A General Hippocampal Computational Model

Combining Episodic and Spatial Memory

in a Spiking Model

Paulo de Castro Aguiar

Doctor of Philosophy

Institute for Adaptive and Neural Computation

School of Informatics

University of Edinburgh

2005

Abstract

The hippocampus, in humans and rats, plays crucial roles in spatial tasks and non-spatial tasks involving episodic-type memory. This thesis presents a novel computational model of the hippocampus (CA1, CA3 and dentate gyrus) which creates a framework where spatial memory and episodic memory are explained together. This general model follows the approach where the memory function of the rodent hippocampus is seen as a "memory space" instead of a "spatial memory".

The innovations of this novel model are centred around the fact that it follows detailed hippocampal architecture constraints and uses spiking networks to represent all hippocampal subfields. This hippocampal model does not require stable attractor states to produce a robust memory system capable of pattern separation and pattern completion. In this hippocampal theory, information is represented and processed in the form of activity patterns. That is, instead of assuming firing-rate coding, this model assumes that information is coded in the activation of specific constellations of neurons. This coding mechanism, associated with the use of spiking neurons, raises many problems on how information is transferred, processed and stored in the different hippocampal subfields. This thesis explores which mechanisms are available in the hippocampus to achieve such control, and produces a detailed model which is biologically realistic and capable of explaining how several computational components can work together to produce the emergent functional properties of the hippocampus. In this hippocampal theory, precise explanations are given to why mossy fibres are important for storage but not recall, what is the functional role of the mossy cells (excitatory interneurons) in the dentate gyrus, why firing fields can be asymmetric with the firing peak closer to the end of the field, which features are used to produce "place fields", among others. An important property of this hippocampal model is that the memory system provided by the CA3 is a palimpsest memory: after saturation, the number of patterns that can be recalled is independent of the number of patterns engraved in the recurrent network.

In parallel with the development of the hippocampal computational model, a simulation environment was created. This simulation environment was tailored for the needs and assumptions of the hippocampal model and represents an important component of this thesis.

Acknowledgements

This thesis greatly benefited from the interaction with many individuals. Among them, my special thanks go to my dear friend Gaurav, with whom I had countless hours of discussions and brainstorms in many interesting topics of neuroscience. His opinions, companionship, patience, help during chapter revisions, were all precious to me.

I would also like to thank Kit, a true example of pure enthusiasm in science. I have learnt a lot from him during my PhD and all our "Hippo-*lunches-in-the*-campus" have been a great source of enlightenment. I also benefited from discussions with Mark van Rossum, David Barber, David Sterratt and Amos Storkey. They were always available when I came to them with questions. I would also like to thank David Willshaw for valuable comments on the first drafts of this thesis, which helped me to improve its organisation and contents.

On the non-academic side, I want to show my gratitude to a group of close friends whose presence made my stay in Edinburgh a great experience. They are Aghlab, Melissa, Georgios and Marco. We were a great team.

From home, I could always count with the support from Sandra, Clotilde and Eduardo, my brothers Ricardo and Rafael, and my parents. Their words of encouragement have been priceless.

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.

(*Paulo de Castro Aguiar*)

To Suzete and Daniel

Table of Contents

List of Figures

List of Tables

List of Definitions, Acronyms and Notation

This list presents the symbols frequently used throughout this thesis. The entry "1" in the "Dimension" column is used for dimensionless variables. Some variables contain their typical value instead of a reference in the "Definition" column.

Symbol	Description	Dimension	Definition
V	membrane voltage	mV	eq. 3.2
V_{rest}	membrane resting potential	mV	$-70mV$
V_{thresh}	membrane firing threshold voltage	mV	$-50mV$
V_{reset}	membrane reset voltage	mV	$-60mV$
E_{syn}	synapse's reversal potential	mV	$0mV$(exc)
			$-90mV$ (inh)
R_m	membrane resistance (or input resistance)	$M\Omega$	$20M\Omega$
C_m	membrane capacitance	nF	$1nF$
I_{in}	total input current	nA	eq. 3.6
I_{syn}	synaptic input current	nA	eq. 3.5
S	post-synaptic conductance	μS	eq. 3.4
S_{peak}	post-synaptic peak conductance	μS	eq. 3.4
S_{max}	maximum potentiated post-synaptic peak conductance	μS	eq. 3.8
S_{thresh}	threshold peak conductance	μS	tab. 6.1
τ_m	membrane time constant	ms	$20ms$
T_{ref}	absolute refractory period	ms	$7ms$
τ_{rise}	post-synaptic conductance rise time constant	ms	$3ms$(exc) $2ms$ (inh)
τ_{decay}	post-synaptic conductance decay time constant	ms	$5ms$(exc) $20ms$ (inh)
N_{PSR}	number of post-synaptic responses	1	eq. 6.1

Symbol	Description	Definition
LTP	long-term potentiation	sec. 3.2.4
ΔLTP	long-term potentiation change	eq. 3.8
LTD	long-term depression	sec. 3.2.4
ΔLTD	long-term depression change	eq. 3.8
H	homosynaptic depression change	eq. 3.9
C^c	convergent connectivity	sec. 2.1.3
C^d	divergent connectivity	sec. 2.1.3
$ECII$ or EC2	entorhinal cortex layer II	fig. 2.3
$ECIII$ or EC3	entorhinal cortex layer III	fig. 2.3
DG	dentate gyrus	fig. 2.3
$CA3$	*cornus ammonis* region 3	fig. 2.3
$CA1$	*cornus ammonis* region 1	fig. 2.3
Sub	subiculum	fig. 2.3
$p.p.$	perforant path	fig. 2.3
$m.f.$	mossy fibres	fig. 2.3
$r.c.$	recurrent collaterals	fig. 2.3
$s.c.$	Schaffer collaterals	fig. 2.3
HCM	Hippocampal Computational Model	fig. 2.3

Chapter 1

Introduction

Adaptation, in the form of a mechanism to increase survival expectancy, appears at two distinct levels. It appears at the species level, with the interplay of mutations and natural selection that progressively produces species more adapted to the host environment. At a much smaller time scale, it appears at the organism level, where specialised systems use previous experiences to increase the organism's life expectancy by affecting its future interactions with the environment. For the organism's time scale, the best solution, according to evolution, was to merge these specialised systems with the systems that control the behaviour of the organism, forming a single, complex, dynamic and hence adaptable system - the nervous system.

The ability to extract from experience relevant information that enables a particular organism to increase its or its progenies' chances of survival, by either adapting to the environment or adapting the environment to its needs, is certainly one of the most fascinating capabilities of living organisms. Slowly, our attempts to understand how biological systems are able to efficiently adapt, and to create artificial systems with similar capabilities, are being rewarded (Holland, 1992).

Two main processes need to be considered in this adaptation process: memory and learning.

Throughout this thesis, the word *memory* is used to refer to the ability to store information directly acquired from input channels (for example from sensory input). This information is stored in the form of internal representations without being subject to complex manipulations.

The word *learning* is used to describe the ability to extract information from present

and past input data, using some form of computation, and allowing the creation of internal models of the interaction of the organism with its environment.

Independently, these processes are very limited. Memory alone is inefficient: in complex environments, the sensory configuration space, i.e. all the possible combinations of contents in the input sensory channels, is infinite, demanding an equally infinite memory storage place. The extreme importance of *generalisation* can therefore be seen: instead of storing a representation of a particular configuration of the environment (with a zero probability of re-happening), each representation stored has to refer to a class of environmental configurations. Also, learning processes alone are inadequate since they have to work on cause-effect relations, or other interacting information acquired or spread in time. This information has to be stored somewhere and available for processing and later construction of correlations.

Throughout evolution, living organisms have developed subsystems capable of integrating these two processes. The mammal *hippocampus*, one of the phylogenetically oldest parts of the mammalian brain (Johnston and Amaral, 1998), is one such subsystem.

Both human and rat hippocampi have been extensively studied experimentally and robust links have been established between the functional role of the hippocampal formation and some forms of memory and learning processes (Squire, 1992; Treves and Rolls, 1994; Burgess *et al*, 2002).

Three points are addressed during this introductory chapter: 1) motivation for the importance of the hippocampus as a crucial structure for certain forms of learning and memory, 2) review our present knowledge about the functional role of the hippocampus in terms of theories and models, and 3) state the objectives of this thesis. Each of these points is presented in a separate section underneath.

1.1   The hippocampus

It is appropriate to begin a study on the hippocampus by mentioning the seminal paper by William Scoville and Brenda Milner (1957). This paper introduced to the neuroscience community the, now famous, patient H. M., who underwent bilateral hippocampal removal (2/3 of the hippocampus and some amygdala) for the treatment of epilepsy. As a consequence of this surgery, H.M. suffered a permanent loss of the abil-

ity to encode new declarative information into long-term memory. This was the first demonstration that the hippocampus, and other temporal lobe structures, is crucial for the storage of declarative memories.

Hundreds of papers followed, narrowing down hypotheses and progressively revealing more and more details about the morphological, physiological and functional properties of the hippocampal formation.

Nevertheless, instead of obtaining robust answers, more doors were opened to the vast complexity of this area of the mammalian brain. The discovery of *place cells* by John O'Keefe and John Dostrovsky (1971), and the discovery of *long-lasting potentiation* by Tim Bliss and Terje Lømo (1973), are undoubtedly among the best examples of that.

Most of these advances in the understanding of the mammal hippocampus resulted from experiments with rat hippocampi. Only after the introduction of non-invasive techniques such as functional magnetic resonance imaging, was there an increase in the number of studies targeted on human hippocampus. Previously, these studies had been totally dependent on pathological cases (Stern and Hasselmo, 1999; Maguire, 2001).

It is still the rat hippocampus that gathers a larger volume of data and therefore an equally larger number of theories and models. It is possible though that many results may be extrapolated from one organism to the other. So far the results are consistent with the idea that, although there are differences in the way that rat and human hippocampus handle specific tasks, some principles that direct the functional properties of these two systems seem to be the same (Squire, 1992).

1.1.1   Roles of the hippocampus

The key functional roles of the primate and rodent hippocampus are summarised here (see also Eichenbaum *et al*, 1999; Eichenbaum, 2000).

Human hippocampus

Instead of one memory system, humans possess several types of interconnected but independent memory systems (Squire and Zola, 1996; Giovanello and Verfaellie, 2001). Damage in one of those systems does not imply total impairment in information storage and retrieval in the others.

Accordingly to Squire and Zola (1996), the long-term memory systems can be divided into two major groups (see figure 1.1): declarative (explicit) memory and non-declarative (implicit) memory. Examples of non-declarative memory types are *habituation*, *sensitisation*, *classical conditioning* and *procedural memory* (see Squire and Kandel, 2002). Declarative memory is responsible for dealing with autobiographical information and is commonly believed to be further broken down into two components (Tulving and Schater, 1990): *episodic memory* (memory for past and personally experienced events) and *semantic memory* (knowledge for the meaning of words and how to apply them).
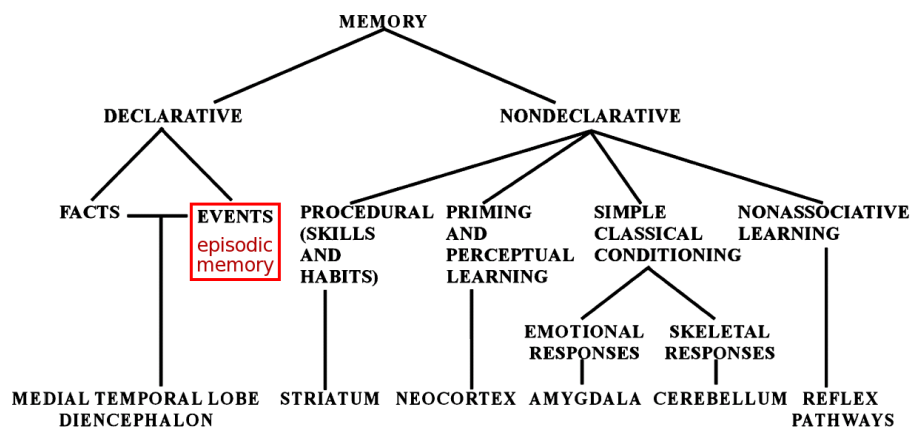
Figure 1.1: Memory systems in the human brain. Adapted from Squire (2004).

This thesis focuses on episodic memory. Personally experienced incidents, events or any information on the form of *what*, *where*, *when*, are all defined as episodic memories.

Strong experimental evidence shows that the hippocampus is fundamental for the encoding of episodic information. Damage to the hippocampus above a certain level leads to anterograde amnesia, associated with the loss of ability to encode new information into long-term memory (Zola-Morgan *et al*, 1986; Tulving and Markowitsch, 1998). The hippocampus is not necessary for other forms of memory such as working memory or procedural memory. Damage to the hippocampus will only affect the formation of new episodic memories. The hippocampus is also not required for the acquisition of semantic memory (Vargha-Khadem *et al*, 1997).

It has also been found that there is a strong link between the hippocampus and nav-

igational tasks, in which spatial information, locations, paths and events are merged together. With the aid of functional magnetic resonance imaging it has also been possible to detect functional asymmetries within the hippocampus: while the right hippocampus appears particularly involved in memory for environmental locations, the left hippocampus is more involved in context-dependent episodic or autobiographical memory (Burgess *et al*, 2002).

Even a structural dependence has been found in hippocampi of humans that systematically perform navigational tasks, specifically taxi drivers in a complex city (Maguire *et al*, 1998).

Rat hippocampus

While in humans the data is more at the cognitive level, in rats detailed information is available at the cellular level.

Generally the rat hippocampus seems to be functionally equivalent to the human hippocampus, serving the same purposes, with crucial roles in spatial and non-spatial tasks that involve episodic type memory (Squire, 1992).

Nevertheless, many support the idea that, in contrast with the humans where the hippocampus seems to play a critical role in memory formation for a broad range of information (Eichenbaum *et al*, 1999), the rat hippocampus seems to have evolved to be more dedicated to space-related tasks. The discovery of place cells by O'Keefe and Dostrovsky (1971), which form the base of the *cognitive map* theory (O'Keefe and Nadel, 1978), started a new wave of enthusiastic research into the functional role of the hippocampus in spatial learning and navigation.

The firing rate of certain types of cells in the rat hippocampus is highly correlated with the location of the rat in its environment. Each of these so called *place cells* fire maximally in defined regions of the environment called *place fields*. Specific lesions on the hippocampus lead to substantial impairment is solving simple navigational tasks (Morris *et al*, 1982). Experiments where rats have to retain some sort of associations of non-spatial nature, also revealed an involvement of specific areas of the hippocampus (Wood *et al*, 1999; but see also Dudchenko *et al*, 2000).

Due to the higher amount of data and theoretical work on the rat than humans, this present study concentrates on the rat hippocampus.

1.2   Current theories and models

The hippocampus is an important structure in the mammal brain because of its pivotal role in encoding, consolidation and retrieval of associations responsible for episodic memory, as well as its well established role is spatial memory and navigation. Understanding how the hippocampus works offers not only a great achievement to the field of neuroscience but also to machine learning, where insights can be given on how to build efficient artificial learning and memory systems, both in terms of generalisation and capacity.

The architecture of the hippocampus has to be considered in order to produce detailed models of its dynamics. For the interest of making the discussion on the current hippocampal theories clearer, a brief description of the anatomy of the hippocampus is given. A more detailed description is the focus of chapter 2.

The hippocampus consists of structurally dissimilar regions (subfields) serially connected and communicating directly with the entorhinal cortex which acts, simultaneously, as the major input and output of the hippocampus. Sensory information from the cerebral association cortex (neocortex) arrives at the superficial layers of the entorhinal cortex (EC) by way of the parahippocampal gyrus (PHG) and/or perirhinal cortex (PR). All hippocampal subfields, namely the dentate gyrus (DG), CA3 and CA1, receive direct input from the superficial layers (II-III) of the entorhinal cortex. The CA1 subfield is the last region of tri-synaptic circuitry carrying out the signals from the hippocampus to the subiculum (Sub) and to the deep layers of the entorhinal cortex (V-VI).

Figure 1.2 shows the major connections in the corticohippocampal network as well as the most relevant hippocampal pathways.

The modular architecture of the hippocampus suggests that each subfield may subserve specific computational functions. This has led to theories for the functional task of each subfield (Treves and Rolls, 1994): CA3, with its extensive recurrent connections, possibly acts as an auto-associative memory and DG's possible task is to produce internal representations that are sparse and orthogonal.

Following the debate on the central functional purpose of the (rat) hippocampus, detailed models have nevertheless been divided into two separate groups: 1) models focused on the storage and recall of episodic memories and 2) models focused on
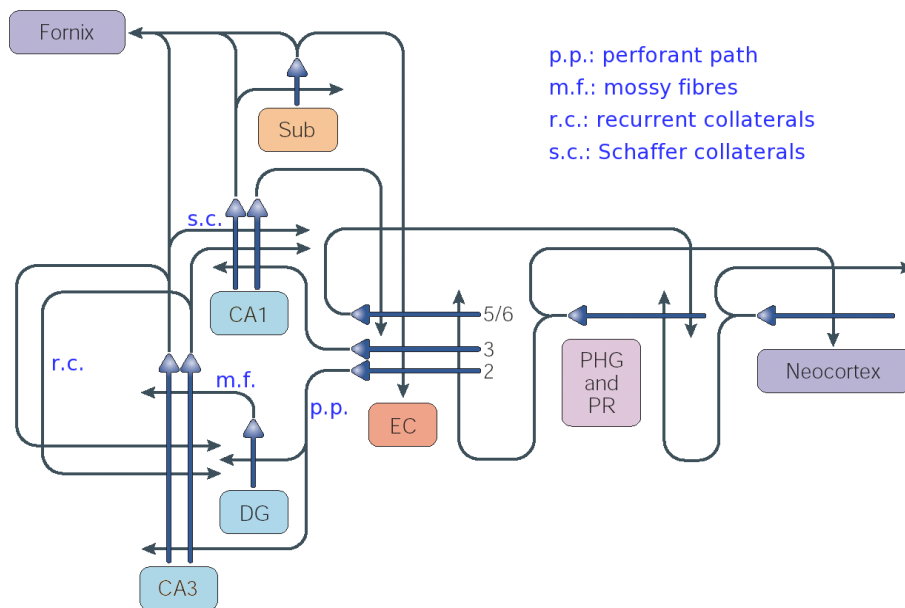
Figure 1.2: Diagram of the hippocampus. The hippocampal subfields are coloured in light blue. Adapted from Nakazawa *et al* (2004).

spatial memory/learning and navigation. Models for episodic, or more specifically associative memory, are typically based on auto-associative networks (Willshaw *et al*, 1969) and discrete attractor networks. Examples include Marr (1971), Treves and Rolls (1994), O'Reilly and McClelland (1994), McClelland *et al* (1995), Hasselmo *et al* (1995), Shastri (2002). On the other side, models for spatial learning often rely on continuous attractor networks to recreate experimental results such as place field formation (Samsonovich and McNaughton 1997; Brunel and Trullier, 1998; Redish and Touretzky, 1998; Arleo and Gerstner, 2000; Káli and Dayan, 2000).

In attractor networks (discrete or continuous), learning and memory are related to the topological properties of the network's energy landscape: a stored memory corresponds to a local minimum on the energy landscape (the reciprocal is not true since a stable point may correspond to spurious memories); storing a new memory corresponds to digging in the energy landscape a local minimum at the state representative of the memory to be stored (Hopfield, 1982). The correct recall of a stored memory requires that the initial state of the network is in the basin of attraction of the correct minimum.

Only few attempts have been made in creating a holistic hippocampal model that

is in accordance with experimental results from both spatial and episodic memory paradigms. Examples of hippocampal unified models of spatial *and* associative (episodic) type memory are: Misják, Lengyel and Érdi (2001) used rate units to developed a model for CA3 that, using the Bienenstock-Cooper-Munro (Bienenstock *et al*, 1982) learning rule, can exhibit place cell properties and can be used for storing episodic memories; Rolls, Stringer and Trappenberg (2002) developed a model based on a combined attractor network, also using rate units, that can store both discrete and continuous representations; Hasselmo's group, Eichenbaum's group and Robert Cannon (CATACOMB), are also working on the project "A Spiking Model of Hippocampus for Guiding Behaviour" (Cannon *et al*, 2002).

In terms of a complete, global model for the hippocampus, most models cited above (unified or not) have, in my opinion, the following limitations:

- they do not recreate, using both morphological and physiological constraints, all hippocampal fields (in fact, most models concentrate on a single field);

- they are developed for a specific task, either related to spatial memory or associative/episodic type memory;

- they use rate-based or binary units to describe individual neurons;

- they require stable states to specify a memory; learning is based on the establishment of new point attractors in the state space landscape;

- they do not allow specific manipulations in order to recreate experimental results (e.g. removal of a specific subfield, blockage of plasticity on a specific path) or predict novel behaviour;

- they make strong restrictions on the inputs the model receives (for example, the input units have to be already tuned to specific features).

In addition, in some of the models cited above, many proposed roles for specific components of the hippocampal circuitry are discussed only in qualitative terms. They are never tested in a full model under neuro-physiological constraints. Examples are the proposed role for the mossy fibres as a mean to create sparse representations in CA3, or the role of theta rhythms.

Overcoming these limitations means achieving a hippocampal model that is more biologically plausible and able to give a deeper understanding on the control mechanisms and computations of the hippocampus. In particular, developing a model for learning and memory in the hippocampus that does not depend upon stable states of attractor networks is of considerable importance. It seems implausible to me that computations in the hippocampus, or anywhere else in the nervous system, require stable configurations. Most signals in the brain are transient and immersed in noise and stochasticity. A proper model for learning and memory in the hippocampus should account for these properties. Following this line of thought, a criticism has been made as to the use of firing-rate models to describe the information transfer in the hippocampus (Rieke *et al*, 1996). Again, in my opinion, firing-rates hide a deeper reality which is firing probability modulation. As it will be shown in this thesis, assuming the approach of firing probabilities (in the context of spiking neurons) instead of the deterministic firing-rate, allows a clear view of the computations occurring in the hippocampus.

1.3 Thesis objectives and goals

This thesis' main goal was to develop a spiking neuron model of the hippocampus (CA1, CA3 and dentate gyrus) that exhibits all the relevant functional properties of this system in a behaving rat. In other words, to produce a holistic computational model for the hippocampus that can recreate hippocampal episodic memory and spatial memory by reproducing key experimental results. Another goal of this thesis was to show that the hippocampal memory system can be recreated and explained without the use of firing-rate coding and memories based in stable attractor states. In addition, and associated with the use of spiking units, this thesis had the objective of investigating the processes of information storage and activity control in spiking networks (following the architecture of the hippocampus).

In parallel with the development of the hippocampal model, a simulation environment was created. This simulation environment was tailored to the needs and assumptions of the hippocampal model.

A model that uses spiking units introduces difficulties that are less problematic in rate-based units. The two most important ones are:

- stability control of the levels of neuronal activity in each subfield and;

- storage and recall of spiking patterns.

When creating a holistic model for the hippocampus, these two initially independent issues become inter-related in the networks introduced here. Specific sections of this thesis have been dedicated to address the stability problem and the storage and recall dichotomy (see chapters 6 and 8).

It should be emphasised that expanding the time dimension using a spiking model allows precise analysis to be done on time related questions. For example, the time delay between the activation of a specific pattern at the entorhinal cortex layer II and the activation of the population of CA1 neurons that signal the recognition and recall can be measured and compared with experimental data.

The hippocampal formation is a complex system with many complex mechanisms. Naturally, only the mechanisms with direct impact on the functional behaviour studied here were modelled. This simplification process is, by itself, a first statement of the assumptions of the model. Thus, in order to create a clearer view of this hippocampal model, a short description of the major assumptions in which it is grounded are presented below:

1. **Neuron Representation**. Each neuron is represented as an *Integrate-and-Fire* unit (Lapicque, 1907). The vast diversity of ionic channels and respective effects are disregarded. Action potentials are stereotyped waves with the single purpose of signalling to synapses.

2. **Synapse Representation**. The synapse is the basic computation element. Each synapse unit is in fact the representation of a small population of correlated synapses sharing the same presynaptic and postsynaptic neuron.

3. **Dendritic non-linearities**. The neuron unit is point-wise and although an approximation for dendritic filtering (spatial and temporal) is considered, the richness of dendritic signal interactions is not accounted for. In other words, all the possible complex non-linear computations that *may* take place in the dendritic tree are disregarded.

4. **Interneurons**. The function of interneurons is activity level control. They do not directly take part in computations and their synapses are not plastic.

5. **Storage**. Information is stored in the network in synaptic efficacies (Hebb, 1949; Morris, 1989).

6. **Activity patterns**. Throughout this thesis the concept of *activity patterns* will be extensively used. It is defined in the network's spatial domain but, instead of a time snapshot of the neural assembly activation configuration, an activity pattern has an associated temporal length on the order of the neuron's membrane time constant τ_m, i.e., the neuron's integration window. The activity pattern at time t is then represented by all the neurons in the population that reached threshold in the interval $[t - \tau_m/2, t + \tau_m/2]$. This concept is related, but not identical, to the notion of simple representations described in Marr (1971).

7. **Memory states**. All information is transferred and manipulated in the form of activity patterns. Stored states, or memories, are also represented by specific activation patterns on neural assemblies, that *do not* need to be stable. This means that temporal coding *at a population level* is used (as opposed to firing-rate coding).

8. **Activity gating**. The gating mechanisms introduced in this thesis are one of the pillars of the computational model presented here. In neural pathways subject to activity gating, only meaningful patterns of activity at the pre-synaptic population are able to excite beyond threshold the post-synaptic population. The gating mechanism is produced through learning.

9. **Plasticity**. Learning in the network is based in associative plasticity (Hebb, 1949), spike-time dependent plasticity (Bi and Poo, 1998) and short-term dynamics (Fuhrmann *et al*, 2002).

10. **Inputs and Outputs**. The properties of the input to the hippocampal formation and which information is available at the output is certainly a critical point. In terms of output, it is considered that CA1 provides the last stage of the computations required for episodic and spatial learning. Most spatial tests on the model are performed with CA1 units. In terms of inputs, strong assumptions are made and described in detail next.

The entorhinal cortex provides the major input of the hippocampal model. It is assumed that the activity patterns formed at entorhinal cortex, layers II (ECII) and

III (ECIII), are functions of the activity of a high number of sensorial channels (e.g. visual, proprioceptive, auditive, olfactive), as well as a function of previously stored information.

A mapping function, $\mathcal{M}$, projects the high dimensional sensory space to a lower (but still high) dimensional space: the entorhinal cortex activity configuration space. Although unknown, this mapping function has to satisfy the following key requirements: be stable over time, be well behaved and be segmented (segregated).

Being stable over time means that across the time-scale of the simulations the mapping function does not change; being well-behaved means that small changes on the input leads to small changes in the output. Finally, a segmented (segregated) projection means that each component in the output space depends on a small number of components in the input space. This is required for extracting features more efficiently from the sensorial signals.

The spatial model by Káli and Dayan (2000) is a particular case of this approach in which they assume that the level of activity of each principal neuron of ECII is broadly tuned to a specific location of space. In fact, experimentally, some results show that some ECII principal cell also act as place cells but with a lower level of selectivity (Frank *et al*, 2000).

One consequence of this mapping hypothesis in the framework of spatial learning is that, once all the model functional parameters have been defined, one can concentrate on manipulations of the environment by affecting the activity patterns at EC. For example, the consequences of geometric manipulations (Burgess, 2002) can be tested and compared with experimental results and with other models.

A less obvious consequence from this mapping assumption, but a core point in this thesis, is that it stops making sense to distinguish *spatial* inputs from *non-spatial* inputs, and equivalently spatial memory from non-spatial memory: the hippocampus computes over streams of sensorial inputs blended together, modulated or not by prior knowledge, and presented at entorhinal cortex.

1.3.1 Thesis outline

This thesis is divided into eleven chapters where the first, current, chapter provides an introduction to the thesis. Chapter 2 briefly discusses morphological and physiologi-

cal properties of the mammalian hippocampus providing relevant information for the construction of the computational model.

The models used for the dynamics of single neurons are presented in chapter 3. Topics regarding the computational properties of single neurons (e.g. dendritic filtering, signal integration, coding) are discussed in detail in this chapter.

The next four chapters discuss important information that is fundamental for the computational model analysed in chapters 7, 8, 9 and 10.

Chapter 4 describes the simulation environment created to build and analyse the holistic computational model for the hippocampus. The simulation environment is itself a component of the computational model: many of the simplifications produced in order to increase the efficiency of the simulations carry with them strong beliefs about the computational role of many components of the system (e.g. synapses, dendrites, action potentials). The simulation environment represents one of the major contributions of this thesis and satisfies all the requirements to work as a useful tool for future research.

Chapter 5 is dedicated to dynamical synapses as a mechanism of decoding information conveyed in the temporal domain. Dynamical synapses play a central role in the computational model for the hippocampus.

The topic of activity level control is discussed in chapter 6. The control of levels of activity on several hippocampal fields and how activity is propagated from field to field, is a also of great importance in the hippocampal model.

Chapter 7 combines the presented information to build a computational model for the functional roles of the hippocampus in the creation of an episodic memory system. The core concepts of the computational model are introduced and explained here.

While chapter 7 presents the theory and assumptions, chapter 8 presents the implementation and the results of the hippocampal computational model in the context of episodic memory. This chapter is the backbone of this thesis where results of several simulations are used to support the theory presented in chapter 7.

Chapter 9 moves in the direction of extending the computational model to incorporate the properties associated with spatial memory. The model discussed in this chapter is a holistic model for the hippocampus that is in close agreement with data arising from both spatial and non-spatial experiments. Results from simulation experiments regarding fundamental concepts from spatial memory are presented in this chapter.

These favourable results pave the way to a more detailed analysis of spatial memory in this novel hippocampal theory and present an important step forward against the view of the rat hippocampus as a specialised structure for spatial memory.

Chapter 10 is used to critically assess the computational model.

A discussion of the points that could be improved in the model, as well as interesting future directions to take, is presented in the last chapter of this thesis, chapter 11.

1.3.2   Testing the new hippocampal spiking model

The computational model for the hippocampus has been subjected to several tests or simulation experiments, presented in the appropriate chapters. This tests were:

- *basic simulations* - fundamental results are tested. The correct storage and recall of spiking patterns is analysed, and as well as place field formation and transformations;

- *validation experiments* - in such simulation experiments, electrophysiological results that were not directly used to constrain the model serve to test the compliance of the artificial system with the biological hippocampus.

- *predictions* - some predictions are made as a result of simulations using conditions not yet experimentally tested.

Chapter 2

The Mammalian Hippocampal Formation

The purpose of this chapter is to provide the reader with the most relevant properties of the hippocampus, in the context of this thesis.

Far from being exhaustive, the information presented here is intended to create a clearer view of the simplifications used to build the computational model discussed in this thesis. This information is therefore highly summarised and targeted to validate the choices made in the construction of the model. For a more complete discussion on the mammalian hippocampus see, for example, Amaral and Witter (1989) and Johnston and Amaral (1998).

Organisation of the Chapter

This chapter is divided into two sections. Section 2.1 addresses the neuro-anatomical information and is organised in subsections dealing with the different anatomical structures that have been modelled. Experimental data regarding population sizes and connectivity are provided.

The major sources for this data are Johnston and Amaral, 1998; Patton and Mc-Naughton, 1995; Freund and Buzsaki, 1996.

Section 2.2 discusses the electrophysiological properties of hippocampal neurons. This section is also organised in subsections where topics such as population activity levels, firing properties and synaptic plasticity are addressed.

2.1   Hippocampal morphology

The hippocampus is part of the limbic system and is located inside the medial temporal
lobe. Figure 2.1 shows the location of the hippocampus in both rat and human brain.
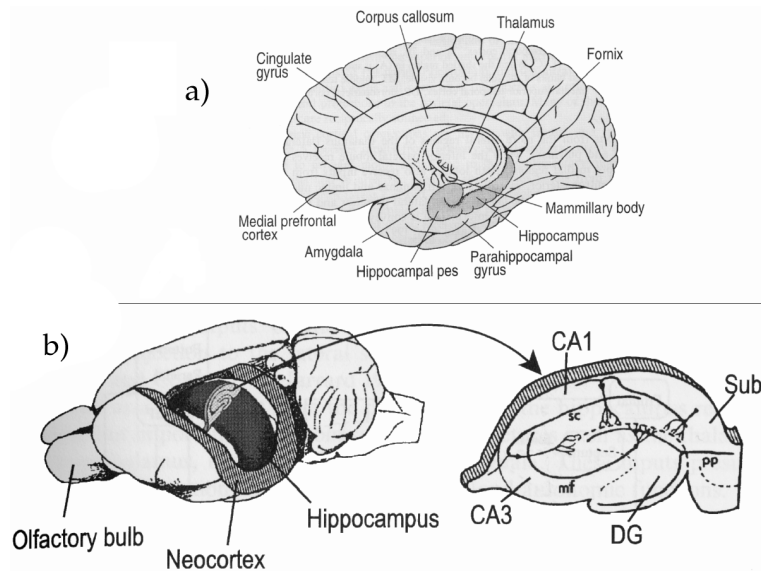
Figure 2.1: Location of the hippocampus in the mammalian brain: a) human hippocampus; b) rat hippocampus. Adapted from Burgess *et al* (1999) ©Oxford University Press.

The hippocampus belongs to a group of structures called the hippocampal formation which includes the dentate gyrus, the hippocampus proper, the subiculum, presubiculum and parasubiculum, and the entorhinal cortex(Johnston and Amaral, 1998).
The hippocampus proper comprises regions CA1, CA2 and CA3 (**C**ornus **A**mmonis).

The word *hippocampus* is used to refer to the dentate gyrus (DG), CA1, CA2 and CA3.

Region CA2 is very reduced in size and is considered to be a transition zone between CA1 and CA3 (Johnston and Amaral, 1998). No functional roles are typically assigned to CA2 and therefore this area is often ignored in theoretical studies. The hippocampal model described in this thesis incorporates the dentate gyrus, CA3 and CA1.

Through the entorhinal cortex and parahippocampal gyrus, the hippocampus receives inputs from virtually all association areas in the neocortex. This means that the hippocampus has direct access to processed information acquired from a panoply of
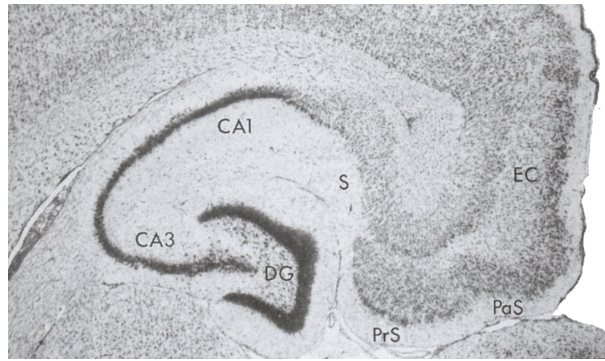
Figure 2.2: Coronal section of the hippocampal formation: dentate gyrus (DG), sub-fields CA1 and CA3, subiculum (S), presubiculum (PrS), parasubiculum (PaS), and entorhinal cortex (EC). Adapted from Johnston and Amaral (1998).

sensory channels (Rolls and Treves, 1998).

Although this multi-modal information is, to some extent, blended together, there is still some topographical segregation of the sources (Rolls and Treves, 1998). This segregation that persists in the entorhinal cortex is of particular relevance for the computational model for spatial memory.

The number of principal cells in the entorhinal cortex is estimated to be $200,000$ in the rat, and its activity level at 7% (O'Reilly and McClelland, 1994).

The hippocampus also receives regulatory cholinergic inputs from the septal nuclei, via the precommissural branch of the fornix. The projection to CA3 is bilateral and at least some of the neurons that project from CA3 to the septal region are GABAergic (Johnston and Amaral, 1998).

The major output of the hippocampus is CA1, sending connections back to the entorhinal cortex (layers III-V, but predominately layer V) and to the subiculum. The connections to the subiculum are again topographically organised (Johnston and Amaral, 1998).

Since the computational model developed in this thesis is addressed to explain the functional behaviour of the hippocampus, this discussion will continue focusing on the properties of dentate gyrus, CA3 and CA1.
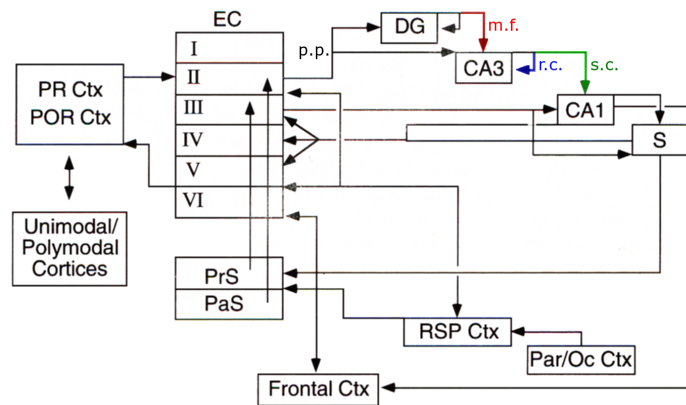
Figure 2.3: Diagram of hippocampal formation connections. New abbreviations: PR, perirhinal; POR, postrhinal; Par/Oc Ctx, pariental occipital cortices; RSP Ctx, retrosplenial cortex. The fornix, not represented in this diagram, connects the hippocampus with the hypothalamus. Pathways: p.p. perforant path; m.f. mossy fibres; r.c. recurrent collaterals; s.c. Schaffer collaterals. Adapted from Johnston and Amaral (1998).

