Rosene, D. L., Killiany, R. J., and Ross, M. B. (1999). Hippocampal formation lesions produce memory impairment in the Rhesus monkey. *Hippocampus*, **9**, 562–574.

Bennett, A. T. D. (1996). Do animals have cognitive maps? *Journal of Experimental Biology*, **199**(1), 219–224.

Berardi, N., Pizzorusso, T., and Maffei, L. (2000). Critical periods during sensory development. *Current Opinion in Neurobiology*, **10**, 138–145.

Berger, T., Rinaldi, P., Weisz, D. J., and Thompson, R. (1983). Single unit analysis of different hippocampal cell types during classical conditioning of the rabbit nictating membrane response. *Journal of Neurophysiology*, **50**, 1197–1219.

Bernabeu, R., Cammarota, M., Izquierdo, I., and Medina, J. (1997). Involvement of glutamate AMPA receptors and a cAMP/protein kinase A/CREB-P pathway in memory consolidation of an aversive learning task in rats. *Brazilian Journal of Medical and Biological Research*.

Binder, J. R., Frost, J. A., Hammeke, T. A., Bellgowan, P. S. F., Rao, S. M., and Cox, R. W. (1999). Conceptual processing during the conscious resting state: A functional MRI study. *Journal of Cognitive Neuroscience*, **11**, 80–93.

Bizon, J. L., Han, J., Hudon, C., and Gallagher, M. (2003). Effects of hippocampal cholinergic deafferentation on learning strategy selection in a visible platform version of the water maze. *Hippocampus*, **13**, 000 –000.

Blum, K. I. and Abbott, L. F. (1996). A model of spatial map formation in the hippocampus of the rat. *Neural Computation*, **8**(1), 85–93.

Bolhuis, J. J., Stewart, C. A., and Forrest, E. M. (1994). Retrograde-amnesia and memory reactivation in rats with ibotenate lesions to the hippocampus or subiculum. *Quarterly Journal of Experimental Psychology: Section B - Comparative and Physiological Psychology*, **47**(2), 129–150.

Bontempi, B., Laurent-Demir, C., Destrade, C., and Jaffard, R. (1999). Time-dependent reorganization of brain circuitry underlying long-term memory storage. *Nature*, **400**(6745), 71–675.

Bookheimer, S., Zeffiro, T., Blaxton, T., Gaillard, W., and Theodore, W. (1995). Regional cerebral blood flow during object naming and word reading. *Human Brain Mapping*, **3**, 93–106.

Bramham, C. R. and Srebro, B. (1989). Synaptic plasticity in the hippocampus is modulated by behavioral state. *Brain Research*, **493**(1), 74–86.

Brasted, P. J., Bussey, T. J., Murray, E. A., and Wise, S. P. (2003). Role of the hippocampal system in associative learning beyond the spatial domain. *Brain*, **126**, 1202–1223.

Breese, C., Hampson, R., and Deadwyler, S. (1989). Hippocampal place cells: stereotypy and plasticity. *Journal of Neuroscience*, **9**, 1097–1111.

Brizzolara, D., Casalini, C., Montanaro, D., and Posteraro, F. (2003). A case of amnesia at an early age. *Cortex*, **39**(4-5), 605–625.

Broman, M., Rose, A. L., Hotson, G., and McCarthy Casey, C. (1997). Severe anterograde amnesia with onset in childhood as a result of anoxic encephalopathy. *Brain*, **120**, 417–433.

Brown, A. S. (2002). Consolidation theory and retrograde amnesia in humans. *Psychonomic Bulletin and Review*, **9**(3), 403–425.

Buhusi, C. and Schmajuk, N. (1996). Attention, configuration and hippocampal function. *Hippocampus*, **6**, 621–642.

Buhusi, C. V. and Meck, W. H. (2000). Timing for the absence of a stimulus: The gap paradigm reversed. *Journal of Experimental Psychology: Animal Behavior Processes*, **26**(3), 305–322.

Bunsey, M. and Eichenbaum, H. (1996). Conservation of hippocampal memory function in rats and humans. *Nature*, **379**, 255–257.

Bures, J., Fenton, A. A., Kaminsky, Y., and Zinyuk, L. (1997). Place cells and place navigation. *Proceedings of the National Academy of Sciences of the USA*, **94**(1), 343–350.

Burgess, N. and O'Keefe, J. (1996). Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus*, **7**, 749–762.

Burgess, N., Becker, S., King, J. A., and O'Keefe, J. (2001). Memory for events and their spatial context: models and experiments. *Philosophical Transactions of the Royal Society of London: B*, **356**, 1493–1503.

Burgess, N., Maguire, E., and O'Keefe, J. (2002). The human hippocampus and spatial and episodic memory. *Neuron*, **35**(4), 625–641.

Buzsáki, G. (1989). 2-stage model of memory trace formation: A role for noisy brain states. *Neuroscience*, **31**(3), 551–570.

Buzsáki, G. (1998). Memory consolidation during sleep: a neurophysiological perspective. *Journal of Sleep Research*, **7**, 17–23 Suppl.

Cassel, J.-C., Cassel, S., Galani, R., Kelche, C., Will, B., and Jarrard, L. (1998). Fimbria-fornix vs selective hippocampal lesions in rats: effects on locomotor activity and spatial learning and memory. *Neurobiology of Learning and Memory*, **69**, 22–45.

Cave, B. C. and Squire, L. R. (1992). Intact and long-lasting repetition priming in amnesia. *Journal of Experimental Psychology: Learning, memory and cognition*, **18**, 509–520.

Cermak, L. and Butters, N. (1972). The role of interference and encoding in the short-term memory deficits of Korsakoff patients. *Neuropsychologia*, **10**, 89–96.

Cermak, L., Butters, N., and Gerrein, J. (1973). The extent of verbal coding ability of Korsakoff patients. *Neuropsychologia*, **11**, 85–94.

Cho, Y., Friedman, E., and Silva, A. (1999). Ibotenate lesions of the hippocampus impair spatial learning but not contextual fear conditioning in mice. *Behavioral Brain Research*, (1), 77–87.

Chrobak, J. J. and Buzsáki, G. (1994). Selective activation of deep layer (V-VI) retrohippocampal cortical-neurons during hippocampal sharp waves in the behaving rat. *Journal of Neuroscience*, **14**(10), 6160–6170.

Chrobak, J. J., Lörincz, A., and Buzsáki, G. (2000). Physiological patterns in the hippocampo-entorhinal cortex system. *Hippocampus*, **10**(4), 457–465.

Chun, M. M. and Phelps, E. A. (1999). Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. *Nature Neuroscience*, **2**(9), 844–847.

Cipolotti, L., Shallice, T., Chan, D., Fox, N., Scahill, R., Harrison, G., Stevens, J., and Rudge, P. (2001). Long-term retrograde amnesia ... the crucial role of the hippocampus. *Neuropsychologia*, **39**(2), 151–172.

Clark, R., West, A., Zola, S., and Squire, L. (2001). Rats with lesions of the hippocampus are impaired on the delayed nonmatching-to-sample task. *Hippocampus*, **11**(2), 176–186.

Clark, R. E. and Squire, L. R. (1998). Classical conditioning and brain systems: a key role for awareness. *Science*, **280**, 77–81.

Clayton, N. S. and Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, **395**, 272–274.

Clayton, N. S. and Krebs, J. R. (1995). Memory in food-storing birds: from behaviour to brain. *Current Opinion in Neurobiology*, **5**(2), 149–154.

Clegg, B. A., DiGirolamo, G. L., and Keele, S. W. (1998). Sequence learning. *Trends in Cognitive Science*, **2**(8), 275–281.

Clifton, P., Vickers, S., and Somerville, E. (1998). Little and often: Ingestive behavior patterns following hippocampal lesions in rats. *Behavioral Neuroscience*, **112**(3), 502–511.

Cobos-Zapiain, C. G., Salado-Castillo, R., Sánchez-Alavez, M., Quirarte, G. L., Roldán-Roldánd Díaz del Guante, M., and Prado-Alcalá, R. A. (1996). High level of footshock during inhibitory avoidance training prevents amnesia induced by intranigral injection of GABA antagonists. *Neurobiology of Learning and Memory*, **65**, 202–206.

Cohen, J. J. and Squire, L. R. (1981). Preserved learning and retention of pattern-analysing skill in amnesia: Dissociation of knowing how and knowing that. *Science*, **210**, 207–210.

Cohen, N. and Eichenbaum, H. (1993). *Memory, amnesia and the hippocampal system*. MIT Press, Cambridge, MA.

Cohen, N., Ryan, J., Hunt, C., Romine, L., Wszalek, T., and Nash, C. (1999). Hippocampal system and declarative (relational) memory: Summarizing the data from functional neuroimaging studies. *Hippocampus*, **9**, 83–98.

Compton, D. M., Griffith, H. R., McDaniel, W. F., Foster, R. A., and Davis, B. K. (1997). The flexible use of multiple cue relationships in spatial navigation: A comparison of water maze performance following hippocampal, medial septal, prefrontal cortex, or posterior parietal cortex lesions. *Neurobiology of Learning and Memory*, **68**, 117–132.

Conway, M. (2001). Sensory-perceptual episodic memory and its context: autobiographical memory. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences*, **356**(1413), 1375–1384.

Conway, M., Turk, D.J.and Miller, S., Logan, J.and Nebes, R., Meltzer, C., and Becker, J. (1999). A positron emission tomography (PET) study of autobiographical memory retrieval. *Memory*, **7**(5-6), 679–702.

Corkin, S. (1968). Acquisition of a motor skill after bilateral medial temporal-lobe excision. *Neuropsychologia*, **6**, 255–265.

Crick, F. H. C. and Mitchison, G. (1983). The function of dream sleep. *Nature*, **304**, 111–114.

Crusio, W. E. (1996). Natural selection on hippocampal circuitry underlying exploratory behaviour in mice: Quantitative-genetic analysis. In E. Alleva, A. Fasola, H. P. Lipp, L. Nadel, and L. Ricceri, editors, *Behavioural Brain Research in Naturalistic and Semi-naturalistic Settings*, NATO Advanced Study Institutes Series D, pages 323–342. Kluwer Academic Press, Dordrecht, The Netherlands.

Curran, T. (1995). On the neural mechanisms of sequence learning. *Psyche*, **2**(12).

Curran, T. (1997). Higher-order associative learning in amnesia: evidence from the serial reaction time task. *Journal of Cognitive Neuroscience*, **9**(4), 522–533.

Czurkó, A., Cźeh, B., Seress, L., Nadel, L., and Bures, J. (1997). Severe spatial navigation deficit in the Morris water maze after single high dose of neonatal x-ray irradiation in the rat. *Proceedings of the National Academy of Sciences of the United States of America*, **94**, 2766–2771.

Damasio, A. (1989a). The brain binds entities and events by multiregional activation from convergence zones. *Neural Computation*, **1**, 123–132.

Damasio, A. (1989b). Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition*, **33**, 25–62.

Dash, P. K., Hebert, A. E., and Runyan, J. D. (2004). A unified theory for systems and cellular memory consolidation. *Brain Research Reviews*, **45**, 30–37.

Davidson, T. and Jarrard, L. (1993). A role for hippocampus in the utilization of hunger signals. *Behavioral and Neural Biology*, **59**, 167–171.

Day, L. B. (2003). The importance of hippocampus-dependent non-spatial tasks in analyses of homology and homoplasy. *Brain, Behaviour and Evolution*, **62**, 96–107.

de Biran, M. (1804). *The influence of habit on the faculty of thinking, 1929 edition*. Williams and Wilkins, Baltimore.

de Hoz, L. (2000). *Memories along the longitudinal axis of a rodent hippocampus: acquisition and consolidation of variants of a spatial task*. Ph.D. thesis, Centre for Neuroscience, University of Edinburgh.

de Hoz, L., Know, J., and Morris, R. G. M. (2003). Longitudinal axis of the hippocampus: both septal and temporal poles of the hippocampus support water maze spatial learning depending on the training protocol. *Hippocampus*, **13**, 587–603.

de Kloet, E. R., Oitzl, M. S., and Joëls, M. (1999). Stress and cognition: are corticosteroids good or bad guys? *Trends in Neurosciences*, **22**(10), 422–426.

Debiec, J., LeDoux, J. E., and Nader, K. (2002). Cellular and systems reconsolidation in the hippocampus. *Neuron*, **36**, 527–538.

Desmond, N. and Levy, W. (1997). Ovarian steroidal control of connectivity in the female hippocampus: An overview of recent experimental findings and speculations on its functional consequences. *Hippocampus*, **7**, 239–245.

Diamond, D. M., Park, C. R., Heman, K. L., and Rose, G. M. (1999). Exposing rats to predators impairs spatial working memory in the radial arm water maze. *Hippocampus*, **9**, 542–552.

Douglas, R. (1967). The hippocampus and behaviour. *Psychological Bulletin*, **67**(6), 416–442.

Dudchenko, P. A., Wood, E. R., and Eichenbaum, H. (2000). Neurotoxic hippocampal lesions have no effect on odour span and little effect on odour recognition memory but produce significant impairments on spatial span, recognition, and alternation. *Journal of Neuroscience*, **20**(8), 2964–2977.

Dusek, J. A. and Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences of the USA*, **94**, 7109–7114.

Dusek, J. A. and Eichenbaum, H. (1998). The hippocampus and transverse patterning guided by olfactory cues. *Behavioral Neuroscience*, **112**(4), 762–771.

Eichenbaum, H. (1994). The hippocampal system and declarative memory in humans and animals: Experimental analysis and historical origins. In D. L. Schacter and E. Tulving, editors, *Memory Systems*. MIT Press, Cambridge, MA.

Eichenbaum, H. (1997). How does the brain organize memories? *Science*, **277**(5324), 330–332.

Eichenbaum, H. (1999). Conscious awareness, memory and the hippocampus. *Nature Neuroscience*, **2**(9), 775–776.

Eichenbaum, H., Fagan, A., and Cohen, N. J. (1986). Normal olfactory discrimination-learning set and facilitation of reversal-learning after medial temporal damage in rats: implications for an account of preserved learning-abilities in amnesia. *Journal of Neuroscience*, **6**(7), 1876–1884.

Eichenbaum, H., Kuperstein, M., Fagan, A., and Nagode, J. (1987). Cue-sampling and goal-approach correlates of hippocampal unit activity in rats performing an odor-descrimination task. *Journal of Neuroscience*, **7**, 716–732.

Eichenbaum, H., Stewart, C., and Morris, R. G. M. (1990). Hippocampal representation in spatial learning. *Journal of Neuroscience*, **10**, 331–339.

Eichenbaum, H., Otto, T., and Cohen, N. J. (1992). The hippocampus - what does it do? *Behavioral and Neural Biology*, **57**, 2–36.

Eichenbaum, H., Otto, T., and Cohen, N. (1994). Two component functions of the hippocampal memory system. *Behavioural and Brain Sciences*.

Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., and Tanila, H. (1999). The hippocampus, memory, and place cells: Is it spatial memory or a memory space. *Neuron*, **23**, 209–226.

Eriksson, P., Perfilieva, E., Bjork-Eriksson, T., Alborn, A.-M., Nordborg, C., Peterson, D., and Gage, F. (1998). Neurogenesis in the adult human hippocampus. *Nature Medicine*, **4**(11), 1313–1317.

Etchamendy, N., Desmedt, A., Cortes-Torrea, C., Marighetto, A., and Jaffard, R. (2003). Hippocampal lesions and discrimination performance of mice in the radial maze: Sparing or impairment depending on the representational demands of the task. *Hippocampus*, **13**(2), 197–211.

Eustache, F., Rioux, P., Desgranges, B., andM.C. Petit-Taboue, G. M., Dary, M., Lechevalier, B., and Baron, J. (1995). Healthy aging, memory sub-systems and regional cerebral oxygen consumption. *Neuropsychologia*, **33**, 867–887.

Evans, J. J., Graham, K. S., Pratt, K. H., and Hodges, J. R. (2003). The impact of disrupted cortico-cortico connectivity: a long-term follow-up of a case of focal retrograde amnesia. *Cortex*, **39**, 767–790.

Fernandez, G., Klaver, P., Fell, J., Grunwald, T., and Elger, C. E. (2002). Human declarative memory formation: Segregating rhinal and hippocampal contributions. *Hippocampus*, **12**(4), 514–519.

Floresco, S., Seamans, J. K., and Phillips, A. G. (1996). A selective role for dopamine in the nucleus accumbens of the rat in random foraging but not delayed spatial win-shift-based foraging. *Behavioural Brain Research*, **80**(1-2), 161–168.

Frankland, P., O'Brien, C., Ohno, M., Kirkwood, A., and Silva, A. (2001). alpha-—CaMKII-dependent plasticity in the cortex is required for permanent memory. *Nature*, **411**(6835), 309–313.

Frankland, P. W., Cestari, V., Filipkowski, R. K., McDonald, R. J., and Silva, A. J. (1998). The dorsal hippocampus is essential for context discrimination but not for contextual conditioning. *Behavioural Neuroscience*, **112**(4), 863–874.

Frey, U. and Morris, R. (1998). Synaptic tagging: implications for late maintenance of hippocampal long-term potentiation. *Trends in Neurosciences*, **21**, 181–188.

Gabrieli, J. D. E., Brewer, J. B., Desmond, J. E., and Glover, G. (1997). Separate neural bases of two fundamental memory processes in the human medial temporal lobe. *Science*, **276**, 264–266.

Gadian, D. G., Aicardi, J., Watkins, K. E., Porter, D. A., Mishkin, M., and Vargha-Khadem, F. (2000). Developmental amnesia associated with early hypoxic-ischaemic injury. *Brain*, **123**(3), 499–507.

Gaffan, D. (1974). Recognition impaired and association intact in the memory of monkeys after transection of the fornix. *Journal of Comparative and Physiological Psychology*, **86**, 1100–1109.

Gaffan, D. (1994a). Dissociated effects of perirhinal cortex ablation, fornix transection and amygdalectomy: Evidence for multiple memory systems in the primate temporal lobe. *Experimental Brain Research*, **99**, 411–422.

Gaffan, D. (1994b). Scene-specific memory for objects: a model of episodic memory impairment in monkeys with fornix transection. *Journal of Cognitive Neuroscience*, **6**, 305–320.

Gaffan, D. and Parker, A. (1996). Interaction of perirhinal cortex with the fornix-fimbria: memory for objects and "object in place" memory. *Journal of Neuroscience*, **16**(18), 5864–5869.

Gaffan, D., Saunders, R. C., Gaffan, E. A., Harrison, S., Sheilds, C., and Owen, M. J. (1984). Effects of fornix transection upon associative memory in monkeys: role of the hippocampus in learned action. *Q J Experimental Psychology B*, **36**, 173–221.

Gaffan, E. A., Bannerman, D. M., and Healey, A. N. (2003). Learning associations between places and visual cues without learning to navigate: neither fornix nor entorhinal cortex is required. *Hippocampus*, **13**, 445–460.

Gall, F. J. (1835). *The influence of the brain on the form of the head*. Marsh, Capen and Lion.

Gallistel, C. R. (1990). *The organization of learning*. MIT Press, Cambridge, MA.

Gaskin, S., Tremblay, A., and Mumby, D. (2003). Retrograde and anterograde object recognition in rats with hippocampal lesions. *Hippocampus*.

Ghaem, O., Meller, E., Crivello, F., Tzourio, N., Mazoyer, B., Berthoz, A., and Denis, M. (1997). Mental navigation along memorized routes activates the hippocampus, precuneus and insula. *Neuroreport*, **8**, 739–744.

Gilbert, P., Kesner, R., and DeCoteau, W. (1998). Memory for spatial location: Role of the hippocampus in mediating spatial pattern separation. *Journal of Neuroscience*, **18**(2), 804–810.

Gilboa, A., Winocur, G., Grady, C., Hevenor, S., and Moscovitch, M. (2004). Remembering our past: Functional neuroanatomy of recollection of recent and very remote personal events. *Cerebral Cortex*, **14**(11), 1214–1225.

Gilissen, E. and Simmons, R. M. T. (2001). Brain evolution: a matter of constraints and permissions? *Behavioural and Brain Sciences*, **24**, 284 – 286.

Giurfa, M. and Calpaldi, E. A. (1999). Vectors, routes and maps: new discoveries about navigation in insects. *Trends in the Neurosciences*, **22**(6), 237–242.

Glisky, E. L., Schacter, D. L., and Tulving, E. (1986a). Computer learning by memory-impaired patients: acquisition and retention of complex knowledge. *Neuropsychologia*, **24**, 313–328.

Glisky, E. L., Schacter, D. L., and Tulving, E. (1986b). Learning and retention of computer-related vocabulary in memory-impaired patients: method of vanishing cues. *Journal of Clinical Experimental Neuropsychology*, **1986**, 292–312.

Gluck, M. (1996). Computational models of hippocampal function in memory. *Hippocampus*, **6**, 565–566.

Gluck, M. and Myers, C. (1997). Psychobiological models of hippocampal function in learning and memory. *Annual Review of Psychology*, **48**, 481–514.

Gluck, M., Meeter, M., and Myers, C. (2003). Computational models of the hippocampal region: linking incremental learning and episodic memory. *Trends in Cognitive Sciences*, **7**(6), 269–276.

Good, M., de Hoz, L., and Morris, R. M. (1998). Contingent versus incidental context processing during conditioning: Dissociation after excitotoxic hippocampal plus dentate gyrus lesions. *Hippocampus*, **8**, 147–159.

Goodale, M. A. and Milner, D. A. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, **15**, 20–25.

Goodhill, G., Simmen, M. W., and Willshaw, D. J. (1995). An evaluation of the use of multidimensional scaling for understanding brain connectivity. *Philosophical Transactions of the Royal Society: B*, **348**, 265–280.

Goshen-Gottstein, Y., Moscovitch, M., and Melo, B. (2000). Intact implicit memory for newly formed verbal associations in amnesic patients following single study trials. *Neuropsychology*, **14**(4), 570–578.

Gothard, K., Skaggs, W., and McNaughton, B. (1996). Dynamics of mismatch correction in the hippocampal ensemble code for space: Interaction between path integration and environmental cues. *Journal of Neuroscience*, **16**(24), 8027–8040.

Gould, E., Beylin, A., Tanapat, P., Reeves, A., and Shors, T. J. (1999). Learning enhances adult neogenesis in the hippocampal formation. *Nature*, **2**(3), 260–265.

Graf, P., Squire, L. R., and Mandler, G. (1984). The information that amnesic patients do not forget. *Journal of Experimental Psychology: Learning, memory and cognition*, **10**, 164–178.

Graham, K. S. and Hodges, J. R. (1997). Differentiating the roles of the hippocampal complex and the neocortex in long-term memory storage: evidence from the study of semantic dementia and Alzheimer's disease. *Neuropsychology*, **11**, 1–13.

Graham, K. S., Patterson, K., and Hodges, J. R. (1999). Episodic memory: new insights from the study of semantic dementia. *Current Opinion in Neurobiology*, **9**, 245–250.

Granger, R., Wiebe, S., Taketani, M., and Lynch, G. (1996). Distinct memory circuits composing the hippocampal region. *Hippocampus*, **6**, 567–578.

Gray, J. (1982). *Precis* of The neuropsychology of anxiety: an enquiry into the functions of the septo-hippocampal system. *Behavioral and Brain Sciences*, **5**, 469–484. See also commentary and response, p484-534.

Greenough, W. and Bailey, C. (1988). The anatomy of memory: convergence of results across a diversity of tests. *Trends in Neurosciences*, **11**(4), 142.

Griffiths, D., Dickinson, A., and Clayton, N. (1999). Episodic memory: what can animals remember about their past? *Trends in Cognitive Sciences*, **3**(2), 74–80.