2.1.1   Dentate gyrus

The principal cells in the dentate gyrus (DG) are the granule cells with a population size of 1000×10^3 in the rat (Boss *et al*, 1987), and 9000×10^3 in the human (Squire *et al*, 1989). Granule cells receive inputs from layer II of the entorhinal cortex (ECII) through the perforant path.

The dentate gyrus has a rich variety of interneurons which include mossy cells, GABAergic polymorphic cells, basket cells, axoaxonic chandelier cells and molecular layer perforant path cells (Freund and Buzsaki, 1996).

Mossy cells stand out among these types of interneurons: besides being one of the most abundant types of interneurons (3×10^4) in DG, these glutamergic cells provide positive feedback to granule cells. Each granule cell receives around 1000 inputs from mossy cells (Patton and McNaughton, 1995).

It has been reported that around 400 synapses (7% of excitatory input) need to be activated in order to trigger an action potential in the granule cell's soma (Patton and McNaughton, 1995).

The axons of granule cells, called mossy fibres (m.f.), establish very sparse connections with CA3. Each mossy fibre connects to around 16 CA3 pyramidal cells

(Johnston and Amaral, 1998). The mossy fibre synapses, also called boutons, are very specialised structures that can have more than 30 release sites (Chicurrel and Harris, 1992). Its efficacy is furthermore increased due to the fact that they establish connections with the soma or proximal dendrites of CA3 pyramidal neurons (Johnston and Amaral, 1998).

2.1.2 CA3 region

The principal cells in the CA3 region are pyramidal neurons, with a population size of 160×10^3 in the rat (Boss *et al*, 1987), and 2300×10^3 in the human (Squire *et al*, 1989).

Pyramidal neurons receive excitatory inputs from (Urban *et al*, 2001):

- perforant path (p.p.) - from entorhinal cortex layer II - 4000 synapses;

- mossy fibres (m.f.) - from dentate gyrus - 50 synapses;

- recurrent collaterals (r.c.) - recurrent connections from CA3 - 12,000 synapses.

Action potentials produced in CA3 are propagated through the Schaffer collaterals to CA1. In terms of projections to other subcortical regions, the only sizable connection is to the lateral septal nucleus (Swanson and Cowan, 1977). This connection is bilateral.

The existence of this projection, together with the fact that acetylcholine is a prominent excitatory neurotransmitter in the hippocampus, stimulated the view of the lateral septal nucleus as modulatory system in the complementary tasks of learning and recall in the associative memory model for CA3 (Hasselmo *et al*, 1995).

2.1.3 CA1 region

The principal cells in CA1 are again pyramidal cells with a population size of 250×10^3 in the rat (Boss *et al*, 1987), and 4600×10^3 in the human (Squire *et al*, 1989). CA1 pyramidal cells receive inputs from layer III of the entorhinal cortex (ECIII) and from CA3 pyramidal neurons, through the Schaffer collaterals. The pyramidal neurons in CA1 also receive a septal projection. Although similar to CA3, this projection is substantially lighter (and unidirectional).

In terms of outputs, CA1 is the first field in the hippocampus originating return projections to the deep layers of entorhinal cortex, especially layer V (Rolls and Treves, 1998). In addition, CA1 also projects to the subiculum.

The following table summarises the population sizes (PS) of the principal neurons in the hippocampus:

region	cell type	PS$\times 10^3$ (rat*)	PS$\times 10^3$ (human)
DG	granular	1000	9000
CA3	pyramidal	160	2300
CA1	pyramidal	250	4600

Table 2.1: Population sizes of the principal neurons in the hippocampus. *These estimates are for the Sprague-Dawley and Wistar strains (Boss *et al*, 1987).

Brief note on connectivity figures

Two connectivity types are used in this thesis. The definitions are:

- divergent connectivity C^d: counts the number of output connections that a neuron type establish with a specific post-synaptic population;

- convergent connectivity C^c: counts the number of input connections that a neuron type receives from a specific pre-synaptic population;

The correct conversion from convergent to divergent connectivity (and vice-versa) requires the knowledge on the connectivity profile spatial distribution. If one ignores the topographic properties present in almost all hippocampus pathways and assume uniform spatial distributions, it is possible to relate C^c with C^d:

$$C^c = \frac{C^d N_A}{N_B} \tag{2.1}$$

The constants N_A and N_B are the sizes of pre and post-synaptic populations, respectively.

2.2   Hippocampal electrophysiology

This section presents relevant properties regarding the functional behaviour of the hippocampus.

2.2.1 Activity levels

The topic of activity levels in the hippocampus is subject to strong debate. The importance of specific levels of activity for the functional behaviour of the hippocampus is unclear and precise experimental measurements of the activity levels, across different conditions, are unavailable. Each field in the hippocampus, as well as the entorhinal cortex, exhibits constrained activity levels. However, instead of constant across time, the activity levels oscillate within controlled bounds (Barnes *et al*, 1990; Jung and McNaughton, 1993). The average activity level is specific to each field.

Experimentally, the available estimates for the levels of activity of a hippocampal field can be computed as the mean firing rate of the neurons belonging to this field, divided by maximum firing rate of the measured sample (Barnes *et al*, 1990). This gives the percentage of neurons firing at maximum rate.

The average activity level values obtained using this method are presented in the table 2.2.

Area	n. neurons	activity level [%]
EC	200,000	7.0
DG	1,000,000	0.5
CA3	160,000	2.5
CA1	250,000	2.5

Table 2.2: Activity levels at EC, DG, CA3 and CA1 fields. Data from Barnes *et al* (1990); Jung and McNaughton (1993).

It is my belief that specific values of activity levels are necessary for the correct functional behaviour of the hippocampus. The hippocampal model constructed in this thesis incorporates this belief. Constrained activity levels are required for the type of population coding assumed in this thesis: since information is represented as the activation of specific constellations of neurons, it is important that the total number of active neurons in the hippocampal populations be constrained.

Throughout this thesis, values for *instantaneous activity levels* will be used instead of the rate based activity levels. The instantaneous activity level is defined as the ratio of neurons that are above threshold over the total number of neurons in the population.

These two quantities, instantaneous and rate-based activity level, are numerically

similar and therefore the figures presented above will be used to provide estimates for the instantaneous activity levels. A justification that both quantities numerically similar is given below:

Let us assume that each neuron's firing follows a Poisson distribution with parameter λ_i, $i = 1, ..., N$ where N is the total number of neurons on the population.

By the properties of the Poisson distribution, the firing of the whole population is also given by a Poisson distribution with parameter $\sum_i \lambda_i = N < \lambda_i >$. From this we take that the number of spikes in the population in a time window Δt is given by $N < \lambda_i > \Delta t$.

The maximum firing rate $Max_i(\lambda_i)$ sets the appropriate time window to convert spike counts to average firing.

This way, we calculate that the amount of spikes in a time snapshot of the population is given by $N < \lambda_i > /Max_i(\lambda_i)$. Furthermore, the fraction of neurons in the population that are firing is $< \lambda_i > /Max_i(\lambda_i)$, which is the definition used to obtain the data in 2.2.

These mathematical steps do not have the intention of showing that the two definitions, instantaneous activity level and rate-base activity level, are equivalent. In fact, strong approximations are used in these steps such as the assumption that each neuron's firing follows a Poisson distribution and that each λ_i is constant during the time window of $1/Max_i(\lambda_i)$. The important thing is that it shows that both quantities are related and within the same order of magnitude. The hippocampal model is not invalidated if the exact values for the activity levels in the hippocampus regions are not used. As long as the true biological activity levels are low, all the properties and implications of the hippocampal model hold.

2.2.2   Firing properties of CA3 and CA1 neurons

Instead of isolated action potentials, both CA3 and CA1 pyramidal neurons tend to fire in short bursts. The properties of the burst are different between the two types of neurons (see figure 2.4):

- CA3 neurons tend to fire in bursts of action potentials with declining amplitudes (Johnston and Amaral, 1998);

- CA1 neurons fire repetitively but show accommodation and fast and slow after-hyperpolarisations (Johnston and Amaral, 1998).
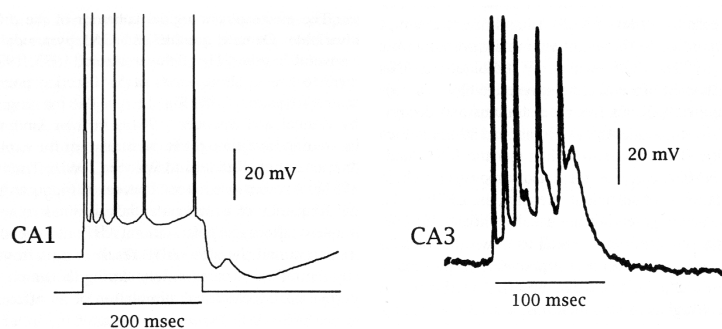
Figure 2.4: Firing properties of CA3 and CA1 pyramidal neurons. Adapted from Johnston and Amaral (1998).

2.2.3 Synaptic plasticity

Several pathways in the hippocampus show plasticity in terms of the excitation level that they provide to the post-synaptic populations. The properties of this plasticity differ nevertheless from pathway to pathway.

Mossy fibres

Mossy fibres have been shown to display a variety of long-term synaptic plasticity including associational and non-associational (NMDA independent) forms of long-term depression and potentiation (LTD and LTP) (see Bliss and Lømo, 1973). Nevertheless experiments with rats in which mossy fibre LTP/LTD was eliminated, had no significant effects on learning task performance (Urban *et al*, 2001).

Mossy fibres also display short-term dynamics, mostly driven by facilitation properties (Urban *et al*, 2001). These short-term dynamics allow mossy fibres to work in a large spectrum of synaptic strengths. Within the input stimuli range from 0 to 100Hz, mossy fibre boutons can increase their post-synaptic response by 4000% (Urban *et al*, 2001).

Perforant path

The perforant path synapses exhibit long-term potentiation. Of particular interest is the associative nature of the perforant path connections with CA1 and CA3 which are thought to have significant functional importance (Resmondes and Schuman, 2004).

Recurrent collaterals and Schaffer collaterals

Both recurrent collaterals and Schaffer collaterals exhibits associative LTP (Johnston and Amaral, 1998). That is, in order for the synapse to get potentiated, the post-synaptic membrane has to be depolarised at the time of arrival of the pre-synaptic signal. In general this depolarisation is the consequence of post-synaptic neuron firing.

2.2.4 Rhythms

The hippocampus is known to show several types of potential oscillations. The two most important in the rodent hippocampus are theta and gamma:

Theta

This is probably the rhythm in the hippocampus which is more studied for its functional role. The theta oscillations are characterised by a frequency on the band of 4-12 Hz (Buzsáki, 1989; Stewart and Fox 1990). These oscillations are often temporally nested with faster gamma-frequency oscillations (Buzsáki, 1983; Bragin et al. 1995).

The theta oscillation is delivered by the septo-hippocampal pathway. More precisely, it originates from the medial septum-diagonal band of Broca (MS-DBB). This oscillation is thought to have a functional role in the hippocampal representation of spatial information (O'Keefe and Recce 1993; Skaggs *et al*, 1996), in time-locking cell activities (Buzsáki, 2002) and regulating leaning by facilitating the induction of synaptic plasticity (Buzsáki, 1989; Huerta and Lisman 1993)

The medial septum modulation driving the theta rhythm has two components, cholinergic and GABAergic:

- The cholinergic input targets principal neurons and inhibitory interneurons in the hippocampus causing an increase in intrinsic cell excitability. This increase in excitability is achieved through suppression of various potassium currents. It also leads to a decrease in Schaffer collaterals synaptic transmission achieved by presynaptic inhibition.

- The inhibitory GABAergic input targets interneurons only. It has the effect of synchronising (phasing) the interneurons leading to population activity at theta frequency.

Theta waves are most present during REM sleep (Jouvet, 1969) and during various types of locomotor activities described by the subjective terms "voluntary", "preparatory", "orienting" or "exploratory" (Vanderwolf, 1969). In the immobile animal theta waves are generally absent but epochs of theta trains can be elicited by noxious conditioned stimuli (Bland, 1986).

Gamma

Gamma oscillations are faster than theta, ranging from 40 to 100 Hz (Buzsáki *et al*, 1983; Bragin *et al*, 1995). Several mechanisms can give rise to these oscillations such as mutual excitation, excitatory-inhibitory loops or mutual inhibition (see Traub *et al*, 1998). While the theta rhythm is driven by an external action, the gamma oscillations are a result of internal mechanisms.

Chapter 3

Spiking Neuron Models

The assumptions required to build a model often imply that the system under study is constrained to a sub-domain of its dynamics. This means that different research goals about a specific system involve the use of different models. In general terms, a good model represents a balance between the complexity of its dynamics, which reproduce reliably many of the system's properties, and the complexity of its mathematical analysis. Another point to take into account is that the degree of complexity of a model, expressed in terms of amount of parameters involved, should always be minimal. It is this process of "variables reduction" that leads to the explanation at the heart of the system's dynamics. Naturally, the computational model for the hippocampus discussed in this thesis also disregards many well known and well described properties of hippocampal morphology and physiology. This chapter's objective is to compare, at the level of single neurons, the most important models and present the arguments for the choice of the *integrate-and-fire* model to represent the neuron units in the hippocampal model.

Organisation of the Chapter

Section 3.1 briefly introduces and compares the two main families of single compartmental models for single neurons: the firing-rate and spiking models. Section 3.2 starts by laying down the arguments for the use of spiking units in the hippocampal model. The characteristics of the type of spiking model used, the integrate-and-fire model, are then described in detail.

The discussion then focuses on the properties of synapses in spiking networks. The topics addressed are post-synaptic currents, synaptic interactions and plasticity.

The chapter concludes with a discussion of the major limitations of the single neuron model used, in the context of the hippocampal model.

3.1   Models for single neurons

At the level of single neurons, several models have been proposed to cope with different questions or domains of analysis. Most of them fall into one of two categories, defined in terms of the neuron's output form: in *firing-rate* models (Dayan and Abbott, 2001) the output is continuous and represents the instantaneous firing frequency of the neuron unit while *spiking* neuron models (or *pulse* models) produce discrete (binary) impulses representing the presence of action potentials (Gerstner and Kistler, 2002).

Detailed conductance-based neuron models, as well as multi-compartmental models, are not discussed. Although they reproduce electrophysiological measurements with a high degree of accuracy (Hines and Carnevale, 1997; Koch and Burke, 1998), their intrinsic complexity and extensive amount of free parameters makes them unsuitable for the creation and analysis of large networks (which are required for the construction of the present hippocampal model).

3.1.1   Firing-rate models

The reason behind using firing rates instead of action potentials to define the output of a neuron goes beyond mathematical simplification. It is grounded on the deep belief that precise timing of spike production/arrival may not contain relevant information (Koch, 1999). This is motivated by the fact that, in most neural systems, responses to the presentation of the same stimuli are not based on identical spike trains (see figure 3.1). This means that all relevant information is coded in terms of neurons' firing rates.

One example with particular interest for this thesis refers to hippocampal place cells in which the information for a place is considered to be coded as firing rates.

Defining v as the firing rate of the output unit and $\vec{u}$ the vector of input firing rates, the general equation for firing rate units is (Dayan and Abbott, 2001):

$$\tau_r \frac{dv}{dt} = -v + f(\vec{w} \cdot \vec{u}) \tag{3.1}$$

The time constant τ_r measures how closely v can follow rapid fluctuations on time-dependent inputs (it does not represent the neuron's membrane time constant). Vector
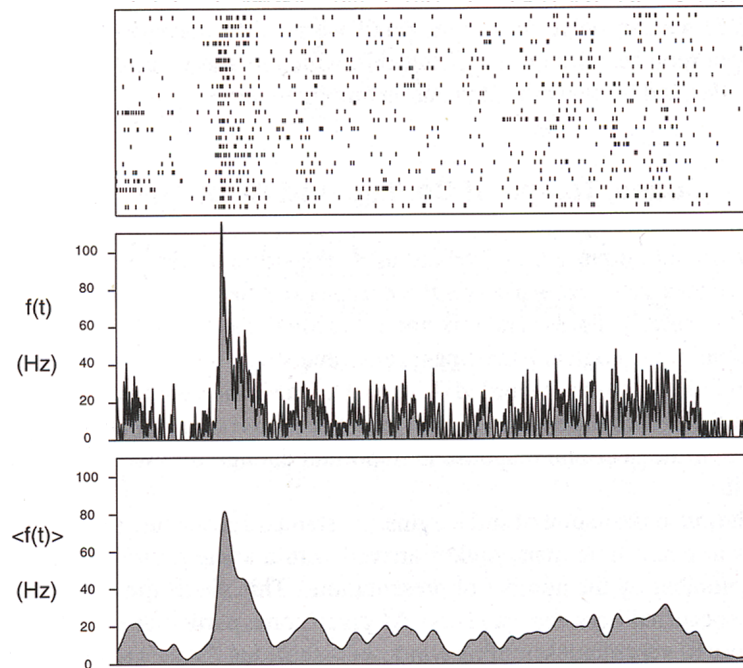
Figure 3.1: The neural spiking responses (upper panel) to the same stimulus vary, indicating that information may not be coded in the precise timing of action potentials. The lower panels correspond to filtered averages of spikes per unit of time. The second panel uses a Gaussian filter with 2 ms standard deviation (producing what we can call an *instantaneous firing rate* $f(t)$) while the lower panel uses a standard deviation of 20 ms. Adapted from Koch, 1999.

$\vec{w}$ represents the coupling (or efficacy weight) between each input unit u_i and the output unit v. The function f is called an activation function and represents the firing rate of unit v in steady-state. The activation function is typically taken to be a threshold linear function or saturating function, such as a sigmoid. Saturating functions are particularly useful for stabilising network dynamics and are therefore extensively used in recurrent network models (Dayan and Abbott, 2001).

In terms of analysis, firing-rate models present several advantages over spiking models. By neglecting very short time scales, firing-rate models can be simulated more efficiently on computers. Also, their tendency to have a smaller number of free parameters in comparison with spiking models makes their parametrisation easier (Dayan and Abbott, 2001). But definitely, the most important advantage of firing-rate models is their simpler mathematical formulation which allows analytical analysis of several

aspects of network dynamics (Hertz *et al*, 1991) that cannot be perform in spiking models due to their non-linearities. One particular domain in which the simple mathematical formulation of firing-rate models becomes important is stability analysis.

Networks described by dissipative non-linear differential equations (or any non-linear dynamical system) generally contain equilibrium points in the state space, also called fixed points. The stability of these equilibrium points can be assessed using Lyapunov's stability theory (Hopfield, 1982). Of particular interest for neural network dynamics are the *uniformly stable* equilibrium points. An equilibrium state $\bar{\mathbf{x}}$ is said to be uniformly stable if for small displacements in the state space, the system converges[1] back to $\bar{\mathbf{x}}$ (see Haykin, 1999).

The network models that exploit uniformly stable equilibrium points are called attractor networks, of which the most well known type is the Hopfield network (Hopfield, 1982). In these networks, uniformly stable equilibrium states are called attractor states. Just as in Hopfield networks, it is possible to calculate and, more importantly, manipulate the attractor states of recurrent networks that employ firing-rate units (Káli and Dayan, 2000).

Attractor networks offer a very attractive way of modelling memory systems for several reasons such as their robustness to noise, correct retrieval with partial cues and graceful degradation. As a result, networks of firing-rate units have been *the* elected model to explain biological memory systems such as the CA3 subfield and to create functional models of the hippocampus (e.g. Treves and Rolls, 1994; Káli and Dayan, 2000).

3.1.2   Spiking models

In spiking models, the discrete nature of a biological neuron's output in terms of stereotyped impulses is preserved and no assumption is made, *a priori*, on how information is coded.

The original model of a neuron in McCulloch and Pitts (1943), the spike response model, SRM, (Gerstner and Kistler, 2002) and the widely used integrate-and-fire model (first investigated by Lapicque, 1907), are all examples of spiking models.

---

[1]Formally, if $x(t)$ is the state vector, an equilibrium state $\bar{\mathbf{x}}$ is said to be uniformly stable if for any given positive ε there is a positive δ such that the condition $\|\mathbf{x}(0) - \bar{\mathbf{x}}\| < \delta$ implies $\|\mathbf{x}(t) - \bar{\mathbf{x}}\| < \varepsilon$ for all $t > 0$.

By including short time-scale dynamics, spiking models provide information about spike timings. Action potentials are still stereotyped impulses but are well defined individual entities in the time domain.

Due to this increase in complexity, most types of spiking models offer a considerable challenge in terms of their analysis, in particular with their stability analysis. In fact, in most cases, an analytical study is not possible without constraining the spiking model to a smaller domain of its dynamics.

These mathematical difficulties arise mainly from the discontinuous transition from sub-threshold to supra-threshold dynamics. It is nevertheless this property, the non-linearity of the dynamics, that gives networks of spiking units their high computational power.

3.2 Hippocampal model's neuron units

Behind the hippocampal computational model there are fundamental assumptions that are incompatible with single unit firing rate models. One such assumption regards how information is coded and transferred in the hippocampus. Therefore, spiking neuron models were used to represent single neurons.

In this thesis it is proposed that information in the hippocampus is coded neither in terms of firing rates nor precise spike timings. Instead, information is transferred, manipulated and produced in the form of spatial activity patterns (section 1.3), i.e. sub-sets of the neural population that produced an action potential within a time window of about 20ms (the order of the neurons' integration time). This form of population coding requires each action potential to be treated as an individual entity.

The *integrate-and-fire* neuron was the spiking model of choice to simulate and analyse the functional properties of the hippocampus. It is a model that has been extensively used and successfully applied in many studies.

In the words of Koch (1999), the "*venerable* integrate-and-fire model comes in many flavours". The characteristics of the particular version used in this thesis are now described in detail.

3.2.1   Integrate-and-fire neuron model

The sub-threshold dynamics of hippocampal model's neuron units are defined by the general integrate-and-fire equation which relates the voltage state variable V ([mV]), with the input current $I_{in}(t)$ ([nA]):

$$\tau_m \frac{dV}{dt} = V_{rest} - V + R_m I_{in}(t) \tag{3.2}$$

Once the voltage exceeds a predefined threshold value V_{thresh}, the time evolution of V is substituted by a stereotyped wave. The duration of the wave form is 5 ms, after which the time evolution of the variable V follows once again the dynamics of equation 3.2 with initial condition $V = V_{reset}$. The shape of the wave is in accordance with experimental data in Spruston *et al* (1995).

The absolute refractory period T_{ref} is coincident with the duration of the stereotyped wave during which all current inputs to the soma are disregarded.

All the constants used in the model take biologically plausible values (Nicholls *et al*, 1992; Koch and Segev, 1998; Koch, 1999; Dayan and Abbott, 2001) and, unless stated, are defined as:

- $R_m = 20M\Omega$: membrane resistance (or input resistance);

- $\tau_m = 20ms$: membrane time constant

- $C_m = 1nF$: membrane capacitance;

- $T_{ref} = 5ms$: absolute refractory period (maximum firing frequency is $\sim 200Hz$);

- $V_{rest} = -70mV$: membrane resting potential;

- $V_{thresh} = -50mV$: membrane firing threshold voltage;

- $V_{reset} \sim -60mV$: membrane reset voltage;

The only atypical value, in the framework of integrate-an-fire models, is the membrane reset voltage. While typically $V_{reset} = V_{rest}$ (Dayan and Abbott, 2001), it has been decided to set V_{reset} at a voltage closer to threshold. The reason for that is that, once a first action potential is fired, it is easier to fire a second. This promotes a bursty behaviour that, as will be shown in chapters 7, 8 and 9, is important for the functional

model of the hippocampus. It should be noted that both CA3 and CA1 pyramidal neurons fire in short bursts (see figure 2.4).

It remains now to discuss the input current term. The simplest case is to consider a constant input current. In this situation, the integrate-and-fire model dynamics acquire an oscillatory behaviour with constant period. The example in which the input current $I_{in}(t) = 1.01nA$ is depicted in figure 3.2. The threshold current is $1.0nA$.
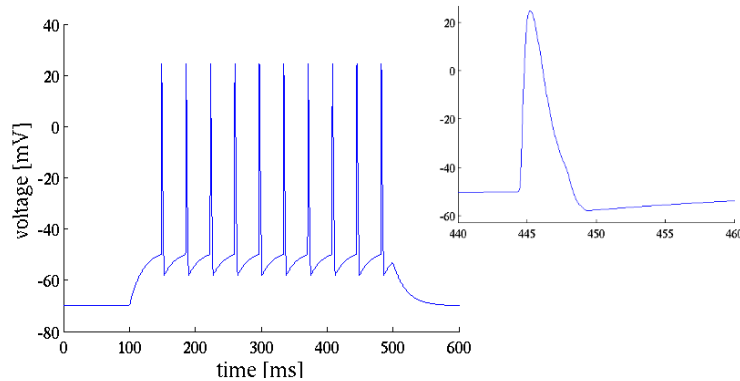
Figure 3.2: Integrate-and-fire neuron. A constant input current $I_{in} = 1.01nA$ is injected between 100 and 500 ms.

With constant input current I_{in} above the threshold current, the neuron unit fires periodically with an inter-spike interval T_f that can be calculated directly from equation 3.2:

$$T_f = T_{ref} - \tau_m \ln \left(\frac{R_m I_{in} + V_{rest} - V_{thresh}}{R_m I_{in} + V_{rest} - V_{reset}} \right) \tag{3.3}$$

The more interesting case in which $I_{in}(t)$ reflects the signal communication with other neuron units is discussed below.

3.2.2 Synaptic currents

When a spike reaches the neuron unit, it generates a synaptic current I_{syn} that contributes to the integrate-and-fire input current I_{in}. This synaptic current is modelled using a post-synaptic conductance profile described by a dual exponential function.

Setting t_{sp} as the spike arrival time, the post-synaptic conductance $S(t)$ is defined for $t \geq t_{sp}$ as:

$$S(t) = S_{peak} \times norm \times \left(e^{-\frac{t}{\tau_{decay}}} - e^{-\frac{t}{\tau_{rise}}} \right) \tag{3.4}$$

For $t < t_{sp}$ the post-synaptic conductance is zero. The normalisation constant *norm* puts the conductance peak at S_{peak} and has the value:

$$norm = \left(\left(\frac{\tau_{rise}}{\tau_{decay}} \right)^{\frac{\tau_{rise}}{\tau_{decay} - \tau_{rise}}} - \left(\frac{\tau_{rise}}{\tau_{decay}} \right)^{\frac{\tau_{decay}}{\tau_{decay} - \tau_{rise}}} \right)^{-1}$$

The dual exponential time constants are specific for the type of synapse. For excitatory synapses, and unless otherwise stated, they take the following values:

- $\tau_{rise} = 2ms$: conductance rise time constant;

- $\tau_{decay} = 5ms$: conductance decay time constant;

The values typically used for inhibitory synapses are $\tau_{rise} = 1ms$ and $\tau_{decay} = 25ms$.

Instead of reproducing the post-synaptic response at the synaptic loci in the dendrite, the conductance time constants are set to mimic the synaptic input current profile at the soma. In this way, the effect of static dendritic filtering is approximated. The dendrite-fibre distance considered for the dendritic filtering corresponds to the average synaptic distance for hippocampal pyramidal cells ($\sim 150\mu m$). More details on this are given in chapter 4. The conductance time constants, for both excitatory and inhibitory synapses, are biologically plausible and in accordance with experimental data (Cowan *et al*, 2001).

The value S_{peak} should also be regarded as an efficacy measure instead of an efficiency. While "efficency" refers to the size of the post-synaptic response, "efficacy" accounts for the true contribution of the synapse to the firing of the neuron.

The synaptic current that affects the soma is given by:

$$I_{syn}(t) = (V - E_{syn}) \times S(t) \tag{3.5}$$

The constant E_{syn} represents the synapse's reversal potential and, again, depends on the nature of the synapse. Throughout this thesis, the synapse's reversal potential takes the value $0mV$ for excitatory synapses and $-90mV$ for inhibitory synapses (Cowan *et al*, 2001).

In order to produce the threshold current of $1.0nA$, a single standardised excitatory synapse ($\tau_{rise} = 2ms$, $\tau_{decay} = 5ms$) has to produce a peak conductance $S_{peak} = 62.6nS$

The dual exponential parameters described above are able to recreate the essential properties of the synaptic responses in the hippocampus (Johnston and Amaral, 1998). In the hippocampal model they are used to recreate excitatory post-synaptic responses and both fast ($GABA_A$) and slow ($GABA_B$) inhibitory post-synaptic responses

So far we have discussed the case of an isolated synapse. The interactions between synapses and how each synaptic current contributes for the total input current of the integrate-and-fire neuron will now be discussed.

3.2.3 Synaptic interactions

In single compartmental models, such as the integrate-and-fire model, all neuronal topological information is disregarded: the whole neuron, including dendritic tree, soma and axon, is collapsed into a point. In other words, all the richness of dendritic synaptic interactions is lost (Koch, 1999).

Another strong simplification taken here is that synaptic currents are assumed to interact linearly. The input current I_{in} that drives the integrate-and-fire neuron is obtained by the linear sum of all arriving synaptic currents:

$$I_{in}(t) = \sum_{syn} I_{syn}(t) \tag{3.6}$$

This reduces greatly the complexity of the model's analysis and simulation, but reduces also its computational power. Non-linearity is a fundamental property for information processing systems (MacKay, 2003) and, in this neuron model, exists only in the firing mechanism. This nevertheless proves to be enough to model and explain the functional behaviour of the hippocampus.

Although many local computations can be performed at dendrites by using specific synaptic spatial configurations and non-linear interactions (Mel, 1994), we assume in this hippocampal model that they do not play an essential role in hippocampal functional behaviour.

Moreover, there is no direct experimental evidence for the existence in the hippocampus of learning mechanisms able to produce such dendritic computational units, or selection mechanisms to enhance advantageous local configurations. It is important to bear in mind that the dimensions of the space in which all possible synaptic configurations are represented (for a specific dendritic tree) is immense. Very effi-

cient mechanisms have to exist in order to create and select these (speculative) local computational units.

3.2.4   Synaptic plasticity

This subsection presents the plasticity dynamics, or learning rules, that have been used with the spiking neuron model presented. A very important point to keep in mind is that these rules are not optimal in the sense of information storage. Rather they reflect mechanisms that 1) are available to hippocampal pyramidal neurons and 2) are supported by biological experimental data.

A synapse can change its efficacy, represented in terms of its peak conductance S_{peak}, as a result of two independent processes:

1. Short-term plasticity changes the synapse's peak conductance temporarily (in the time scale from hundreds of *ms* to a few seconds).

2. Long-term plasticity produces permanent changes on S_{peak}.

Since chapter 5 is entirely dedicated to short-term plasticity in the hippocampus (particularly on its computational repercussions), the following discussion is centred around long-term plasticity in spiking neurons.

Long-term plasticity can be of the form of long-term potentiation (LTP) or long-term depression (LTD).

Two types of long-term potentiation (LTP) exist in hippocampal fibres: associational and non-associational. The simplest type is the non-associational, that is, the potentiation of the synapse occurs independently of post-synaptic neuron's state. On the other hand, associational LTP requires that the pre-synaptic signal arrives at a time at which the post-synaptic membrane loci is depolarised. In modelling studies this is typically regarded as the consequence of post-synaptic firing. There are nevertheless other mechanisms, such as local synaptic interactions, that may locally depolarise dendrite fibres. These other mechanisms, nevertheless, have to be disregarded in point-wise spiking neuron models.

Two forms of associational long-term potentiation are considered in the hippocampal computational model: simple associative (Hebbian) plasticity and spike-timing dependent plasticity (STDP). In simple associative plasticity, a synapse is modified only

when both pre and post-synaptic neurons fire within a short time window (independently of the relative firing times). In these circumstances, the synapse is potentiated using the rule for potentiation shown in equation 3.8 (explained in detail below). The dynamics for the spike-timing dependent plasticity (STDP) follow the results reported by Bi and Poo (1998) in cultured hippocampal cells. Some modifications are nevertheless introduced.

In spike-timing dependent plasticity the quantity t_{STDP} is defined as the difference between the pre-synaptic spike time arrival t_{sp}^{pre} and the post-synaptic spike time t_{sp}^{post}. That is, $t_{STDP} = t_{sp}^{pre} - t_{sp}^{post}$. In the canonical STDP model (e.g. Song *et al*, 2000), LTD is induced if $t_{STDP} < 0$, otherwise LTP is induced. The amount of potentiation, and depression, decays exponentially with the absolute value of t_{STDP}. The time constants for the decay measured by Bi and Poo (1998) were 34 ± 13 ms for depression and 17 ± 9 ms for potentiation. Nevertheless, in the model used here instead of an exponentially decreasing modification magnitude, constant values were used within time windows of 40 ms. A comparison between the canonical STDP model and the model used is depicted in 3.3. This rectangular STDP rule enables a simpler interpretation of the results from the hippocampal computational model.
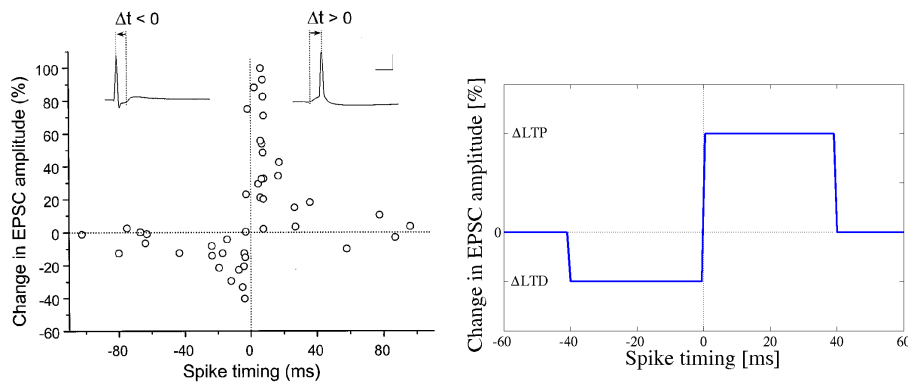
Figure 3.3: Spike-timing dependent plasticity. Left panel: canonical STDP (taken from Bi and Poo, 1998). Right panel: STDP rule used in this thesis (the quantities ΔLTP and ΔLTD are explained in the text). EPSC: excitatory post-synaptic current.

Following experimental data (Bi and Poo, 1998), the modification size induced on a synapse undergoing LTP is dependent on the synapse's initial strength. On a synapse undergoing LTD, the modification size is independent of its initial strength.

The constant ΔLTP represents the maximum potentiation that a synapse can experience. This corresponds to the conductance increase in a synapse with zero initial conductance. The constant ΔLTD corresponds to the fixed conductance decrease in synapses subject to LTD.

The amplitude of the modification size is given by:

$$\textbf{Potentiation}: \quad \Delta S_{peak} = \Delta LTP \left(1 - \frac{S_{peak}}{S_{max}} \right) \qquad (3.7)$$

$$\textbf{Depression}: \quad \Delta S_{peak} = \Delta LTD \qquad (3.8)$$

Soft-bound dynamics (van Rossum *et al*, 2000) are used to describe the potentiation increase. The quantity S_{max} represents the biological upper bound for the synaptic peak conductances.

Spike-timing dependent plasticity rules say nothing about the situation in which a pre-synaptic neuron continuously fails to fire a post-synaptic neuron. Since there is experimental evidence that support the efficacy reduction of synapses establishing such connections, the plasticity model used here was modified to accommodate this property. Thus, if a synapse is activated it always undergo some form of modification.

The depression caused in synapses that systematically fail in leading the post-synaptic neuron to threshold is nevertheless different from STDP depression. These synapses undergo homosynaptic depression and are depressed by a constant, small, percentage H.

$$\textbf{Homosynaptic Depression}: \Delta S_{peak} = H \, S_{peak} \qquad (3.9)$$

The homosynaptic depression constant H is typically chosen from 0% to 5%.

The reason to include this homosynaptic depression becomes clear when the hippocampal computational model is discussed in chapter 6. Mainly this form of depression secures a soft-degradation of information stored that allows the creation of a palimpsest type memory.

The complete set of learning rules applied to the spiking neurons network is summarised below:

pre-synaptic neuron	post-synaptic neuron	synaptic modification
silent	silent	none
silent	spike	none
spike	silent	homosynaptic depression (3.9)
spike	spike	LTP or LTD (3.8)

Table 3.1: Learning rules applied to spiking neuron units.

3.3 Final remarks

This chapter has presented the properties of the neuron model used to describe each unit in the hippocampal computational model.

The major limitation on the spiking neuron model used is undoubtedly its reduced computational power due to removal of local synaptic interactions. However, it is my belief that such eventual local computations are not necessary for the main functional properties of the hippocampus.

All the relevant details on how the integrate-and-fire dynamics were simulated will be given in chapter 4 where the simulation environment created to analyse the hippocampal model will be discussed.

Chapter 4

Hippocampal Simulation Environment

This chapter is dedicated to the description of the simulation environment, SE, which was used for all the numerical experiments and simulations of the hippocampal computational model (HCM).

Specific details about the implementation of the spiking neuron units are explained. The approximations, suppositions and limitations of the simulation environment are themselves a component of the hippocampal computational model that will be later introduced. Any model can only be considered fully described when all its simplifications and assumptions are stated.