Haist, F., Gore, J., and Mao, H. (2001). Consolidation of human memory over decades revealed by functional magnetic resonance imaging. *Nature Neuroscience*, **4**(11), 1139–1145.

Hamilton, D. A., Driscoll, I., and Sutherland, R. J. (2002). Human place learning in a virtual morris water task: some important constraints on the flexibility of place navigation. *Behavioural Brain Research*, **129**, 159–170.

Hampson, R. E., Jarrard, L. E., and Deadwyler, S. A. (1999). Effects of ibotenate hippocampal and extrahippocampal destruction on delayed-match and nonmatch-to-sample behavior in rats. *Journal of Neuroscience*, **19**, 1492–1507.

Harley, C. W. and Martin, G. M. (1999). Open field motor patterns and object marking, but not object sniffing, are altered by ibotenate lesions of the hippocampus. *Neurobiology of Learning and Memory*, **72**, 202–214.

Hasselmo, M., Wyble, B., and Wallenstein, G. (1996). Encoding and retrieval of episodic memories: Role of cholinergic and GABAergic modulation in the hippocampus. *Hippocampus*, **6**, 693–708.

Hasselmo, M. E. (1995). Neuromodulation and cortical function: modeling the physiological basis of behaviour. *Behavioural Brain Research*, **67**, 1–27.

Hebb, D. (1949). *The organization of behavior*. Wiley, New York.

Henke, K., Buck, A., Weber, B., and Wieser, H. G. (1997). Human hippocampus establishes associations in memory. *Hippocampus*, **7**(3), 249–256.

Henke, K., Kroll, N. E. A., Behniea, H., Amaral, D. G., Miller, M. B., Rafal, R., and Gazzaniga, M. S. (1999). Memory lost and regained following bilateral hippocampal damage. *Journal of Cognitive Neuroscience*, **11**(6), 682–697.

Hetherington, P. and Shapiro, M. (1993). A simple network simulates hippocampal place fields: 2. Computing goal-directed trajectories and memory fields. *Behavioral Neuroscience*, **107**(3), 434–443.

Hirsh, R. (1974). The hippocampus and contextual retrieval of information from memory: A theory. *Behavioral Biology*, **12**, 421–444.

Hock, B. J. and Bunsey, M. D. (1998). Differential effects of dorsal and ventral hippocampal lesions. *Journal of Neuroscience*, **18**(17), 7027–7032.

Hoh, T., Beiko, J., Boon, F., Weiss, S., and Cain, D. P. (1999). Complex behavioral strategy and reversal learning in the water maze without NMDA receptor-dependent long-term potentiation. *Journal of Neuroscience*, **19**(RC2), 1–5.

Holdstock, J., Mayes, A. R., Cezayirli, E., Isaac, C. L., Aggleton, J., and Roberts, N. (2000). A comparison of egocentric and allocentric spatial memory in a patient with selective hippocampal damage. *Neuropsychologia*, **38**(4), 410–425.

Holdstock, J., Mayes, A. R., Isaac, C. L., Gong, Q., and Roberts, N. (2002). Differential involvement of the hippocampus and temporal lobe cortices in rapid and slow learning of new semantic information. *Neuropsychologia*, **40**(7), 748–768.

Honey, R. C., Watt, A., and Good, M. (1998). Hippocampal lesions disrupt an associative-mismatch process. *Journal of Neuroscience*, **18**, 2226–2230.

Hopkins, R., Waldram, K., and Kesner, R. (2004). Sequences assessed by declarative and procedural tests of memory in amnesic patients with hippocampal damage. *Neuropsychologia*, **42**(14), 1877–1886.

Horn, G. (1998). Visual imprinting and the neural mechanism of recognition memory. *Trends in Neurosciences*, **21**(7), 300–305.

Huerta, P. T., Sun, L. D., Wilson, A. A., and Tonegawa, S. (2000). Formation of temporal memory requires NMDA receptors within CA1 pyramidal neurons. *Neuron*, **25**, 473–480.

Hunt, M. E., Kesner, R. P., and Evans, R. B. (1994). Memory for spatial location: functional dissociation of entorhinal cortex and hippocampus. *Psychobiology*, **22**, 186–194.

Ioalè, P., Gagliardo, A., and Bingman, V. P. (2000). Hippocampal participation in navigational map learning in young homing pigeons is dependent on training experience. *European Journal of Neuroscience*, **12**, 742–750.

Isaacs, E., Vargha-Khadem, F., Watkins, K., Lucas, A., Mishkin, M., and Gadian, D. (2003). Developmental amnesia and its relationship to degree of hippocampal atrophy. *Proceedings of the National Academy of Sciences of the USA*, **100**(22), 13060–13063.

Isaacson, R. (1974). *The Limbic System*. Plenum Press, New York.

Ishizuka, N., Weber, J., and Amaral, D. G. (1990). Organisation of intrahippocampal projections originating from CA3 pyramidal cells in the rat. *Journal of Comparative Neurology*, **295**, 580–623.

Ivanco, T. L. and Racine, R. J. (2000). Long-term potentiation in the reciprocal corticohippocampal and corticocortical pathways in the chronically-implanted, freely moving rat. *Hippocampus*, **10**, 143–152.

Izquierdo, I. and Medina, J. H. (1997). Memory formation: the sequence of biochemical events in the hippocampus and its connection to activity in other brain structures. *Neurobiology of Learning and Memory*, **68**, 285–316.

Izquierdo, I., Schröder, N., Netto, C. A., and Medina, J. H. (1999). Novelty causes time-dependent retrograde amnesis for one-trial avoidance in rats through NMDA receptor- and CaMKII-dependent mechanisms in the hippocampus. *European Journal of Neuroscience*, **11**, 3323–3328.

James, W. (1890). *The principles of Psychology*. Henry Holt.

Jeffrey, K. J. and Anderson, M. I. (2003). Dissociation of the geometric and contextual influences on place cells. *Hippocampus*, **13**, 1–5.

Jenkins, J. G. and Dallenbach, K. M. (1924). Obliviscence during sleep and waking. *American Journal of Psychology*, **35**, 605–612.

Jonasson, Z., Ballantyne, J. K., and Baxter, M. G. (2004). Preserved anterograde and retrograde memory of rapidly acquired olfactory discriminations after neurotoxic hippocampal lesions. *Hippocampus*, **14**, 28–39.

Judge, M. E. and Quartermain, D. (1982). Characteristics of retrograde amnesia following reactivation of memory in mice. *Physiology and Behaviour*, **28**, 585–590.

Jung, M. W., Wiener, S. L., and McNaughton, B. L. (1994). Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *Journal of Neuroscience*, **14**, 7347–7356.

Kandel, E., Schwartz, J., and Jessell, T. (1995). *Essentials of Neural Science and Behaviour*. Prentice Hall International, international edition.

Kapur, N. (2000). Focal retrograde amnesia and the attribution of causality: an exceptionally benign commentary. *Cognitive Neuropsychology*, **17**(7), 623–637.

Kapur, N. and Brooks, D. J. (1999). Temporally-specific retrograde amnesia in two cases of discrete bilateral hippocampal pathology. *Hippocampus*, **9**, 247–254.

Karni, A., Tanne, D., Rubenstein, B. S., Askenasy, J. J. M., and Sagi, D. (1994). Dependence on REM-sleep of overnight improvement of a perceptual skill. *Science*, **265**(5172), 679–682.

Kartsounis, L. D., Rudge, P., and Stevens, J. M. (1995). Bilateral lesions of CA1 and CA2 fields of the hippocampus are sufficient to cause a severe amnesic syndrome in humans. *Journal of Neurology, Neurosurgery and Psychiatry*, **59**, 95–98.

Kesner, R. (1998). Neurobiological view of memory. In J. Martinez and R. Kesner, editors, *Neurobiology of Learning and Memory*, pages 361–416. Academic Press.

Kesner, R. and DiMattia, B. (1987). Neurobiology of an attribute model of memory. In *Progress in Psychology and Physiological Psychology*. Academic Press, New York.

Kesner, R., Gilbert, P., and Barua, L. (2002). The role of the hippocampus in memory for the temporal order of a sequence of odors. *Behavioural Neuroscience*, **116**(2), 286–290.

Kesner, R. P., Bolland, B. L., and Dakis, M. (1993). Memory for spatial locations, motor responses, and objects: Triple dissociation among the hippocampus, caudate nucleus, and extrastriate visual cortex. *Experimental Brain Research*, **93**, 462–470.

Kessels, R., de Haan, E., Kappelle, L., and Postma, A. (2001). Varieties of human spatial memory: a meta-analysis on the effects of hippocampal lesions. *Brain Research Reviews*, **35**(3), 295–303.

Kim, J. J. and Fanselow, M. S. (1992). Modality-specific retrograde amnesia of fear. *Science*, **256**, 675–677.

Kim, J. J., Clark, R. E., and Thompson, R. F. (1995). Hippocampectomy impairs the memory of recently, but not remotely, acquired trace eyeblink conditioned responses. *Behavioural Neuroscience*, **109**(2), 195–203.

Kimble, D. (1968). Hippocampus and internal inhibition. *Psychological Bulletin*, **70**(5), 285–295.

Kimble, D. P. and Pribram, K. H. (1963). Hippocampectomy and behavior sequences. *Science*, **139**, 824–825.

Kirkwood, A., Lee, H. K., and Bear, M. F. (1995). Co-regulation of long-term potentiation and experience-dependent synaptic plasticity in visual cortex by age and experience. *Nature*, **375**, 328–331.

Kitchener, E. G., Hodges, J. R., and McCarthy, R. (1998). Acquisition of post-morbid vocabulary and semantic facts in the absence of episodic memory. *Brain*, **121**, 1313–1327.

Knierim, J., Kudrimoti, H., and McNaughton, B. (1995). Place cells, head direction and the learning of landmark stability. *Journal of Neuroscience*, **15**, 1648–1659.

Knight, R. (1996). Contribution of human hippocampal region to novelty. *Nature*, **383**, 256–259.

Knowlton, B. J. and Fanselow, M. S. (1998). The hippocampus, consolidation and online memory. *Current Opinion in Neurobiology*, **8**, 293–296.

Knowlton, B. J. and Squire, L. R. (1993). The learning of categories: Parallel brain systems for item memory and category knowledge. *Science*, **262**(5140), 1747–1749.

Knowlton, B. J. and Squire, L. R. (1995). Remembering and knowing: Two different expression of declarative memory. *Journal of Experimental Psychology*, **21**(3), 699–710.

Knowlton, B. J. and Squire, L. R. (1996). Artificial grammar learning depends on implicit acquisition of both abstract and exemplar-specific information. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **22**(1), 169–181.

Knowlton, B. J., Ramus, S. J., and Squire, L. R. (1992). Intact artificial grammar learning in amnesia - dissociation of classification learning and explicit memory for specific instances. *Psychological Science*, **3**(3), 172–179.

Kogan, J. H., Frankland, P. W., and Silva, A. J. (2000). Long-term memory underlying hippocampus-dependent social recognition in mice. *Hippocampus*, **10**, 47–56.

Kopelman, M, D. and Kapur, N. (2001). The loss of episodic memories in retrograde amnesia: single-case and group studies. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences*, **356**(1413), 1409–1421.

Kopelman, M. D. (2000). Focal retrograde amnesia and the attribution of causality: an exceptionally critical review. *Cognitive Neuropsychology*, **17**(7), 585–621.

Kopelman, M. D. (2002). Disorders of memory. *Brain*, **125**, 2152–2190.

Kubie, J. L., Sutherland, R. J., and Muller, R. U. (1999). Hippocampal lesions produce a temporally graded retrograde retrograde amnesia on a dry version of the Morris swimming task. *Psychobiology*, **27**, 313–330.

Kudrimoti, H. S., Barnes, C. A., and McNaughton, B. L. (1999). Reactivation of hippocampal cell assemblies: Effects of behavioral state, experience, and EEG dynamics. *Journal of Neuroscience*, **19**, 4090–4101.