The simulation environment presented in this chapter is, by itself, a considerable achievement of this thesis. It not only permits a detailed analysis of the hippocampal computational model but also creates a general framework where many other hypothesis can be tested. This simulation environment can be an extremely useful tool for future research in the area of spiking networks.

Organisation of the Chapter

Section 4.1 is concerned with the decisions taken about which language to code the hippocampal model (HCM). Section 4.2 exposes the general structure of the simulation environment and explains the elemental computational units. The functional behavioural of each component is presented and discussed.

The situations where specific approximations on the SE may have a relevant impact on HCM's results, are discussed in the last section.

4.1   Coding language

A model with a few thousands integrate-and-fire units, each with a couple of hundred synapses, demands a considerable amount of space to store all the state variables and requires millions of operations to compute a single time step propagation.

The limiting factor in a spiking model network simulation is the number of synapses. A model with 5,000 units, each with 200 synapses (totalising 1,000,000 synapses) requires about 80MBytes to store all the variables needed (where each synapse is characterised by ~ 20 floating point variables occupying 4 bytes each).

The number of operations required to update the state variables at each time step rules out the use of any high level programming language such as *Matlab* or *Mathematica*. Nevertheless, these languages were of great use for the treatment of results and confirmation of specific simulation calculations. *Mathematica* was of great use in all analytical calculations.

An attempt was made to use *NEURON* (Hines and Carnevale, 1997), taking advantage of the newly introduced event driven mechanisms. Some tests were made with the help of Michael Hines. Unfortunately, the difficulty in tuning specific algorithms and simulation properties, became a considerable drawback.

The decision was made to use *C* programming language. A program tuned to the task in hand was then written, with execution speed greatly improved compared with most languages, and full control over all the algorithms. This required full programming of almost all the algorithms used in the simulation environment.

4.2   Simulation environment core ideas

A biological neuron can be divided into 4 structural entities: soma, axons, dendrites and synapses. At the simulation environment, a neuron is represented by a soma unit and a list of synapse units.

The simulation environment's central structure is the synapse unit, which contains information about:

1. transmission delay - axonic and synaptic delay constants;

2. dendritic filtering - dendritic length constant and variable for electrotonic length;

3. connectivity - identification of input and output soma units;

4. arriving spikes - times of arrival of the spikes produced at the input soma unit;

5. produced post-synaptic currents (PSC) - times of initiation of PSC's at the output soma unit;

6. post-synaptic current shape - dual exponential time constants and normalisation factor;

7. synaptic strength - connection strength variable, expressed in terms of peak conductance;

8. short-term dynamics - facilitation and depression constants and variables;

9. long-term dynamics - LTP and LTD constants and variables required for spike time dependent plasticity.

Each synapse unit has its own "private" axon fibre and dendritic fibre. This means that dendritic signals interact only through the soma. This simplification is motivated by two reasons: 1) it is my belief that dendritic computations are not necessary to establish the core functional properties of hippocampal neurons and 2) this simplification greatly enhances the computational performance of the simulation environment. My belief that dendritic computations are not necessary to obtain the most relevant hippocampal functional behaviours is based in the fact that highly complex mechanisms are required to establish and maintain specific computations in the dendritic tree. Such mechanisms have not been described in the literature and, in addition, the most relevant functional properties of the hippocampus can be recreated without the use of dendritic computations.

The SE's neuron equivalent can be seen as a central soma compartment surrounded by a shell of synapse compartments connected through filters. There are no branching points at the neuron unit's dendritic tree. It is important to emphasise that spatially the neuron unit is a point: the size of the dendritic fibre associated with each synapse unit is only used for the purpose of post-synaptic current (PSC) amplitude filtering. Spatially all synapses are directly connected to the soma's surface.
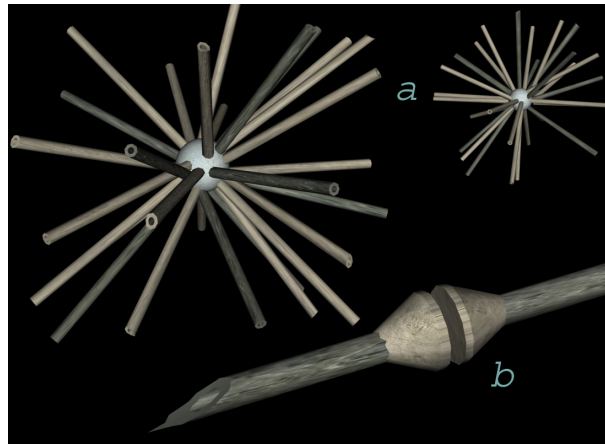
Figure 4.1: a) Illustrative shape of SE neuron's equivalent. b) Each synapse unit has its own "private" axon and dendrite fibres. The size of the axon represent the transmission delay while the size of the dendrite represents the attenuation. Spatially both these fibres have zero length.

Building a neuronal network in the simulation environment is a two step procedure: first, soma units are defined and placed; then synapse units, which are axon-synapse-dendrite wires, are placed connecting specific pairs of soma units.

Some of the simulation environment core concepts are developed below.

4.2.1   Transmission delay

Axons in this model are seen simply as introducing transmission delays. In the SE's framework, a biological axon fibre that bifurcates and ends in two synapses corresponds to two independent axons. Each of these, now independent, axons is characterised by the total time that a spike would take in travelling between the common soma and the corresponding synapse (figure 4.2).

This view does not introduce any relevant approximation, in terms of signal processing, between model and biological system. The same does not happen with the dendrite conversion.
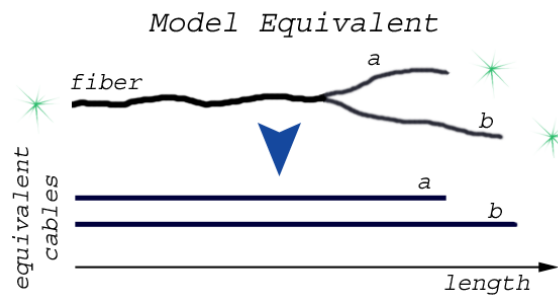
Figure 4.2: The SE neuron fibres always join the soma with a single synapse. The SE counterpart for a fibre that bifurcates and ends in one synapse each, is two independent cables. Each of these new fibres contain information about signal delay or filtering depending on the nature of the fibre.

4.2.2 Dendritic filtering

For dendrites, the conversion is identical but involves more complex approximations. Each synapse unit contains information about the dendritic distance to the soma (figure 4.2). Introducing this distance is not for delay calculation purposes; the significant signal delay is accounted for by the "transmission delay" constant mentioned before. Instead, the dendritic fibre length is required for calculating dendritic amplitude attenuation due to the presence of resistive and capacitive elements.

The low-pass filtering effect on the post-synaptic current profile is not continuously calculated. Instead, the post-synaptic current wave form (defined as a dual exponential function) already corresponds to the associated steady-state profile created at the soma level (see section 4.2.5 for further details).

The effect of frequency on the wave profile in not taken into account. For the purpose of the hippocampal model, the error associated with this approximation is negligible. On the other hand, the current attenuation produced with typical hippocampal firing frequencies is considered to be relevant (see figure 4.3). This mechanism can be a considerable ally to the inhibitory interneuron populations in the control of the activity levels at principal neuronal cells (chapter 6 is dedicated to the control of activity levels).

According to data from Brown *et al.* (1981), Spruston and Johnston (1992), Thurbon *et al.* (1994), the dendritic space constants in hippocampal neurons, including
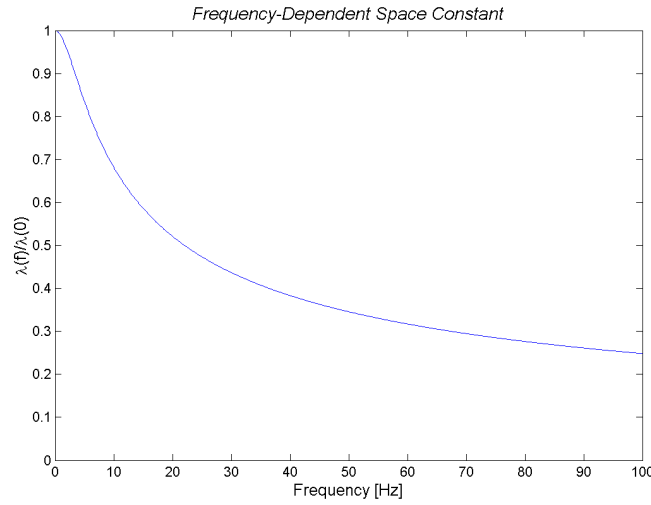
Figure 4.3: Frequency dependent space constant ($\tau_m = 50ms$).

dentate granule cells, CA1 and CA3 pyramidal neurons and also interneurons, are quite long, in the order of $2000\mu m$. In addition, the dendritic membrane time constants are quite slow, in the order of $50ms$.

This implies that steady-state potentials suffer almost no attenuation as opposed to signals propagating at higher frequencies.

The frequency-dependent space constant $\lambda(f)$ can be calculated using Fourier transforms, and is given in terms of the steady-state space constant $\lambda(0)$ by:

$$\lambda(f) = \frac{\lambda(0)}{Re\left\{\sqrt{1 + i\,2\pi f\,\tau_m}\right\}} \tag{4.1}$$

Using cable theory, the attenuation on the post-synaptic current originating at a distance x from the soma, in a dendritic fibre of total length L (see figure 4.4), is given by:

$$\frac{I_{soma}}{I_{syn}} = \frac{\cosh\left(\frac{L-x}{\lambda}\right)}{\cosh\left(\frac{L}{\lambda}\right)} \tag{4.2}$$

Since the neuron unit model assumes each dendrite containing a single synapse unit placed at the end of the fibre, therefore $x = L$. The dendritic filtering attenuation factor F_{atten}, as a function of the fibre length L and stimulus frequency f, is then given by:
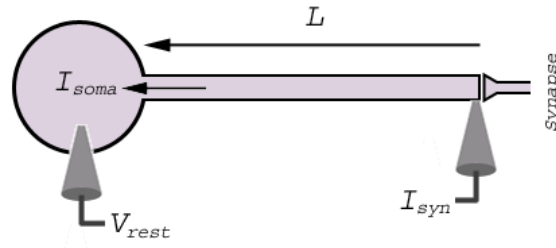
Figure 4.4: Current amplitude attenuation of a stimulus created at the dendrite, as seen at the soma.

$$F_{atten}(f,L) = \frac{1}{\cosh\left(\frac{L}{\lambda(0)} Re\left\{\sqrt{1+i\,2\pi f\,\tau_m}\right\}\right)} \qquad (4.3)$$
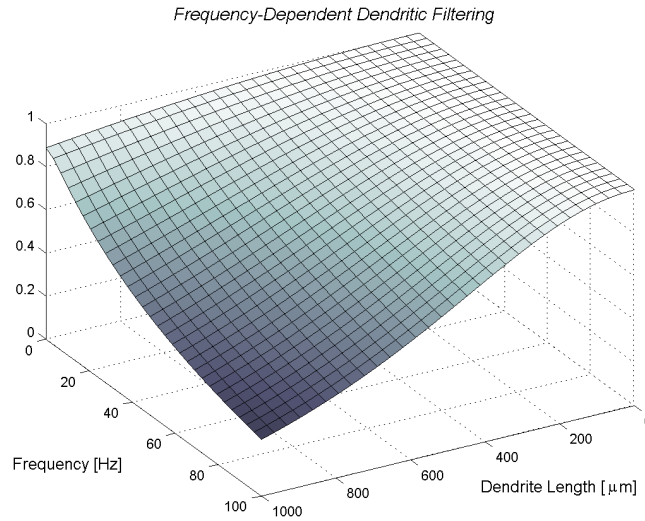
Figure 4.5: Dendritic filtering attenuation factor F_{atten} as a function of the fibre length L and stimulus frequency f.

The dendritic filtering attenuation factor is extensively used throughout the simulations. It is unwise to use this analytically calculated expression due to computational efficiency. An approximation of the form $\tilde{F}_{attenuation}(f,L) = \exp\left(k_1\left(k_2+f\right)L^2\right)$ is used instead, using the experimental values $\lambda(0) = 2000\mu m$ and $\tau_m = 50ms$. It should be noticed that the constraint $F_{atten}(f,0) = 1$ is met. The constants $k_{1,2}$ are calculated

using a least-squares fit.

This new expression, $\tilde{F}_{attenuation}(f,L)$, can be calculated using fewer computational steps. This expression is invoked around 2x per synapse unit, for several thousands synapse units, at each time step.

The approximated surface is presented in figure 4.6 together with a graphical visualisation of the approximation errors. The error is only bigger when both frequency and fibre length values are high. In this region the approximation underestimates the amplitude attenuation. This error leads nevertheless to small change in the post-synaptic currents and therefore does not pose a problem.

The attenuation in a dendritic tree with arbitrary branching can be analytically calculated using, for example, the *recursive algorithm* (Rall and Agmon-Snir, 1998). Unfortunately this method is computationally intensive. Besides that, and most importantly, it is part of the assumptions of the hippocampal model that the main functional role of extended branched dendritic trees is to collect and deliver signals to the soma; specific pre-computations at the dendritic level are not required to achieve the functional behaviour of the hippocampus.

4.2.3   Soma units and connectivity

The soma units gather all incoming synaptic currents, which are summed. Only one type of response is considered for the neuron units: spikes. Other possible forms of signalling, such as modulators like neuro-active peptides known to be present at the hippocampus (Johnston and Amaral, 1998) are absent in the HCM.

As described in the previous chapter, a linear integrate-and-fire equation (Gerstner and Kistler, 2002) is used to describe the dynamics of the only variable that defines the soma's state: the voltage V. For under-threshold conditions, the spiking model used was:

$$\tau_m \frac{dV}{dt} = V_{rest} - V + R_m I_{in} \qquad (4.4)$$

The current I_{in} is the sum of all current inputs, namely all synaptic currents with respective complex time dependent profiles.

The numerical method used to propagate in time the dynamic system was the fourth order Runge-Kutta approximation method (Press *et al.*, 1994; Wilson, 1999).
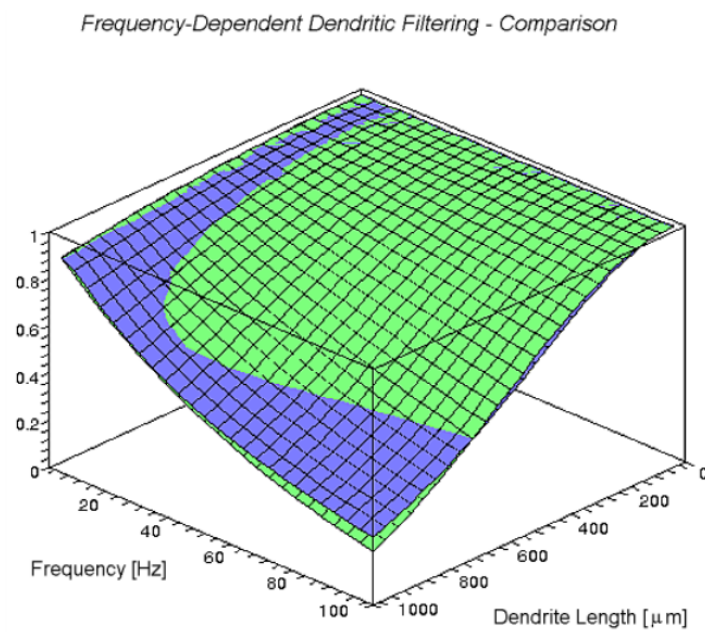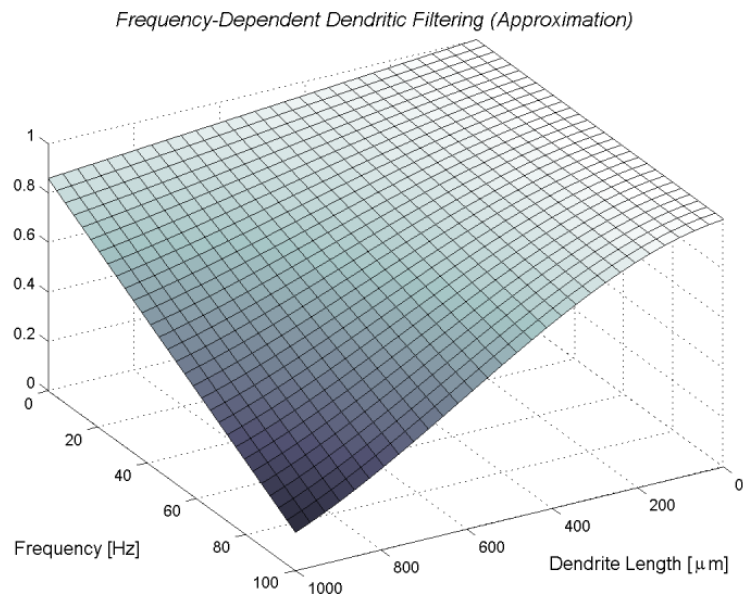
Figure 4.6: Upper graph: dendritic filtering attenuation approximation factor $\tilde{F}_{atten}$ as a function of the fibre length L and stimulus frequency f. Lower graph: $\tilde{F}_{atten}(f,L)$ in green and $F_{atten}(f,L)$ in blue.

The 'luxury' of using a fourth order method to deal with the soma's differential equations is reasonable as the SE's bottleneck is at the computations associated with synapse units (Rotter and Diesmann, 1999). This way, more accurate calculations are possible for the soma units without decreasing significantly the SE's performance.

As in every integrate-and-fire model, whenever the state variable V reaches a pre-defined threshold value V_{thresh}, a non-linearity is introduced into the dynamics. In this soma unit model, when voltage V reaches V_{thresh}, a stereotyped voltage wave is generated, for a duration equal to the neuron's *absolute refractory period* (5 ms).

The stereotyped voltage wave's role is not of mimicking an action potential. As mentioned before, spikes are fully characterised in the model by their arrival time at the corresponding synapse. That is, the propagation of action potentials is not simulated.

Instead, the stereotyped spike voltage wave follows the profile of a back-propagated action potential (BPAP) as seen at distance of $\sim 150 \mu m$ in the dendritic tree. The BPAP profile used in the SE was extracted from simulations with *NEURON* software using hippocampal neurons with physiologically realistic properties (Poirazi *et al*, 2003) which, in turn, is in accordance with the experimental data from Spruston *et al* (1995).

The voltage wave amplitude attenuation is calculated again using cable theory, assuming only membrane passive properties, and following the same approximations used in dendritic filtering:
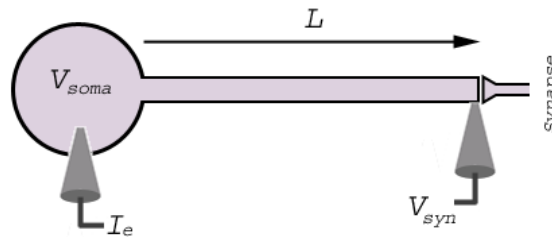
Figure 4.7: Current amplitude attenuation of a stimulus created at the soma, seen at the end of a dendritic fibre.

The attenuation on the back-propagated action potential at a distance L from the soma, in a dendritic fibre of total length L (see figure 4.2.3), is given by:

$$\frac{V_{syn}}{V_{soma}} = \frac{1}{\cosh\left(\frac{L}{\lambda}\right)} \qquad (4.5)$$

This equation leads to the same attenuation factor defined in equation 4.3. The voltage calculated at each synapse loci has impact over two mechanisms: post-synaptic current amplitude (product of driving voltage and conductance) and learning (affecting posterior conductance amplitudes through LTP and LTD).

A third and last computation is associated to the soma units: measuring firing rates. A correct, and continuous, estimation of each soma's unit firing frequency is important for correct calculation of frequency-dependent dendritic attenuation and to estimate neuron's firing-rate allowing overall comparison with firing-rate models. Several methods have been proposed in the literature to produce firing-rate estimates. The method used in the SE uses linear filters and linear kernels (Dayan and Abbott, 2001). The kernel, or window term, has to be causal which has the consequence of typically delaying estimations. A kernel is called *causal* when the firing rate estimate at time t depends only on spikes fires before t.

Two components are used to produce the continuous firing rate estimates:

- a record of the previous 50 spike times for each soma unit;

- a causal kernel;

The value 50 is chosen as the minimal number of spike times that lead to a very-good frequency estimation (within the range of hippocampal frequencies).

The window function adopted in the SE is the commonly used α function (Baddeley *et al.*, 1997):

$$\omega(\tau) = \left[\alpha^2 \tau \exp\left(-\alpha\tau\right)\right]_+$$

The half-wave rectification operation $[\cdot]_+$ is defined as:

$$[x]_+ = \begin{cases} x & : & x \geq 0 \\ 0 & : & x < 0 \end{cases}$$

The constant α^{-1} determines the temporal resolution of the resulting firing-rate estimate. Two factors influence the choice of α^{-1}: the range of frequencies in which the system is likely to work and the maximum changing rate of frequencies. A value of $\alpha^{-1} = 100ms$ was used to account for an empirical maximum changing rate of 100 Hz/s.

Explicitly, the approximated, instantaneous, firing-rate at any time t is given by:

$$\tilde{r}(t) = \sum_{i=1}^{50} \left[\alpha^2 (t - t_i) \exp\left(-\alpha(t - t_i)\right) \right]_+$$

The values $t_i, i = 1, ..., 50$ are the previous 50 spike times. In the beginning of a simulation, when 50 previous spike times are unavailable, the estimated firing-rate is calculated assuming spike times at $-\infty$.

The quality of the firing-rate estimation algorithm can be accessed applying the kernel convolution to a spike train produced by a non-homogeneous Poisson process (Ross, 1996) with known time dependent arrival time. Figure 4.8 below shows the results using both $N = \infty$ (A) and $N = 50$ (B).
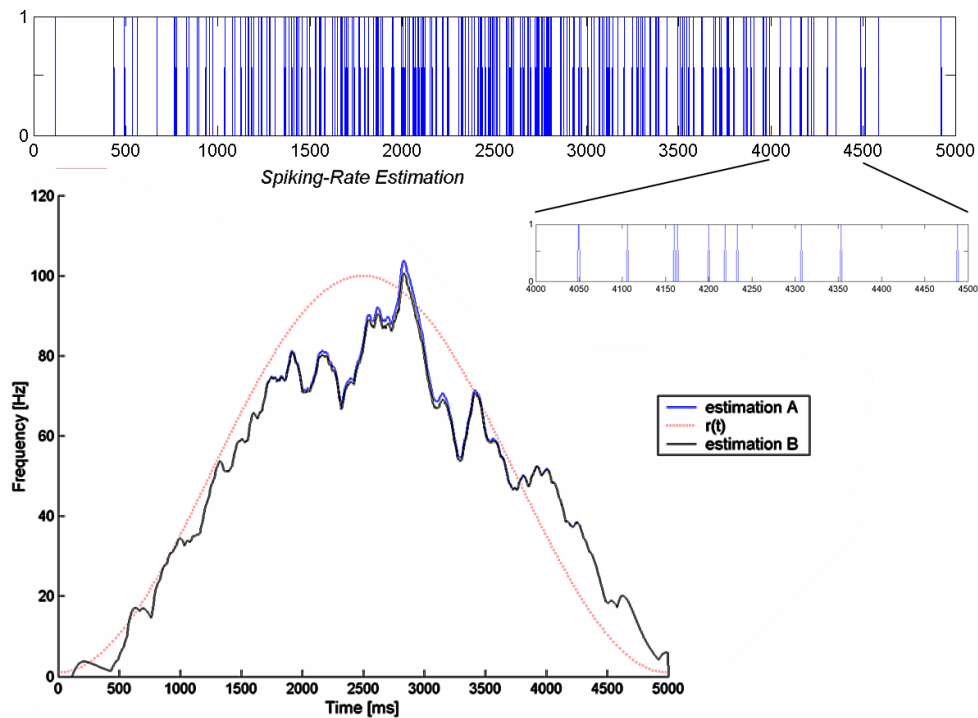
Figure 4.8: Upper graph: spikes produced by a non-homogeneous Poisson process. Lower graph: original firing rate used to produce the spike train; estimated firing-rate using all previous spike times (A), or up to 50 spike times (B), for the convolution operation.

The delay in the frequency estimator (due to the use of an asymmetrical kernel) is of the order of $\alpha^{-1} = 100ms$. This is a point that offers little concern. The frequency

dependent processes at the neuron that receives this spike train may be ruled by first order kinetic equations, to which the α-function profile is an approximate solution. In fact, more than the value that minimises the firing-rate estimation errors, α^{-1} should be representative of the time constants of the frequency-dependent internal processes.

4.2.4 Spikes

Each synapse unit keeps a list of all spikes present at its afferent axon. The information is in the form of remaining time for post-synaptic conductance activation. This counter includes both remaining time for synaptic terminal arrival and synaptic transmission delay ($\sim 1ms$).

Only in the case of considerable axonic delays and very high frequencies does the spike list contain more that 1-2 records (counters). The synaptic unit's spike list is updated at each time step, according to the numerical integration time dt.

A spike record is created each time the afferent neuron unit fires, and is removed, and converted to a post-synaptic current (PSC) record, once its (negative) counter reaches zero, i.e. the signal reached the post-synaptic neuron.

4.2.5 Post-synaptic currents

Just as with incoming spikes, each synapse unit keeps a list of all post-synaptic currents present at its associated dendrite fibre. The synaptic unit list is formed by the elapsed times since the PSC's initiation. These (positive) counters are updated at each SE time step, according to the numerical integration time dt.

As mentioned before, each PSC counter record is created and set to zero each time a spike record is eliminated (denoting spike arrival). Removal occurs when the current injected at the soma unit falls below a predefined, negligible, value.

The currents are calculated assuming a dual exponential model for the post-synaptic conductance response $S(t)$. See equation 3.4.

The voltage amplitude required to obtain the PSC from the conductance is calculated as the difference between the synapse's reversal potential and the attenuated soma's voltage (accordingly to the fibre length).

The total current injected by a specific synaptic unit into the soma is calculated as the convolution of the post-synaptic response function $\sum \delta(t - t_{PSC})$, with the synapse's

PSC profile[1].

4.3   Final remarks and critical summary

The approach taken to implement the dendritic trees, where they are temporally and spatially collapsed into a point, obviously removes all eventual local non-linear interactions between post-synaptic currents (Mel, 1994). There is still an indirect interaction, through the soma unit since the PSC generated by each synapse is dependent on the membrane driving force. The integration of synaptic signals is linear.

This is indeed a big compromise between a greater complexity in the model and computation time in the SE. The consideration of dendritic geometries and post-synaptic currents interactions requires an enormous amount of computations. The theory for the hippocampus that will be discussed later, does not require computations happening at the dendritic level in order to recreate the functional behaviour seen in episodic and spatial memory. Controlled computations happening at dendritic branches require a much larger number of synapses[2] that would even further load the SE. In some simulations, the dendritic fibre length is neglected: dendritic democracy, active properties and back-propagated action potentials, make the steady-state attenuation very small. Frequency dependent attenuation is nevertheless considered due to its implications on stability control.

A final remark goes to the efficiency of the simulation environment. Using a PC with a AMD64 3.5GHz and 1024MB of memory, an experiment of 60s (with $dt = 0.1ms$) in a network with 100000 synapses can be simulated in around 3.5 hours.

4.3.1   Snapshots of the simulation environment

Throughout this thesis snapshots of the simulation environment's graphical output are presented. A sample is presented in figure 4.9.

This simulation environment's graphical output is a schematic representation of the network being simulated. Each neuron unit is shown as a filled circle and all units

---

[1]In the case where the synaptic strength changes as a consequence of short and long-term plasticity, the post-synaptic response function is weighted accordingly

[2]For specific configurations of synapses, capable of useful/meaningful computations, to be formed through a continuous adaptation and selection process, a huge population of synapses have to be used.

belonging to the same population (or field) are aggregated in rectangular matrices. The voltage on each neuron unit is colour-coded:

- bright blue: $-90mV$; darker blues toward V_{rest}.

- black: V_{rest}.

- bright red: just below V_{thresh}; darker reds toward V_{rest}.

- yellow: all voltages above V_{thresh}.

All fields are identified by their initials (see Definitions page). The subscript p is used for identifying the principal population, while i stands for inhibitory interneurons. The activity level is also shown beneath the representation of each population with the initials "a.l.".
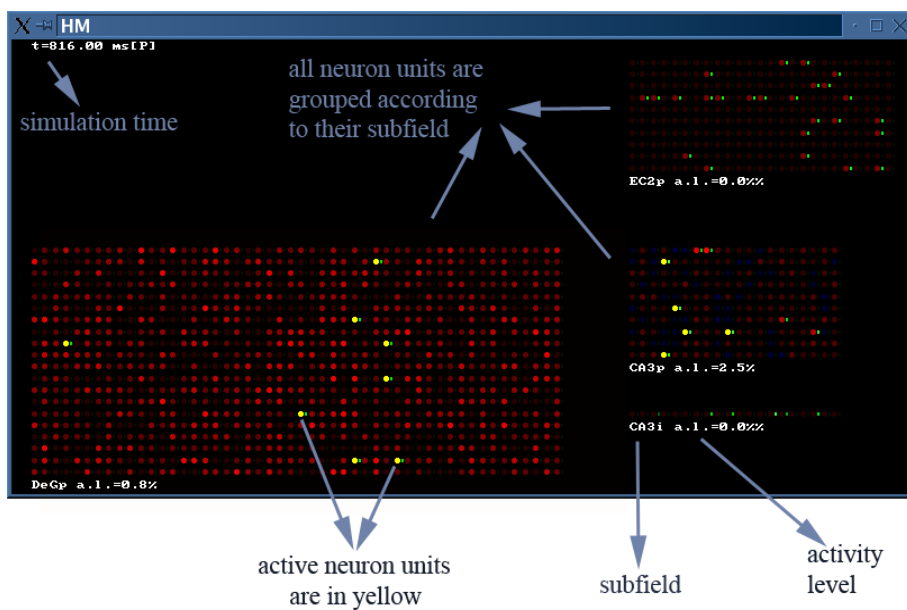
Figure 4.9: Snapshot of the simulation environment's graphical output. In this case there are 250 entorhinal cortex II principal neurons (EC2p), 1000 dentate gyrus principal neurons (DeGp), 200 CA3 principal neurons (CA3p) and 20 CA3 inhibitory interneurons (CA3i)

Chapter 5

Temporal Decoding and Dynamical Synapses

The notion of activity patterns was briefly introduced in chapter 1 (section 1.3). According with this form of coding, the information regarding a certain feature is complete at each instant t, and is independent of previous times. In other words, the activity patterns are purely spatial: in terms of transmission of information, coding in the temporal domain is non-existent.

Nevertheless the hypothesis of information being conveyed in the temporal domain should not be disregarded. Part of the hippocampus, namely the dentate gyrus, receives sensorial information conveyed in the temporal domain (O'Keefe, 1976; McNaughton *et al.*, 1983; Eichenbaum *et al.*, 1989) and it is important to analyse which processes are available at the hippocampal fields to access and use this type of information. There is controversy on the effective use of these temporal correlations for relevant computations at the hippocampal fields.

It will be argued that, at least at some stages of the hippocampal processing, analysis on the temporal domain is performed. At these stages, the decoding of temporal correlations between input patterns of activity is required for the correct functional behaviour of the hippocampus. This chapter provides the information required for the introduction of dynamical synapses into the hippocampal model.

For a cellular mechanism to be able to extract information from the temporal domain it needs to have a form of memory, i.e. it needs to be a system in which the current state depends not only on the inputs but also on the previous state(s). Several
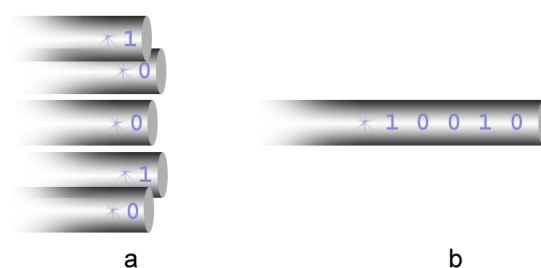
Figure 5.1: In spatial coding, the set of states of multiple channels at each time *t* codes the information (*a*). In temporal coding, a single channel transfer the information as a sequence of states (*b*). The two methods combined are called spatio-temporal coding.

mechanisms satisfy this condition and form reasonable candidates for decoding temporal information. Examples include the panoply of neuronal ionic currents found in hippocampal cells, such as I_{AHP}, I_A and I_M.

This chapter focuses on one such history dependent mechanism: *dynamical synapses*. Dynamical synapses exhibit short-term plasticity expressed as *facilitation* and *depression*, changes which are dependent on the neurotransmitter release process. Instead of being static at short-time scales, the post-synaptic peak currents in dynamical synapses can vary dynamically according to the synapses' activity history (see figure 5.2).

These synapses are widely distributed throughout several types of neurons in the hippocampal formation (Liaw and Berger, 1996; Buonomano, 1997) and offer experimentally confirmed mechanisms to condition the post-synaptic currents to the temporal properties of the input spike trains.

Organisation of the Chapter

In the first part of this chapter, the unreliability of synaptic signal transmission is discussed. Dynamical synapses are introduced and the properties of *facilitation* and *depression* behind the short-term dynamics will be described briefly in section 5.2, together with their biophysical basis.

In section 5.3 some dynamical synapse models are reviewed. A derived probabilistic model for synapses with a single release site is introduced and compared in section 5.4.

Section 5.5 deals with the following topics:

- Synaptic Switch - the synaptic reliability is tuned to a specific value of a spike
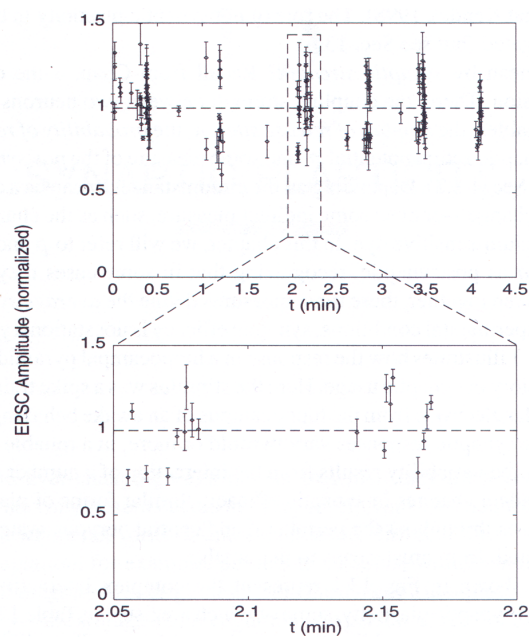
Figure 5.2: The upper graph shows the normalised excitatory post-synaptic currents (EPSCs) recorded from CA1 pyramidal neurons in a hippocampal slice in response to stimulation of the Schaffer collateral input. The stimulus is a spike train recorded *in vivo* from the hippocampus of an awake behaving rat. The EPSCs mean and standard deviation after 4 repetitions of the same stimulus are shown. The normalised strength of the EPSC varies in a deterministic manner depending on the prior usage of the synapse. For a constant synaptic efficiency, the normalised amplitudes should all fall in the dashed line. The bottom graph shows a blown-up version of part of the top graph. Reproduced from Koch (1999). ©The MIT Press.

train's temporal property, for example its frequency;

• Temporal to Spatial Conversion - a population of neurons with dynamical synapses all sharing the same input convert, in a small time window, the temporal code into spatial code;

The chapter concludes in section 5.6 with some final remarks on the role of dynamical synapses in the hippocampus.
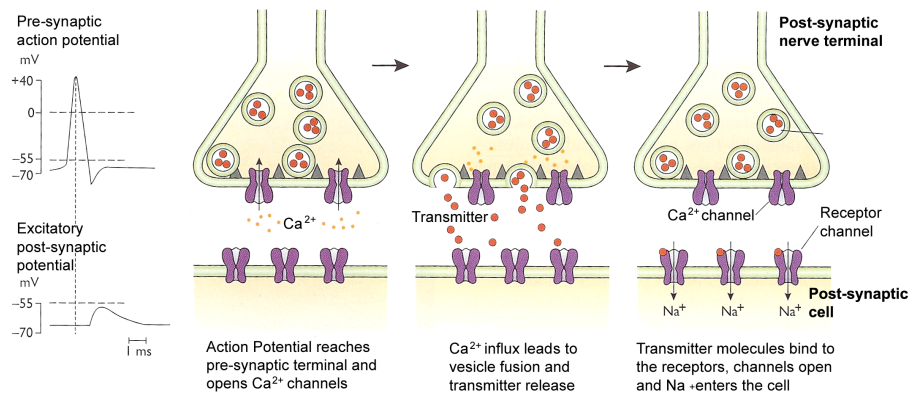
Figure 5.3: Diagram of the processes behind transmission at a chemical synapse, showing the level of complexity. Adapted Squire and Kandel (2002). ©W. H. Freeman and Company

5.1 Synaptic transmission

The synapse is a complex system, rich in subprocesses that mediate communication between pre and post-synaptic neurons (Shepherd, 1994). When an action potential, travelling though the axon, reaches the synapse and depolarises it, voltage dependent calcium channels open, leading to an increase of the intracellular concentration of Ca^{2+} ions. In turn, this calcium increase initiates a sequence of events that *may* culminate in the release of one neurotransmitter vesicle per release site. These neurotransmitters will then excite, or inhibit, the post-synaptic neuron, depending on the nature of the synapse (see figure 5.3).

The release of vesicles, where each vesicle is considered a *quantum*, is a stochastic process. A very simple but successful model for the time-averaged post-synaptic response size *PSR* is given by (Nicholls *et al*, 1992):

$$PSR = n\,p\,q \qquad (5.1)$$

In this expression n is the number of release sites, p the probability of vesicle release for each independent site, and q a measure of post-synaptic response to a single vesicle. While n and p are quantities dependent on pre-synaptic mechanisms, the post-synaptic response q is mostly dependent in post-synaptic factors such as the number of post-synaptic receptors. The post-synaptic response takes the units of the q variable,

which can be a voltage, a current or, as used throughout this thesis, a peak conductance S_{peak}. The peak conductance is the maximum conductance that the post-synaptic membrane reaches as a consequence of synaptic activation.

Although q is associated with a fixed unitary quantity, a *quantum*, it is nevertheless subject to small variations mainly due to Gaussian fluctuations on the amount of neurotransmitter molecules present in the vesicle (Nicholls *et al*, 1992).

The probability of k quanta being released in a synapse with n release sites, each with release probability p, follows a binomial distribution:

$$p(n,k) = \binom{n}{k} p^k (1-p)^{(n-k)}$$

It can be seen that the successful transmission probability is:

$$p(k > 0) = 1 - (1-p)^n$$

A typical feature of the hippocampal synapses is that each bouton contains only one or a few active release zones (Dobrunz and Stevens, 1997; Jonhston and Amaral, 1998). Hippocampal mossy fibre boutons are the exception (Chicurel and Harris, 1992) that will be discussed later.

As a consequence, the arrival of an action potential to a synapse often leads to the release of a single vesicle or simply to release failure (Korn and Faber, 1993). That is, the *one or none* assumption in the number of released quanta applies to most hippocampal synapses. Since the fluctuations in the post-synaptic response size are inversely proportional to $\sqrt{n}$, the level of variability is maximal in these single release site synapses.

In dynamical synapses, the probability of vesicle release p depends directly on the history of prior use of the synapse. This makes dynamical synapses an excellent mechanism to decode temporal information. Details of the key factors behind this dependency will now be given.

5.2 Dynamical synapses - facilitation and depression

The transmission probability p depends on several factors. One such factor is the calcium concentration - $[Ca^{2+}]$: the arrival of an action potential makes the synaptic internal calcium concentration increase by hundreds of micromolars (Lando and Zucker, 1994). This rapid increase triggers release of neurotransmitters into the synaptic cleft.

It has been observed in many synapses, including Schaffer collateral synapses, that the close arrival of input spikes leads to a progressive increase in the release probability (Dittman *et al*, 2000). The rise in the post-synaptic response can be of one order of magnitude. Depending on the number of stimulus and its frequency, this form of potentiation takes the designation *paired-pulse facilitation* (PPF), *augmentation* or *tetanic potentiation*.

Consecutive arrival of action potentials leads to an increase in the residual calcium concentration on the order of a 1 micromolar (Delaney and Tank, 1994). This goes against the empirical view of facilitation simply being the result of the additive effect of residual calcium together with the transient increase generated by the arrival of an action potential. The latter is two orders of magnitude bigger than the former.

Although the degree of complexity of this form of potentiation is obviously higher, models for the dynamics of facilitation based simply on additive effects have been applied with success.

The fact that the vesicles pool is finite, together with the fact that the full vesicle regeneration process is not instantaneous, leads to another property of synapses called *depression* (Tsodyks and Markram, 1997). Depression is characterised by a strong decrease in the release probability due to the depletion of available vesicles.
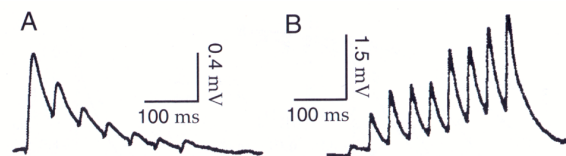
Figure 5.4: Experimental evidence of average decrease due to depression (**A**) and average increased in post-synaptic responses due to facilitation (**B**). (Taken from Markram and Tsodyks, 1996 and from Markram *et al*, 1998.

Depression is typically present in synapses between pyramidal neurons. Facilitation is more widely present and possibly more relevant in synapses between pyramidal neurons and interneurons (Thomson and Deuchars, 1994; Markram *et al*, 1998). In the hippocampus, Schaffer collaterals typically respond to a train of inputs with facilitation followed by depression (Dobrunz and Stevens, 1997). Mossy fibres are an exception to pyramidal-pyramidal usual synapse types showing no depression at natural input spike

rates, and being extremely facilitating with increases in PSR of up to 40 times (Urban *et al*, 2001).

Table 5.1 summarises the short-term dynamics that, in my opinion are relevant for the functional behaviour of the hippocampus. Discussion about their functional implications will be presented in the chapters 7, 8 and 9.

Synapse	Depression	Facilitation	Proposed functional role
pyramidal-interneuron	weak	strong	activity level control
Schaffer collaterals	moderate	moderate	temporal correlations between spatial patterns
mossy fibres	absent	very strong	temporal correlations between spatial patterns

Table 5.1: Short-term dynamics with impact in the functional behaviour of the hippocampus.

5.3 Models for dynamical synapses

Slightly different approaches have been taken in the literature to create models for dynamical synapses. These models can be directed to explain the collective post-synaptic response caused by multiple synapses (Varela *et al*, 1997) or the post-synaptic response caused by a single synapse.

Most models are directly or indirectly related to the original phenomenological models used in Tsodyks and Markram (1997) and Abbott *et al* (1997).

In Tsodyks *et al* (1998), a kinetic model was created in which three synaptic states are considered: recovered, active and inactive. The synapse is supposed to have a finite amount of resources and each pre-synaptic spike activates a fraction of those resources. After activation, this fraction will then quickly inactivate with a time constant of a few milliseconds, and recover with a time constant of about one second (for review of synaptic vesicle pool dynamics see Rizzoli and William, 2005). A facilitation effect is incorporated by increasing the fraction of resources used by a certain amount, which is elevated by each spike arrival, and decays between spikes. The full model is expressed

by a system of four differential equations. Simplified, non-kinetical versions of this model, based on the same principles, have been successfully applied in different studies (Dittman *et al*, 2000).

The version used here is the one analysed in Fuhrmann *et al* (2002). The dynamics are expressed by two coupled differential equations with variables R, the amount of available resources ([0,1]), and U_{SE}, the amount of resources *used* upon action potential arrival ([0,1]). Three constants define the dynamics of the synapse: $U1$, amount of resources used when resources are at maximum ([0,1]); τ_{fac}, facilitation decay time (ms); and τ_{rec}, the resources recovery time constant.

$$\frac{dR}{dt} = \frac{1-R}{\tau_{rec}} - U_{SE} \cdot R \cdot \delta(t - t_{sp}) \qquad (5.2)$$

$$\frac{dU_{SE}}{dt} = -\frac{U_{SE}}{\tau_{fac}} + U1 \cdot (1 - U_{SE}) \cdot \delta(t - t_{sp}) \qquad (5.3)$$

The variable t_{sp} represents a spike arrival time. The dynamical synapse is fully modelled by setting its post-synaptic response $PSR = W R(t_{sp})$, where W is a constant representing the maximal PSR obtained if all the synaptic resources are released simultaneously.

In the case of stimulation with constant period T and 100% reliable synapses, the steady-state values of the variables (denoted with a *ss* subscript) are:

$$R_{ss}(T) = \frac{1 - \exp\left(-\frac{T}{\tau_{rec}}\right)}{1 - (1 - U_{ss}(T)) \exp\left(-\frac{T}{\tau_{fac}}\right)} \qquad (5.4)$$

$$U_{ss}(T) = \frac{U1}{1 - (1 - U1) \exp\left(-\frac{T}{\tau_{fac}}\right)} \qquad (5.5)$$

Typical experimental values for dynamical synapses' constants are presented in table 5.2. All values for the dynamical synapses used in the hippocampal computational model were taken from these experimental ranges.

This model was chosen to describe dynamical synapses in the hippocampal model, and in the study of mossy fibre boutons (see appendix) due to its good results shown in several studies where it was applied. The use of a well established model for dynamical synapses is preferable for the discussion and analysis of the hippocampal model.

Synapse Type	$U1$	$\tau_{fac}\,[ms]$	$\tau_{rec}\,[ms]$
depressing†			500-1,500
(pyramidal- pyramidal)			
facilitating‡	0.012-0.086	550-3,044	104-694
(pyramidal - interneuron)			

Table 5.2: Experimental values for dynamical synapses. † Markram (1997) and ‡ Markram *et al* (1998).

A model that, is my opinion, describes better the *one or none* type of synapses in the hippocampus will be introduced and discussed in the next section. This modified model is totally probabilistic. It has some similarities with the probabilistic model in Fuhrmann *et al* (2002).

5.4 Dynamical synapses probabilistic model

The models discussed above are best applied to synapses with multiple vesicle release sites or populations of synapses, where they reproduce the average response to an input stimuli. The fact that many synapses in the hippocampus contain at most a single active zone (Schikorski and Stevens, 1997) motivated the modifications toward a novel probabilistic model. Another motivation was the creation of a model that fits in the equation for the post-synaptic response $PSR = n\,p_{rel}\,q$. The subscript *rel* has been added to the probability variable p in equation 5.1 to emphasise that it represents a vesicle release probability.

This modified model is targeted to the Schaffer collateral synapses. The model by Tsodyks and Markram (1997) is still more appropriate in the case of mossy fibre boutons where the number of release sites can be up to 37 (Chicurel and Harris, 1992). It will be shown that different dynamics are achieved with these modifications.

The following model is entirely probabilistic. The probability for the release of a vesicle is a function of facilitation, identified by the subscript F, and of the amount of resources available at the synapse, R.

Defining it as a jointly distributed random variable, the probability of release p_{rel} can be written as:

$$p_{rel} = p(F,R) = p_F(F)\,p_R(R) \tag{5.6}$$

The two independent marginalised terms $p_F(F)$ and $p_R(R)$ are developed below.

Facilitation

The probability of release, subject only to facilitation, is assumed to be proportional to a *facilitation precursor* which increases at each spike arrival and decays between spikes. A difference exists in the dynamics of p_F and its homologous U_{SE}: the upper bound is removed from the differential equation and applied externally. This way the dynamics assume that the transient increase in the facilitation precursor (which may not be directly the calcium concentration) is constant. After a limiting level though, the effect is considered to be saturated, at which $p_F = 1$.

The dynamics for p_F are therefore expressed in terms of the facilitation precursor decrease time constant τ_F and the step increment Ic_F (dimensionless) that is either a predefined constant, or zero to keep $p_F \leq 1$:

$$\frac{dp_F}{dt} = -\frac{p_F}{\tau_F} + Ic_F\delta(t - t_{spike}) \tag{5.7}$$

When saturation is reached $p_F(t_{sp}) = 1$, otherwise the solution is:

$$p_F(t) = p_F(0)e^{-\frac{t}{\tau_F}} + Ic_F e^{-\frac{t}{\tau_F}} \sum_n e^{\frac{t_n}{\tau_F}} UnitStep(t - t_n) \tag{5.8}$$

Each value t_n is the time of arrival of the n^{th} spike and the function $UnitStep$ is defined as:

$$UnitStep(x) = \begin{cases} 0 & : & x < 0 \\ 1 & : & x \geq 1 \end{cases}$$

In the presence of a spike train with constant inter-spike interval T, the value p_F, immediately after each spike arrival, converges to:

$$p_F(nT) = \frac{Ic_F}{1 - e^{-\frac{T}{\tau_F}}} \tag{5.9}$$

with $n \in \mathcal{N}$ and $n \gg 1$.

Depression

The notion of resources is also used in the probabilistic model. The probability for release conditioned to depression suffers a constant decrement Ic_R (dimensionless) each time a vesicle is *effectively* released, and relaxes continuously to a maximal value with time constant τ_R, as a consequence of replenishment of resources.

In the previous models (Tsodyks and Markram, 1997; Fuhrmann *et al*, 2002), the decrement was proportional to the resources used, which assumed the existence of multiple release zones in the synapse as opposed to the single release zone assumption behind this probabilistic model.

The step increment Ic_R, applied at each release time t_{rel}, is either constant or zero to keep $p_R \geq 0$.

$$\frac{dp_R}{dt} = -\frac{1 - p_R}{\tau_R} + Ic_R \delta(t - t_{rel}) \tag{5.10}$$

As long as the resource pool is not depleted, $p_R > 0$ and the system is described by:

$$p_R(t) = (p_R(0) - 1) e^{-\frac{t}{\tau_R}} + e^{-\frac{t}{\tau_R}} \left(e^{\frac{t}{\tau_R}} + Ic_F \sum_k e^{\frac{t_k}{\tau_R}} UnitStep(t - t_k) \right) \tag{5.11}$$

where t_k are the *vesicle release* times.

This is the major difference when compared with the previous model which used, instead, the action potential arrival time. For these stochastic synapses, the probability of signal transmission failure is quite high. In terms of the dynamics, a failure leads to a faster recovery of the synapse while facilitation is unaffected. This point is not taken into account by the previous model.

If a vesicle is released with a period T, the value p_R immediately after each release converges to:

$$p_R(kT) = \frac{Ic_R}{1 - e^{\frac{T}{\tau_R}}} + 1 \tag{5.12}$$

where $k \in \mathcal{N}$ and $k \gg 1$. Please note that $Ic_R < 0$.

Model comparison

This model is purely probabilistic and targeted to synapses with a single release site. In addition, it differs from Tsodyks and Markram's (1997) seminal model in the

increment/decrement values which are now considered either constant (independent of the state variables) or zero (once saturation was reached).

However the most important difference lies in the use of effective release times t_{rel} instead of spike arrival times t_{sp} for the dynamics associated with depression. A low p_R leads to a decrease in p_{rel} that, in turn, leads to an increase in p_F due to release failures. This negative feedback forces the *PSR* to oscillate around a non-zero value, once p_F has stabilised.

An experiment consisting of sending high frequency, vesicle pool depleting, stimuli through Schaffer collaterals, would be useful for comparing the two models: Tsodyks and Markram's (1997) model leads to $PSR \approx 0$ as opposed to a higher value predicted by the probabilistic model (figure 5.5).

The numerical steps for simulating the probabilistic model are described below:

When $t = t_{spike}$:

- $p_F(t_{sp}) = Min[p_F(t_{sp} - \Delta t) - \frac{p_F(t_{sp} - \Delta t)}{\tau_F} + Ic_F, 1]$

- sample X, where $X \sim Uniform(0, 1)$

- $p_{rel}(t_{sp}) = p_F((t_{sp}) \, p_R((t_{sp} - \Delta t))$

- If $p_{rel}(t_{sp}) > X$ then: $p_R(t_{sp}) = Max[p_R(t) + \frac{1 - p_R}{\tau_R} + Ic_R, 0]$

- If $p_{rel}(t_{sp}) < X$ then: $p_R(t_{sp}) = p_R(t) + \frac{1 - p_R}{\tau_R}$

When $t \neq t_{spike}$, the dynamics are driven by the exponential decay and relaxation with time constants τ_F and τ_R, respectively.

The values used in the simulations for all the constants where: $Ic_F = 0.2$, $\tau_F = 1500ms$, $Ic_R = -0.125$, $\tau_R = 400ms$ (see Fuhrmann *et al*, 2002). Although the marginal probabilities are defined for every instant, the vesicle probability release is only meaningful at each spike arrival time.

While previous models have the depression dynamics deterministic and fully conditioned by spike arrivals, this model draws probable paths for p_R values (figure 5.6).

Most of the graphs showing synaptic short-term plasticity are produced by averaging the post-synaptic currents produced by a specific spike train over several trials. Figure 5.7 is produced by averaging 100 responses to a Poisson spike train and convolving the p_{rel} values with a dual exponential function describing the post-synaptic
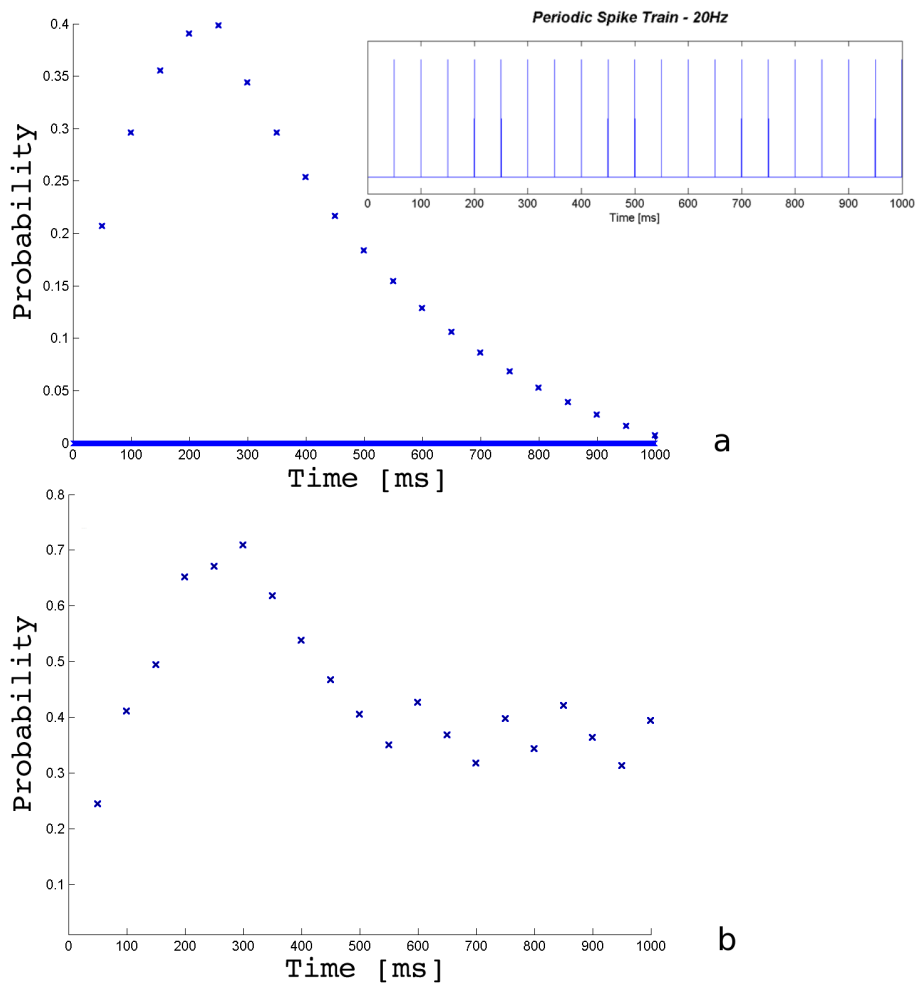
Figure 5.5: Release probability profile for a constant 20Hz spike train using Tsodyks and Markram's (1997) reference model *(a)* and using the presented novel probabilistic model *(b)*.
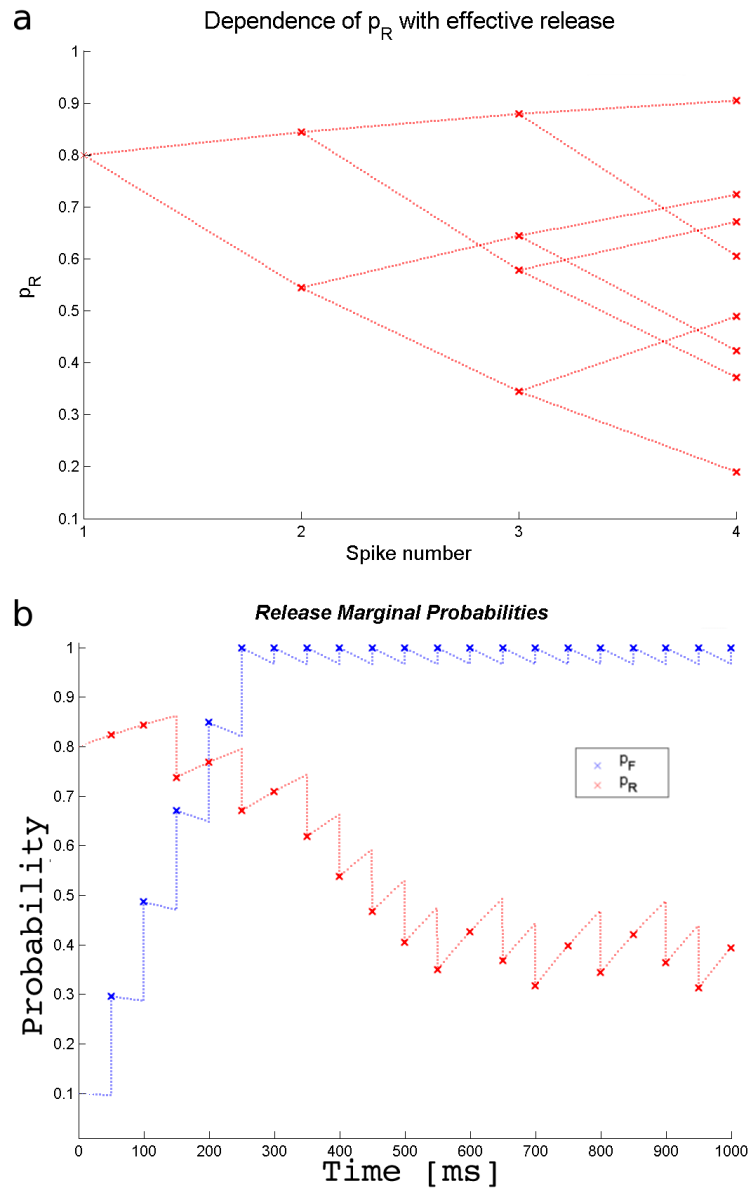
Figure 5.6: The value for p_R is updated accordingly to the effective release or not of a vesicle. The possible paths of p_R after arrival of four spikes with constant inter-spike interval are shown in *(a)*. Each upper trace in a bifurcation represents a release failure. Figure *(b)* shows the release probability associated at each spike arrival in a high frequency spike train with a possible course for p_R depending on effective release of vesicles. The facilitation factor saturates very quickly in this example.

current. Parameters for the dual exponential are: rise time constant = 3ms; decay time constant 5ms.

5.5 Synaptic switches and temporal code to spatial code conversions

Our knowledge on how neural assemblies *could* decode or process temporal information is quite small compared to the possible of strategies and mechanisms to process spatial code.

For small time windows ($\sim 100ms$), the times of consecutive spikes in many neural sub-systems are not entirely independent (see for example Richmond *et al*, 1990; Mechler *et al*, 1998). Therefore some effort has been made towards finding possible neuronal strategies to convert temporal code into spatial code that can then be extracted and processed using the available mechanisms. It as been shown theoretically that, as long the conversion from a temporal code to a spatial code satisfies some specific constraints, *any* real-time computation can be performed on time-varying inputs (Jaeger, 2001a, 2001b; Maass *et al*, 2002).

However it is still a subject of debate if, except for very specialised systems, neural assemblies *do* use complex temporal coding or temporal correlations: the fact that spike sequences contain temporal correlations as a result of sensorial transduction, or any other process, does not necessarily mean that this information is used subsequently.

Among the properties used to explain how this temporal to spatial conversion can be achieved are *paired pulse facilitation and depression* and *slow inhibitory post-synaptic potentials* (Buonomano *et al*, 1995 1997, 2000), resonance properties at the synaptic and cellular level (Izhikevich *et al*, 2003), post-inhibitory rebound (Hooper, 1998), and facilitation and depression (Markram *et al*, 1998; Aguiar and Willshaw, 2004).

It is again on dynamical synapses, though, that this section focuses its attention since none of the other processes were included in the theory for the functional behaviour of the hippocampus.

Two properties make dynamical synapses suitable for this temporal to spatial code conversion. First, the time scale of short-term plasticity is appropriate for decoding
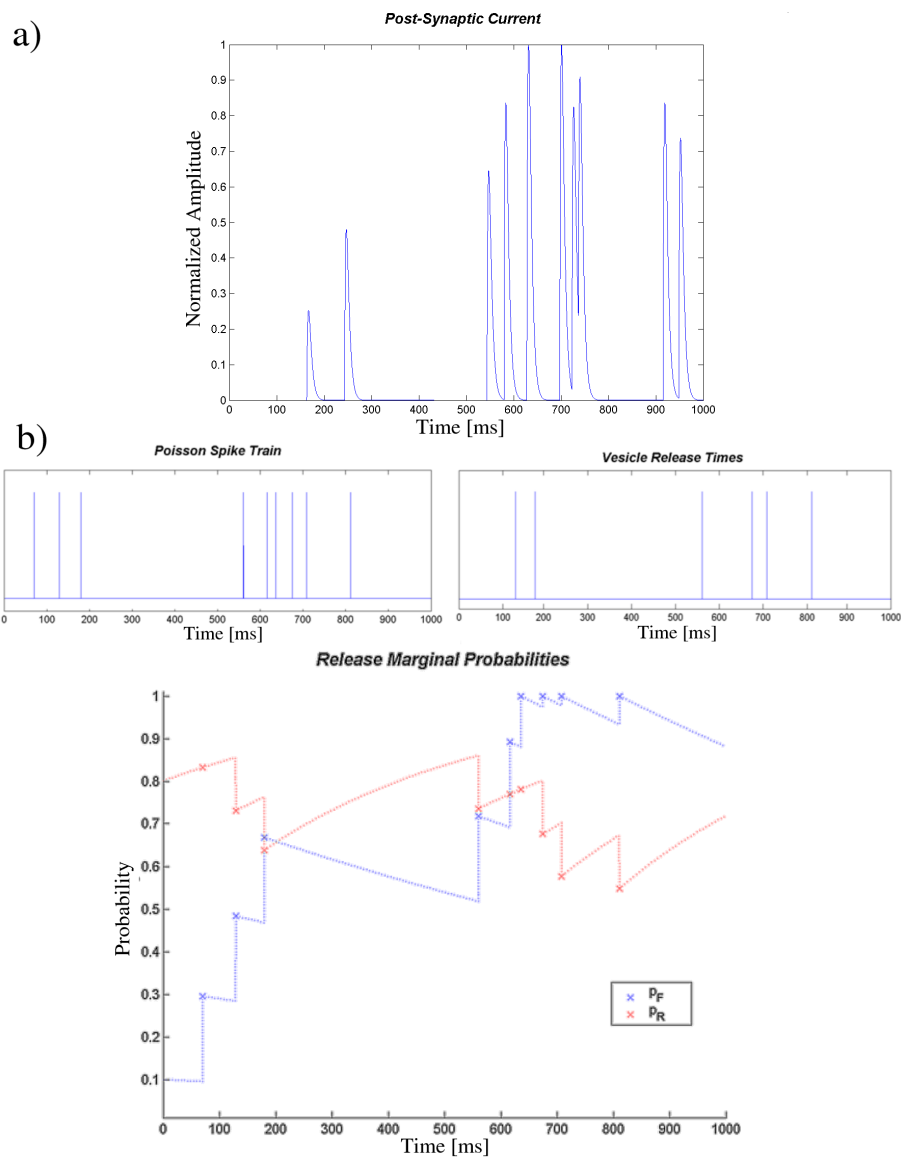
Figure 5.7: a) Typical averaged response of a dynamic synapse to a Poisson spike train. b) Example of a pair input spike train/output release times. The paths followed by the dynamic probabilities are shown in the lower graph.

eventual information conveyed in the temporal domain. Secondly, it has been shown that *PSRs* show a considerable heterogeneity that arrives from differences in facilitation/depression properties (Dobrunz and Stevens, 1997; 1999). Even synapses from the same axon can exhibit different short-term dynamics. These two properties have stimulated theoretical work on the simpler form of temporal to spatial conversion: synaptic "gates" or "switches".

Synaptic switches are tuned to a particular class of spike trains. Such synapses act as a switch, only releasing neurotransmitters if the spike train satisfies certain temporal constraints. In Natschläger and Maass, 2001, a technique is presented to compute the temporal pattern which maximises signal transfer for a synapse, given its dynamic parameters. The hippocampal synapses that gather the best characteristics to work as dynamical switches are the mossy fibre boutons.

In terms of generic temporal to spatial conversion, Maass and Markram (2002) have shown that temporal information can be efficiently decoded, in biologically plausible ways, using dynamic synapses and spiking units.

Although emphasis is being given to dynamical synapses, other mechanisms, as stated before, can be used for decoding temporal information. Figure 5.8 refers to a study by Buonomano (1995) in which paired pulse facilitation (PPF) and slow inhibitory post-synaptic potentials where used to convert temporal coding to spatial coding.

5.6 Discussion and final remarks

Two major questions motivated and directed this chapter. The first was concerned with the existence in the hippocampal fields, of mechanisms sensitive to temporal features. Through several reported experimental and theoretical studies we have seen that such mechanisms do exist. Among them, dynamical synapses combine several properties adequate to temporal decoding. The second, regarding the *effective* use of temporal information contained in reasonably sized time windows, was nevertheless left without a conclusive answer. Mechanisms for temporal decoding do exist, certain pathways on the hippocampus do have their synaptic efficacy dependent on previous use, but to which extent are these characteristics relevant, in functional terms, for the hippocampal computations remains unknown. I will nevertheless claim that dynamical synapses
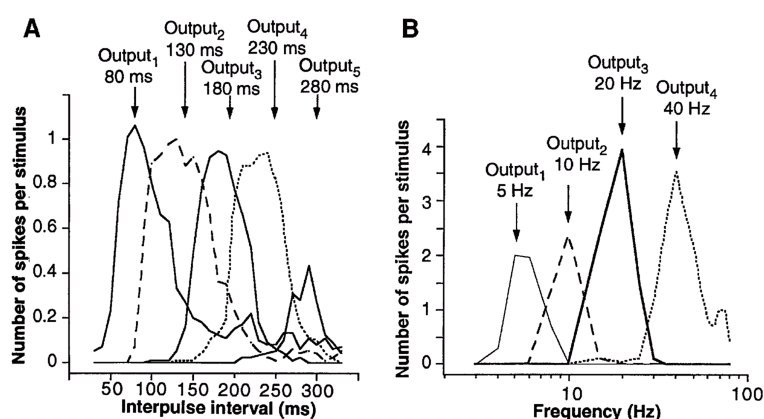
Figure 5.8: A network of integrate-and-fire elements with PPF and slow IPSPs was trained to discriminate among four different spike trains corresponding to specific short stimuli. Four output units were tuned to each of the four stimuli (*A*). Another network with the same properties was trained to discriminate between different inter-spike intervals (*B*). Adapted from Buonomano, 1995.

play an important role in the following three pathways: pyramidal-interneuron fibres, Schaffer collaterals and mossy fibres.

Pyramidal-Interneuron Fibres

The experimentally observed dynamical synapses between pyramidal and inhibitory interneurons lead to a better control in the activity levels. Facilitation on these connections creates an important non-linearity on the feedback control mechanism produced by inhibitory populations. This subject is analysed in detail in chapter 6, section 6.5.

A robust control of activity levels in each field is fundamental in the proposed computational model for the hippocampus.

Schaffer Collaterals

As stated before, several studies have reported the existence of well defined short-term dynamics on Schaffer collaterals. These fibres transfer information from the CA3 region to the CA1 region. Their role would be primarily in coding transitions between two consecutive spatial activity patterns at CA3.

With short-term plasticity, Schaffer collaterals can produce CA1 activation patterns that depend not only on the present but also on the preceding CA3 activation profiles. The dependence on previous CA3 activation patterns must, however, remain small,

that is, the activation profile produced in CA1 should be mostly driven by the present CA3 pattern of activity. Such ability creates the conditions for time-series learning, e.g. sequence learning and navigation.

Mossy Fibres

The importance of mossy fibres in creating the required CA3 sparse representations has already been brought out by many previous hippocampal models. This property results from the sparse connections between mossy fibre synapses and CA3 pyramidal neurons. I will nevertheless claim, in the theory presented in chapter 7, that the functional purpose of mossy fibres is significantly wider than that. As a result of the short-term dynamics (facilitation), mossy fibres are able to selectively propagate only patterns of activity that are statistically relevant across time. Mossy fibres actively chose which features should undergo storage.

One important topic mentioned before but worth further discussion, regards synaptic failures. In the hippocampal computational model, individual synapses are substituted by aggregates of coherent synapses. This means that each synaptic unit in the hippocampal computational model (see chapter 4) is the representation of a small population of synapses in a hippocampal neuron sharing the same dendritic branch and same pre-synaptic neuron. Synaptic failure probabilities are therefore accommodated into the parameters of the model's synaptic units, specifically into the conductance efficacies, which account for the average number of failures. The mechanisms of facilitation and depression, so important for the hippocampal computational model, are intimately related with the effects of synaptic failure. Therefore, synaptic failure is not just accounted for in the hippocampal model but rather is a fundamental component of its functional mechanisms.

Chapter 6

Activity Level Control in Spiking Networks

Control mechanisms that keep the levels of activity bounded are a key component in excitatory neural networks where signal amplification can occur. Signal amplification can result from the properties of specific network architectures or from changes in synaptic efficacy that lead to an increase in excitation levels. The hippocampus, as a learning system and containing recurrent networks, is highly susceptible to uncontrolled increases in excitation levels that could lead to catastrophic activity (e.g. temporal lobe epilepsy). It is important to emphasise that the correct functional behaviour of excitatory networks depends on a tight control of its levels of activity.

It is therefore important to understand how the control of activity levels is achieved in the hippocampus, and how to recreate the mechanisms behind this control in the hippocampal model.

In the hippocampus, several mechanisms are thought to be responsible for this control. In particular, the variety of interneurons that are systematically merged with the populations of principal neurons are thought to be dedicated to this control (Jonhston and Amaral, 1998).

In models based on firing-rate units, the control of (rate-based) activity levels can be established using one or more inhibitory units. On the other hand, creating efficient control mechanisms in spiking networks is quite challenging, especially if one wants to mirror biological constraints instead of applying rules like k-winners-take-all (O'Reilly and Munakata, 2000).

The goal of this chapter is to describe the control mechanisms that were implemented in the hippocampal computational model and discuss how they could work in the biological hippocampus. Novel techniques are introduced for the parameterisation of hippocampal connectivity properties that lead to constrained activity levels.

Organisation of the Chapter

Section 6.1 introduces the concept of activity level control and describes the stability problem which is the subject of this chapter. Section 6.2 briefly states which activity level control mechanisms are available in the hippocampus. It then moves to the description of the stability problem in the context of spiking neuron networks with dual exponential synaptic conductances. The notion of threshold peak conductance in discussed in section 6.3 which is then used in section 6.4 to describe how to obtain reliable inhibitory feedback control in spiking networks.

Attention is transferred from the single neuron level to the population level in section 6.5 where the mechanisms to produce precise levels of activity in post-synaptic populations are described. Simulation results from activity level control in stereotypical network architectures are discussed in section 6.6. The objective is to validate the claims presented in the previous sections. Some final remarks are presented in section 6.7.

6.1 Stability problem

Non-linear dynamics have to be used in order to model the important computational features of a neuron (Koch, 1999). Unfortunately, the non-linear properties that give neural networks their strong computational power also lead, typically, to analytically intractable dynamics. The number of biologically realistic models in which it is possible to study, analytically and in detail, particular aspects of their dynamics is limited.

One particular aspect of crucial importance in any model is the assessment of its stability. Given the properties of a neuron network model, it is important to be able to test the stability of the network and devise mechanisms for its control.

For spiking neuron models, the analysis of stability is quite difficult. An approach typically taken is to calculate the firing probability density functions as a function of time, according to the network connectivity properties. These differential equa-

tions, which describe the temporal evolution of a probability distribution (often called Fokker-Planck equations), offer a back door to the dynamics of the system that avoids the problems of the threshold non-linearity (Fourcaud and Brunel, 2002).

The analysis with firing-rate models is substantially simpler, and has motivated their use in many theoretical studies. Nevertheless, it is my strong belief that most of the information in the hippocampus is conveyed, and processed, in the form of (spatial) activity patterns (see section 1.3). Therefore, a study on stability of spiking networks, in the context of this novel hippocampal model, was required.

The methods presented in the following sections are devised for the neuron unit model used, the integrate-and-fire model (see chapter 3). The main problem addressed in this chapter can be stated in the following way:

> Consider a population B that receives inputs from a population A with a defined activity level $\mathcal{A}_A$. The objective is to relate the activity level of population B, denoted as $\mathcal{A}_B$, with network properties such as connectivity, number of units on each population and post-synaptic response strengths.

Note that population A can be an ensemble of several populations. Also, population A can be replaced by a sub-set of population B, thereby modelling the case of recurrent connections.

Since in the hippocampal model the connectivity and population sizes are constrained by biological data, the discussion is centred around the dependence between activity levels and synaptic strengths.

There are two ways of achieving stability in the excitation levels:

1. having precise synaptic strengths which set amplification close to zero (leads to unstable equilibrium);

2. having compensatory (inhibitory) feedback control;

Combining the two strategies leads to a robust control of activity levels. The following sections show how to implement both methods in network of integrate-and-fire neurons satisfying hippocampal constraints.

6.2 Activity level control mechanisms

It is worth repeating that, in this thesis, the quantity *activity level* is defined for a neural assembly, for each time t, as the fraction of units in which the voltage is above

threshold level. It is therefore an instantaneous measure as opposed to the rate-related measure used in O'Reilly and McClelland, 1994.