Land, C., Bunsey, M., and Riccio, D. C. (2000). Anomalous properties of hippocampal lesion-induced retrograde amnesia. *Psychobiology*, **28**(4), 476–485.

Lang, C. and Bastian, A. (2001). Additional somatosensory information does not improve cerebellar adaptation during catching. *Clinical Neurophysiology*, **112**(5), 895–907.

Lashley, K. S. (1950). In search of the engram. *Symposia for the Society of Experimental Biology*, **4**, 454–482.

Lathe, R. (2000). The enteroceptive hippocampus. Unpublished.

Lavenex, P. and Amaral, D. G. (2000). Hippocampal-neocortical interaction: A hierarchy of associativity. *Hippocampus*, **10**(4), 420–430.

Lemaire, V., Aurousseau, C., Le Moal, A., and Abrous, D. N. (1999). Behavioural trait of reactivity to novelty is related to hippocampal neurogenesis. *European Journal of Neuroscience*, **11**, 4006–4016.

Leonard, B. J., McNaughton, B. L., and Barnes, C. A. (1987). Suppression of hippocampal synaptic plasticity during slow-wave sleep. *Brain Research*, **425**(1), 174–177.

Levine, B., Black, S., Cabeza, R., Sinden, M., Mcintosh, A., Toth, J., Tulving, E., and Stuss, D. (1998). Episodic memory and the self in a case of isolated retrograde amnesia. *Brain*, **121**(10), 1951–1973.

Levy, W. B. (1996). A sequence predicting CA3 is a flexible associator that lear ns and uses context to solve hippocampal-like tasks. *Hippocampus*, **6**(6), 579–590.

Liaw, J. and Berger, T. W. (1996). Dynamic synapse: A new concept of neural representation and computation. *Hippocampus*, **6**, 591–600.

Lisman, J. E. (1999). Relating hippocampal circuitry to function: recall of memory sequences by reciprocal dentate-CA3 interactions. *Neuron*, **22**(2), 233–242.

Lorente de Nó, R. (1934). Studies on the structure of the cerebral cortex: II. Continuation of the study of the ammonic system. *Journal of Physiological Neurology*, **46**, 113–177.

Luo, Q., Perry, C., Peng, D., Jin, Z., Xu, D., Ding, G., and Xu, S. (2003). The neural substrate of analogical reasoning: an fMRI study. *Cognitive Brain Research*, **17**(3), 527–534.

Lynch, G. and Granger, R. (1992). Variations in synaptic plasticity and types of memory in corticohippocampal networks. *Journal of Cognitive Neuroscience*, **4**(3), 189–199.

Mactutus, C. F., Riccio, D. C., and Ferek, J. M. (1979). Retrograde amnesia for old reactivated memory: Some anomalous characteristics. *Science*, **74**, 1319–1320.

Mactutus, C. F., Ferek, J. M., George, C. A., and Riccio, D. C. (1982). Hypothermia-induced amnesia for newly acquired and old reactivated memories: commonalities and distinctions. *Physiological Psychology*, **10**, 79–95.

Maguire, E., Henson, R., Mummery, C., and Frith, C. D. (2001a). Activity in prefrontal cortex, not hippocampus, varies parametrically with the increasing remoteness of memories. *Neuroreport*, **12**(3), 441–444.

Maguire, E., Vargha-Khadem, F., and Mishkin, M. (2001b). The effects of bilateral hippocampal damage on fMRI regional activations and interactions during memory retrieval. *Brain*, **124**, 1156–1170.

Maguire, E. A. and Frith, C. D. (2003). Lateral asymmetry in the hippocampal response to the remoteness of autobiographical memories. *The Journal of Neuroscience*, **23**(12), 5302–5307.

Maguire, E. A. and Mummery, C. J. (1999). Differential modulation of a common memory retrieval network revealed by positron emission tomography. *Hippocampus*, **9**, 54–61.

Maguire, E. A., Frackowiak, R. S. J., and Frith, C. D. (1996). Learning to find your way: A role for the human hippocampal formation. *Proceedings of the Royal Society of London Series B-Biological Sciences*, **263**, 1745–1750.

Maguire, E. A., Frackowiak, R. S. J., and Frith, C. D. (1997). Recalling routes around London: Activation of the right hippocampus in taxi drivers. *J. Neuroscience*, **17**, 7103–7110.

Maguire, E. A., Burgess, N., Donnett, J. G., Frackowiak, R. S. J., Frith, C. D., and Okeefe, J. (1998a). Knowing where and getting there: A human navigation network. *Science*, **280**, 921–924.

Maguire, E. A., Frith, C. D., Burgess, N., Donnett, J. G., and Okeefe, J. (1998b). Knowing where things are: Parahippocampal involvement in encoding object locations in virtual large-scale space. *J. Cognitive Neuroscience*, **10**, 61–76.

Mandler, G. (1980). Recognizing: The judgement of previous occurrence. *Psychological Review*, **87**, 252–271.

Manns, J. R. and Squire, L. R. (1999). Impaired recognition on the Doors and People Test after damage limited to the hippocampal region. *Hippocampus*, **9**, 495–499.

Manns, J. R., Clark, R. E., and Squire, L. R. (2000). Awareness predicts the magnitude of single-cue trace eyeblink conditioning. *Hippocampus*, **10**, 181–186.

Maren, S., Aharonov, G., and Fanselow, M. S. (1997). Neurotoxic lesions of the dorsal hippocampus and Pavlovian fear conditioning in rats. *Behavioral Brain Research*, **88**, 261–274.

Markus, E., Barnes, C., McNaughton, B., Gladden, V., and Skaggs, W. (1994). Spatial information content and reliability of hippocampal CA1 neurons: Effects of visual input. *Hippocampus*, **4**, 410–421.

Markus, E., Qin, Y., Leonard, B., Skaggs, W., McNaughton, B., and Barnes, C. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *Journal of Neuroscience*, **15**, 7079–7094.

Marr, D. (1970). A theory for cerebral neocortex. *Procedings of the Royal Society*, **176**, 161–234.

Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London*, **262**(841), 23–81.

Martin, A. (1999). Automatic activation of the medial temporal lobe during encoding: lateralized influences of meaning and novelty. *Hippocampus*, **9**, 62–70.

Martin, A., Wiggs, C. L., Ungerleider, L. G., and Haxby, J. V. (1996). Neural correlates of category-specific knowledge. *Nature*, **379**, 649–652.

Maviel, T., Durkin, T., Menzaghi, F., and Bontempi, B. (2004). Sites of neocortical reorganization critical for remote spatial memory. *Science*, **305**(5680), 96–99.

Mayes, A., Isaac, C. L., Holdstock, J. S., Hunkin, N. M., Montaldi, D., Downes, J. J., MacDonald, C., Cezayirli, E., and Roberts, J. (2001). Memory for single items, word pairs, and temporal order of different kinds in a patient with selective hippocampal lesions. *Cognitive Neuropsychology*, **18**(2), 97–123.

Mayes, A., Montaldi, D., Spencer, T., and Roberts, N. (2004). Recalling spatial information as a component of recently and remotely acquired episodic or semantic memories: an fMRI study. *Neuropsychology*, **18**(3), 426–441.

Mayes, A. R., Isaac, C. L., Holdstock, J. S., Cariga, P., Gummer, A., and Roberts, N. (2003). Long-term amnesia: a review and detailed illustrative case study. *Cortex*, **39**, 567–603.

McClelland, J. L. and Goddard, N. H. (1996). Considerations arising from a complementary learning systems perspective on hippocampus and neocortex. *Hippocampus*, **6**(6), 654–665.

McClelland, J. L., McNaughton, B. L., and O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, **102**(3), 419–457.

McDonald, R. and White, N. (1995). *Behavioural Neuroscience*, **109**.

McDonald, R. J. and Hong, N. S. (2000). Rats with hippocampal damage are impaired on place learning in the water task when overtrained under constrained conditions. *Hippocampus*, **10**, 153–161.

McDonald, R. J. and White, N. (1993). A triple dissociation of memory systems: Hippocampus, amygdala, and dorsal striatum. *Behavioural Neuroscience*, **107**, 3–22.

McDonald, R. J., Murphy, R. A., Guarraci, F. A., Gortler, J. R., White, N. M., and G.Baker, A. (1997). Systematic comparison of the effects of hippocampal and fornix-fimbria lesions on acquisition of three configural discriminations. *Hippocampus*, **7**, 371–388.

McEchron, M. D. and Disterhoft, J. F. (1997). Sequence of single neuron changes in CA1 hippocampus of rabbits during acquisition of trace eyeblink conditioned responses. *Journal of Neurophysiology*, **78**, 1030–1044.

McEchron, M. D. and Disterhoft, J. F. (1999). Hippocampal encoding of non-spatial trace conditioning. *Hippocampus*, **9**, 385–396.

McEchron, M. D., Bouwmeester, H., Tseng, W., Weiss, C., and Disterhoft, J. F. (1998). Hippocampectomy disrupts auditory trace fear conditioning and contextual fear conditioning in the rat. *Hippocampus*, **8**, 636–646.

McEwen, B. S. (1999). Stress and hippocampal plasticity. *Annual Review of Neuroscience*, **22**, 105–122.

McKee, R. D. and Squire, L. R. (1993). On the development of declarative memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **19**, 397–404.

McNaughton, B. (1989). Neuronal mechanisms for spatial computation and information storage. In L. Nadel, L. Cooper, P. Culicover, and R. Harnish, editors, *Neural connections, Mental computation*, chapter 9, pages 285–350. MIT Press, Cambridge, MA.

McNaughton, B. and Morris, R. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences*, **10**(10), 408–415.

McNaughton, B., Barnes, C., Meltzer, J., and Sutherland, R. (1989). Hippocampal granule cells are necessary for spatial learning but not for spatially-selective pyramidal cell discharge. *Experimental Brain Research*, **76**, 485–496.

McNaughton, B., Chen, L., and Markus, E. (1991). "Dead reckoning," landmark learning and the sense of direction: A neurophysiological and computational hypothesis. *Journal of Cognitive Science*, **3**(2), 190–202.

McNaughton, B., Barnes, C., Gerrard, J., Gothard, K., Knierim, M., Kudrimoti, J., Qin, Y., Skaggs, W., Suster, M., and Weaver, K. (1996). Deciphering the hippocampal polyglot: The hippocampus as a path integration system. *Journal of Experimental Biology*, **199**(1), 173–186.

Means, L. W., Alexander, S. R., and O'Neal, M. F. (1992). Those cheating rats: male and female rats use odor trails in a water-escape "working memory" task. *Behav Neural Biol*, **58**, 144–151.

Meeter, M. and Murre, J. M. J. (2004). Consolidation of long-term memory: Evidence and alternatives. *Psychological Bulletin*, **130**, 843–857.

Mehta, M., Barnes, C., and McNaughton, B. (1997). Experience-dependent, asymmetric expansion of hippocampal place fields. *Proceedings of the National Academy of Sciences of the United States of America*, **94**(16), 8918–8921.

Merigan, W. H. and Maunsell, J. H. R. (1993). How parallel are the primate visual pathways? *Annual Review in the Neurosciences*, **16**, 369–402.

Meunier, M., Hadfield, J., Bachevalier, J., and Murray, E. A. (1996). Effects of rhinal cortex lesions combined with hippocampectomy on visual recognition memory in rhesus monkeys. *Journal of Neurophysiology*, **75**, 1190–1205.

Milekic, M. H. and Alberini, C. M. (2002). Temporally graded requirement for protein synthesis following memory reactivation. *Neuron*, **36**, 521–525.

Milner, B. (1962). Les troubles de la mémoire accompagnant des lésions hippocampiques bilatérales. *Physiologie de l'hippocampe*, pages 257–270.

Milner, P. (1989). A cell assembly theory of hippocampal amnesia. *Neuropsychologia*, **27**(1), 23–30.

Minai, A. A. and Levy, W. B. (1994). Sequence learning in a single trial. In *INNS World Congress on Neural Networks*, pages I:582–586.

Misanin, J. R., Miller, R. R., and Lewis, D. J. (1968). Retrograde amnesia produced by electroconvulsive shock after reactivation of consolidated memory trace. *Science*, **160**, 554–555.

Mishkin, M. and Petri, H. (1984). Memories and habits: Some implications for the analysis of learning and retention. In L. R. Squire and N. Butters, editors, *Neuropsychology of memory*. Guilford Press, New York.

Mishkin, M., Ungerleider, L. G., and Macko, K. A. (1983). Object vision and spatial vision: two cortical pathways. *Trends in the Neurosciences*, **6**, 414–417.

Mishkin, M., Vargha-Khadem, F., and Gadian, D. G. (1998). Amnesia and the organization of the hippocampal system. *Hippocampus*, **8**, 212–216.

Miyashita, Y., Kameyama, M., Hasegawa, I., and Fukushima, L. (1998). Consolidation of visual associative long-term memory in the temporal cortex of primates. *Neurobiology of Learning and Memory*, **70**(1-2), 197–211.

Miyashita *et al*, Y. (1994). Society for Neuroscience Abstracts, 20:248.

Moll, M., Miikkulainen, R., and J.Abbey (1994). The capacity of convergence-zone episodic memory. In *Proceedings of the Twelth National Conference on Artificial Intelligence*.

Morris, R. and Frey, U. (1997). Hippocampal synaptic plasticity: role in spatial learning or the automatic recording of attended experience? *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, **352**(1360), 1489–1503.

Morris, R., Garrud, P., Rawlins, J., and O'KeefeE, J. (1982). Place navigation impaired in rats with hippocampal-lesions. *Nature*, **297**(5868), 681–683.

Morris, R. G. M. (1999). Topographical knowledge survives hippocampal damage. *Current Biology*, **9**, R890–892.

Morris, R. G. M., Schenk, F., Tweedie, F., and Jarrard, L. E. (1990). Ibotenate lesions of hippocampus and/or subiculum: Dissociating components of allocentric learning. *European Journal of Neuroscience*, **2**, 1016–1028.

Morton, J., Hammersley, R. H., and Bekerian, D. A. (1985). Headed records: a model for memory and its failures. *Cognition*, **20**, 1–23.

Moscovitch, M. and Nadel, L. (1998). Consolidation and the hippocampal complex revisited: in defense of the multiple-trace model. *Current Opinion in Neurobiology*, **8**, 297–300.

Moscovitch, M. and Nadel, L. (1999). Multiple-trace theory and semantic dementia. *Trends in Cognitive Sciences*, **3**(3), 87–89.

Moser, M.-B. and Moser, E. (1998). Functional differentiation in the hippocampus. *Hippocampus*, **8**, 608–619.

Moser, M.-B. and Moser, E. I. (2000). Pretraining and the function of hippocampal long-term potentiation. *Neuron*, **26**, 559–561.

Moser, M. B., Moser, E. I., Forrest, E., Andersen, P., and Morris, R. G. M. (1995). Spatial learning with a minislab in the dorsal hippocampus. *Proceedings of the National Academy of Sciences of the United States of America*, **92**, 9697–9701.

Moss, H. E., Kopelman, M. D., Cappelletti, M., de Mornay Davies, P., and Jaldow, E. (2003). Lost for words or loss of memories? autobiographical memory in semantic dementia. *Cognitive Neuropsychology*, **20**(8), 703–732.

Moyer, J. R., Deyo, R. A., and Disterhoft, J. F. (1990). Hippocampectomy disrupts trace eye-blink conditioning in rabbits. *Behavioral Neuroscience*, **104**, 243–252.

Muller, R. (1996). A quarter century of place cells. *Neuron*, **17**, 813–822.

Muller, R., Kubie, J., and Ranck, J. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *Journal of Neuroscience*, **7**, 1935–1950.

Muller, R., Kubie, J., and Saypoff, R. (1991a). The hippocampus as a cognitive graph. *Hippocampus*, **1**(3), 243–246.

Muller, R., Kubie, J., Bostock, E., Taube, E., and Quirk, G. (1991b). Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In J. Paillard, editor, *Brain and Space*, chapter 17, pages 296–333. Oxford University Press, New York.

Muller, R., Stead, M., and Pach, J. (1996). The hippocampus as a cognitive graph. *Journal of General Physiology*, **107**(6), 663–694.

Muller, R. U. and Stead, M. (1996). Hippocampal place cells connected by Hebbian synapses can solve spatial problems. *Hippocampus*, **6**, 709–719.

Mumby, D. G. (2001). Perspectives on object-recognition memory following hippocampal damage: lessons from studies in rats. *Behavioural Brain Research*, **127**(1), 159–181.

Mumby, D. G., Wood, E. R., Duva, C. A., Kornecook, T. J., Pinel, P. J., and Phillips, A. G. (1996). Ischaemia-induced object-recognition deficits in rats are attenuated by hippocampal ablation before or soon after ischemia. *Behavioral Neuroscience*, **110**, 266–281.

Murray, E., Gaffan, D., and Mishkin, M. (1993). Severe tactual as well as visual memory deficits follow combined removal of the amygdala and hippocampus in monkeys. *Journal of Neuroscience*, **4**, 2565–2580.

Murray, E. A. (1996). What have ablation studies told us about the neural substrates of stimulus memory? *Seminars in the Neurosciences*, **8**, 13–22.

Murray, E. A. and Bussey, T. J. (2001). Consolidation and the medial temporal lobe revisited: Methodological considerations. *Hippocampus*, **11**(1), 1–7.

Murray, E. A. and Mishkin, M. (1998). Object recognition and location memory in monkeys with excitotoxic lesions of the amygdala and hippocampus. *Journal of Neuroscience*, **18**, 6568–6582.

Murray, E. A. and Wise, S. P. (1996). Role of the hippocampus plus subjacent cortex but not amygdala in visuomotor conditional learning in rhesus monkeys. *Behavioral Neuroscience*, **110**, 1261–1270.

Murray, E. A., Bussey, T. J., Hampton, R. R., and Saksida, L. M. (2000). The parahippocampal region and object identification. *Annals of the New York Academy of Sciences*, **911**, 166–174.

Murre, J. M. and Sturdy, D. P. (1996). The connectivity of the brain: multi-level quantitative analysis. *Biological Cybernetics*, **73**(6), 529–545.

Murre, J. M. J. (1996). Tracelink: A model of amnesia and consolidation of memory. *Hippocampus*, **6**, 675–684.

Murre, J. M. J. (1999). Interaction of cortex and hippocampus in a model of amnesia and semantic dementia. *Reviews in the Neurosciences*, **10**, 267–278.

Musen, A. and Squire, L. R. (1993a). Implicit learning of colour-word associations in a Stroop paradigm. *Journal of Experimental Psychology: Learning, memory and cognition*, **19**, 789–798.

Musen, A. and Squire, L. R. (1993b). On the implicit learning of novel associations by amnesic patients and normal subjects. *Neuropsychology*, **119-135**.

Myers, C. E., Ermita, B., Harris, K., Hasselmo, M., Solomon, P., and Gluck, M. A. (1996). A computational model of the effects of septohippocampal disruption on classical eyeblink conditioning. *Neurobiology of Learning and Memory*, **66**, 51–66.

Nadel, L. (1994). Multiple memory systems: What and why, an update. In D. L. Schacter and E. Tulving, editors, *Memory Systems*. MIT Press, Cambridge, MA.

Nadel, L. and Bohbot, V. (2001). Consolidation of memory. *Hippocampus*, **11**(1), 56–60.

Nadel, L. and Moscovitch, M. (1997). Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology*, **7**, 217–227.

Nadel, L., Samsonovich, A., Ryan, L., and Moscovitch, M. (2000). Multiple trace theory of human memory: Computational, neuroimaging, and neuropsychological results. *Hippocampus*, **10**(4), 352–368.

Nader, K., Schafe, G., and Le Doux, J. E. (2000). Fear memories required protein synthesis in the amygdala for reconsolidation after retrieval. *Nature*, **206**, 722–726.

Neisser, U. and Harsh, N. (1992). Phantom flashbulbs: False recollections of hearing the news about challenger. In E. Winograd and U. Neisser, editors, *Accuracy and affect in recall: Studies of flashbulb memories*, pages 9–31. Cambridge University Press, Cambridge.

Nestor, P. J., Graham, K. S., Bozeat, S., Simons, J. S., and Hodges, J. R. (2002). Memory consolidation and the hippocampus: further evidence from studies of autobiographical memory in semantic dementia and frontal variant frontotemporal dementia. *Neuropsychologia*, **40**, 633–654.

Niki, K. and Luo, J. (2002). An fMRI study on the time-limited role of the medial temporal lobe in long-term topographical autobiographic memory. *Journal of Cognitive Neuroscience*, **14**(3), 500–507.

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

Nitz, D. A., Pauer, M., and McNaughton, B. L. (1997). Divergent modulation of CA1 and DG interneurons by spatial novelty. *Society of Neuroscience Abstracts*, **23**, 196.4.

Nyberg, L., Tulving, E., Habib, R., Nilsson, L., Kapur, S., Houle, S., Cabeza, R., and McIntosh, A. R. (1995). Functional brain maps of retrieval mode and recovery of episodic information. *Neuroreport*, **6**, 249–252.

Odegard, T. N. and Lampinen, J. M. (2004). Memory conjunction errors for autobiographical events: More than just familiarity. *Memory*, **12**(3), 288–300.

Ogden, J. A. (1993). Visual object agnosia, prosopagnosia, achromatopsia, loss of visual imagery and autobiographical amnesia following recovery from cortical blindness. *Neuropsychologia*, **31**, 571–589.

Ohl, F. and Fuchs, E. (1999). Differential effects of chronic stress on memory processes in the tree shrew. *Cognitive Brain Research*, **7**, 379–387.

Okaichi, H. (1996). Effects of fimbria-fornix lesions on door discrimination and route choice in a lattice maze by rats. *Neurobiology of Learning and Memory*, **66**, 155–166.

O'Kane, G., Kensinger, E. A., and Corkin, S. (2004). Evidence for semantic learning in profound amnesia: an investigation with patient H.M. *Hippocampus*, **14**, 417–425.

O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Experimental Neurology*, **51**, 78–109.

O'Keefe, J. (1996). The spatial prepositions in English, vector grammar and the cognitive map theory. In P. Bloom, M. Peterson, L. Nadel, and M. Garrett, editors, *Language and Space*, pages 277–316. MIT Press, Cambridge, MA.

O'Keefe, J. (1999). Do hippocampal pyramidal cells signal non-spatial as well as spatial information. *Hippocampus*, **9**(4), 352–364.

O'Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely moving rat. *Experimental Brain Research*, **34**, 171–175.

O'Keefe, J. and Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford University Press.

Olton, D. (1979). Mazes, maps and memory. *Amer. Psychol.*, **34**, 583–596.

Olton, D. (1983). Memory functions and the hippocampus. In W. Seifert, editor, *Neurobiology of the hippocampus*, pages 335–374. Academic Press, London.

Olton, D., Becker, J., and Handelmann, G. (1979). Hippocampus, space amd memory. *Behav. Brain Sci.*, **2**, 313–322.

Orbach, J., Milner, B., and Rasmussen, T. (1960). Learning and retention in monkeys after amygdala-hippocampus resection. *Archives of Neurology*, **3**, 230–251.

O'Reilly, R. and Norman, K. A. (2002). Hippocampal and neocortical contributions to memory: advances ion the complementary learning systems framework. *Trends in Cognitive Sciences*, **6**(12), 505 –510.

Ostergaard, A. L. (1987). Episodic, semantic and procedural memory in a case of amnesia at an early age. *Neuropsychologia*, **25**, 341–357.

Owen, A., Milner, B., Petrides, M., and Evans, A. (1996). A specific role for the right parahippocampal gyrus in the retrieval of object-location: A positron emission tomography study. *Journal of Cognitive Neuroscience*, **8**(6), 588–602.