Several biological mechanisms exist to ensure the control of activity levels in all hippocampal fields. Examples are frequency adapting ion channels, dendritic filtering, refractory periods, synaptic vesicle depletion and, most importantly, inhibitory interneurons.

In the hippocampus, the average levels of activity in freely moving rats seem to be constrained to well defined values: $\sim 0.5Hz$ in DG, $\sim 6Hz$ in EC and $\sim 2.5Hz$ in CA3 (Barnes *et al*, 1990; Jung and McNaughton, 1993). However, the average firing rates can change considerably with changes in the behavioural state of the animal, or with changes in the environmental conditions (Wilson and McNaughton, 1993). A correct control of the levels of activity on the hippocampus is nevertheless crucial: failure to control them is often the reason behind several hippocampal pathological behaviours such as epilepsy. The hippocampus is the brain region with the lowest seizure threshold (Johnston and Amaral, 1998).

Synaptic strength, expressed here as peak conductance S_{peak}, is an important factor for post-synaptic activity levels. Even small deviations on the average synaptic strength can lead to catastrophic activity or complete shutdown of the network. However, assigning tight bounds to the average synaptic strength is not enough to achieve proper activity level control.

In the hippocampal model, the activity level control in each field is mainly achieved through inhibitory interneurons. Interneurons are not directly involved in any form of learning. My hypothesis is that their unique role is to keep excitatory units' activity levels bounded and promote competition between excitatory units. The only other mechanism included in the hippocampal model that contributes to activity control is dendritic filtering. This form of signal attenuation has a strong effect on high frequency signals (see figure 4.5). Depression on dynamical synapses have also a strong contribution in damping high frequency signals. However the dynamical synapses included in the hippocampal model were not targeted to this purpose.

The discussion focuses on how to produce good estimates for synaptic strengths that lead to specific post-synaptic activity levels, and how to use inhibitory populations as a negative feedback control mechanism. The equations required for this analysis are the integrate-and-fire model (3.2) and the dual exponential post-synaptic response

(3.4) that are rewritten below:

$$\tau_m \frac{dV}{dt} = V_{rest} - V + R_m I_{in}(t)$$

The dual exponential post-synaptic conductance is zero for $t < 0$, and for $t \geq 0$ it is defined as:

$$S(t) = S_{peak} \times norm \times \left(e^{-\frac{t}{\tau_{decay}}} - e^{-\frac{t}{\tau_{rise}}} \right)$$

Defining E_{syn} as the synapse's reversal potential, the soma input current associated with the post-synaptic conductance is then:

$$I_{syn}(t) = (V - E_{syn}) \times S(t)$$

The variables that affect synaptic strength estimates for the interaction between two populations A and B are:

- populations sizes, N_A and N_B;

- populations activity levels, $\mathcal{A}_A$ and $\mathcal{A}_B$;

- connectivities, C_{AB}^d, C_{BA}^d, C_{BB}^d and C_{AA}^d.

- dual exponential post-synaptic responses, τ_{rise} and τ_{decay};

The subscripts in C_{AB}^d indicate that A is the pre-synaptic population and B is the post-synaptic population. Unless stated, the connectivity values always refer to divergent connectivity (see 2.1.3).

In the following sections, the number of active inputs N_{PSR} that a neuron from population B receives from population A will be systematically required for calculations. The subscript *PSR* stands for Post-Synaptic Responses.

The properties of the number of active inputs N_{PSR} (also the number of Post-Synaptic Responses initiated) are discussed below.

The number of active inputs N_{PSR} that a neuron from population B receives from population A follows a binomial distribution with parameters C_{AB}^c (number of "trials") and $\mathcal{A}_A$ (probability of "success"). That is (see figure 6.1):

$$N_{PSR} \sim Binomial\left(C_{AB}^c, \mathcal{A}_A\right) \tag{6.1}$$

In addition, the convergent connectivity C_{AB}^c is given by:

$$C_{AB}^c \sim Binomial\left(N_A, \frac{C_{AB}^d}{N_B}\right) \tag{6.2}$$
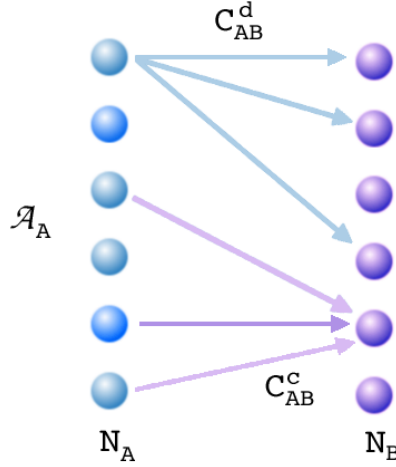
Figure 6.1: Conversion from divergent to convergent connectivity.

In order to calculate the probability of a unit from B receiving x active inputs, $P(N_{PSR} = x)$, the following expression is used:

$$P(N_{PSR} = x) = \sum_{C_{AB}^c} P(N_{PSR} = x | C_{AB}^c) \, P\left(C_{AB}^c | N_A, N_B, C_{AB}^d\right) \tag{6.3}$$

The two conditional probabilities inside the sum are given by expressions 6.1 and 6.2.

6.3 Threshold peak conductance

Consider a single unit u, at the resting potential, which receives N_{PSR} active inputs from homogeneous synapses defined by equal τ_{rise}, τ_{decay} and S_{peak}. The active inputs may be received simultaneously or spread through the unit's integration time ($\sim \tau_m$).

The threshold S_{peak}, the minimal peak conductance that leads to post-synaptic firing, can be calculated by solving numerically the integrate-and-fire equation. Since linearity is assumed for current summing, the total current $I_{in}(t)$ is obtained by adding

all N_{PSR} post-synaptic current profiles. Assuming that the post-synaptic responses are initiated at each synapse at times t_j, the total input current $I_{in}(t)$ is then:

$$I_{in}(t) = S_{peak} \times norm \times (V - E_{syn}) \times \sum_{j=1}^{N_{PSR}} \left(e^{-\frac{t_j}{\tau_{decay}}} - e^{-\frac{t_j}{\tau_{rise}}} \right) \qquad (6.4)$$

The linearity on current summation means that a new S_{peak} can be estimated for a different activity level (different N_{PSR}) without having to recalculate the numerical solution of the differential equation.

For default neuron parameters ($Rm = 20M\Omega$, $\tau_m = 20ms$) and post-synaptic dual exponential response with $\tau_{decay} = 5ms$ and $\tau_{rise} = 2ms$, the result for the threshold peak conductance (using equations 3.2 and 6.4) in the situation of simultaneous initiation of five *PSR* is $S_{peak} = 12.52nS$ (figure 6.2). In other words, a default neuron receiving 5 simultaneous, identical, synaptic inputs will only reach threshold if the peak conductance of each response is, at least, 12.52 nS.
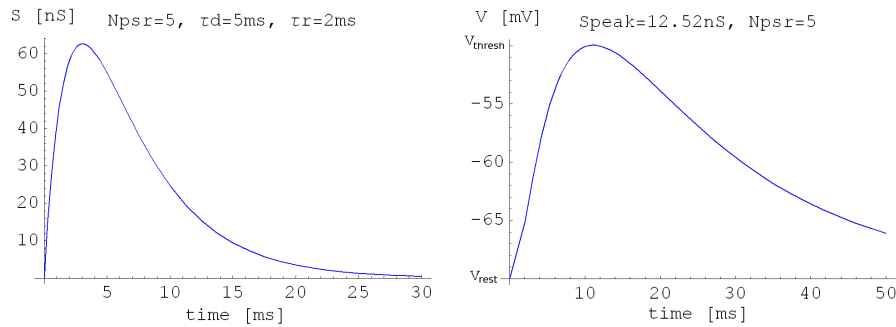
Figure 6.2: Numerical solution for I&F model unit receiving 5 simultaneous post-synaptic responses with $\tau_{rise} = 2ms$, $\tau_{decay} = 5ms$ and average $S_{peak} = 12.52nS$.

For under-threshold conditions, the solution for the voltage profile is a low-pass filter of the current profile.

Using the simulation environment with the default time step (0.1ms) and recreating the conditions used in the previous calculation, the estimation is $S_{peak} = 12.46nS$. This represents an error of about 0.5%. This confirmation of the precision of the SE is required to export confidently the theoretical results to the model.

Non-simultaneous initiation of the PSRs, within the membrane's integration time, affects only slightly the threshold S_{peak}. For the same parameters, the estimate for the minimum peak conductance under jittering condition (all five post-synaptic responses

uniformly spread over $\tau_m/2ms$) is $S_{peak} = 13.05$. This represents a variation of less than 5% (figure 6.3).
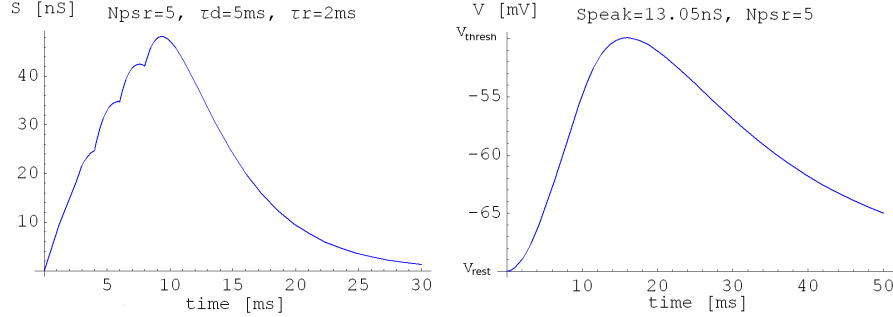
Figure 6.3: Numerical solution for I&F model unit receiving 5 non-simultaneous post-synaptic responses with $\tau_{rise} = 2ms$, $\tau_{decay} = 5ms$ and $S_{peak} = 13.05nS$. The initiation of the post-synaptic responses were spread over $\tau_m/2$ ms.

Table 6.1 contains some threshold peak conductances for different configurations of the post-synaptic conductance profile. The number of post-synaptic responses is set to 1 ($PSR = 1$) in order to give the overall threshold peak conductances, or S_{thresh}.

$\tau_{rise}[ms]$	$\tau_{decay}[ms]$	$S_{thresh}[nS]$
2	5	62.6
3	5	55.4
0.5	30	38.7

Table 6.1: Threshold peak conductances.

The information contained in table 6.1 is very important. It tells what is the minimal sum of all post-synaptic responses that leads the default neuron to fire. For example, a neuron receiving synapses with parameters $\tau_{rise} = 2ms$ and $\tau_{decay} = 5ms$, will fire if the sum of the peak conductances of the activated synapses is at least $62.6nS$.

By simply using the threshold S_{peak} values calculated using this method, it is already possible to parameterise a reliable inhibitory feedback control to keep activity levels bounded. This will be the topic of discussion in the next section.

6.4 Reliable inhibitory feedback control

In this section, the general technique used in the hippocampal model to parameterise the non-plastic inhibitory synapses is presented. As stated before, the unique role of the inhibitory populations associated with each field is assumed to be of bounding the activity levels of excitatory populations.

A good example of the importance of reliable inhibitory feedback control is given by the CA3 subfield. The principal population of the CA3 subfield receives excitatory input from the perforant path, mossy fibres and recurrent collaterals. In order to produce proper computations and provide the post-synaptic populations with the correct profile of activations, the net excitation level in CA3 must be constrained. This means that the CA3 system must be resilient to fluctuations in the input excitation levels and to the result of possible feedback amplification produced by the recurrent collaterals. Even transient excesses of excitation must be avoided in order to prevented propagation of erroneous signals to post-synaptic populations. The inhibitory feedback control mechanisms acting in the CA3 principal population are therefore extremely important. These mechanisms, provided by the inhibitory populations, secure the constrained levels of excitation required for the proper functional behaviour of this system. Failure or absence of the inhibitory feedback control in CA3 leads to abnormal storage and recall of information.

After this introductory example, let us consider now the general case of one excitatory population (E) and its associated inhibitory population (I). The following properties are defined:

- population sizes, N_E and N_I;

- population activity levels, $\mathcal{A}_E$ and $\mathcal{A}_I$;

- connectivities, C_{EI}^d, C_{IE}^d;

- dual exponential post-synaptic responses, τ_{rise} and τ_{decay}.

The population E receives excitatory inputs from other populations and/or from recurrent connections. Although these inputs may vary significantly, the activity levels at population E must remain constrained.
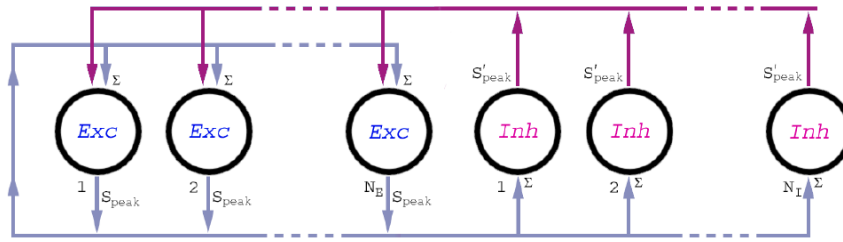
Figure 6.4: Excitatory (blue) and inhibitory (red) populations.

In the hippocampal model, the strategy used to keep the activity level at excitatory populations bounded (and centred around a pre-assigned value) makes use of the non-linearity of integrate-and-fire units: the average peak conductances of excitatory synapses into interneurons are set so that interneurons only fire when the number of active excitatory units is statistically above the pre-assigned activity level for the excitatory population. The higher the deviation from the pre-assigned activity level is, the larger will be the number of interneurons recruited to fire.

According to the connectivity properties, it is possible to infer the average number of excitatory connections each inhibitory neuron unit receives. In order to do that the (probabilistic) convergent connection measure is required.

Each unit from population E makes C_{EI}^d connections with random target units. A uniform distribution for the choice of targets is assumed (biologically, physical constraints and topographic connections typically generate complex distributions). Using the variables defined above, the average number of active inputs $< N_{PSR} >$ can be calculated using expression 6.3.

The strength of the excitatory connections to the inhibitory neuron units are then calculated as the threshold S_{peak} for $< N_{PSR} > +1$ active inputs. This means that inhibitory units will only fire when the activity level of the excitatory population is statistically above the predefined level. The higher the pre-synaptic activity level is, the higher will be the probability for a inhibitory unit to fire, which, in population terms, means that more inhibitory units are recruited.

Figure 6.5 shows the recruitment of inhibitory units according to the activity level of the controlled excitatory population. Different levels of activity are imposed, in sequence, to the excitatory population and the response of the inhibitory population

is measured. Feedback connections have been removed here to create a clear figure. The activity level of the excitatory population is normalised to the preassigned activity level ($\mathcal{A}_E = 5\%$). Although not relevant, the inhibitory population's activity level is also normalised to $\mathcal{A}_I = 100\%$.
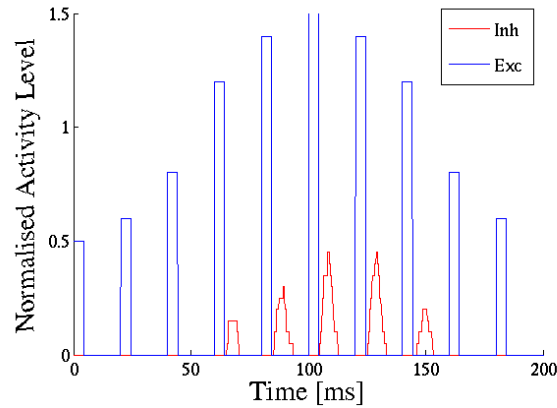
Figure 6.5: Recruitment of inhibitory units in activity control.

It can be seen that only when $\mathcal{A}_E > 1.0\%$ interneurons start firing. The amount of interneurons recruited increases with the deviation size. The connections C_{IE} from the inhibitory population I to population E provide the feedback control mechanism. The amount of inhibition provided should be enough to bring $\mathcal{A}_E$ just below the predefined value. It is important that the inhibition is graded according to the amplitude of the deviation in $\mathcal{A}_E$. If the inhibition is too strong for small deviations, the activity level at the controlled population can be completely shut down.

In practice, using only the non-linearity in the firing of the inhibitory units is not enough to control the activity levels in some architectures, or under abnormal excitation levels. It should to be noticed that, once above threshold, the control mechanism is based on a linear balance of excitatory and inhibitory inputs. It is well known that linear mechanisms do not provide a robust control. While for short deviations the inhibitory feedback perfectly compensates for the increase in excitation the same does not happen with *abnormal* initial levels of excitation.

A solution to increase the strength of the inhibitory feedback is to increase the population of (inhibitory) interneurons. Nevertheless, increasing the inhibitory population provides only a slight increase in the amplitude of the deviations that the network can

cope with.

This problem is particularly strong in recurrent excitatory networks with high con-
nectivity *and* very narrow distribution of peak conductances, in which the amplification
mechanism leads to catastrophic activity very rapidly. If the connectivity is very high
but the synaptic peak conductances are widely distributed around the threshold peak
conductance (very high variance), then the amplification process can more easily be
balanced.

The solution to cope with abnormal levels of excitation is to take advantage of
another non-linearity available in the biological hippocampal system. Two candidate
mechanisms are: 1) long lasting inhibitory post-synaptic conductances and 2) facilita-
tion dynamics that exist in pyramidal-inhibitory connections.

Facilitating synapses are able to build up their inhibitory effect if the excitatory
population increases its activity. This way, if the activity level on the excitatory pop-
ulation systematically increases, not only the number of interneurons recruited for the
control increases but also their negative feedback strength.

The non-linearity provided by long-lasting inhibitory post-synaptic conductances
works in a similar way. In the situation of activity above the pre-defined level, suc-
cessive waves of inhibition lead to an increase in inhibitory currents, overcoming am-
plification: when a second wave of inhibition is initiated, it builds up with the effects
of the previous inhibition. The presence of slowly inactivating $GABA_B$ channels in
hippocampal interneurons offers a strong support for this hypothesis. Since long last-
ing inhibitory post-synaptic conductances are enough to make the control mechanism
robust to abnormal increases in activity levels, the facilitating dynamics were not intro-
duced in the simulations of the hippocampal model. Introduction of facilitating dynam-
ics in interneurons would greatly reduce the performance of the simulations without
contributing for the functional behaviour of the system. Long lasting inhibitory post-
synaptic conductances are implemented just by changing the properties of the dual
exponential conductance profile.

6.5   Activity levels dependence on peak conductances

So far the analysis was conducted to check, at the individual level, if a neuron unit will
fire according to the input network activity and connectivity properties. Now a step

further is taken and the analysis is conducted at the population level, i.e. a probability is estimated for the firing of each unit in the population. At the population level, the firing probability of a unit and the network activity level value are equivalent[1]. This means that relating the unit's firing probability with the pre-synaptic activity level and connectivity properties, is solving the stability problem stated in section 6.1. That is, I will show how to produce a specific activity level at a post-synaptic population given the activity level of the pre-synaptic population and the connectivity properties.

Let us assume that excitation arrives in volleys, with each wave limited to a time window of the order of the membrane time constant. This corresponds to assuming that the inputs to this system are in the form of patterns of activity. This follows the assumption of this thesis that, at the hippocampus, information is carried as activity patterns.

Although this analysis can be applied in any architecture, for simplicity let us again consider a population B receiving excitatory inputs from a population A, and that all B units initially are at resting potential (this condition can be relaxed).

The total current that reaches the soma unit is a function of the following variables:

- population sizes, N_A and N_B;

- population A activity level, $\mathcal{A}_A$;

- connectivities, C^d_{AB} (for simplicity $C^d_{BA} = C^d_{BB} = C^d_{AA} = 0$).

- dual exponential post-synaptic responses, τ_{rise}, τ_{decay} and S_{peak};

In order to calculate the activity level at population B as a function of the synaptic peak conductance, $\mathcal{A}_B\left(S_{peak}\right)$, one has to define the total input current. A unit will reach voltage firing threshold when its inputs generate a threshold current. It is therefore important to analyse the statistical properties of the total input current that each post-synaptic neuron receives. Equation 6.3, which gives the probability of a unit from B receiving x active inputs, provides the means to calculate the probability distribution for total input current exciting each neuron.

For a given S_{peak}, and predefined τ_{rise} and τ_{decay}, it is possible to calculate numerically (using the integrate-and-fire equation and the dual exponential model for the

---

[1]In a frequentist approach, the probability of a unit firing is simply the ratio of above-threshold units over the size of the population. This is precisely the definition for the activity level used here.

post-synaptic conductance) what is the minimal number of post-synaptic signals psr_0 that leads to threshold voltage. For example, assuming $\tau_{rise} = 3ms$ and $\tau_{decay} = 5ms$ and consulting table 6.1, if $S_{peak} = 12nS$ then a neuron requires at least five active inputs to go beyond threshold voltage. In other words, $psr_0 = 5$.

It is the value of psr_0 that, in the distribution of the number of active inputs that each neuron receives, provides the estimate to the fraction of the post-synaptic population that receives enough excitation to reach threshold. The $\mathcal{A}_B$, or equivalently the firing probability of population B, is then given by:

$$
\begin{aligned}
\mathcal{A}_B\left(S_{peak}\right) &= 1 - P\left(N_{PSR} < psr_0\left(S_{peak}\right)\right) \\
&= 1 - \sum_{x=0}^{psr_0} P\left(N_{PSR} = x\right)
\end{aligned}
\tag{6.5}
$$

Figure 6.6 provides a visual idea of the method.
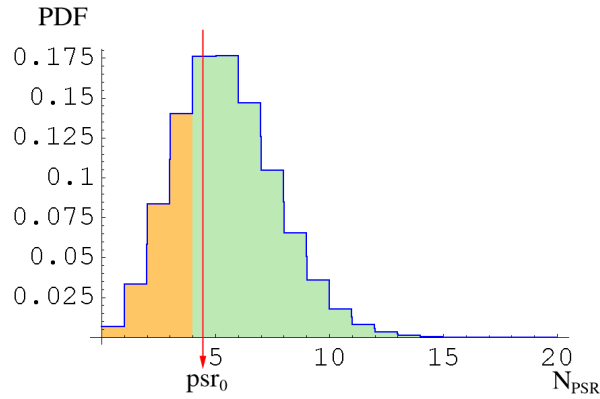
Figure 6.6: Probability density function (PDF) for N_{PSR}. The minimal number of post-synaptic signals that leads to firing is marked as psr_0. The area of the orange region, on the left, is numerically the same as the fraction of the population that received sub-threshold excitation. Notice that psr_0 value is not included in the orange region. The green area, on the right, is numerically the same as $\mathcal{A}_B$.

Using this method, one can calculate the activity level in a spiking network as a function of S_{peak}. Due to the discrete nature of the binomial distribution, the possible post-synaptic activity levels are quantised if S_{peak} is the same for all synapses (or takes only discrete values).

Graph 6.7 refers to a spiking network with default neuron unit properties, $N_A = 200$, $N_B = 100$, $\mathcal{A}_A = 0.05$, $C_{AB} = 50$, $\tau_{rise} = 3ms$ and $\tau_{decay} = 5ms$. The red line represents the expected values using the equations described above. Each point in the graph is the average of five simulations using the simulation environment. The standard deviation for each data point is presented as error bars.
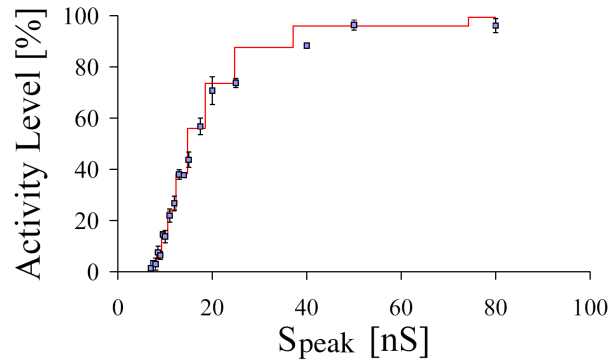
Figure 6.7: Activity level as a function of S_{peak}. Parameters: $N_A = 200$, $N_B = 100$, $\mathcal{A}_A = 0.05$, $C_{AB} = 50$, $\tau_{rise} = 3ms$ and $\tau_{decay} = 5ms$.

It is important to emphasise that in case population *A* would be heterogeneous (or even a sub-set of population *B*) this method would still be applicable. The main difference is that equation 6.4 had to accommodate the different post-synaptic responses respecting their relative contributions. In addition, probability distributions can be incorporated for the other dependent variables. This however leads to a much more complicated discussion that would not benefit this thesis. It is important to mention that in the case where the variables are normally distributed with small variances, this method still provides good results.

6.6 Activity level control in the hippocampal model

This section shows examples of how the activity level control mechanisms were effectively implemented in the computational model of the hippocampus.

The method presented in section 6.4 is used to parameterise the connections between excitatory and inhibitory populations. The results presented in section 6.5 provide the estimates for the connection strengths that create the correct activity levels at

post-synaptic field, given the pre-synaptic field's activity level.

In the computational model for the hippocampus, besides sampling the activity at the associated field, the inhibitory populations also sample the activity of fields that provide inputs to the controlled population. For example, interneurons in CA3 region not only receive excitation from CA3 pyramidal neurons but also from DG granule cells (see hippocampal diagram in figure 2.3).

This feature in the connectivity profile, in complete agreement with experimental data on hippocampal connectivity, provides an even more efficient control since strong deviations in the inputs can be accounted for before they affect the post-synaptic activity levels. In abnormal input excitation levels, inhibitory interneurons are highly recruited preventing even transient increases in the activity level of the controlled population.

It may also be important to clarify that the hippocampal model does not assume perfectly stable activity levels. On the contrary, the levels oscillate and many computations depend on these fluctuations. What is being emphasised is that abnormal increases in activity are neither propagated to subsequent fields, nor lead to catastrophic amplification (in case of existence of recurrent connections).

The simulations presented here refer to a simplified version of CA3 in which only specific pathways were considered. The simplifications are described in detail below.

A snapshot of the network is presented in figure 6.6 (refer to the 4.3.1 for colour codes).
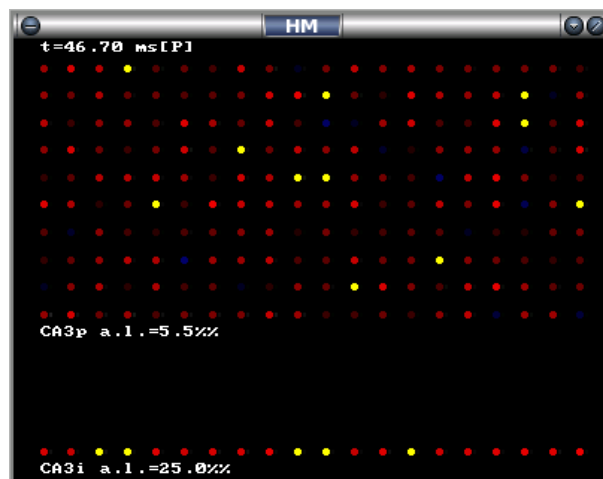
Figure 6.8: Snapshot of CA3 type network.

6.6.1 Generating the correct activity levels

In here, the connections from a hypothetical entorhinal cortex to this CA3 type field are parameterised. The following quantities are defined:

- $N_{EC} = 250$ and $N_{CA3} = 200$;

- $\mathcal{A}_{EC} = 0.1$;

- $C^d_{EC:CA3} = 43$;

- dual exponential post-synaptic responses, $\tau_{rise} = 2ms$ and $\tau_{decay} = 5ms$;

The objective is to set the synaptic peak conductance so that an activity level proximal to 5% (a predefined value) is induced in the CA3 field by the EC connections. By examining the probability distribution for the number of active inputs that each CA3 unit receives (equation 6.3), we calculate that only 4.6% of the CA3 population receives more than nine active inputs.

Numerically solving the integrate-and-fire equation for the parameters of the post-conductance profile, we get a threshold conductance of $62.6nS$ (see table 6.1). Due to linearity of input summation, we set $S_{peak} = 6.26nS$ in order to produce an activity level of 4.6% in the CA3 population. The synaptic strengths may be normally distributed around this value. Figure 6.9 presents the activity level produced in CA3 using $< S_{peak} > = 6.26nS$ with variance of $2nS$. The activity level at the hypothetical EC field is set to 10% at time 0. In order to show consistency, the CA3 activity level wave shown in blue is the average of the waves generated by 10 random connectivity profiles.

6.6.2 Controlling the levels of activity

In order to test the robustness of the activity level control, several simulations were performed. In these simulations, recurrent connections were added to the CA3 field and the activity level in entorhinal cortex was allowed to deviate from the predefined value of 10%.

A population of interneurons was associated with the CA3 excitatory population. The inhibitory connections were defined to keep the excitatory population's activity
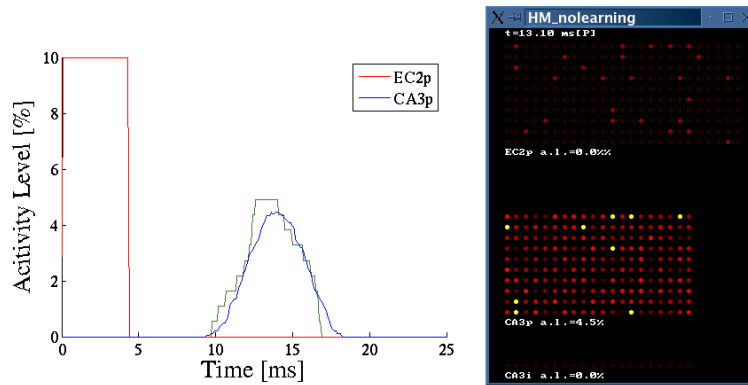
Figure 6.9: Producing predefined activity levels. The wave shown in green represents a single activity wave as opposed to the average in blue. It emphasises that the variability in the activity waves generated by different connectivity profiles is not big.

level at 5%, using the method described in section 6.4. The inhibitory post-synaptic conductance used $\tau_{rise} = 0.5ms$ and $\tau_{decay} = 30ms$ (this is just a possibility).

Two configurations were used:

In a first configuration, a sequence of patterns were induced in EC region (with a periodicity of 100ms) leading all to an activity level of 10%. The results of the control established by the interneurons (CA3i) are shown in figure 6.10. All activity levels are normalised to the preassigned values (with inhibitory $\mathcal{A}_{CA3i}$ normalised to 100%). Notice that the inhibitory control properties create a slow decay of the CA3's activity level. In the hippocampal model, where CA3 acts as an auto-association network, the 2-3 cycles of sustained activity are enough to induce the complete recall of a stored pattern (details in chapters 7 and 8). This means that a stored activity pattern can be recalled in less than 100 ms.

In the second configuration, a sequence of increasing activity levels were induced in EC region (with a periodicity of 100ms). The EC activity levels ranged from 10% to 20%. Values for $\mathcal{A}_{EC} < 10\%$ were not included since they are not able to raise $\mathcal{A}_{CA3}$ significantly. When the deviations in $\mathcal{A}_{EC}$ are too big, a significant number of interneurons start firing preventing the high excitation levels of being transfered to the CA3p field. This ensures that noise is not propagated.

A last comment goes to the situation in which the source, this case EC, continuously provides excitation. That is, instead of periodically raising its activity level, $\mathcal{A}_{EC}$
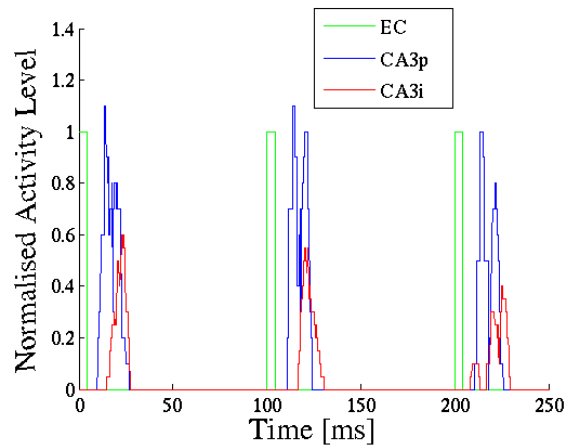
Figure 6.10: Inhibitory control with several exposures. Parameters: $N_{EC} = 250$, $N_{CA3p} = 200$ and $N_{CA3i} = 20$.

is kept at 10%. Inhibition increases rapidly leading to a complete shut down of the target population.

In the context of the hippocampal model in which information is considered to be transfered as patterns of activity, if the exposure rate would be too high, this easily leads to an overlap of the patterns. This overlap reduces the information content of each excitation volley and the consequence is inhibition completely shutting down the activity at the controlled network so that errors do not propagate. In the computational model for the hippocampus the rate of the excitation volleys corresponds to the theta cycle. That is, excitation and consequently levels of activity, raise with a periodicity of 80-100 ms (theta rhythm: 10-12Hz).

6.7 Final remarks

The method presented in this chapter is essential for the parametrisation of the connections in the hippocampal model which is discussed in the next chapter. This is a general method that can be applied in many other networks of spiking neurons.

As a final remark, the same way inhibitory interneurons were shown here to be able to bound from above activity levels, using the same methods one can show that excitatory interneurons can bound from below activity levels. In fact, I propose that this mechanism effectively happens in the dentate gyrus: on one hand, inhibitory interneu-
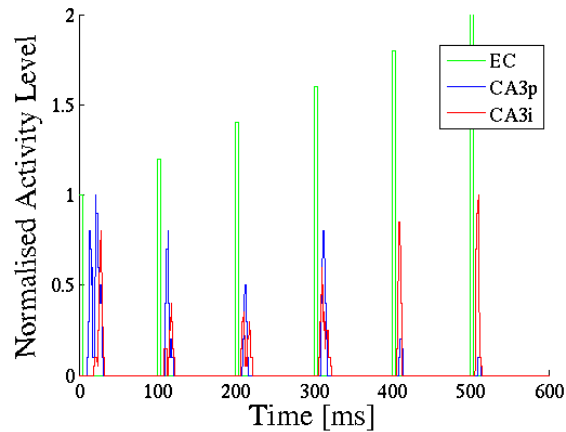
Figure 6.11: Inhibitory control with several increasing exposures. Parameters: $N_{EC} = 250$, $N_{CA3p} = 200$ and $N_{CA3i} = 20$.

rons (such as basket cells and GABAergic polymorphic cells) control the activity level from above; on the other hand, excitatory interneurons (mossy cells), merged with the principal (granule) cells in a recurrent feedback loop, control the activity level from below. The activity level in the dentate gyrus is very important for the process of storing information in CA3 recurrent collaterals and perforant path (from EC to CA3). In the storage process, the size of the constellation of CA3 units chosen to form the internal representation of an environmental configuration is crucial. As it will be argued in chapter 7, it is important that different memories be represented by constellations with roughly the same size. The size of the CA3 constellations are intimately related with the activity level at the dentate gyrus.

Computer simulations were carried in order to show that the described activity level control can be achieved in the DG. Three populations were used to represent the dentate gyrus (see figure 6.12-**B**): a population of principal units representing the granule cells, a population of excitatory interneurons representing the mossy cells and a population of assorted inhibitory interneurons. The size of the principal (P) population was of 1000 units while both excitatory (E) and inhibitory (I) population were composed of 50 units each.

The parameters of the network were set, using the methods described in this chapter, so that an activity level of 1.0% was produced in the principal population. The behaviour of the network was assessed by analysing its response to an initial level of

activity. Three regions were clearly produced and their interpretation should be done taking into account the effect of the mossy fibres in CA3:

1. For an initial activity level below 0.5% no amplification occurred. In this case the activity would be considered to be noise and no transmission would occur to the CA3 region since there was not enough excitation to bring CA3 neurons above threshold.

2. When the initial activity level is above 0.5%, the excitatory interneurons amplify the level of activity and activate further units. The amplification is supervised by the inhibitory interneurons which assure that the optimal level of 1.0% is not exceeded.

3. When the initial activity level is close to 1.0%, the excitatory interneurons are completely silenced by the inhibitory interneurons. Levels much higher than 1.0% are equally not transfered to the CA3 region since the contacts that filopodial extensions of mossy fibres establish with CA3 interneurons (Acsády *et al*, 1998) completely block the excessive excitation to the pyramidal neurons in CA3.

Figure 6.12-**A** shows these three regions of control. The blue bars represent the activity added to the initial activity level shown in the horizontal axis. This added activity results from units that are brought above threshold due to the amplification process. The total level of activity, as seen by CA3, is represented in figure 6.12-**A** by the red line, which is simply the sum of the initial activity level with the added activity. For example: when an activity level of 0.7% is induced in the dentate gyrus, the control mechanisms (recurrent excitation and recurrent inhibition) bring more 0.37% of the principal population to the threshold level in less then 5 ms. This way, in integration time window ($\sim \tau_m$) of the CA3 region, the dentate gyrus supplies an activity level of 1.07. The data presented in 6.12-**A** is an average of 3 simulations.
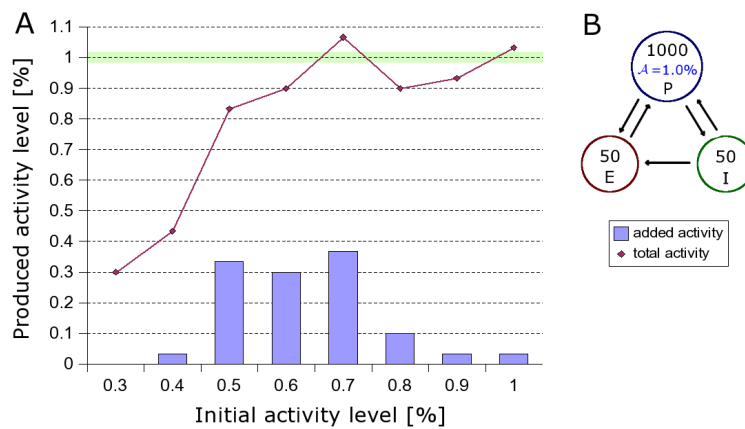
Figure 6.12: Activity level control in the dentate gyrus. See text for detailed description of the figures **A** and **B**. Parameters: $N_P = 1000$, $N_E = 50$ and $N_I = 50$; $C_{PE}^d = 25$, $C_{EP}^d = 100$, $C_{PI}^d = 40$, $C_{IP}^d = 200$ and $C_{IE}^d = 100$.

Chapter 7

Episodic Memory: Assumptions and Theory

The hippocampus is believed to be essential for the encoding of episodic memories (Squire, 1992). In this chapter, a theory for the hippocampus' functional behaviour in the formation of episodic memories will be presented.

This theory is based on Marr (1971), Treves and Rolls (1992; 1994), O'Reilly and McClelland (1994) and Hasselmo *et al* (1995; 1996). It places CA3, and its extensive recurrent connections, in the centre of the computational machinery required to perform information storage and recall. It differs, nevertheless, in some important functional roles for particular fields and pathways.

All the central assumptions of this theory can be effectively implemented and tested in a biologically plausible spiking network with properties that follow closely the hippocampally most relevant morpho-physiological constraints. In fact, this is the purpose of chapter 8 where a computational model of the hippocampus' contribution to the formation of episodic memories is created. This computational model allows detailed analysis of the quantitative aspects of the theory. This represents a significant step forward in the understanding of the hippocampus' functional behaviour. It is worth mentioning that many hypotheses raised in previous theories - directly or indirectly related to the original theories alluded above - have not been tested in quantitative terms. In fact, as shown later, many specific computational issues only become evident after effective implementation of a detailed model.

Organisation of the Chapter

This chapter is divided into two sections. The first section describes the theory behind the hippocampal computational model for episodic-type memory, giving special attention to the functional roles of each field. The second section addresses assumptions regarding inputs and outputs of information and plasticity in the hippocampal computational model.

7.1   Theory of episodic memory in the hippocampus

Before going into the details of the computational model, it is important to describe my beliefs for the functional roles of the hippocampus.

The central idea is that the hippocampus provides a limited storage space where relevant multi-modal information is temporarily stored. The storage process is triggered whenever the environment configuration is significant. The information stored is an internal representation of the environmental configuration (an episode), thanks to the multi-modal nature of the inputs to the hippocampus. The level of significance of a particular environmental configuration depends on novelty and on emotional response to the configuration. It is important to keep in mind that the hippocampus is part of the limbic system and one of the oldest phylogenetic regions of the brain. It is directly connected with equally old phylogenetic structures: the septum and the amygdala. In general terms, the septum enhances or suppresses all motivations for acquisition while the amygdala enhances or suppresses all motivations for active avoidance (retreat). These two regions, the septum (Hasselmo *et al*, 1995; 1996) and the amygdala (Richter-Levin, 2004), act as memory modulators for the hippocampus.

Although it was said that the hippocampus provides a temporary storage space, it may keep information for very long periods of time as long as the information, in the form of a memory, is periodically revived.

7.1.1   Neocortex, hippocampus and declarative memory

In mammals, the neocortex provides another storage space, with a much higher capacity, where memories can be solidly engraved. There is experimental evidence that the hippocampus and the neocortex work together (Norman and O'Reilly, 2003) in establishing long-term memories. Particularly in humans, where a declarative memory

system exists, the hippocampus plays an essential role in the encoding of new declarative memories. The hippocampus is nevertheless not required for the retrieval of old, solid memories, which are stored in the neocortex.

Albeit having a much higher capacity, the repository for time-resilient memories provided by the neocortex is still finite. It is therefore important to select which new information is effectively relevant to be stored. In addition, engraving new information in the neocortex, which already contains a lot of stored information, is a delicate procedure: if the learning process is not properly done, previously stored episodes can be corrupted. A critical variable in this process is the learning rate which has to be slow. The storage has to be mediated through small changes on the network, accommodating the new information without disrupting previously stored episodes. These two processes, selection and slow transfer, require a space where information is temporarily stored.

The information transfer from the hippocampus to other cortical areas (neocortex), where the information is consolidated, is often argued to happen during sleep (see for example Káli and Dayan, 2004). The hypothesis of slow information transfer from the hippocampus to neocortex is supported by experimental results in which hippocampal damage impairs information retrieval not only regarding experiences posterior to seizure but also experiences that occurred during a short period prior to seizure.

Since the hippocampus is not required to retrieve consolidated memories, it is reasonable to assume that, somewhere, a short-cut is formed in the information flow between the (invocation) input cortical areas and the long-term storage areas. That is, information does not go through the hippocampal formation in order to produce memory retrieval.

7.1.2 Functional roles of the hippocampus

This thesis looks at the hippocampus independently of the neocortex's declarative system. The purpose is to analyse the functional properties of the hippocampus alone in storing episodic information through a fast learning process. It is worth repeating that, in this thesis, "episodic" information has the meaning of multi-modal information representing an environmental configuration. An environmental configuration is defined by the state in the sensory space (which includes somesthetic, proprioceptive

and other high level sensory information). In a more general and abstract view, an environmental configuration may also include cognitive or behaviour states in addition to the perceptual states.

In this hippocampal theory, information is assumed to be processed in the form of patterns of activity. These activity patterns are defined, in a population, as the constellation of neurons that have reached threshold within a time window of a few tenths of milliseconds (see definition in section 1.3). The size of the constellations are roughly constant and related with the average activity level of the population. Due to the constraints in the synaptic strengths, the activity produced in a post-synaptic population depends directly on the number of active neurons in the pre-synaptic population. While small assemblies of active pre-synaptic neurons fail to effectively raise post-synaptic groups of cells, large assemblies of active pre-synaptic neurons may bring too many post-synaptic neurons to threshold level. It is therefore important that the size of the coding constellations (activity patterns) is roughly similar and dependent of the connectivity properties and maximum synaptic strengths.

I hypothesise that the functional roles of the hippocampus in order to establish a system capable of storing episodic-type memories are:

1. Choose, from its input space, the patterns of activity which are statistically relevant. This task requires history-dependent dynamics.

2. Store the relevant information in the form of internal representations. A proper balance between pattern completion and pattern separation is required.

In addition, for the selected stored patterns, the hippocampus is required to produce an output activity pattern. This new activity pattern is the one that will in fact be available for use by the other cortical areas.

The selection process (point **1**) of which information to store carries considerable responsibility: the storage capacity is limited and only statistically relevant information should be stored. That is, only features of the environmental configuration that are consistent across time, or that evoke strong emotional responses (modulated by the septum and, maybe, the amygdala), should undergo the storage process. Moreover, the way information is stored is also an important factor, with orthogonal representations taking the maximal profit of the storage capacity. It is then also a task of the

hippocampus to create internal episodic memory representations that improve storage and recall.

In order to provide more detail to point **2** above, the following definitions are introduced.

Consider the set Π defined as the set of all pure input patterns for which the hippocampus has to store an internal representation. A pure pattern corresponds to the exact constellation of active neurons that motivated the storage process.

For each pure pattern $\pi^k \in \Pi$, we can define a class $\tilde{\Pi}_\varepsilon$ which incorporates all perturbed versions $\tilde{\pi}^k$ of the pure patterns with an error $\leq \varepsilon$. Depending on the discussion, the error can be the percentage of active units that are missing (partial cue) or that are misplaced (noisy pattern), when compared with the pure activity patterns.

Using this notation, point **2** above expresses that for any input pattern $\tilde{\pi}^k$, the hippocampus should provide the same output, which corresponds to its internal representation of π^k. On the other hand, different pure patterns π^k (or patterns referring to different information) should have different internal representations. Ideally these representations should be as distant as possible from each other.

The distance between two activity patterns is defined here as the sum of the square differences of the states (either active or inactive) of all neurons in the population that holds the patterns. If the population contains N neurons and s_n^k represents the state (0 for inactive, 1 for active: spike produced) of neuron n in pattern k, then the distance d between patterns i and j can be written as:

$$d(\pi^i, \pi^j) = \frac{1}{N \mathcal{A}} \sum_{n=1}^{N} \left(s_n^i - s_n^j \right)^2 \tag{7.1}$$

The term $\frac{1}{N \mathcal{A}}$ is a normalisation factor that accounts for the number of active neurons in the patterns.

The more active neurons are shared between the patterns, the lower will be the distance measure (see also section 8.7).

7.1.3 Inputs and outputs of information

The two fundamental computational tasks of pattern selection and pattern storage, described above, take place in two fields of the hippocampus: dentate gyrus and CA3.

The inputs to these fields are provided by layer II of the entorhinal cortex which in turn receives connections from association neocortex via the para-hippocampal gyrus and perirhinal cortex. The ECII activity patterns therefore reflect the combination of information from different sources and of different natures (see figure 2.3)

Particularly relevant to this view is that entorhinal cortex is a major gateway for sensory information into the hippocampal formation. It has been shown that neurons in rat's lateral entorhinal cortex react to auditory, visual, somesthetic, gustative and olfactory stimulations (Vayssettes-Courchay and Sessler, 1983).

Equally important is that the entorhinal cortex, and dentate gyrus, maintain some segregation of inputs originating from different parts of the cerebral cortex (Suzuki and Amaral, 1994b; Rolls and Treves, 1998). That is, different regions of EC reflect the influence of different sources.

Taking into account this convergence of multi-channel information, one can say that the activity patterns formed at entorhinal cortex layer II represent by themselves episodic-type information: sensorial inputs such as visual, olfactory, proprioceptive (and maybe previously stored information) are all blended together providing a *what-where-when* structure (episodic structure).

By storing an internal representation of a ECII pattern of activity, the hippocampus is storing the combined information representing a particular environmental configuration, or, in other words, an episode in the sensorial space.

From now on, we consider the patterns of activity at the entorhinal cortex to be the representation of an episode[1]. The underlying assumption is quite strong: it is being argued that the hippocampus, instead of actively combining information from different sources to produce "episodic information", will be continuously monitoring the already combined patterns at ECII for statistically relevant activation patterns, expressed in terms of consistency across time.

At most, the hippocampus will combine different input patterns that are systematically correlated (e.g. one pattern always precedes the other) by creating a single internal representation for both input patterns. This way, future exposures to any of the two patterns is interpreted by the hippocampus as a partial cue for the combined representation that should be recalled.

---

[1]The nature of the information expressed in the activation patterns of ECII will be discussed in more detail in the next chapter in which spatial learning is addressed.

This concludes the discussion of the inputs to the hippocampal system. In terms of outputs, both CA1 and subiculum, the two layers that succeed CA3, send connections to entorhinal cortex layers III and IV but predominantly to layer V. While the superficial layer II in the entorhinal cortex provides the input to the hippocampus, the deep layers correspond to the output.

For each pure pattern $\pi^k \in \Pi$, the hippocampus produces, in the output layers, an associated activity pattern which results from its stored internal representation of π^k. Since the crucial computational tasks are performed at DG and CA3, we will assume for now that CA3 represents the output. The CA1 stage involves mainly a translation process which is functionally important but not computationally demanding.

We now move on into the discussion on how the two main tasks assumed above are implemented in the hippocampus architecture. That is, we will discuss the mechanisms that implement the computational tasks described in points **1** and **2**. The next two subsections discuss in detail the functional roles of dentate gyrus and CA3 respectively.

7.1.4 Functional roles of dentate gyrus

Taking into consideration the experimentally measured average activity levels for ECII, DG and CA3, as well as their population numbers, we can see that both ECII and CA3 configurations space are much smaller than the size of the intermediate space formed by the DG.

Assuming that each activity pattern contains roughly the same number of active neurons - defined as the product of the population size N by the activity level $\mathcal{A}$ - we can verify the previous statement by calculating the total number of different patterns available on each field. Calculating the combinations we get that (see table 8.1 for numerical values):

$$\begin{pmatrix} N_{ECII} \\ N_{ECII}\,\mathcal{A}_{ECII} \end{pmatrix} \ll \begin{pmatrix} N_{DG} \\ N_{DG}\,\mathcal{A}_{DG} \end{pmatrix} \gg \begin{pmatrix} N_{CA3} \\ N_{CA3}\,\mathcal{A}_{CA3} \end{pmatrix}$$

The fact that the dentate gyrus works as an intermediate higher dimensional space has supported the idea that this field's functional role is to create sparse and separate (more distant) representations on CA3, thus increasing the hippocampal storage capacity (Treves and Rolls, 1992).

The DG can be seen working as a support vector machine by taking advantage of its higher dimensional space to linearly separate the input patterns (Baker, 2001).

A more plausible way by which DG can increase the separation between input patterns is discussed in O'Reilly and McClelland, 1994. The core idea is that separation (distance) is increased simply through sparse random connections.

Although I agree that a major functional role of the dentate gyrus is to increase the separation of CA3's representation, I do not think that this represents its only purpose. The dentate gyrus has the very important roles of (1) actively selecting which input patterns are relevant (meaning that their representation should be engraved in CA3) and (2) controlling the activity level of the representations in CA3 so that the number of active neurons in any stored patterns is roughly constant (i.e. independent of the ECII input activity pattern).

The hippocampus is continuously receiving inputs in the form of patterns of activity at the entorhinal cortex. But only a small fraction of these patterns should trigger storage and creation of associated internal representations. Two reasons support this statement: first the hippocampus possesses a limited capacity as a storage system; secondly the hippocampus is required to perform generalisation in the mapping from input patterns to internal representations (therefore should not produce a different output for every single input pattern). Two mechanisms execute this selection process.

At a first basic level, only relevant information will be able to raise the EC's activity levels (Hasselmo *et al*, 2000) to the point that allows propagation of activity for subsequent levels, i.e. dentate gyrus and CA3. Low activity levels at the entorhinal cortex do not bring a statistically relevant number of post-synaptic neurons (in DG and CA3) to the threshold level.

The second selection stage occurs at mossy fibre synapses. In a naive (without stored information) hippocampus, all perforant path connections to CA3 are extremely weak and are therefore incapable of directly transferring activity from ECII to CA3 by themselves. Nevertheless, these synapses exhibit associative, Hebbian-like, plasticity which means that, if a constellation of neurons in the CA3 field is activated - through the mossy fibres pathway - then all the synapses between active ECII and CA3 neurons will be enhanced. The same will happen to the recurrent synapses between active CA3 neurons which also exhibit associative plasticity. It is this later potentiation that leads to the storage of the CA3 patterns of activity (the hippocampus internal representation).

Mossy fibres play a critical role in selecting which CA3 neurons should become active not because of its static strength but as a consequence of its high variability. In other words, it is the variance of the efficacy of mossy fibres synapses, instead of the mean value, that truly represent their effect in CA3 (O'Reilly and McClelland, 1994).

Mossy fibres are unique in the central nervous system for their dynamical properties (discussed in detail in chapter 5). The unitary mossy fibre EPSC can vary a lot, due to short-term plasticity dynamics such as facilitation (depression seems to be small at natural stimuli frequencies). Mossy fibre facilitation can lead to a 2-fold potentiation for frequencies as low as 0.2Hz and to more than 40-fold potentiation of EPSCs for 100Hz stimulation (Urban *et al*, 2001). The strong short-term facilitation that can act upon mossy fibre synapses means that this structures can be reliable input frequency estimators (see equations 5.4 and 5.5): mossy fibre synapses that are systematically activated consistently increase their post-synaptic effect.

The direct implication of this facilitation process is that the activity level at CA3 will only significantly rise when the same constellation of ECII neurons, and associated DG constellation, has been repeatedly activated. In other words, fluctuations and noisy patterns are filtered out.

Mossy fibres also show LTP and LTD, both Hebbian and non-Hebbian. Interestingly, knockout mice with impaired LTP and LTD at mossy fibre synapses do not show any deficit in learning tasks (Huang *et al*, 1995). It seems that short-term plasticity is of much more importance for the role of dentate gyrus than long-term plasticity. The dynamical properties grant these synapses the ability to act as a frequency discriminator. This point has been studied in Aguiar and Willshaw (2004) (see Appendix).

An effective experiment able to shutdown only the short-term plasticity at mossy fibres has not been achieved yet. Such experiment would possibly uncover a correlation between learning and mossy fibres short-term plasticity. Nevertheless, in Lassalle *et al* (2000), it has been shown, through an excellent experiment where the mossy fibres are inactivated in a reversible manner, that mossy fibre synapses play a crucial role in acquiring new memories: inactivating mossy fibre synapses impairs spatial learning in mice doing the Morris navigation task. Consolidation and retrieval are not affected. Reactivating the m.f. synapses allows the mice to perform spatial learning. This result follows the now canonical role for the dentate gyrus introduced by some models (Treves and Rolls 1992; 1994; McNaughton and Smolensky 1991) which stated that

mossy fibres are needed for the process of learning but not for memory retrieval.

Another important hypothesis that I make is that all stored patterns in CA3 are composed of roughly the same amount of active neurons. The reasons to assume, and require this property, will be described when discussing the functional roles of CA3 (section 7.1.5). This assumption has a direct impact on the properties of the mossy fibres. It means that in order to create a proper internal representation, the dentate gyrus has to bring above threshold roughly the same amount of CA3 neurons, independently of the ECII input pattern. Therefore specific mechanisms to control the amount of excitation delivered to the CA3 pyramidal population must exist. It should be noticed that the activity level control mechanism that results from the CA3 interneuron population receiving merely inputs from the CA3 pyramidal cells, does not avoid transients in the activity level.

This hypothesis is in accordance with the complex connections that mossy fibres establish with CA3 interneurons. Each mossy fibre establishes contacts with perhaps no more than 14 CA3 pyramidal cells. But on the other hand, each mossy fibre forms filopodial extensions and small terminals contacting a number of interneurons 10 times larger (Acsády *et al*, 1998).

To avoid abnormal excess excitation to CA3 pyramidal cells (as a result, for example, of a higher dentate gyrus' activity level), the mossy fibres connect directly to CA3 interneurons through this small terminal, correcting for any excess of excitatory input before the CA3 activity rises.

It is an experimental fact that increased activity of granule cells suppresses the overall excitability of the CA3 recurrent system (Acsády *et al*, 1998). As it will be shown later, this is a key point in this theory for the role of hippocampus in episodic memory formation. If abnormal increases in the dentate gyrus' activity level propagates to CA3, storage and retrieval of internal representations will most probably be corrupted.

Besides the complex feed-forward inhibition system that controls from above the activity level induced in CA3, I believe that another system exists to cope with lower excitation levels. It is essential that the spatial coding created by the dentate gyrus and mossy fibres produces a representations of roughly constant size in CA3. Proper information recall depends on this.

Therefore, besides constraining from above the amount of excitation that mossy fi-

bres provide to CA3, it is also important to constrain this excitation from below. Since the activity on mossy fibres reflect the activity levels on granule cells, this means that a mechanism should exist to increase excessively low levels of activity in the DG. I purpose that this is the role of mossy cells, the dominant type of interneurons in DG (see section 2.1.1). This glutamergic, excitatory cells, merged with the population of granule cell can act as a single network with the recurrent connections allowing feedback amplification. The other type of (inhibitory) interneurons guarantee nevertheless the upper bound control of the activity level.

The result is a network that can produce the behaviour depicted in figure 6.5, that is, increase its activity level (when above minimum value) until the interneurons come into play controlling the level of activity. The amount of excitation transfered to CA3 is, in this way, compensated avoiding the creation of small representations that would lead to poor recall.

7.1.5 Functional roles of CA3

The task of associating the relevant activity patterns with internal representations falls into the domain of CA3.

The key property of CA3 field that allows the execution of this task is its extensive plastic recurrent connectivity. The recurrent connections allow CA3 to behave as an auto-associative memory. This means that CA3 is able to recall stored patterns through the presentation of noisy or incomplete versions of these patterns. Although the idea that this is the core property that drives the recall process in the hippocampus is shared in this thesis, I present an alternative to the form in which information is assumed to be stored in CA3.

Most models for CA3 are based in attractor networks. In such dynamical systems, it is possible to create stable states, with large basis of attraction, by manipulating the system's energy landscape[2]. Initial conditions that fall inside a basis of attraction, will converge, in the state space, to the local minimum (for more details on the subject of neurodynamics see Haykin, 1999). The core idea behind attractor networks is therefore viewing each stored pattern (network state) as an attractor point.

---

[2]In fact, such dynamical systems (dissipative systems) may contain several types of attracting sets or manifolds. Only *point attractors* are used in the context of auto-associators. Others like limit cycles are behind many models for spike generation mechanisms.

The literature that discusses CA3 as an attractor network focuses then in the biologically plausible ways to manipulate the energy landscape. Many interesting problems arises like how to create large and deep basins of attraction (allowing increased robustness to noisy patterns), how to avoid spurious states (natural local minima), among many others.

There is another key assumption in the CA3 attractor network literature which is that of rate coding. This means that the state of each independent unit on these, either discrete or continuous, attractor networks represents a firing frequency. This goes against my view on how information is propagated and computed in the hippocampus, i.e. through patterns of activity.

The assumption of temporal coding at a population level has motivated a different hypothesis on how episodic memories are temporarily stored in the hippocampus. Please note that the *precise* time of each spike is not important. An activity pattern is defined within a time window of the order of the membrane's time constant (see definition in 1.3). According to this hypothesis, the transfer of activity between hippocampal fields is gated; that is, only in proper conditions will a pre-synaptic field be able to bring above threshold a considerable number of neurons on the post-synaptic field.

Two examples of this gating have already been described in this thesis: first, dentate gyrus will lead to an increase in the activity level of CA3 only when the same constellation is continuously activated; and, secondly, ECII will only gain access to rise the CA3 levels when a proper translation between input pattern and internal representation has been learnt (mediated by mossy fibres).

According to this gating hypothesis, CA3 will only be able to bring a CA1 population to fire if its activity is close to the desired level (above 2.5% in rat's CA3). While CA3 activity level that is above the desirable level triggers strong inhibitory inputs to CA1, small CA3 fluctuations are incapable of providing enough excitation to bring enough CA1 neurons above threshold. Nevertheless, small activity at CA3 that fits a partial cue of a stored pattern can, through the auto-associative properties of this field, rebuild the whole pattern, raise the activity level and then transfer information control to CA1.

If the activity level at CA1 goes beyond the level of fluctuations, this *means* that a pattern has been recalled. The pattern recalled is represented by the CA1 constellation

of active units.

This translation from CA3 internal representations to CA1 patterns is mediated by entorhinal cortex layer III connections. Naive Schaffer collaterals provide only weak excitation to CA1 neurons which require the combined excitation from the ECIII perforant path to exceed threshold. The word "naive" is again used to denote connections that have not been modified through learning.

Nevertheless, once a specific constellation of CA1 neurons has been activated due to combined excitation from naive Schaffer collaterals and ECIII, plasticity in the form of associative LTP can act upon Schaffer collateral synapses stabilising the translation code. Without plasticity (long-term plasticity) on Schaffer collaterals, the translation code that converts the CA3 stored representations to CA1 patterns is not consistent across time (fluctuations on ECIII select different CA1 constellations of active neurons). The notion of translation connections is one of the pillars of this hippocampal computational model. This translation occurs in several fields always using the same scheme: an activity profile in one field is translated to an activity profile in another field through the mediation of a third population of neurons that chooses the spatial code.
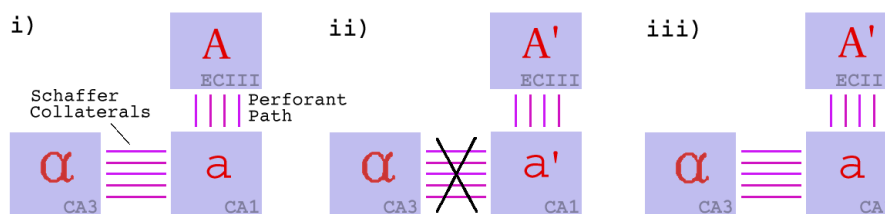
Figure 7.1: Translation code between CA3 and CA1 provided by the Schaffer collaterals. In **i)**, a pattern **A** in the entorhinal cortex generates a pattern α in CA3 and a pattern **a** in CA1. Schaffer collaterals store the translation from α to **a**. In **ii)**, in the absence of Schaffer collaterals, a noisy version of pattern **A** would generate in CA1 a different pattern **a'**. The Schaffer collaterals, in situation **iii)**, stabilises the pattern in CA1 using the pattern completion properties of CA3. Using the connections from CA1 back to the EC, the activity pattern in EC can be corrected using the stored information in the hippocampus.

Interneurons in CA1 have again the fundamental task of bounding the activity lev-

els at the field by securing that Schaffer collaterals, combined with ECIII connections, do not activate too much CA1 pyramidal neurons. If the excitation is too strong, the interneurons enforce a *k*-winners-take-all situation (O'Reilly and McClelland, 1994) where *k* represents the number of active units in the peak activity level of CA1 field.

Taking into account the hypothesis described above, it can be seen that the episodic memory spiking model introduced here does not require stable states. Storage and recall is performed without point attractors and stable patterns of activity. In fact I argue that convergence to stable firing configurations, in a biological network subject to oscillations, fluctuations and current noise, is not required for storage and recall. Transient states are sufficient to induce recollection.

In order to complete this discussion about the individual tasks of the hippocampus, supporting the establishment of episodic-type memories, a word must be said on the functional roles of CA1 and subiculum, the two fields that close the loop with the entorhinal cortex.

It has been suggested in McClelland and Goddard, 1996, that Schaffer collaterals together with the bi-directional pathways between EC and CA1, support a memory decoding system through which a retrieved CA3 representation (coding) of an EC pattern can reinstate that pattern on EC. This way, corrections can be performed in the EC present pattern using the information about the pure pattern stored in CA3. This nevertheless removes the possibility of the hippocampus having a role in producing (different) representations of the EC input that benefits storage and recall in the areas for long-term storage.

Not going against this suggestion, we consider that another possible role to CA1, and subiculum, is to direct the establishment of connections between entorhinal cortex superficial layers with the deep layers. According to my gating hypothesis, the activity level at CA1 only rises significantly when a pattern has been retrieved, i.e. when an ECII activity pattern has been recognised. Although without any experimental evidence to support this claim, the EC intra-layer connections could be strengthened through learning, using the gating support (see hippocampus connections diagram 2.3) from CA1 field (and subiculum). The feed-forward connections between layers of EC could then provide a bypass to the hippocampus relay station. It should be emphasised that these EC intra-layer connections are only subject to learning when CA1 activity denotes recognition of a pattern. It is also important to keep in mind that (although in a

much less efficient way) feed-forward connections are also capable of pattern comple-tion. This is, as long as post-synaptic neurons have access to the pure pattern, which is the case here, if we accept the memory decoding proposal of McClelland and Goddard (1996).

The hypothesis that the hippocampus also has the functional role of creating proper[3] representations for long-term storage (as opposed to correcting the EC pattern) fits into this framework since the input patterns in layer II can be associated with the sparse representations produced by the hippocampus in EC layers III, IV and V. These feed-forward intra-layers connections could then produce a hetero-associative system robust to some fluctuations in the input space.

It is worth mentioning again that a direct storage of the multi-channel information forming the episodic-like memories is not computationally appropriate. The space for long-term storage is limited, therefore requiring a proper selection of which infor-mation is relevant to store and also requiring the construction of representations that maximise pattern separation and completion (recall).

The hippocampus memory system provides a way to select and temporarily store new memories. The internally efficient representations of this memories, enhancing capacity values, can then be retrieved and "played back" (McClelland and Goddard, 1998) to the neocortex where they are accommodated with other memories for long-term storage.

7.2 Important Assumptions

This section discusses important assumptions in the hippocampal computational model with strong impact on its functional behaviour (see chapter 8).

7.2.1 Inputs and outputs to the computational model

The entorhinal cortex II spiking neurons provide the input to the episodic memory model. The activity profiles at ECII are set under the assumption of the existence of a mapping function $\mathcal{M}$ with the properties stated in section 1.3. This function projects to the high dimensional space, in which sensorial, vestibular and previously stored infor-

---

[3]Maximising capacity and balancing pattern completion with separation.

mation is represented, into the lower-dimensional space of ECII's patterns of activity. For simplicity let us disregard vestibular and previously stored information (consolidated memories) from the EC's patterns construction. In other words, let us assume for now, and without loss of generality, that ECII's activity pattern reflects solely the configuration in the sensorial space.

What is in fact represented in the EC's patterns, is not important for the computational model for episodic memory. All that is important is that it expresses a current state in the whole system. The same does not happen in the computational model for spatial memory presented in chapter 9 in which the features represented in the EC's patterns are important.

What is expected for this episodic memory computational model is to be able to:

- decide which patterns (which information) are relevant;

- store and recall the relevant patterns in a network with high capacity;

- cope with the situation in which a noisy version or partial cue is presented.

Let us define C^k as the set of patterns that represent noisy versions or substantial fragments of pattern k. Pattern p is considered to be stored when the same output state is produced for any input pattern belonging to class C^k.

There is a significant difference in the way information is stored in this model: while other models for associative memory based in recurrent networks are grounded on the creation and management of stable states (Hopfield, 1982; Treves and Rolls, 1992), the existence of long lasting patterns of activity is not required for the computational model presented here. In this model, each memory is represented by the activation of a specific neural ensemble with the capacity to propagate activity to subsequent layers. Stored patterns of activation in CA3 have the ability to raise significantly the CA1's activity level while unstored patterns do not provide enough excitation to bring CA1 neurons above threshold level. This gating of the activity is achieved through Hebbian learning that selectively enhances the connections formed with the ensembles of neurons that represent stored configurations in the input space. The details about all this information storage and propagation will now be discussed.

7.2.2 Plasticity

Following the main dogma in computational neuroscience, information is assumed here to be stored in the form of modifications in synaptic efficacy. Both short-term and long-term plasticity are present in the hippocampal model.

Of all the connections discussed previously, only the recurrent collaterals and perforant path connections to CA3 exhibit long-term plasticity. Mossy fibres exhibit short-term plasticity and all the other connections are not plastic.

Mossy fibres - short-term plasticity

The large boutons that mossy fibres create in the very proximal dendrites of CA3 neurons have led to the view that mossy fibres can act as "detonators". Although it is incorrect that a mossy fibre alone can bring a CA3 pyramidal neuron to threshold, it is commonly agreed that the mossy fibre bouton's unique morpho-physiological properties are behind equally unique computational properties (Urban *et al*, 2001, Aguiar and Willshaw, 2004).

It has been suggested by Treves and Rolls (1992), that mossy fibre input creates efficient information storage through their sparse and strong influence on the CA3 cell firing rates. This hypothesis has been verified experimentally (in a spatial memory framework) by Lassalle *et al*, 2000, in which reversible inactivation of the mossy fibre' synapses were shown to impair learning, but neither consolidation nor retrieval. A detailed quantitative model on how mossy fibres affect the storage process is nevertheless inexistent (but see analysis present in O'Reilly and McClelland, 1994).

From the variety of forms of plasticity that mossy fibres have been shown to display (Urban *et al*, 2001), including both Hebbian and non-Hebbian forms of LTP and LTD, only short-term plasticity was included in the present computational model. This choice is based on the fact that in all studies done so far, impairment in any form of long-term plasticity has failed to produce effects on the performance of a learning task (see chapter 2 for more on this). It is then assumed here that the functional role of mossy fibres is deeply bounded with their short-term properties. Mossy fibre short-term dynamics are modelled using the dynamics from Fuhrmann *et al*, 2002, discussed in chapter 5. The properties of strong potentiation even at low excitation levels, as well as hardly measured depression, have led to the choice of the values $U1 = 0.4$, $\tau_{fac} = 2000ms$ and $\tau_{rec} = 10ms$ (see equations 5.2 and 5.3).

The maximal post-synaptic response, expressed in terms of the synapse's maximal

conductance, is set in order to avoid high activity levels in CA3 (i.e. above 5%).

Perforant path (to CA3) - long-term plasticity

The plasticity in the perforant path synapses with CA3 was modelled using spike timing-dependent plasticity. The dynamics follow equations 3.8 with potentiation and depression time windows of equal size, 40*ms*. The dynamics with potentiation/depression strength decaying exponentially with the spike time interval (see figure 3.3) were not used since this increase in the complexity did not have a particular role in the computational model.

An upper limit for the peak conductances, defined as S_{max}, is imposed. This maximum peak conductance prevents synaptic conductances from going beyond biologically plausible values, as a result of continuous potentiation. The soft-bound dynamics attempts to avoid the peak of the synaptic conductance distribution at S_{max} (van Rossum *et al*, 2000). The distribution is bimodal due to the considerable amount of synapses with negligible peak conductances. In the biological systems these synapses could degenerate and disappear through a selection process that keeps only significant synapses (and therefore saving resources in the maintenance of this structures). In a computational model that does not include synaptogenesis, all synapses with zero or negligible conductance should be interpreted as absent.

The upper-bound S_{max} value is set in order to promote cooperation of neuron units. That is, S_{max} is kept small enough so that a significant number of active perforant path inputs will always be required to raise a CA3 unit's voltage above threshold. This way a neuron unit will not reach threshold with just few synapses potentiated to the limit.

Recurrent collaterals' long-term plasticity

The long-term plasticity in CA3 recurrent collaterals follows simple associative learning dynamics with a time window of 40*ms*. The rule for simple associative potentiation is the same as the STDP potentiation rule shown in equations 3.8. An upper-bound S_{max} is imposed on the synapses' peak conductances. This value, while promoting again the cooperation of units, is nevertheless bigger, in relative terms, when compared with the perforant path synapses. This way, less active inputs are required to drive CA3 neuron units above threshold level. The idea is that a small sub-set of active units are then able to rebuild the complete stored pattern.

From equation 6.1 the average number of active recurrent collateral inputs can

be calculated. This value provides then an estimate for S_{max}. For the connectivity values presented, the average number of active recurrent collaterals is 6. The maximum conductance is then set so that, at least 6 active recurrent collateral inputs are required to fire a CA3 neuron unit. Due to the soft-bounds in the potentiation dynamics, in practice, a minimum of 7 active recurrent collaterals are required to fire a CA3 neuron unit. Considering the average size of CA3 activity patterns, in order to rebuild a stored pattern, a partial cue must contain at least half of neuron units in the pure activity pattern above threshold.

Chapter 8

Episodic Memory: Implementation and Results

This chapter is fundamental in making this work different from many other proposed theories for the role of the hippocampus in episodic memory formation: in order to complement the qualitative description of the tasks taking place in the hippocampal fields (see chapter 7), I create a computational model to analyse, in quantitative terms, all the hypotheses behind the theory. Through the creation and simulation of the computational model, many specific functional details emerge which provide a deeper understanding of the strengths and limitations of the theory. In addition, results from the simulations provide information to create refinements in the model. The computational model focuses on the two tasks that have been described as the main functional roles of the hippocampus: selection of relevant information in the multi-modal input space and storage of this information in the form of an internal representation of the combined presence of particular multi-modal features.

This chapter focuses on the results of several simulations of the hippocampal computational model for episodic memory. The simulations performed had the objective of comparing the dynamics of the hippocampal model with results well established in the hippocampus literature.

Organisation of the Chapter

This chapter is divided into eight sections. The first section discusses the assumptions behind the scaled down version of the hippocampus used in the simulations. The

results from the simulations refer to the following assessments, which are all addressed in separate sections:

- Pattern separation: the role of the perforant path and mossy fibres in increasing the distance between internal representations referring to distinct environmental contexts, or configurations;

- Feature selection: the contribution of the short-term plasticity of mossy fibres to the storage process;

- Storage and recall of an internal representation: complete time course of the functional events that mediate the creation and recall of an internal representation;

- Firing properties: firing profile of neuron units involved in the storage and recall of an internal representation ;

- Theta rhythm: the importance of periodic levels of activity in order to separate activity patterns containing different information.

- Capacity: empirical analysis of the capacity of the memory system.

Section 8.8 presents some final remarks and concludes the chapter.

8.1   Scaling down the hippocampus

Three spiking networks, representing entorhinal cortex II, dentate gyrus and CA3, compose the episodic memory computational model. A population of interneurons is integrated into the CA3 spiking network.

Although the created simulation environment (chapter 4) allows the use of an extremely high number of spiking neuron units, the biological network numbers have to be scaled down in order to efficiently simulate and analyse the computational model. While the rat's hippocampus number of units in each field is of the order of 10^5 (see table 2.1), the number of neuron units in each of the computational model's field is of the order of 10^2.

The network properties can be scaled-down in several different ways, depending on which constraints are chosen to direct the scaling process. Since the hippocampal

theory proposed here is based in patterns of activity, the average activity levels are chosen as the central constraint for the creation of network counterparts for ECII, DG and CA3.

Table 8.1 presents both biological and computational model numbers for the neurons in each area as well as the activity levels $\mathcal{A}$. All these values are averages but for legibility the $< . >$ notation is dropped.

Area	rat hippocampus		computational model	
	n. neurons	$\mathcal{A}(\%)$	n. neurons	$\mathcal{A}(\%)$
EC	200,000	7.0	250	10.0
DG	1000,000	0.5	1000	1.0
CA3	160,000	2.5	200	5.0

Table 8.1: Population numbers in the rat hippocampus and in the computational model.