Oxbury, S., Oxbury, J., Renowden, S., Squier, W., and Carpenter, K. (1997). Severe amnesia: An unusual late complication after temporal lobectomy. *Neuropsychologia*, **35**(7), 975–988.

Packard, M. G. (1999). Glutamate infused posttraining into the hippocampus or caudate-putamen differentially strengthens place and response learning. *Proceedings of the National Academy of Sciences of the USA*, **96**(22), 12881–12886.

Packard, M. G. and Cahill, L. (2001). Affective modulation of multiple memory systems. *Current Opinion in Neurobiology*, **11**, 752–756.

Packard, M. G. and McGaugh, J. L. (1996). Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiology of Learning and Memory*, **65**, 65–72.

Packard, M. G. and Teather, L. A. (1999). Dissociation of multiple memory systems by post-training intracerebral injections of glutamate. *Psychobiology*, **27**, 40–50.

Packard, M. G., Hirsh, R., and White, N. (1989). Differential effects of fornix and caudate nucleus lesions on two radial maze tasks: Evidence for multiple memory systems. *Journal of Neuroscience*, **9**, 1465–1472.

Papanicolaou, A., Simos, P., Castillo, E., Breier, J., Katz, J., and Wright, A. (2002). The hippocampus and memory of verbal and pictorial material. *Learning and memory*, **9**(3), 99–104.

Papez, J. (1937). A proposed mechanism of emotion. *Archives of Neurology and Psychology*, **38**, 725–743.

Parkin, A. J. (1997). Human memory: Novelty, association and the brain. *Current Biology*, **7**, R768–769.

Parkinson, J. K., Murray, E. A., and Mishkin, M. (1988). How spatial memory and episodic memory are related to each other in humans and primates remains a high-priority issue for research. *Journal of Neuroscience*, **8**, 4159.

Parron, C., Poucet, B., and Save, E. (2001). Re-evaluation of the spatial memory deficits induced by hippocampal short lasting inactivation reveals the need for cortical co-operation. *Behavioural Brain Research*, **127**, 71–79.

Pearce, J. M., Roberts, A. D. L., and Good, M. (1998). Hippocampal lesions disrupt navigation based on cognitive maps but not heading vectors. *Nature*, **396**, 75–77.

Penttonen, M., Kamondi, A., Sik, A., Acsady, L., and Buzsaki, G. (1997). Feed-forward and feed-back activation of the dentate gyrus in vivo during dentate spikes and sharp wave bursts. *Hippocampus*, **7**(4), 437–450.

Phillips, R. G. and Le Doux, J. E. (1992). Differential contribution of amygdala and hippocampus to cued and contextual fear conditioning. *Behavioral Neuroscience*, **106**(2), 274–285.

Phillips, R. G. and LeDoux, J. E. (1994). Lesions of the dorsal hippocampal formation interfere with background but not foreground contextual conditioning. *Learning and Memory*, **1**, 34–44.

Piefke, M., Weiss, P., Zilles, K., Markowitsch, H., and Fink, G. (2003). Differential remoteness and emotional tone modulate the neural correlates of autobiographical memory. *Brain*, **126**, 650–668.

Piolino, P., Belliard, S., Desgranges, B., Perron, M., and Eustache, F. (2003a). Autobiographical memory and autonoetic consciousness in a case of semantic dementia. *Cognitive Neuropsychology*, **20**(7), 619–639.

Piolino, P., Desgranges, B., Belliard, S., , Matuszawski, V., Lalevee, C., de la Sayette, V., and Eustache, F. (2003b). Autobiographical memory and autonoetic consciousness: triple dissociation in neurodegenerative diseases. *Brain*, **126**, 2203–2219.

Piolino, P., Giffard-Quillon, G., Desgranges, B., Chetelat, G., and Baron, JC.. Eustache, F. (2004). Re-experiencing old memories via hippocampus: a PET study of autobiographical memory. *Neuroimage*, **22**(3), 1371–1383.

Poldrack, R., Clark, J., Pare-Blagoev, E., Shohamy, D., Moyano, J., Myers, C., and Gluck, M. (2001). Interactive memory systems in the human brain. *Nature*, **414**(6863), 546–550.

Poliakoff, E. and Meudell, P. R. (2000). New learning and remote memory in the same and different domains of experience: implications for normal memory and amnesia. *Cortex*, **36**, 195–211.

Poucet, B., Hermann, T., and Buhot, M. C. (1991). Effects of short-lasting inactivations of the ventral hippocampus and medial septum on long-term and short-term acquisition of spatial information in rats. *Behavioural Brain Research*, **44**, 53–65.

Price, C. J., Wise, R. J. S., Watson, J. D. G., Patterson, K., Howard, D., and Frackowiak, R. S. J. (1994). Brain activity during reading - the effects of exposure duration and task. *Brain*, **117**, 1255–1269.

Przybyslawski, J. and Sara, S. J. (1997). Reconsolidation of memory after its reactivation. *Behavioural Brain Research*, **84**, 241–246.

Przybyslawski, J., Roullet, P., and Sara, S. J. (1999). Attenuation of emotional and nonemotional memories after their reactivation: role of β adrenergic receptors. *Journal of Neuroscience*, **19**(15), 6623–6628.

Qin, Y.-L., McNaughton, B. L., Skaggs, W. E., and Barnes, C. A. (1995). Reactivation during sleep of cortico-cortical and hippocampo-cortical correlation states from preceding behaviour. *Society for Neuroscience Abstracts*, **21**, 941.

Qin, Y. L., McNaughton, B. L., Skaggs, W. E., and Barnes, C. A. (1997). Memory reprocessing in corticocortical and hippocampocortical neuronal ensembles. *Philos Trans R Soc Lond B Biol Sci*, **352**(1360), 1525–1533.

Quillfeldt, J. A., Zanatta, M. S., Schmitz, P., Quevedo, J., Schaeffer, E., DeLima, J. B., Medina, J. H., and Izquierdo, I. (1996). Different brain areas are involved in memory expression at different times from training. *Neurobiology of Learning and Memory*, **66**(2), 97–101.

Quirk, G., Muller, R., and Kubie, J. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *Journal of Neuroscience*, **10**, 2008–2017.

Ramírez-Amaya, V., Escobar, M. L., Chao, V., and Bermúdez-Rattoni, F. (1999). Synaptogenesis of mossy fibers induced by spatial water maze overtraining. *Hippocampus*, **9**(6), 631–636.

Ramón y Cajal, S. (1893). Estructura del asta de Ammon y fascia dentata. *Ann Soc Esp Hist Nat*, **22**.

Ramos, J. M. J. (1998). Retrograde amnesia for spatial information: a dissociation between intra and extramaze cues following hippocampal lesions in rats. *European Journal of Neuroscience*, **10**, 3295–3301.

Ranck, J. B. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. I: behavioral correlates and firing repertoires. *Experimental Neurology*, **41**, 461–535.

Rawlins, J. (1985). Associations across time: The hippocampus as a temporary memory store. *Behavioral and Brain Sciences*, **8**(3), 479–497.

Reber, P. J. and Squire, L. R. (1994). Parallel brain systems for learning with and without awareness. *Learning and memory*, **1**, 217–229.

Recce, M. and Harris, K. (1996). Memory for places : A navigational model in support of Marr's theory of hippocampal function. *Hippocampus*, **6**, 735–748.

Redish, A. (1997). *Beyond the cognitive map: Contributions to to computational neuroscience theory of rodent navigation*. Ph.D. thesis, School of Computer Science, C.M.U. Pittsburgh.

Redish, A. (1999). *Beyond the cognitive map: From place cells to episodic memory*. A Bradford Book, MIT Press, Cambridge Massachusetts.

Reed, J. M. and Squire, L. R. (1998). Retrograde amnesia for facts and events: findings from four new cases. *Journal of Neuroscience*, **18**(10), 3943–3954.

Reed, J. M. and Squire, L. R. (1999). Impaired transverse patterning in human amnesia is a special case of impaired memory for two-choice discrimination tasks. *Behavioral Neuroscience*, **113**(1), 3–9.

Reed, J. M., Hamann, S. B., Stefanacci, L., and Squire, L. R. (1997). When amnesic patients perform well on recognition memory tests. *Behavioral Neuroscience*, **111**(6), 1163–1170.

Reid, I. C. and Morris, R. G. M. (1992). Smells are no surer: rapid improvement in olfactory discrimination learning is not due to the acquisition of a learning set. *Proceedings of the Royal Society of London: B series*, **247**(1319), 137–143.

Rempel-Clower, N. L., Zola, S. M., Squire, L. R., and Amaral, D. G. (1996). Three cases of enduring memory impairment after bilateral damage limited to the hippocampal formation. *Journal of Neuroscience*, **16**(16), 5233–5255.

Ribot, T. (1881). *Diseases of Memory*. Appleton, New York.

Richmond, M. A., Yee, B. K., Pouzet, B., Veenman, L., Rawlins, J. P., Feldon, J., and Bannerman, D. M. (1999). Dissociating context and space within the hippocampus: effects of complete, dorsal, and ventral excitotoxic hippocampal lesions on conditioned freezing and spatial learning. *Behavioral Neuroscience*, **113**(6), 1189–1203.

Rickard, T. C. and Grafman, J. (1998). Losing their configural mind: Amnesic patients fail on tranverse patterning. *Journal of Cognitive Neuroscience*, **10**(4), 509–524.

Ridley, R. and Baker, H. (1997). Evidence for a specific information processing deficit in monkeys with lesions of the septo-hippocampal system. *Cortex*, **33**(1), 167–176.

Robbins, M. J. and Meyer, D. R. (1970). Motivational control of retrograde amnesia. *Journal of Experimental Psychology*, **84**, 220–225.

Roesler, R., Vianna, M., Sant'Anna, M. K., Kuyven, C. R., Kruel, A. V. S., Quevedo, J., and Ferreira, M. B. C. (1998). Intrahippocampal infusion of the NMDA receptor antagonist AP5 impairs retention of an inhibitory avoidance task: Protection from impairment by pretraining or preexposure to the task apparatus. *Neurobiology of Learning and Memory*, **69**, 87–91.

Rolls, E. (1996). A theory of hippocampal function in memory. *Hippocampus*, **6**, 602–620.

Rolls, E. and Treves, A. (1998). *Neural Networks and Brain Function*. Oxford University Press, Oxford, UK.

Rose, S. (1994). *The making of memory*. Bantam Press, Transworld Publishers, London, Great Britain.

Rosenbaum, R., Priselac, S., Kohler, S., S.E., B., Gao, F., Nadel, L., and Moscovitch, M. (2000). Remote spatial memory in an amnesic person with extensive bilateral hippocampal lesions. *Nature Neuroscience*, **3**(10), 1044–1048.

Rosenbaum, R., Winocur, G., and Moscovitch, M. (2001). New views on old memories: re-evaluating the role of the hippocampal complex. *Behavioural Brain Research*, **127**, 183–197.

Rosenbaum, R, S., McKinnon, M. C., Levine, B., and Moscovitch, M. (2004). Visual imagery deficits, impaired strategic retrieval, or memory loss: disentangling the nature of an amnesic person's autobiographical memory deficit. *Neuropsychologia*, **In press**.

Rosenzweig, M. R. (1998). Historical perspectives on the development of the biology of learning and memory. In J. Martinez and R. Kesner, editors, *Neurobiology of Learning and Memory*, pages 1–53. Academic Press.

Rothblat, L. A. and Kromer, L. F. (1991). Object recognition memory in the rat: The role of the hippocampus. *Behavioural Brain Research*, **42**, 25–32.

Rubin, D. C. and Greenberg, D. L. (1998). Visual memory-deficit amnesia: A distinct amnesic presentation and etiology. *Proceedings of the National Academy of Sciences of the USA*, **95**, 5413–5416.

Rudick, C. N. and Woolley, C. S. (2000). Estradiol induces a phasic Fos response in the hippocampal CA1 and CA3 regions of adult female rats. *Hippocampus*, **10**, 274–283.

Rudy, J. W. and Sutherland, R. J. (1995). Configural association theory and the hippocampal formation: An appraisal and reconfiguration. *Hippocampus*, **5**(5), 375–389.