Roughly a factor of 1000/1 is used in scaling of the number of units. The activity levels are increased in order to enhance robustness to noise: fluctuations on the number of active units produce smaller changes on the activity levels. Take the example of CA3 in which, for the biological activity level of 2.5%, 5 units, on average, would be active in a network of 200. A variation of a single active unit means a change of 20% on the activity level. Duplicating the activity levels halves this effect.

The discussion now moves to the connectivity numbers. Instead of being set simply to maintain the relative excitatory contribution from each area, the connection properties are chosen to address other aspects which will now be discussed.

While many models make CA3 the backbone of the associative memory network, in the computational model presented here, many crucial roles are taken by the dentate gyrus. Some of these roles are dependent on the dentate gyrus' levels of activity. The connectivity parameters that secure the correct bounded values for DG's activity are therefore fundamental. These parameters were estimated using the methods described in chapter 6 with the requirement that a granule cell would not reach threshold with less than two active inputs. This measure is required to increase the robustness to fluctuations in the entorhinal cortex's inputs.

A divergent connection value $C^d = 52$ is chosen for the connectivity between ECII and DG. This means that each ECII neuron unit establishes a connection with 52 granule cell units. Setting the average strength of these synapses in such a way that every

granule cell requires at least 4 active inputs to fire means that, on average, only 1.0% of the whole DG population will be above threshold (see chapter 6 for the mathematics behind this estimate) - which is the required activity level for this field. This value is robust to fluctuations in the synapse's efficacy (peak conductance) and transmission delays, if these properties are assumed to follow normal distributions.

This previous statement can be clarified using the central limit theorem. Let us consider the case of the synapse's peak conductance. Let N_{PSR} be the minimum number of active inputs required to fire a granule cell (in the case described above $N_{PSR} = 4$). The variability in the average S_{peak} has to guarantee that the probability of the granule cell firing with just $N_{PSR} - 1$ active inputs is very small. In other words, the probability that the average S_{peak} in N_{PSR} signals is greater than $S_{thresh}/(N_{PSR} - 1)$ has to be very small. In addition, the probability that the average S_{peak} in N_{PSR} signals is less than $S_{thresh}/(N_{PSR} + 1)$ has to be very small. The standard normal distribution is used with the following two expressions in order to compute the maximum standard deviation allowed for S_{peak}:

$$Z_{up} = \frac{\frac{S_{thresh}}{N_{PSR}-1} - \frac{S_{thresh}}{N_{PSR}}}{\frac{\sigma}{\sqrt{N_{PSR}}}} \tag{8.1}$$

$$Z_{down} = \frac{\frac{S_{thresh}}{N_{PSR}+1} - \frac{S_{thresh}}{N_{PSR}}}{\frac{\sigma}{\sqrt{N_{PSR}}}} \tag{8.2}$$

As an example, for $S_{thresh} = 62.6nS$ and $N_{PSR} = 4$, the probability that a neuron will fire with less than 4 active inputs, or stay quiet with more than 3 active inputs, is less than 0.1% if the synaptic conductance follows a normal distribution with $\mu = S_{peak}$ and $\sigma < 2.42nS$.

It is important to notice that, as the required number of active inputs N_{PSR} increase, the mechanism becomes more robust to the S_{peak} variability in each synapse. The only requirement is that, for large N_{PSR}, the firing probability transition has to be very sharp: the firing probability has to be close to 0 if the average S_{peak} in N_{PSR} signals is just below $S_{thresh}/(N_{PSR})$, and close to 1 if the average S_{peak} in N_{PSR} signals is just above $S_{thresh}/(N_{PSR})$. Using again the central limit theorem, one can see that the fluctuations in the average S_{peak} are proportional to $1/\sqrt{N_{PSR}}$, allowing a considerable variability (σ) in each synapse's S_{peak}.

The same reasoning applies to the transmission delays and to all the other synaptic

properties following normal distributions that affect, through a combined manner, the neuron's state.

Back to the main topic of this section, now one has to discuss the CA3 connectivity in the hippocampal model. The data presented in table 8.2 is used in order to set the CA3's afferent connection properties.

Input	quantal size [pA]	active zones/syn.	number CA3 cells	mean unitary amplitude [pA]	activity [% of EC]
m.f.	9	20	50	70	3
r.c.	5	1	12,000	7	40
p.p.	-	-	4,000	1	100

Table 8.2: Strength of excitatory inputs to CA3. Adapted from Urban *et al*, 2001.

For each input pathway, the average number of active inputs is estimated accordingly to the activity level of the originating area. This estimation assumes that the connectivity profile follows a uniform distribution. There are no topographic preferences (this point is revised in section 8.2). Table 8.3 summarises these biological estimated (average) values.

pathway	$< n.\,inputs >$	$< n.\,active\,units >$
perforant path	4,000	7.0% → 280
recurrent collaterals	12,000	2.5% → 300
mossy fibres	50	0.5% → 0.25

Table 8.3: Afferent connections to CA3 neurons in the rat hippocampus. Average number of inputs and average number of active inputs.

The average number of input connections to each CA3 neuron unit in the model roughly preserves the original relative contributions.

There are nevertheless other reasons for these chosen connectivity values, mainly related to activity level control (chapter 6), which will now be described.

For the mossy fibres, two requirements are imposed: first, the probability of a CA3 neuron unit not receiving any active mossy fibre has to be high[1] and secondly,

---

[1] In the rat hippocampus this probability is around 77% (assuming a binomial distribution with parameters 50 and 0.005 for the number of inputs and probability of input being active, respectively).

pathway	$< n.\, inputs >$	$< n.\, active\, units >$
perforant path	37	$10.0\% \rightarrow 3.7$
recurrent collaterals	80	$5.0\% \rightarrow 4.0$
mossy fibres	16	$1.0\% \rightarrow 0.16$

Table 8.4: Afferent connections to CA3 neuron units in the computational model. Average number of inputs and average number of active inputs.

a single active mossy fibre should not cause spiking (otherwise fluctuations would be easily propagated). With the parameters stated in table 8.1 above, the probability of CA3 neuron unit not receiving any active mossy fibre input is around 70% and the probability of receiving at least two is around 5%.

In the case of the recurrent collaterals, the opposite condition is required, that is, the probability that each CA3 units does not receive any active recurrent collateral has to be almost zero. In the rat hippocampus this probability is less that 10^{-100} while in the model it is around 10^{-2}. The average number of active recurrent collaterals in both systems is 300 and 6 respectively. The information content of each synapse in the computational model is therefore much higher, making it also significantly less robust to fluctuations and noise.

The same argument is also behind the chosen properties for the perforant path. Besides the fact that the probability of each CA3 unit not receiving any active perforant fibre is very small, the average number of active inputs is similar to the average number of active inputs from the recurrent collaterals. Instead of just trying to mimic the biology, this property is required to allow storage and recall in the CA3 network by levelling the importance - in terms of efficacy - of the perforant and recurrent collaterals pathways.

In order to cope with the sensible CA3 activity level control, a population of 20 inhibitory neuron units was connected to the CA3 pyramidal population units. This number of units represents a balance between the typical biological quantity - around 20% of the pyramidal population size - and the single inhibitory unit used in many models. Although a single *omniscient and omnipotent* unit allows an easier control over the whole network activity level, as has been shown in a model like Káli and Dayan (2000), it may disguise the true biological mechanisms behind activity control and competition enhancement. This solution was therefore rejected in favour of an

interneuron population.

The connectivity between CA3 pyramidal and interneuron units follows the properties presented in chapter 6 regarding the stability control in recurrent networks: the inhibitory population is set to fire only when the number of pyramidal neuron units are statistically above proper activity level. An extra component is nevertheless added: instead of receiving inputs solely from the CA3 pyramidal population, the CA3 interneuron units also receive inputs from the mossy fibres, just as in the biological hippocampus. Again, the interneuron population passes threshold when the mossy fibre active inputs are statistically above their proper activity level value, which is around 1.0%. This connection allows a much more efficient control of CA3 activity level since it prevents CA3 activity from raising improperly in the first place.

The figure 8.1 summarises all the connectivity properties discussed above.
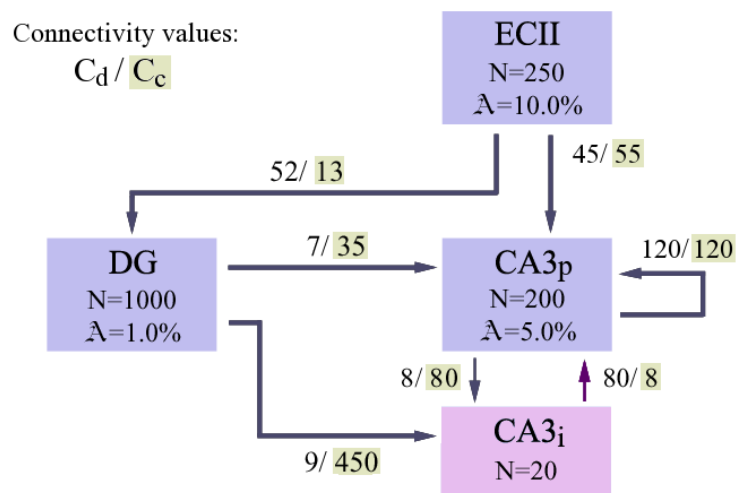
Figure 8.1: Connectivity diagram for the episodic computational model. The convergent connectivity values are average values calculated from the divergent connectivities using equation 6.2

8.2 Pattern separation

In order to increase the efficiency of the hippocampal memory system, it is expected that stored environmental configurations referring to different situations should be associated to completely different internal representations. That is, all stored representa-

tions should have a minimum amount of overlap among them. Minimising the overlap between the CA3 representations, or equivalently maximising their distances (equation 7.1, means that the storage capacity of the system will be significantly increased due to less interference between the representations (Hopfield, 1982).

This problem, know as pattern separation, is solved in the hippocampus by the dentate gyrus (O'Reilly and McClelland, 1994). In the intermediate stage provided by the dentate gyrus, the entorhinal cortex input pattern is projected into a higher dimensional space, with a lower activity level, before producing an activity pattern in CA3 through the sparse connections of the mossy fibres. These two steps greatly reduce the amount of overlap between the CA3 activity patterns.

This mechanism of pattern separation was tested in the hippocampal model by measuring the amount of overlap in the CA3 activity patterns against the amount of overlap in the input patterns of activity placed in the EC. The results are shown in figure 8.2. The overlap in the representations with, and without, the dentate gyrus stage are compared.
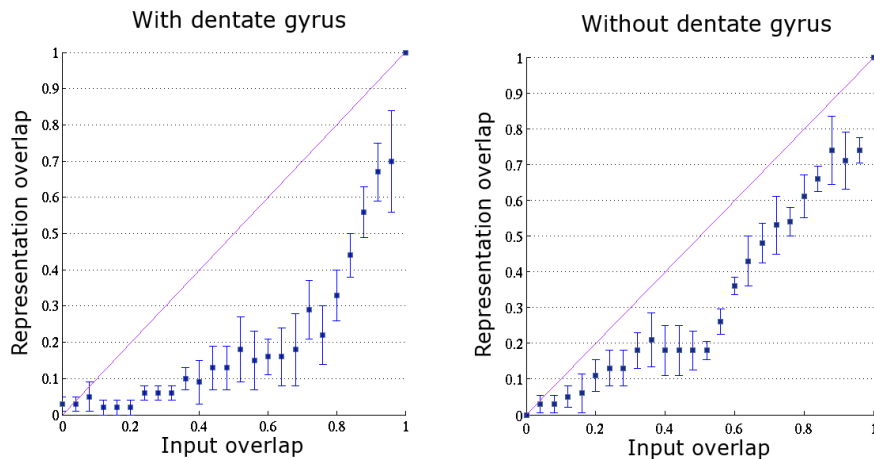
Figure 8.2: Pattern separation. Overlap is measured as the fraction of shared neuron units.

The effect of the dentate gyrus can be clearly seen: even EC input patterns with a substantial amount of overlap produce CA3 activity patterns which share a very small amount of neuron units. In the storage process, the distance between representations associated with similar input patterns can, therefore, be greatly increased.

Although these results are in accordance with previous theoretical studies (e.g

O'Reilly and McClelland, 1994), they are against one property which, I would think, would be desirable in the hippocampal memory system:

The number of degrees of freedom in the input space of the hippocampal memory system is enormous. Even if an environmental configuration is exactly recreated, it is hardly probable that the input to the hippocampus will be exactly the same in the two situations. Internal states, history-dependent mechanisms or simply noise, contribute to this difference. In a memory system with finite storage capacity, it is important to distinguish between noisy versions of a same (known) environmental configuration and a completely new (unknown) configuration. Otherwise redundant information is constantly stored. In addition to this, but related to the finite storage capacity, is the fact that the hippocampal memory system is expected to be capable of generalising: a stored (internal) representation should not be simply a snapshot of all the environmental configuration; it should be a concatenation of features that made that environmental configuration relevant. This way, all environmental configurations which contain the same set of features should be associated with the same internal representation. Due to these two aspects I would expect that very similar EC input patterns should lead to the activation of a similar CA3 constellation. This would allow a more efficient recall of the stored CA3 representation. This way, while maintaining a strong pattern separation for a medium/small degree of overlap, input patterns with a big degree of overlap should be allowed to activate very similar CA3 constellations. It is worth mentioning that this is not the same as pattern completion. This property would nevertheless aid substantially the pattern completion mechanism.

It was found that this different shape in the pattern separation curve can be achieved if the topographical segregation of inputs present in EC and DG (see chapter 2) is considered. The results shown in figure 8.3 were obtained with simulations where there was a topographical arrangement in the connections between EC and DG. Both EC and DG where divided into five sections and each section of the EC established random (uniform) connections only with a paired section in the DG. While testing pattern separation, each EC section always contained the same amount of active units. With this connectivity profile it is very interesting to notice that similar input patterns lead to very similar CA3 constellations. This would allow a more efficient recall, case the input would correspond to a stored memory. This result is even more interesting if we take into account that the segregation on the EC is derived from the convergence

of multi-modal sensory information. Different regions on the EC and DG are mostly driven by information with different nature. This result shows that an episode, described by several features, can still be easily recalled (very similar initial CA3 state) with one or more features changed or absent.
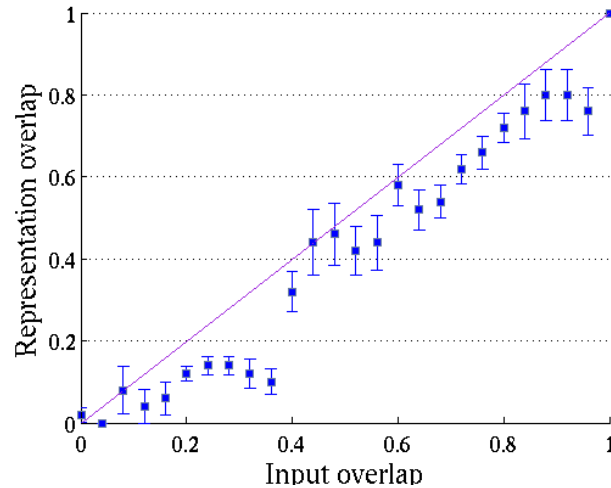
Figure 8.3: Pattern separation in the presence of segregation in the EC and DG. Overlap is measured as the fraction of shared neuron units.

For the parameters used, the transition for strong separation is made at about 40% of overlap. While inputs sharing more than 40% of the units tend to activate CA3 constellations with equal amount of overlap, inputs with less units shared activate CA3 constellations with little or no overlap. The transition is very sharp and its value is dependent on the connectivity profile (topographical properties).

8.3 Feature selection

In the hippocampal model, it has been argued that the short-term plasticity of mossy fibres plays a very important role in the storage process. The overall contribution of mossy fibres for CA3 excitation level is lesser, in relative terms, than the contribution provided by the recurrent collaterals and perforant path. Nevertheless, their great post-synaptic response amplitude, consequence of their short-term dynamics, puts them in a privileged position in leading constellations of CA3 neurons to the threshold level.

As discussed in chapter 5, the post-synaptic response of mossy fibres is greatly enhanced when they are consistently activated. One hypothesis proposed in the hippocampal model is that a constellation of CA3 neurons will only be chosen to form an internal representation, and initiate the storage process, when similar input patterns have been presented at the input level during a short period of time. This way, the CA3 internal representation is associated with the neurons in the input space that have been consistently activated, discarding all the subsets of neurons that fired stochastically or were not relevant in coding the present environmental configuration. The figure 8.4 provides a visual idea of the selection process.

The consistency across time of the activity patterns in the dentate gyrus is a reflection, in a higher dimensional space, of the consistency of the input activity patterns in the entorhinal cortex. I suggest that the time constant of the facilitation mechanism in mossy fibres (2000 ms in these simulations) defines the amplitude for the integration time window, which is in the order of a few seconds (see chapter 5).
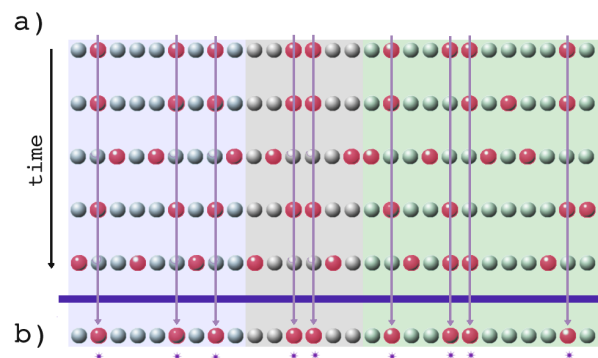
Figure 8.4: Feature selection by the dentate gyrus. In **a)**, a temporal sequence of activity patterns in the dentate gyrus is shown. Circles in red represent active neurons and the three regions with different background colours (light blue, gray and green) illustrates the topographical segregation in the dentate gyrus. The pattern transferred by the mossy fibres to the CA3 region is shown in **b)** and reflect the pattern that has been consistently activated across time. Mossy fibres that have not been systematically activated do not bring CA3 neurons to threshold level.

The reason to call this process feature selection instead of simply pattern selection is based in the topographical segregation in DG and EC. I argue that the activity patterns at the entorhinal cortex, and the correspondent patterns in the dentate gyrus,

should not be seen as a single identity. They are instead a concatenation of information blocks due to the topographical segregation present in these two structures. The subsets of neurons that are consistently activated in a specific environmental configuration should not be seen as encoding simple features, such as a specific isolated sensory feeling. These subsets result from a projection into a higher dimensional space (from EC to DG) and therefore the features they encode will probably represent complex higher order information.

The ability of mossy fibres to perform this selection of features was tested with simulations of the hippocampal model. In a preliminary simulation, the short-term dynamics of mossy fibres were removed and substituted with a static post-synaptic response, with a magnitude similar to the one produced with high frequency stimuli. It was then observed which CA3 constellation was activated in response to a predefined pure pattern presented in EC. Simulations were then run with the mossy fibre short-term dynamics activated and where a sequence of patterns were presented at the input level (EC). The sequence consisted of random patterns alternated with the pure pattern, presented at frequency of 10 Hz. The results are shown in figure 8.5. The CA3 constellations produced by each input in the sequence are compared with the original constellation produced by the pure input pattern.

8.4　Storage and recall of an internal representation

The sequence of functional steps that, according to the hippocampal model, mediate the process of storage and recall of internal representations was extensively tested and analysed using the simulation environment. The simulations shown here demonstrate that all the concepts discussed in the hippocampal model (e.g. management of activity levels, role of mossy fibres, notion of activity patterns) work together harmoniously to form a memory system that mimics, and explains, the functional behaviour of the hippocampus.

Storage

The first results refer to the process of formation of an internal representation in CA3, and its storage in the recurrent collaterals. The experiment performed in the simulation environment consisted of applying repeatedly the same activity pattern in the EC and measuring the propagation of the activity to the DG and CA3. The input
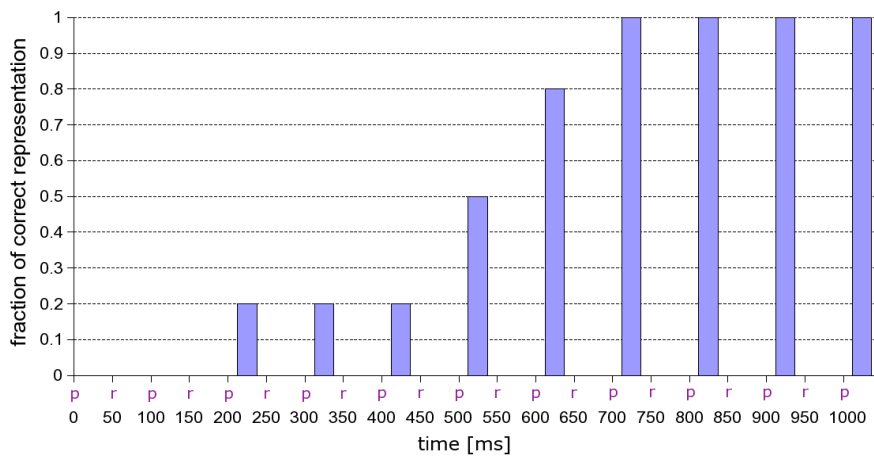
Figure 8.5: Simulation showing the viability of the feature selection process. The horizontal represent the time in ms. The instants when the pure pattern was placed in the EC are marked with a p while the instants when a random pattern was placed instead are marked with a r. The vertical bars measure the fraction of correct units that were activated in response to the input patterns. The random patterns never activated any CA3 neuron. The correct CA3 constellation was invoked after the 8^{th} exposition, after 700 ms. No spurious units were activated.

pattern was applied at a frequency of 10Hz.

The properties of the network recreating EC, DG and CA3 followed the information presented in section 8.2. Both the perforant path and the recurrent collaterals exhibited long-term plasticity (see section 3.2.4). The connectivity properties are shown in figure 8.6. All conductance values were assigned using the methods from chapter 6. For connections exhibiting long-term plasticity the S_{max} values are shown, while for the other connections the value S_{peak} is presented. All these values account for the number of active inputs, N_{PSR}, that each neuron unit requires to reach threshold. The homosynaptic depression was $H = 5\%$. For the perforant path $\Delta LTP = 7.5nS$ and $\Delta LTD = 0.5nS$; for the recurrent collaterals $\Delta LTP = 5.5nS$ and $\Delta LTD = 0.5nS$.

Both ΔLTP were set to very high values: roughly 50% of S_{max}. The reason for that was purely practical: it allowed a much faster learning process thus reducing significantly the simulations time. Nevertheless, the true value should be much smaller so that isolated, stochastic potentiation of spurious units do not affect the performance of the system. This point is discussed in detail with the simulations on the recall process.
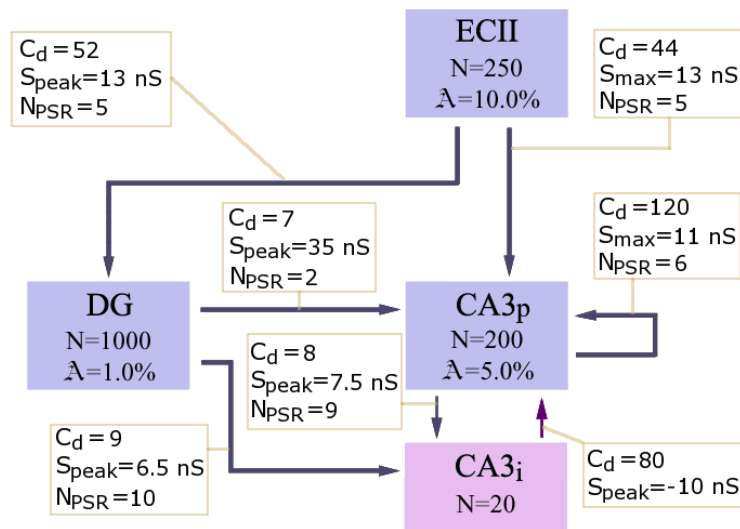
Figure 8.6: Connectivity conductances.

perforant path	recurrent collaterals
$\Delta LTP = 7.5\,nS$	$\Delta LTP = 5.5\,nS$
$\Delta LTD = 0.5\,nS$	$\Delta LTD = 0.5\,nS$
$H = 5.0\,\%$	$H = 5.0\,\%$

Table 8.5: Plasticity in recurrent collaterals and perforant path.

Figure 8.7 shows the time evolution of the system in storing an internal represen-
tation. The memory system had already 10 previously stored representations.

There are four important stages in the storage process:

- **First Stage**. Repeated activation of the same neuron units in the DG, as a result
  of consistent activation in the EC, leads to an enhancement of the mossy fibres
  strength. After the third exposure of the same activity pattern in EC, the facilita-
  tion in the mossy fibres is already sufficient to provide the remaining excitation
  that some CA3 principal neurons require to pass threshold level. Four CA3 neu-
  ron units are activated at around 225 ms (snapshot **A**). These initial units are
  very important because they will become the most informative units in the inter-
  nal representation. The units that were activated in CA3 unleash the long-term
  plasticity dynamics in the perforant path (between the EC's input pattern and
  the CA3 units) and in the recurrent collaterals (among the CA3 activated neuron
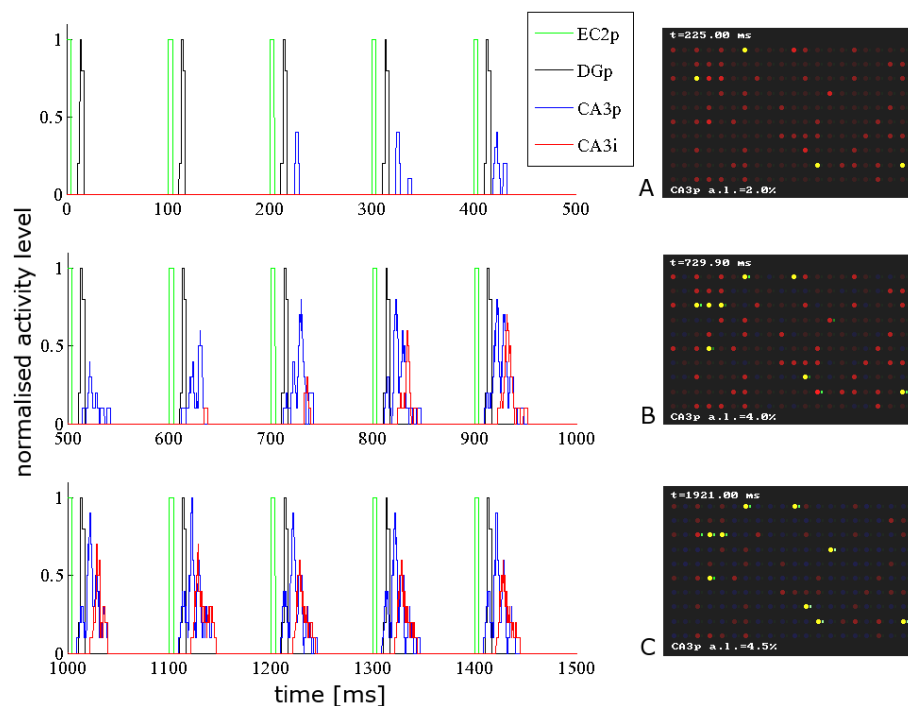
Figure 8.7: Storing a CA3 representation in the hippocampal model: representative sequence of functional steps that mediate the storage of an internal representation. Frames **A**, **B** and **C** contain snapshots of the graphical output of the simulation environment containing the principal population of neuron units from the CA3 field. See text for detailed analysis and description. Parameters from figure 8.6 and table 8.5.

units).

- **Second Stage**. After several exposures of the input activity pattern, the mossy fibres select the complete constellation of CA3 neuron units that will form the internal representation (snapshot **B**). The connections to the entire constellation (ten neuron units) start now to be enhanced (after 700 ms). The inhibitory interneurons in CA3 (which receive connections from the principal neurons in CA3 and DG) start to play a more important role in maintaining the size of the CA3 representation.

- **Third Stage**. The activation time of the CA3 neuron units drifts to earlier times. This is a consequence of the enhancement of the perforant path connections which start to be sufficient to activate the CA3 representation by themselves.

- **Fourth Stage**. Close to 2000 ms, the CA3 representation is activated at roughly the same time as the DG activity pattern. The first CA3 units selected to form the representation have stronger connections and are always the first units to respond to the input pattern. They fire several times and act as promoters in the activation of the CA3 representation. All CA3 neuron units fire more than once.

In light of these results, I suggest that a fifth stage establishes a translation between the CA3 representation and the CA1 activity pattern produced by the perforant path from ECIII (temporoamonic pathway). The hypothesis introduced here is that naive Schaffer collaterals are not enough to activate CA1 neurons. Nevertheless, when the CA1 activity pattern is combined with a CA3 constellation with a statistically relevant size, long-term plasticity mechanisms take place and enhance the Schaffer collaterals connections. The size of the CA3 constellation is only statistically relevant when the CA3 representation is complete (second stage mentioned above). All the CA1 activity patterns that are not paired with activated representations in CA3 do not produce any changes in Schaffer collaterals. In subsequent recall situations, strengthen Schaffer collaterals may dominate over the temporoamonic pathway for the activation of the CA1 representation.

The Schaffer collaterals should be seen, in this hypothesis, as providing a translation code between the CA3 representation and the CA1 representation. The role of the temporoamonic pathway is similar to the mossy fibres in selecting a CA3 constellation. However, since CA1 does not work as an auto-associative but rather as an hetero-associative memory between CA3 and CA1, the requirements for these connections are far different from the ones for the mossy fibres. Briefly, the perforant path from ECIII chooses a CA1 representation and Schaffer collaterals stabilise it.

Simulation results regarding this mechanism are not included here because they do not involve any new computation or mechanism that has not been shown already. They consist mainly of an activity transfer like the one shown in section 6.6.1, figure 6.9. Simulations involving the CA1 field are nevertheless included in chapter 9.

As a final remark to this section, it should be noticed that, in a novel environment, internal representations are developed faster in CA1 than CA3. This property of the HCM is in accordance with experimental results (Leutgeb *et al*, 2004).

Information is only considered to be stored if it is available for future recall. Before

presenting the results regarding the simulations in the context of recalling stored infor-
mation, it is appropriate to discuss the storage/recall dichotomy in the hippocampal
model.

Solving the storage/recall dilemma

In auto-associative networks where the information is stored in the recurrent con-
nections, the recall process should occur without any changes in the connections. A
successful recall is usually composed of a succession of noisy patterns that converge to
the pure, stored pattern and changes that may occur during this process may corrupt the
information related to the stored patterns. Therefore, different mechanisms have been
proposed to commute between the learning phases, where the connections are dynamic
and incorporate the new information, and the recall phases, where the connections are
unchangeable and the stored patterns are retrieved (e.g. Káli and Dayan, 2000).

In this hippocampal model, the work by Hasselmo *et al* (1995; 1996) is taken into
account to solve the storage/recall problem. My hypothesis is that the hippocampus
is, by default, always in recall mode. Nevertheless plastic synapses can suffer small
changes in their efficacies even in recall mode.

Two points are important to keep in mind. Firstly, in the recall process, the activity
levels of all fields fluctuate at low levels and only rise significantly when the activity
patterns involve previously stored constellations of active neurons. It is important to
emphasise that activity levels below the values in table 8.1 do not effectively propagate
activity to subsequent (post-synaptic) fields. Secondly, in order to initiate storage, a
constellation of CA3 neurons, of appropriate size, has to be activated. Only then, the
associative long-term plasticity mechanisms take place.

It has been discussed above that the mossy fibres are required to provide the exci-
tation deficit that the CA3 constellation requires to reach threshold. I now add to this
idea that the facilitated mossy fibres will only be able to provide the excitation deficit
when the excitation level in the CA3 region is *uniformly* raised. This uniform excita-
tion is provided by the medial septum which modulates the learning dynamics in the
hippocampus (Hasselmo *et al*, 1995; 1996). The septum rises the excitation level in the
hippocampus on specific emotional contexts, such as novelty and arousal (see section
8.1). The diagram in figure 8.8 provides a visual idea of the described mechanism.

The high specialisation, in functional terms, of mossy fibres can be further appre-
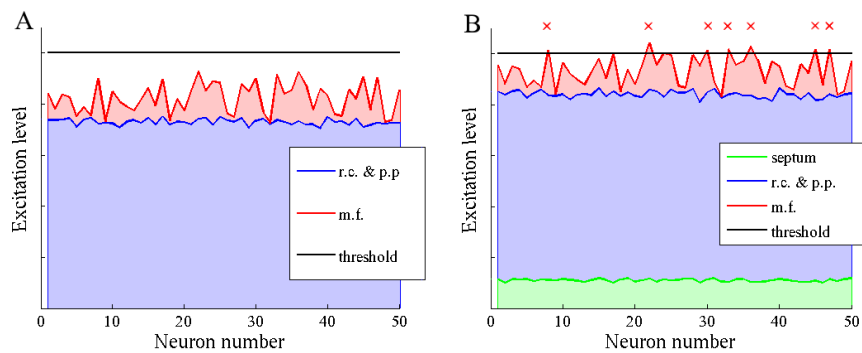
Figure 8.8: Excitation provided to CA3 neurons. Images **A** and **B** provide an idea on how each input connection contributes to the excitation level of 50 CA3 neurons. In order to initiate the storage process with the selection of a CA3 representation, the mossy fibres require the uniform excitation provided by the septum (**B**). Only then a constellation of CA3 neurons reach threshold (marked with red "X").

ciated if we analyse the distribution of voltages in CA3 neurons during the selection of an internal representation. The idea depicted in 8.8-**B** is only possible because the connectivity properties from the mossy fibres guarantee that there is a big voltage gap between the neurons that were activated and the neurons that did not reach threshold. Figure 8.9 shows the distribution of voltages in CA3 neurons when a full constellation is chosen to form an internal representation. With 5% of the units above threshold (-50 mV), the average voltage of the sub-threshold units is around -65 mV.

If, instead, the perforant path would be responsible for choosing the constellation of CA3 neurons (to form the internal representation), many neurons would be very close to threshold level. This means that the process would be much less robust to fluctuations: any perturbation could generate oversized constellations that would corrupt the memory system. The simulation results shown in figure 8.10 corroborate this statement.

Going back to the discussion on the storage/recall dichotomy, it can be shown that the external modulation provided by the medial septum is enough to prevent the corruption of previously stored information. Changes on the recurrent collaterals and perforant path *only* take place when CA3 neurons fire, due to the associative nature of the plasticity dynamics. In turn, CA3 neurons only fire in the following two conditions:

1. When the medial septum raises the excitation level of CA3 and mossy fibres

Figure 8.9: Voltage distribution in CA3 neurons when the mossy fibres create the internal representation. Marked in red is the fraction of the population above threshold (activity level). Average sub-threshold voltage is -65.2 mV. Parameters from figure 8.6 and table 8.5.

detect consistency across time in the input activity patterns (features). In this case the recurrent collaterals and perforant path will be storing new information.

2. When the EC activity pattern resembles a stored episode and the strong perforant path connections, enhanced through learning, are enough to activate fully, or partially, the internal representation in CA3. In this case, the recurrent collaterals and perforant path will neither be storing new information nor storing noise (corrupting the memory system). They will simply be reinforcing the connections associated with the internal representation.

Neuron units in the CA3 field that erroneously may fire (stochastically or due to improbable combination of inputs) will indeed lead incorrectly to the potentiation of some connections. Nevertheless, these same connections will be depressed shortly after, as a result of homosynaptic depression (see section 3.2.4 for the synaptic plasticity dynamics). For this balance to be possible, it is important that the changes produced by each potentiation are small. This way, only consistent activation of the same neurons will produce a relevant rise in synaptic efficacy. This point is not against the idea of fast learning in the hippocampus: if 100 combined, and consistent, activations are required to produce an effective potentiation, and if storage occurs at theta frequency (10 Hz),
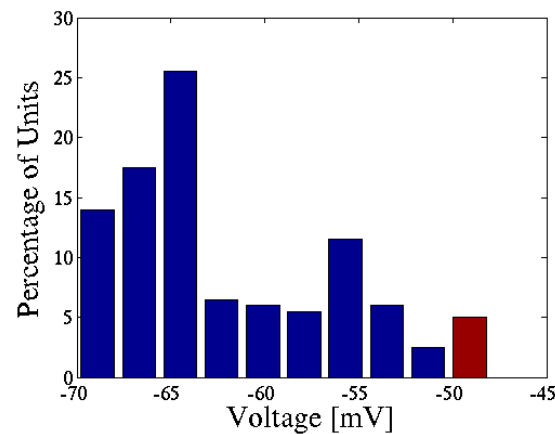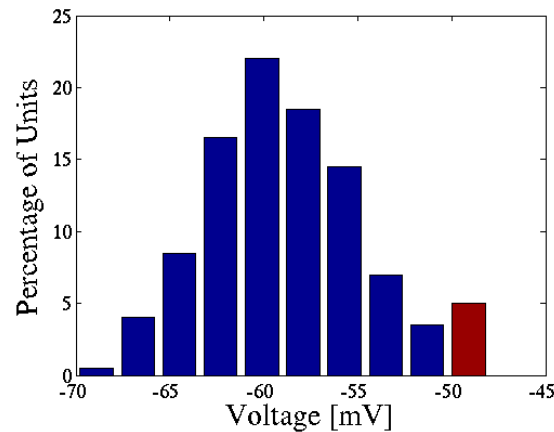
Figure 8.10: Voltage distribution in CA3 neurons when the perforant path chooses the internal representation. Marked in red is the fraction of the population above threshold (activity level). Average subthreshold voltage is -61.9 mV. Parameters from figure 8.6 and table 8.5.

then 10 seconds are enough to complete and stabilise a CA3 internal representation.