Ryan, L., Nadel, L., Keil, K., Putnam, K., Schnyer, D., Trouard, T., and Moscovitch, M. (2001). Hippocampal complex and retrieval of recent and very remote autobiographical memories: evidence from functional magnetic resonance imaging in neurologically intact people. *Hippocampus*, **11**, 707–714.

Sachetti, B., Lorenzini, C. A., Baldi, E., Tassoni, G., and Bucherelli, C. (1999). Auditory thalamus, dorsal hippocampus, basolateral amygdala, and perirhinal cortex role in the consolidation of conditioned freezing to context and to acoustic conditioned stimulus in the rat. *Journal of Neuroscience*, **19**(21), 9570–9578.

Salas, C., Broglio, C., and Rodriguez, F. (2003). Evolution of forebrain and spatial cognition in vertebrates: conservation across diversity. *Brain, Behaviour and Evolution*, **62**, 72–82.

Samsonovich, A. and McNaughton, B. (1997). Path integration and cognitive mapping in a continuous attrator neural network model. *Journal of Neuroscience*, **17**(15), 5900–5920.

Sara, S. J. (2000). Retrieval and consolidation: Toward a neurobiology of remembering. *Learning and Memory*, **7**, 73–84.

Saucier, D. and Cain, D. (1995). Spatial learning without NMDA receptor-dependent long-term potentiation. *Nature*, **378**, 186–189.

Save, E., Buhot, M. C., Foreman, N., and Thinus-Blanc, C. (1992a). Exploratory activity and response to a spatial change in rats with hippocampal or posterior parietal cortical lesions. *Behavioural Brain Research*, **47**(2), 113–127.

Save, E., Poucet, B., Foreman, N., and Buhot, M. C. (1992b). Object exploration and reactions to spatial and nonspatial changes in hooded rats following damage to parietal cortex or hippocampal formation. *Behavioural Neuroscience*, **106**(3), 447–456.

Schacter, D. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **13**, 501–518.

Schacter, D. L. and Dodson, C. S. (2001). Misattribution, false recognition and the sins of memory. *Philosophical Transactions of the Royal Society London: B*, **356**, 1385–1393.

Schmajuk, N. A. and DiCarlo, J. J. (1992). Stimulus configuration, spatial learning, and the hippocampus. *Psychological Review*, **99**, 268–305.

Schnider, A., Regard, M., and Landis, T. (1994). Anterograde and retrograde amnesia following bitemporal infarction. *Behav Neurol*, **7**, 87–92.

Schnider, A., Bassetti, C., Gutbrod, K., and Ozdoba, C. (1995). Very severe amnesia with acute onset after isolated hippocampal damage due to systemic lupus erythematosis. *Journal of Neurology, Neurosurgery and Psychiatry*, **59**, 644–645.

Schoenemann, P. T. (2001). Brain scaling, behavioural ability and human evolution. *Behavioural and Brain Sciences*, **24**, 293295.

Scoville, W. and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery and Psychiatry*, **20**, 11–21.

Seidenbecher, T., Reymann, K. G., and Balschun, D. (1997). A post-tetanic time window for the reinforcement of long-term potentiation by appetitive and aversive stimuli. *Proceedings of the National Academy of Sciences of the USA*, **94**(4).

Sergent, J., Ohta, S., and McDonald, B. (1992). Functional neuroanatomy of faces and and object processing. *Brain*, **115**, 15–36.

Shadmehr, R., Brandt, J., and Corkin, S. (1998). Time-dependent motor memory processes in amnesic subjects. *Journal of Neurophysiology*, **80**, 1590–1597.

Shapiro, M. and Hetherington, P. (1993). A simple network simulates hippocampal place fields: I. Parametric analyses and physiological predictions. *Behavioral Neuroscience*, **107**(1), 34–50.

Shapiro, M. L. and Olton, D. S. (1994). Hippocampal function and interference. In D. L. Schacter and E. Tulving, editors, *Memory Systems*, pages 87–117. MIT Press, Cambridge, MA.

Sharp, P. (1991). Computer simulation of hippocampal place cells. *Psychobiology*, **19**(2), 103–115.

Shen, B. and McNaughton, B. (1996). Modelling the spontaneous reactivation of experience-specific hippocampal cell assembles during sleep. *Hippocampus*, **6**, 685–692.

Shen, J., Barnes, C. A., McNaughton, B. L., Skaggs, W. E., and Weaver, K. L. (1997). The effect of aging on experience-dependent plasticity of hippocampal place cells. *Journal of Neuroscience*, **17**(17), 6769–6782.

Shen, J., Kudrimoti, H. S., McNaughton, B. L., and Barnes, C. A. (1998). Reactivation of neuronal ensembles in hippocampal dentate gyrus during sleep after spatial experience. *Journal of Sleep Research*, **7**, 6–16.

Shepherd, G. (1994). *Neurobiology*. Oxford University Press, New York, 3rd edition.

Sherry, D. F., Jacobs, L. F., and Gaulin, S. J. (1992). Spatial memory and adaptive specialization of the hippocampus. *Trends in Neurosciences*.

Shiflett, M., Smulders, T., Benedict, L., and DeVoogd, T. (2003). Reversible inactivation of the hippocampal formation in food-storing black-capped chickadees (poecile atricapillus). *Hippocampus*, **13**(4), 437–444.

Shimamura, A. P. and Squire, L. R. (1987). A neuropsychological study of fact memory and source amnesia. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **13**, 464–473.

Shroeder, J. A., Wingard, J., and Packard, M. G. (2002). Post-training reversible inactivation of the dorsal hippocampus reveals interference between multiple memory systems. *Hippocampus*, **12**, 280–4.

Silva, A. J., Giese, K. P., Fedorov, N. B., Frankland, P. W., and Kogan, J. H. (1998). Molecular, cellular and neuroanatomical substrates of place learning. *Neurobiology of Learning and Memory*, **70**, 44–61.

Skaggs, W. and McNaughton, B. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, **271**, 1870–1873.

Skaggs, W. E., McNaughton, B. L., Wilson, M. A., and Barnes, C. A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, **6**(2), 149–173.

Smith, C. (1996). Sleep states, memory processes and synaptic plasticity. *Behavioural Brain Research*, **78**(1), 49–56.

Solomon, P. R. and Moore, J. W. (1975). *J. Comp. Physiol. Psychol.*, **89**, 1192–1203.

Solomon, P. R., Vander Schaaf, E. R., and Thompson, R. F. (1986). Hippocampus and trace conditioning of the rabbit's classically conditioned nictitating membrane response. *Behavioural Neuroscience*, **100**, 729–744.

Spiegler, B. J. and Mishkin, M. (1996). *Behavioral Brain Research*, **3**, 303.

Spiers, H. J., Maguire, E. A., and Burgess, N. (2001). Hippocampal amnesia. *Neurocase*, **7**, 357–382.

Squire, L. (1983). The hippocampus and the neuropsychology of memory. In W. Seifert, editor, *Neurobiology of the hippocampus*, pages 491–512. Academic Press, New York.

Squire, L. (1992). Declarative and non-declarative memory: Multiple brain systems supporting learning and memory. *Journal of Cognitive Neuroscience*, **4**(3), 232–243.

Squire, L. and Zola, S. (1998). Episodic memory, semantic memory, and amnesia. *Hippocampus*, **8**(3), 205–211.

Squire, L. and Zola-Morgan, S. (1988). Memory: brain systems and behavior. *Trends in Neurosciences*, **11**(4), 170.

Squire, L. R. and Alvarez, P. (1995). Retrograde amnesia and memory consolidation: a neurobiological perspective. *Current Opinion in Neurobiology*, **5**, 169–177.

Squire, L. R. and Knowlton, B. J. (1995). Learning about categories in the absence of memory. *Proceedings of the National Academy of Sciences of the USA*, **92**, 12470–12474.

Squire, L. R. and Zola, S. M. (1996). Ischemic brain damage and memory impairment: A commentary. *Hippocampus*, **6**(5), 546–552.

Squire, L. R., Clark, R. E., and Knowlton, B. J. (2001). Retrograde amnesia. *Hippocampus*, **11**(1), 50–55.

Stark, C. E. L. and Squire, L. R. (2003). Hippocampal damage equally impairs memory single items and memory for conjunctions. *Hippocampus*, **13**, 281–292.

Steele, R. J. and Morris, R. G. M. (1999). Delay-dependent impairment of a matching-to-place task with chronic and intrahippocampal infusion of the NMDA-antagonist D-AP5. *Hippocampus*, **9**(2), 118–136.

Sutherland, G. R. and McNaughton, B. L. (2000). Memory trace reactivation in hippocampal and neocortical neuronal ensembles. *Current Opinion in Neurobiology*, **10**, 180–186.

Sutherland, R. and Rudy, J. (1989). Configural association theory: The role of the hippocampal formation in learning, memory and amnesia. *Psychobiology*, **17**(2), 129–144.

Sutherland, R. J., Chew, G. L., Baker, J. C., and Linggard, R. C. (1987). Some limitations on the use of distal cues in place navigation by rats. *Psychobiology*, **15**, 48–57.

Sutherland, R. J., Kolb, B. E., and Gibb, R. (1993). Society for Neuroscience Abstracts, 19:362.

Sutherland, R. J., Weisend, M. P., Mumby, D., Astur, R. S., Hanlon, F. M., Koerner, A., Thomas, M. J., Wu, Y., Moses, S. N., Cole, C., Hamilton, D. A., and Hoesing, J. M. (2001). Retrograde amnesia after hippocampal damage: Recent vs. remote memories in two tasks. *Hippocampus*, **11**(1), 27–42.

Suzuki, A., Josselyn, S. A., Frankland, P. W., Masushige, S., Silva, A. J., and Kida, S. (2004). Memory reconsolidation and extinction have distinct temporal and biochemical signatures. *The Journal of Neuroscience*, **24**(20), 4787–4795.

Suzuki, W. and Amaral, D. (1994). The perirhinal and parahippocampal cortices of the monkey: cortical afferents. *Journal of Comparative Neurology*.

Suzuki, W. A. (1996). Neuroanatomy of the monkey entorhinal, perirhinal and parahippocampal cortices: Organization of cortical inputs and interconnections with amygdala and striatum. *Seminars in the Neurosciences*, **8**, 3–12.

Suzuki, W. A., Miller, E. K., and Desimone, R. (1997). Object and place memory in the Macaque entorhinal cortex. *Journal of Neurophysiology*, **78**, 1062–1081.

Swanson, L., Wyss, J., and Cowan, W. (1978). An autoradiographic study of the organisation of intrahippocampal association pathways in the rat. *Journal of Computational Neurology*, **181**, 681–715.

Sziklas, V. and Petrides, M. (2002). Effects of lesions to the hippocampus or the fornix on allocentric conditional associative learning in rats. *Hippocampus*, **12**(4), 543–550.

Takehara, K., Kawahara, S., Takatsuki, K., and Kirino, Y. (2002). Time-limited role of the hippocampus in the memory for trace eyeblink conditioning in mice. *Brain Research*, **951**, 183–190.

Takehara, K., Kawahara, S., and Kirino, Y. (2003). Time-dependent reorganisation of the brain components underlying memory retention in trace eyeblink conditioning. *The Journal of Neuroscience*, **23**(30), 9897–9905.

Taube, J. S. (1999). Some thoughts on place cells and the hippocampus. *Hippocampus*, **9**, 452–457.

Teng, E. and Squire, L. R. (1999). Memory for places learned long ago is intact in hippocampal damage. *Nature*, **400**, 675–677.

Teyler, T. and Discenna, P. (1986). The hippocampal memory indexing theory. *Behavioural Neuroscience*, **100**(2), 147–154.

Thompson, R. (1988). The neural basis of basic associative learning of discrete behavioural responses. *Trends in Neurosciences*, **11**(4), 152.

Thornton, J. A., Rothblat, L. A., and Murray, E. A. (1997). Rhinal cortex removal produces amnesia for preoperatively learned discrimination problems but fails to disrupt postoperative acquisition and retention in rhesus monkeys. *Journal of Neuroscience*, **17**(21), 8536–8549.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, **55**, 189–208.

Treves, A. and Rolls, E. (1994). Computational analysis of the role of the hippocampus in memory. *Hippocampus*, **4**(3), 374–391.

Tsukiura, T., Fujii, T., Okuda, J., Ohtake, H., Kawashima, R., Itoh, M., Fukuda, H., and Yamadori, A. (2002). Time-dependent contribution of the hippocampal complex when remembering the past: a pet study. *Neuroreport*, **13**(17), 2319–2323.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving and W. Donaldson, editors, *Organization of memory*. Academic Press, New York.

Tulving, E. (1983). *Elements of episodic memory*. Clarendon Press, Oxford.

Tulving, E. (1995). *Cognitive Neurosciences*, chapter Organization of memory: Quo vadis?, pages 839–847. MIT Press, Cambridge, MA.

Tulving, E. and Markowitsch, H. (1997). Memory beyond the hippocampus. *Current Opinion in Neurobiology*, **7**, 209–216.

Tulving, E. and Markowitsch, H. (1998). Episodic and declarative memory: Role of the hippocampus. *Hippocampus*, **8**(3), 198–204.

Tulving, E., Hayman, C. A. G., and MacDonald, C. (1991). Long-lasting perceptual priming and semantic learning in amnesia: a case experiment. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **17**, 595–617.

Tulving, E., Markowitsch, H., Kapur, S., Habib, R., and Houle, S. (1994). Novelty encoding networks in the human brain: positron emission tomography data. *Neuroreport*, **5**, 2525–2528.

Vanderwolf, C. H. (1969). Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalogr Clin Neurophysiol*, **26**, 407–418.

Vanderwolf, C. H., Kramis, R., Gillespie, L., and Bland, B. (1975). Hippocampal rhythmical slow activity and neocortical low voltage fast activity: relations to behaviour. In R. Isaacson and K. Pribram, editors, *The hippocampus, II: Neurophysiology and Behaviour*, pages 101–128. Plenum, NY.

Vann, S. D., Brown, M. W., Erichsen, J. T., and Aggleton, J. P. (2000). Fos imaging reveals differential patterns of hippocampal and parahippocampal subfield activation in rats in response to different spatial memory tests. *Journal of Neuroscience*, **10**(7), 2711–2718.

Vargha-Khadem, F., Gadian, D., Watkins, K., Connelly, A., Paesschen, W. V., and Mishkin, M. (1997). Differential effects of early hippocampal pathology on episodic and semantic memory. *Science*, **277**, 376–380.

Verfaellie, M., Reiss, L., and Roth, H. (1995). Knowledge of new english vocabulary in amnesia: as examination of premorbidly acquired semantic memory. *Journal of the International Neuropsychological Society*, **1**, 443–453.

Vertes, R. P. and Eastman, K. E. (2000). The case against memory consolidation in REM sleep. *Behavioral and Brain Sciences*, **23**(6).

Viskontas, I. V., McAndrews, M. P., and Moscovitch, M. (2000). Remote episodic memory deficits in patients with unilateral temporal lobe epilepsy and excisions. *Journal of Neuroscience*, **20**(15), 5853–5857.

Wallenstein, G. V. and Hasselmo, M. E. (1997). GABAergic modulation of hippocampal population activity: sequence learning, place field development, and the phase precession effect. *Journal of Neurophysiology*, **78**, 393–408.

Wallenstein, G. V., Eichenbaum, H., and Hasselmo, M. E. (1998). The hippocampus as an associator of discontiguous events. *Trends in Neurosciences*, **21**(8), 317–323.

Walz, R., Roesler, R., Quevedo, J., Sant'Anna, M., Madruga, M., Rodrigues, C., Gottfried, C., Medina, J., and Izquierdo, I. (2000). Time-dependent impairment of inhibitory avoidance retention in rats by posttraining infusion of a mitogen-activated protein kinase kinase inhibitor into cortical and limbic structures. *Neurobiology of Learning and Memory*, **73**(1), 11–20.

Wan, H., Touretsky, D., and Redish, A. (1994). Towards a computational theory of rat navigation. In M. Mozer, P. Smolensky, D. Touretsky, J. Elman, and A. Weigend, editors, *Proceedings of the 1993 Connectionist Models Summer School*, pages 11–19, Hillsdale, New Jersey. Lawrence Earlbaum Associates.

Wan, H., Aggleton, J. P., and Brown, M. W. (1999). Different contributions of the hippocampus and perirhinal cortex to recognition memory. *Journal of Neuroscience*, **19**(3), 1142–1148.

Warrington, E. and Weiskrantz, L. (1968). New method of testing long-term retention with special reference to amnesia patients. *Nature*, **217**, 972–974.

Warrington, E. and Weiskrantz, L. (1970). Amnesic syndrome: Consolidation or retrieval? *Nature*, **228**, 628–630.

Waters, N. S., Klintsova, A. Y., and Foster, T. C. (1997). Insensitivity of the hippocampus to environmental stimulation during postnatal development. *The Journal of Neuroscience*, **17**(20), 7967–7973.

Weiskrantz, L. and Warrington, E. K. (1979). *Neuropsychologia*, **17**, 187–194.

Westmacott, R., Leach, L., Freedman, M., and Moscovitch, M. (2001). Different patterns of autobiographical memory loss in semantic and medial temporal lobe amnesia: a challenge to consolidation theory. *Neurocase*, **7**, 27–55.

Westmacott, R., Black, S., Freedman, M., and Moscovitch, M. (2003). The contribution of autobiographical significance to semantic memory: evidence from Alzheimer's disease, semantic dementia, and amnesia. *Neuropsychologia*, **42**(1), 25–48.

Whishaw, I. Q. and Jarrard, L. E. (1995). Similarities vs differences in place learning and circadian activity in rats after fimbria-fornix section or ibotenate removal of hippocampal cells. *Hippocampus*, **5**(595-604).

Whishaw, I. Q. and Tomie, J.-A. (1991). Acquisition and retention by hippocampal rats of simple, conditional, and configural tasks using tactile and olfactory cues: implications for hippocampal function. *Behavioral Neuroscience*, **105**(6), 787–797.

Whishaw, I. Q., Cassel, J. C., and Jarrard, L. E. (1995). Rats with fimbria-fornix lesions display a place response in a swimming pool: A dissociation between getting there and knowing where. *Journal of Neuroscience*, **15**, 5779–5788.

Whishaw, I. Q., McKenna, J. E., and Maaswinkel, H. (1997). Hippocampal lesions and path integration. *Current Opinion in Neurobiology*, **7**, 228–234.

White, N. M. and Wallet, P. A. (2000). Dorsal hippocampal function in unreinforced spatial learning. *Hippocampus*, **10**, 226–235.

White, N. M., Holahan, M. R., and Goffaux, P. (2003). Involuntary, unreinforced (pure) spatial learning is impaired by fimbria fornix but not by dorsal hippocampal lesions. *Hippocampus*, **13**, 324–333.

Whitlow, Althoff, and Cohen (1995). *Society for Neuroscience Abstracts*, **21**, 303.5.

Wickelgren, W. A. (1979). Chunking and consolidation: a theoretical synthesis of semantic networks, configuring, S-R versus cogntive learning, normal forgetting, the amnesic syndrome, and the hippocampal arousal system. *Psychological Review*, **86**, 44–60.

Wiener, S., Paul, C., and Eichenbaum, H. (1989). Spatial and behavioural correlates of hippocampal neuronal activity. *Journal of Neuroscience*, **9**, 2737–2763.

Wiggs, C. L. and Martin, A. (1998). Properties and mechanisms of perceptual priming. *Current Opinion in Neurobiology*, **8**, 227–233.

Wiggs, C. L., Martin, A., and Sunderland, T. (1997). Monitoring frequency of occurrence without awareness: evidence from patients with Alzheimer's disease. *Journal of Clinical Experimental Neuropsychology*, **13**, 235–244.

Wilson, M. and McNaughton, B. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, **265**, 676–679.

Winocur, G. (1990). Anterograde and retrograde amnesia in rats with dorsal hippocampal or dorsomedial thalamic lesions. *Behavioural Brain Research*, **38**(2), 145–154.

Winocur, G., McDonald, R. M., and Moscovitch, M. (2001). Anterograde and retrograde amnesia in rats with large hippocampal lesions. *Hippocampus*, **11**(1), 18–26.

Wise, S. and Murray, E. (1999). Role of the hippocampal system in conditional motor learning: Mapping antecedents to action. *Hippocampus*, **9**, 101–117.

Wise, S. and Murray, E. (2000). Arbitrary associations between antecedents and actions. *Trends in Neurosciences*, **23**(6), 271–276.

Witter, M. P., Naber, P. A., van Haeften, T., Machielsen, W. C., Rombouts, S. A., Barkhof, F., Scheltens, P., and Lopes da Silva, F. H. (2000). Cortico-hippocampal communication by way of parallel parahippocampal-subicular pathways. *Hippocampus*, **10**(4), 398–410.

Wood, E. R., Mumby, D. G., Pinel, P. J., and Phillips, A. G. (1993). Impaired object recognition memory in rats following ischaemia-induced damage to the hippocampus. *Behavioral Neuroscience*, **107**, 51–62.

Wood, E. R., Dudchenko, P. A., and Eichenbaum, H. (1999). The global record of memory in hippocampal neuronal activity. *Nature*, **397**, 613–616.

Wood, E. R., Dudchenko, P. A., Robitsek, R. J., and Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron*, **27**, 623–633.

Wu, X., Tyrcha, J., and Levy, W. (1998). A neural network solution to the transverse patterning problem depends on repetition of the input code. *Biological Cybernetics*, **79**, 203–213.

Yamamoto, T. (1993). Neural mechanisms of taste aversion learning. *Neuroscience Research*, **16**(3), 181–185.

Yonelinas, A., Hopfinger, J., Buonocore, M., Kroll, N., and Baynes, K. (2001). Hippocampal, parahippocampal and occipital-temporal contributions to associative and item recognition memory: an fMRI study. *Neuroreport*, **12**(2), 359–363.

Zelkowicz, B., Herbster, A., Nebes, R., Mintun, M., and Becker, J. (1998). An examination of regional blood flow during object naming tasks. *Journal of the International Neuropsychological Society*, **4**, 160–166.

Zilles, K., Wu, J., Crusio, W. E., and Schwegler, H. (2000). Water maze and radial maze learning and the density of binding sites of glutamate, GABA, and serotonin receptors in the hippocampus of inbred mouse strains. *Hippocampus*, **10**(3), 213–225.

Zipser, D. (1985). A computational model of hippocampal place cells. *Behavioral Neuroscience*, **99**(5), 1006–1018.

Zola, S. M., Squire, L. R., Teng, E., Stefanacci, L., Buffalo, E. A., and Clark, R. E. (2000). Impaired recognition memory in monkeys after damage limited to the hippocampal region. *Journal of Neuroscience*, **20**(1), 451–463.

Zola-Morgan, S. and Squire, L. R. (1990). The primate hippocampal formation: Evidence for a time-limited role in memory storage. *Science*, **250**, 288–290.

Zola-Morgan, S., Squire, L. R., and Amaral, D. G. (1986). Human amnesia and the medial temporal region: enduring memory impairment following a bilateral lesion limited to field CA1 of the hippocampus. *Journal of Neuroscience*, **6**, 2950–2967.

Zola-Morgan, S., Squire, L. R., and Amaral, D. G. (1989). Lesions of the hippocampal formation but not lesions of the fornix or the mammillary nuclei produce long-lasting memory impairment in monkeys. *Journal of Neuroscience*, **9**, 898–913.

Zola-Morgan, S., Squire, L. R., Rempel, N. L., Clower, R. P., and Amaral, D. G. (1992). Enduring memory impairment in monkeys after ischaemic damage to the hippocampus. *Journal of Neuroscience*, **12**, 2582–2596.