Although the idea of learning regulation provided by the medial septum is inspired in Hasselmo *et al* (1995; 1996), it should be emphasised that the mechanisms by which the storage/recall problem is solved are specific to the hippocampal model of this thesis. One point of particular interest is that localised modulation on particular pathways is not required: uniform, distributed excitation is sufficient to create the conditions that solve the storage/recall problem. Such precise modulations would nevertheless add robustness to the system.

Recall

The results from the simulations testing the recall mechanism in the hippocampal computational model will now be presented. The simulations followed the approach taken in the simulations for storage process, and the results describe the time evolution of the activity in all relevant fields (figure 8.11).

In the recall simulations the constants which set the synaptic potentiation changes were replaced by the more realistic values: for the perforant path $\Delta LTP = 0.5nS$ and $\Delta LTD = 0.5nS$; for the recurrent collaterals $\Delta LTP = 0.5nS$ and $\Delta LTD = 0.5nS$. This way any potentiation in a spurious unit would be balanced by depression, and the correct units would be slightly reinforced. Storage of a new representation with this values

would just mean that many consistent and systematic enhancement changes would be required. With this low value for ΔLTP, storage and recall can be interchanged without affecting the memory system performance.
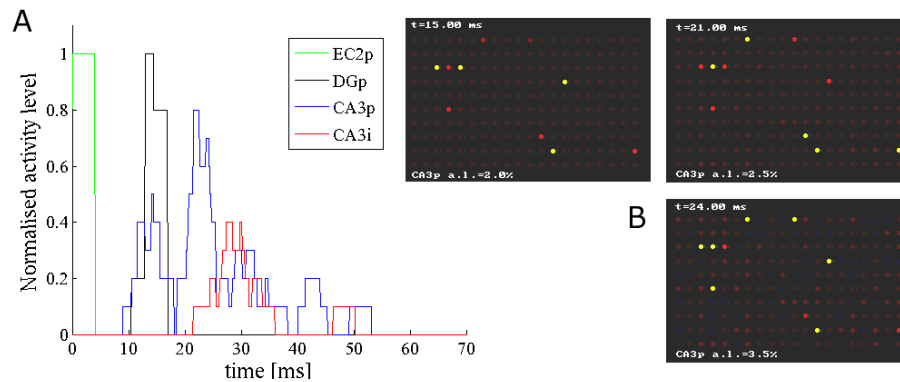
Figure 8.11: Recalling a CA3 representation in the hippocampal model. **A**: time evolution of the activity levels in the principal populations of ECII, DG and CA3, and in the inhibitory population of CA3. **B**: snapshots of the graphical output of the simulation environment at several times. The entire representation is activated within a time window of less than 10 ms. No spurious neuron units are activated. Parameters from figure 8.6.

A very important point to mention is that, no impairment is produced in the recall process in simulations where the mossy fibre connections are removed. This is in accordance with the experimental results in Lassalle *et al* (2000). In fact, even in the simulations where the mossy fibres are present, they are almost silent in response to single input activity patterns that are inconsistent with the previously presented patterns.

The next step in the recall process would be the rise of the activity level in the CA1 field, with the activation of a previously learned constellation of CA1 neuron units. The entire recall process, from exposure of input activity pattern in the EC to the activation of the CA1 constellation, takes less than 50 ms.

Figure 8.12 shows the results of the hippocampal model recall process in the situation where, instead of a pure pattern, a noisy version of a pure pattern is presented at the EC region. Two types of noisy patterns were used:

- partial cues, where some units of the original pure pattern are removed from the input activity pattern;

- corrupted cues, where some units of the original pure pattern are substituted by other, random, units in the input activity pattern;

The fraction of units removed or substituted in the pure pattern define the error of the noisy pattern. For example, a partial cue with 20% noise contains only 80% of the units in the pure activity pattern.
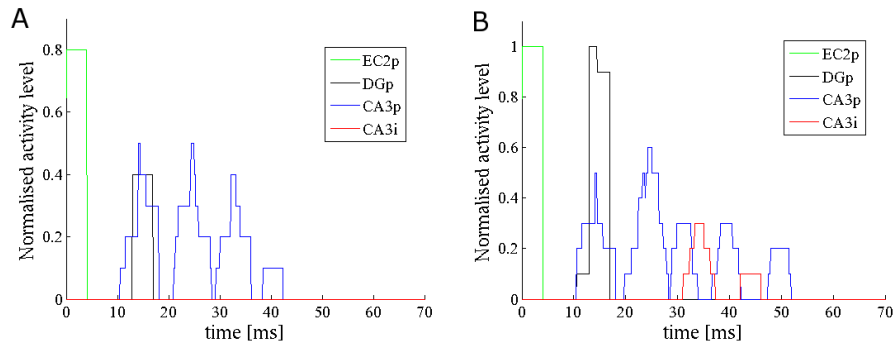
Figure 8.12: Recalling a CA3 representation in the hippocampal model with noisy inputs. Both noisy inputs led to the activation (exclusively) of all the neuron units of the CA3 representation. **A**: Recall with a partial cue with 20% noise. **B**: Recall with a corrupted cue with 20% noise. Parameters from figure 8.6.

Although all the units from the representation are at some point activated, the recall process is slower, i.e., few units are activated on each cycle of the recurrent connections. The representation is therefore completed with more effort and most units fire only once. Nevertheless, for errors under 20% the memory system still performs a perfect recall. On the other hand, for errors above the 20% level, the activated CA3 constellations start becoming corrupted compared with the original stored representation. This 20% level sets the radius of the basins of attraction, the regions in the state space where the system dynamics are able to recall the correct internal representation. Distinct situations occur when the errors are bigger than 20%, i.e. outside the basins of attraction radius, depending on the nature of the input noise:

- For partial cues, the excitation provided to the neuron units of the CA3 representation diminishes and units start to fail reaching threshold. Although the recurrent connections have the role of rebuilding the complete representation, it is important to remember that their maximum peak conductance is assigned

so that a minimal number of active inputs are required to reach threshold level. Input excitation below a critical value fails to initiate the proper feedback amplification and leads to a decay in the activity level, before the complete internal representation is invoked. For partial cues with 40% noise, 70-90% of the internal representation is activated, and for partial cues with more than 60% noise no units are activated in CA3. If one would consider k-winners-take-all mechanisms, correct recalls would be possible at a very high degree of noise.

- For corrupted cues with a high degree of noise, spurious units are activated in the CA3 field. The activated CA3 constellation is completely corrupted for input noises above 30%. This should not be seen as a limitation though, since this level of "corruption" at the EC input pattern would most probably represent a different environmental configuration.

8.5   Firing properties

All the simulation results shown so far refer to the collective behaviour of each hippocampal field. However, the firing characteristics of the neuron units are important for a better understanding of the dynamics of the hippocampal memory system. The voltage profile of all individual neurons involved in the recall process depicted in figure 8.11-**A** are shown in figure 8.13. Although detailed information about each individual voltage profile is obscured due to the high density of traces, the contrast between silent and active neuron units can be appreciated.

Isolated individual traces are presented in 8.14. These traces refer to the simulation that led to figure 8.11. The EC input pattern was presented at a frequency of 10 Hz for 0.5 seconds. The absence of variability in the input pattern during these 0.5 seconds assumes that no relevant changes in the environmental configuration occurred in this period.

It is important to emphasise that although these traces are for one particular recall process, they are representative of the firing characteristics of active CA3 neurons in any recall sequence. Some important observations can be made from the analysis of many traces:

- Most neurons involved in the internal representation fire more that once in re-
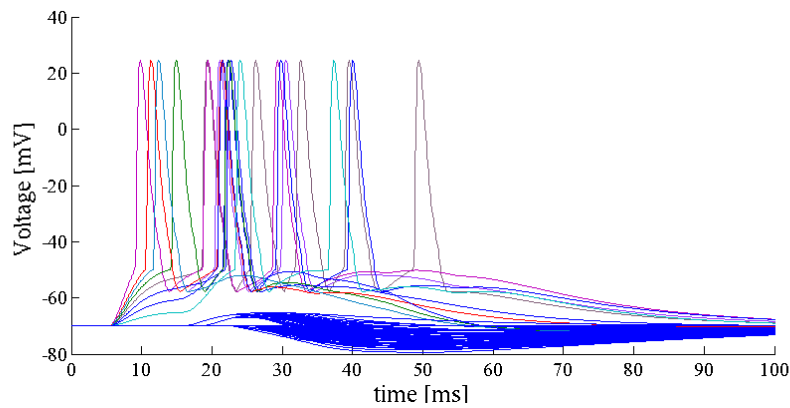
Figure 8.13: Voltage profile of the CA3 population during recall. Parameters from figure 8.6.

sponse to the presentation of a known input activity pattern. The average number of spikes per burst[2] is 2.6 and very rarely a neuron unit contributes with a single spike in the activation of the stored representation (figure 8.14-**H**).

- The response of some neurons is increased after repeated exposures of the same input pattern (figure 8.14-**A**, **E** and **G**). Although long-term dynamics are active, this result is not an effect of learning. What it reflects is change in the state of the network. The membrane time constant and the post-synaptic conductance profiles put the active neuron units at a higher excitation level at the time a second input pattern is presented.

- Neurons that fire more times contribute more actively to the formation of the CA1 activity pattern and are more robust in the situation of recall with partial cues. Their increased amount of input excitation results from a slightly bigger number of inputs connections (see section 6.2 for the connections distribution).

- Some spiking bursts show accommodation (figure 8.14-**B** and **I**). This is not the result of frequency adapting channels since they were simply not included in the neuron model. This effect is created by threshold input currents that decay slowly instead of abruptly.

---

[2]The number of spikes per burst depends on the connectivity parameters that change the excitation levels. This number is, nevertheless, not controllable since it is determined mostly by the conductance values that are constrained by the activity level control.
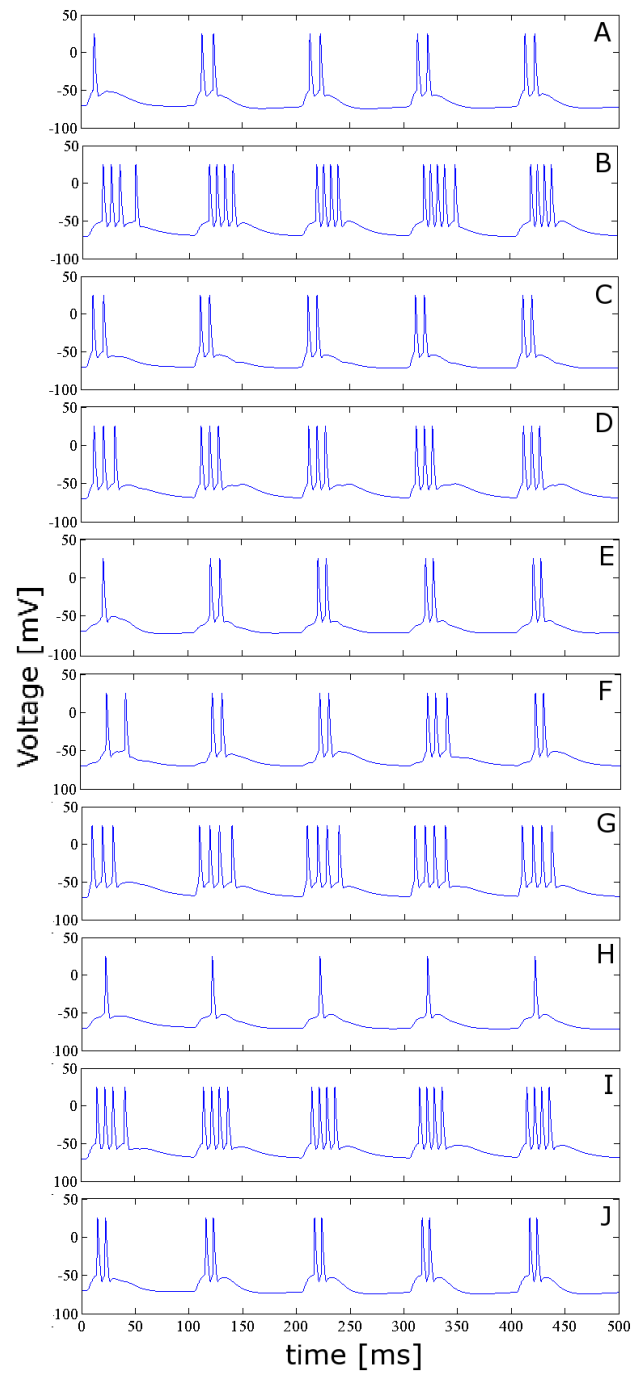
Figure 8.14: Representative voltage profiles of CA3 active neurons during recall. The average number of spikes during each activation of the internal representation is 2.6. Parameters from figure 8.6.

These firing properties will form the core of the arguments that merges episodic memory model with spatial memory in a complete hippocampal model accounting for both experimental paradigms (see chapter 9).

8.6   Hippocampal rhythms

Motivated by observations on many simulations performed with the hippocampal computational model, I include here a brief discussion on my view of the functional roles behind the theta and gamma rhythms in the hippocampal subfields. Two points are addressed:

1. In this hippocampal theory, information is coded, transferred and manipulated in the hippocampal regions in the form of activity patterns. In order to accomplish their tasks, these constellations of active neurons forming the patterns have to be synchronised within a time window of a few tenths of milliseconds, i.e., the neuron's integration time given by τ_m. If the firing time of individual neurons drifts too much, it will corrupt the following, or preceding, activity pattern. In addition, if the activation of the complete constellation is too spread in time it will fail to activate its post-synaptic neurons. Mechanisms that constrain the time frame at which all the neurons in the constellation have to fire are, therefore, of extreme importance.

2. In addition to what has been discussed in section 8.4, the particular time course of the excitation level modulation provided by the medial septum has strong implications in the effectiveness of the storage and recall processes.

If the hypothesis that information is coded as activity patterns is correct, then the firing of all neurons in an active constellation has to be concentrated in time window of τ_m (point **1** above). Segregation between subsequent activity patterns therefore means that there is a periodicity of τ_m between activity peaks. The peaks of activity would result from the higher density of spikes in the "centre" of the time window[3]. With the biologically reasonable value of $\tau_m = 20ms$ used throughout this thesis, the frequency

---

[3]That is, is it reasonable to assume that the distribution of spikes, in the integration time window, would be peaked.

of the segregation mechanism would be of 50 Hz. This is in the range of the measured gamma frequencies (Bragin *et al*, 1995; Buzsáki *et al*, 1983).

With respect to point **2** above, and following the ideas of Buzsáki (1989), it is reasonable to assume that the up-state of the theta modulation is associated with a storage phase. This mechanism would nevertheless not substitute, but rather complement, the mechanism discussed with figure 8.8. The two combined effects means that, in novel environments, storage occurs at theta frequency (see Larson *et al*, 1986). This would be important to avoid interference between activity pattern containing different information. The theta period corresponds roughly to the estimated time that each activity pattern takes to propagate through all fields of the hippocampal memory system.

8.7   Memory storage capacity

A detailed analytical study of the capacity of the episodic memory model was not performed. This novel hippocampal computational model relies on spiking units with complex dynamics and continuous synapses with dual-exponential post-synaptic responses. A theoretical study of its capacity would, by itself, be a theme for a thesis. Some empirical analysis and observations were nevertheless produced and corroborated with results extracted from simulations. Previous capacity studies using neuron units of different natures, such as binary (Willshaw *et al*, 1969; Marr, 1971; Willshaw and Buckingham, 1990) or threshold-linear (Treves and Rolls, 1991), were taken into account in order to provide some insights to the limitations of this novel CA3 spiking memory model.

The number of internal representations that can be stored in the CA3 recurrent network depends on the network's top activity level, $\mathcal{A}$, which defines the size of the internal representations. In the hippocampal model $\mathcal{A}_{CA3}$ was set to 5%, meaning that in a network of 200 units, the size of the representations was 10. If in the rat hippocampus the activity level is roughly 2.5% (table 8.1), then the representation size is on the order of 4000 neurons (in a population of 160000 principal neurons). This means that each episodic memory would be identified by a constellation of 4000 active neurons. The number of different constellations resulting from combinations of 4000 active neurons is colossal: 10^{8121}. Obviously the network is incapable of distinguishing all these different constellations!

After a constellation of $N\mathcal{A}$ neurons is selected to form an internal representation, long-term plasticity takes place and connections between members of the constellation are potentiated. Connections between constellation neurons and external neurons are slightly depressed as a result of homosynaptic depression. As a consequence of the dynamics of the storage process, the distribution of the peak conductances (the synaptic efficacies) is strongly bimodal: synaptic peak conductances tend to concentrate either around zero or very close to their maximum allowed value, S_{max}.

In this hippocampal theory, the capacity of the CA3 memory system depends directly on the choice of S_{max}. The efficacy of a potentiated synapse is related to how informative each presynaptic active neuron is. In turn, the amount of information provided by an active neuron is related to the average number of representations each neuron participates. If R is the number of stored representations, every neuron participates, on average, in $R\mathcal{A}$ representations. In the hippocampal model, the neuron constellations that are selected to form the internal representations are chosen randomly by the dentate gyrus. The participation of CA3 neurons in the internal representations is, therefore, expected to be uniform.

In order to produce a correct recall, a fraction of the constellation which forms the internal representation should be enough to lead to the activation of the remaining neurons, while avoiding the activation spurious neurons. The competition provided by the interneurons, through the activity level control, is of extreme importance for this process.

Let us call F_{min} the minimum fraction capable of activating the whole representation. Given the activity level $\mathcal{A}$ and the connectivity C^c, the minimum number of post-synaptic responses that should lead to threshold is $C^c\mathcal{A}F_{min}$. The synaptic efficacies are parameterised according to this value. The relation is:

$$S_{peak} \approx S_{max} \geq \frac{S_{thresh}}{C^c\,\mathcal{A}\,F_{min}} \qquad (8.3)$$

As a reminder, S_{thresh} is the conductance that leads to the threshold level (see section 6.3).

In a recurrent network with connectivity C^c, the average number of active inputs that a neuron receives is $C^c\,\mathcal{A}\,p$, where p is the proportion of synapses with $S_{peak} \approx S_{max}$. If only a fraction F_{min} of the representation is activated, the average number of active inputs that a neuron receives, drops to $C^c\,\mathcal{A}\,p\,F_{min}$.

Although the efficacy has to be sufficiently big to allow a proper recall with partial constellations it should also be small enough to avoid the activation of spurious neurons. The process of setting S_{max} is equivalent to setting the detection threshold in Willshaw *et al* (1969).

It is important to estimate the maximum number of stored representations that can coexist without interfering with each other. From interference results spurious activations and, consequently, corrupted recall. Let us assume that, after R representations have been engraved in the recurrent connections, the fraction of connections with $S_{peak} \approx S_{max}$ is p. The fraction of synapses with $S_{peak} \approx 0$ is then $1 - p$. On average, the number of synapses which peak conductance $S_{peak} \approx S_{max}$ is $N C^d \mathcal{A}^2$. Since there are $N C^d$ synapses in the spiking network, the fraction of synapses enhanced during the storage of each representation is $\mathcal{A}^2$. Therefore, the fraction of synapses which still have $S_{peak} \approx 0$ after the storage of R representations is:

$$\left(1 - \mathcal{A}^2\right)^R = 1 - p \qquad (8.4)$$

This equation can be approximated to $R \mathcal{A}^2 = -\log(1 - p)$.

In order to avoid the activation of spurious neurons during the recall process, it is important that only the neurons belonging to the representation receive, at least, $C \mathcal{A} F_{min}$ active inputs. The number of active inputs, X, follows a binomial distribution with parameters $N \mathcal{A}$ (repetitions) and $p C^d / N$ (probability of success). The probability of a spurious neuron reaching threshold is represented by $P(X \geq C \mathcal{A} F_{min})$.

The central point in this reasoning is that the expected number of spurious neurons activated should be as small as possible. On the limit, one can allow the activation of one spurious neuron (zero activated spurious neurons implies $p = 0$). This leads to the following equation:

$$N (1 - \mathcal{A}) P (X \geq C^c \mathcal{A} F_{min}) = 1 \qquad (8.5)$$

The probability p, representing the fraction of synapses with $S_{peak} \approx S_{max}$ after R representations are stored, can be calculated using equation 8.5. With this limiting probability calculated, equation 8.4 provides an estimate to the number of representations R that can be stored before spurious neurons start to become activated. The number of representations R define the capacity of the memory system.

For the CA3 spiking memory used in the hippocampal model simulations, where $N = 200$, $C^c = C^d = 120$, $\mathcal{A} = 0.05$ and $S_{max} \approx S_{thresh}/5$ (that is, 5 combined input were required to fire a CA3 neuron), the estimated capacity is $R = 97$. This theoretical result is corroborated by the simulation results shown in figure 8.15. In order to test the capacity in the model of CA3 subfield, several simulations were performed. Each simulation consisted on the following steps:

1. A sequence of R representations with a size of $N \mathcal{A}$ neurons were stored in the CA3 spiking model;

2. A sequence of R noisy versions of the stored representations were placed in the network and the recall of each representation was tested. The activity level was constrained to 5%.

While in discrete time models the recall performance can be tested easily by calculating, for example, the dot product between the stored representation and the recalled constellation, the different nature of this spiking model forced the use of more elaborated methods. The solution adopted was to create a population of CA1 classification neurons. During the storage process, each representation was associated with one, and only one, CA1 classification neuron. During the recall process, the performance was assessed by counting how many partial representations suffered enough reconstruction to lead to the activation of their CA1 classification neuron. This solution solves the problem associated with the continuum nature of the activity patterns and is in agreement with the roles of the hippocampal theory for the CA1 subfield. The only difference is that, in the hippocampal theory, a constellation of CA1 neurons are associated with the CA3 representation, instead of just one neuron.

In simulations used to produced figure 8.15, partial cues with a noise of 30% were used.

From the analysis of figure 8.15 one can see that when the network stores more than 100 representations, its performance decreases sharply. Performance is measured here as the ratio between the number of correct recalls and the number of representations stored. This simulation result is in very close agreement with the theoretical value of 97 representations, calculated with equation 8.5. This agreement increases the confidence in the assumptions used to produce the theoretical estimate for the spiking network storage capacity.
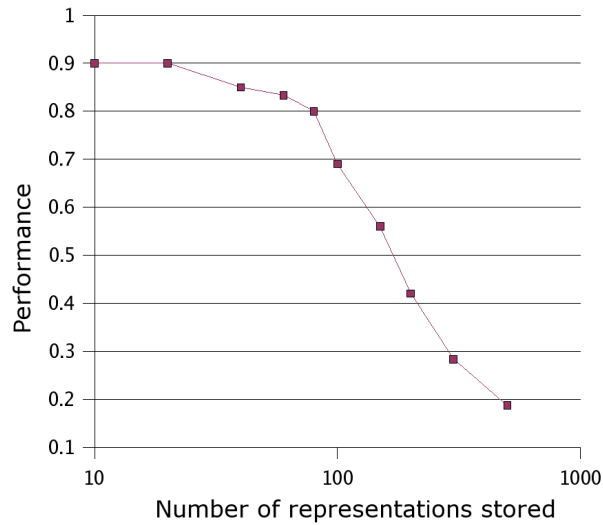
Figure 8.15: Capacity of the CA3 spiking network. Performance is measured as the ratio between the number of correct recalls and the number of representations stored. Parameters $N_{CA3} = 200$, $C^c(CA3) = C^d(CA3) = 120$, $\mathcal{A}_{CA3} = 0.05$ and $S_{max} = 17.0nS$ for the recurrent collaterals.

The simulation results can nevertheless be shown in another format which emphasises a very important property of this model. While many other memory models lose all recall capability when the capacity limit is exceeded, the CA3 network of this hippocampal model maintains the ability to store and recall new representations by erasing less invoked ones. In other words, independently of the number of representations that have been stored, this CA3 spiking model can always recall a number of representations defined by its capacity. This is shown in figure 8.16 where the maximum number of patterns that the network can have simultaneously stored saturates close to 100. This palimpsest property of the CA3 memory system results from the homosynaptic depression dynamics.

In the simulations, the representations that tended to be correctly recalled were typically the most recent. Nevertheless, "old" memories which, by chance, had more or/and stronger connections were also correctly recalled. Since in this hippocampal theory the recall process also produces changes (although small) in the synaptic plasticity, a single recall of a fainting representation is sufficient to reinforce it, making it more robust to the storage of new representations. It is important to notice that, due
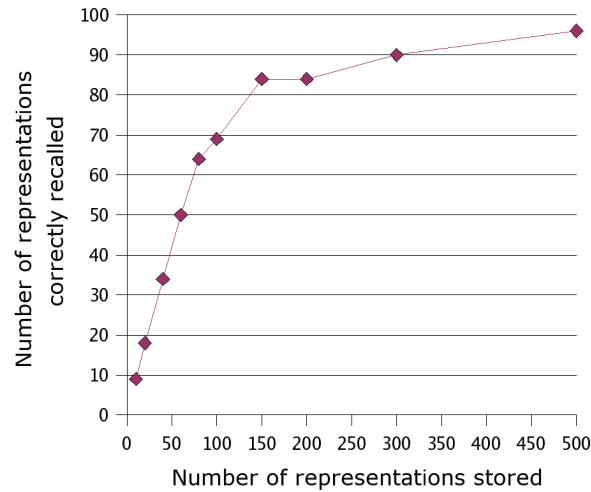
Figure 8.16: The CA3 memory system has palimpsest properties. Independently of the number of representations that are stored, the CA3 spiking model can always recall a number of representations defined by its capacity. Parameters $N_{CA3} = 200$, $C^c(CA3) = C^d(CA3) = 120$, $\mathcal{A}_{CA3} = 0.05$ and $S_{max} = 17.0nS$ for the recurrent collaterals.

to the plasticity dynamics, a representation is only faded when a new representation is stored. In other words, fading is a not a function of time; it depends on the amount of new representations stored.

This novel hippocampal model introduced then a new definition of capacity: instead of being the number of representations that can be stored before the memory system's performance collapses, the capacity of a system is the maximum number of representations that it can be simultaneously stored in the memory system.

Finally in this section, an estimate for the memory capacity of the CA3 subfield in the rat is calculated. In order to do that, it is required to know the number of active inputs necessary for threshold. Following experimental data in which the granule cells are said to fire with few hundred combined excitatory inputs (Patton and McNaughton, 1995), I speculate that CA3 requires 300 post-synaptic responses from r.c. to reach threshold. Using $N = 160000$, $C^c = C^d = 12000$ and $\mathcal{A} = 0.025$ it is possible to calculate that:

- The fraction of synapses with $S_{peak} \approx S_{max}$ is 0.78.

- The capacity is 2400.

Unfortunately the estimated value for the fraction of synapses with $S_{peak} \approx S_{max}$ cannot be tested directly experimentally. In the rat CA3, synapses are created and eliminated according to the information processing requirements. If a synapse is always silent, $S_{peak} \approx 0$, it is expected to be removed. The fraction of synapses with $S_{peak} \approx S_{max}$ in the rat CA3 is therefore, not the same as the fraction in the model were the number of synapses is fixed

8.8 Final remarks

Albeit not specific to the episodic memory model, there are some points worth some final remarks. The first is about the activity level control provided by the CA3 inhibitory population. In the simulations with the hippocampal model it has been found that the precision of the inhibitory control mechanism is degraded if multiple sources are converged to the same population. This is the case of the CA3 inhibitory population which controls the activity level in CA3 in accordance to the activity level in the DG and in the CA3. In these conditions, the inhibitory populations tend to produce lower activity levels due to undifferentiated integration of sources. The solution is to use different inhibitory populations to control the different sources of excitation. This result may be one of the reasons behind the high number of distinct inhibitory interneuron populations that are present in each hippocampal subfield.

It is also worth mentioning that the episodic memory model could have been built on top of k-winners-take-all mechanisms (O'Reilly and McClelland, 1994). Although leading to a complete control on the activity levels, it would have made obscure many interesting points (e.g. temporal computations performed by the mossy fibres, time evolution of storage and recall processes). However, in capacity analysis it is important to eliminate the big fluctuations in activity levels (and constellations sizes) that originate from scaling down the memory system. It was therefore worth to define and parameterise the inhibitory control mechanisms based on spiking neurons. More and more attention should be devoted to the study of the inhibitory populations which seem to be behind many interesting and complex computations.

Chapter 9

Spatial memory

The existence of hippocampal cells that fire in complex bursts when the animal moves through a specific location in an environment were shown by O'Keefe and Dostrovsky (1971) and O'Keefe (1976). These special cells are called place cells and their regions of preference are called place fields. Inside its place field, a place cell may fire at a rate of 20 Hz or more, whereas outside its field, it may fire at a rate of 0.1 Hz, or less.

The existence of place cells offers strong support for the hippocampus' involvement in spatial mapping (O'Keefe and Nadel, 1978) in which a sufficient number of place cells and their respective fields are able to cover and map any given environment. The view presented in this thesis, however, is that place cells are a property, and a consequence, of the more general episodic memory function of the hippocampus. This idea follows Eichenbaum *et al* (1999) and Wood *et al* (1999) where the memory function of the rodent hippocampus is seen as a "memory space" instead of a "spatial memory". This general memory space is capable of encoding different types of information that can occur as "episodes", i.e. multi-modal sensory experiences circumscribed in time.

This chapter shows that the hippocampal model presented in chapter 7, which was developed to explain the functional properties of the hippocampus regarding episodic memory, recreates, alone, key results from spatial learning experiments with rats. Although the results in this chapter are preliminary, they pave the way to a more detailed analysis of spatial memory in the hippocampal theory. The results from the preliminary simulations presented in this chapter provide a high degree of confidence in the ability of the hippocampal model to incorporate the results from spatial memory ex-

perimental paradigms. This is an extremely important step forward against the view of the rat hippocampus as a specialised structure for spatial memory.

Organisation of the Chapter

The first section states which are the key properties that define a place cell and a place field. Section 9.2 then shows that cells taking part in internal representations of episodic-type memories exhibit the properties of place cells. Simulations of the hippocampal computational model are used to corroborate this statement. The notion of "activation fields" is then introduced to substitute the term "place field". Section 9.3 addresses the hypothesis that the nature of the activation fields of entorhinal, perirhinal and postrhinal cortices, reported in the literature, are in fact different from the hippocampal activation fields. Some final remarks are presented in section 9.4.

9.1   Properties that define place cells

Hippocampal place cells fire in complex bursts when the animal moves through a specific location in an environment. The complex, or patterned, spike bursts are comprised of anywhere from 2 to 10 spikes, with the amplitude of successive spikes typically, but not always, decaying. The "intra-burst" frequency ranges from 200 to 250 Hz (4 to 5 ms period) while the "inter-burst" frequency is typically in the range between 0.1 and 1 Hz (seconds to thousands of milliseconds) (Ranck, 1973; Fox and Ranck, 1975; 1981).

Place cells are not limited to the hippocampus though. Already at the input level to the hippocampus, in the medial entorhinal cortex, there are cells showing space selectivity. The spatial firing properties are heterogeneous across the bands of the medial entorhinal cortex (Fyhn *et al*, 2004):

- dorsolateral band - cells show multi-peaked fields (median number of 4) with intervening silent areas, i.e. the fields are clearly delimited against the intervening background; the multi-peaked firing fields are completely dispersed on the spatial environment.

- intermediate band - most cells are spatially modulated but their place fields are broader and less coherent; the place fields contain less peaks (median 1.75) and lacked the characteristic silent areas observed in the dorsolateral band.

- ventromedial band - only very weak spatial modulation is apparent.

This heterogeneity is projected topographically to the hippocampus with a dorsolateral-to-ventromedial axis from the medial entorhinal cortex corresponding to the dorsal-to-ventral axis in the hippocampus (Dolorfo and Amaral, 1998; Fyhn *et al*, 2004). In the dorsal hippocampus, cells in CA3 and CA1 show spatial modulation which is sharper than the spatial selectivity found in EC cells. The number of subfields (local activity peaks) in the hippocampal cells is also smaller than the multi-peaked fields from the dorsolateral band of the medial entorhinal cortex (Fyhn *et al*, 2004).

Figure 9.1 shows what has become the standardised way of presenting place fields. In this style, developed by Muller *et al* (1987), a firing rate map of a place cell is plotted on top of the Cartesian map of the environment. The firing rate map is produced by creating a grid in the environment and counting, for each sector of the grid, the number of spikes fired by the place cell and the amount of time spent. A firing rate for each sector is computed with these values. The firing rate is colour coded with each successive darker colour being represented by 80% fewer pixels than the preceding lighter colour. This way, the "Muller" rate map depicts changes in firing rate tied to the area of the field and not directly to the firing rate itself. A place cell may have a firing rate of 20 Hz, or above, within its place field, and have a firing rate of 0.1 Hz or less, outside. For a more detailed view on the complexities of analysing place cells see Muller *et al* (1987).

Beyond the defining criteria regarding the firing rate map, place cells and their place fields have a set of interesting properties from which the most relevant are:

- Stability

- Responses to environmental manipulations

- Relations to theta rhythm

All these properties will now be discussed in detail.

9.1.1 Place field stability

Place fields are created during exploration of a new environment. The time for formation of place fields upon entry into a novel environment is of the order of 10 minutes
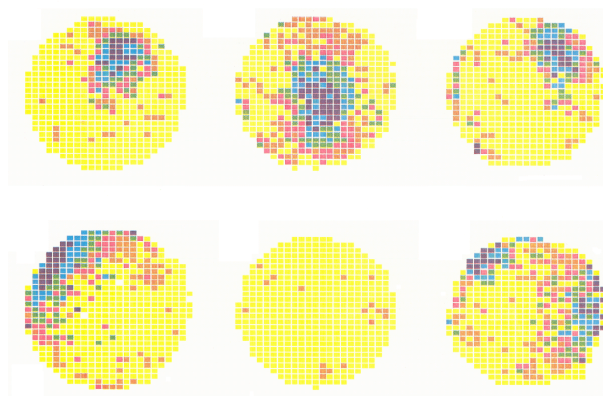
Figure 9.1: Representation of place field data: the "Muller" firing map. Six place cells, with clear place fields, are presented (Muller *et al*, 1987). The colour sequence, from lower firing rate to higher firing rate is yellow, orange, red, green, blue, purple. Yellow represents 0 Hz while purple is around 20 Hz.

(Wilson and McNaughton, 1993). If the environment configuration is static, these place fields are stabilised and remain stable for very long periods of time, up to months (Thompson and Best, 1990). Stability means that place fields remain roughly unchanged in subsequent visits to a previously explored environment.

One important experimental result regarding place field stability is that rats with impaired long-term plasticity in the Shaffer collaterals pathway are incapable of stabilising place fields in the CA1 region (Tsien *et al*, 1996). Although the animal forms naturally place fields each time it explores a specific environment, these place fields are not maintained in subsequent visits to the same environments. That is, on each visit, the animal forms different place fields as if it would be the first time that it was exploring the environment. One can therefore say that, although long-term potentiation in the Schaffer collaterals is not required for the correct formation of place cells, it is necessary for the properties and, most importantly, for the stability of place fields.

9.1.2   Responses to environmental manipulations

Another important characteristic of place cells is their remapping properties, which result from geometric manipulations in the environment. It has been show that manipulations to the environment affect place fields' topographical properties in a similar

way. Possible manipulations include rotation (O?Keefe and Conway, 1978), cue re-
moval (Shapiro *et al*, 1997) and size transformations (O'Keefe and Burgess, 1996).
Rotation and scaling have impressive results, producing equivalent effects in the ge-
ometry of the place fields. More complex manipulations, which may involve changing
the contents of the environment, produce different results.

A very important point is that the effect of environmental manipulations in CA1
and CA3 cell ensembles are distinct. While in CA3 distinct constellations of pyramidal
cells are activated in different environments with similar background features, in CA1
similar features in different environments lead to overlap in activated constellations
(Leutgeb *et al*, 2004).

9.1.3   Relations to theta rhythm

When the temporal properties of single complex-spike bursts are compared with the
theta oscillation of the population one observation can be made: the discharges of
complex cells occur during the positive half of each theta cycle. This relationship
between the complex-spike bursting and the theta oscillation is called "phase-locking"
(O'Keefe and Recce, 1993). In addition, the spike activity of a place cell advances
to earlier phases of the theta cycle as the rat passes through the cell's place field (see
Skaggs *et al*, 1996). This phenomenon, known as the "theta phase precession" or
"theta phase advance", has been suggested as providing additional information (phase-
coding) on the animal's location (Jensen and Lisman, 2000).

9.2   Place cells as components of internal representa-
tions

It is important to notice that filtering and smoothing algorithms are often applied to the
spike counts in order to produce the continuous firing rate maps seen in the literature.
However, it is very important to keep in mind that the raw data regarding spike counts is
far from being smooth. The firing of a place cell is stochastic and the approach taken in
this hippocampal theory is that place fields are defined as a modulated firing probability
instead of a modulated firing rate. Although these two approaches are numerically
the same, phenomenologically they are grounded in completely different assumptions.

While firing rate modulation implies rate coding, the firing probability modulation is in accordance with the form of coding assumed in this model, which is based in activity patterns (constellations of active neurons). Multiple spikes produced by one cell, have, in the firing probability modulation approach, the effect of an increase in the fidelity of the signal transmitted. For an activity pattern to be reliably recognised by a post-synaptic cell, a spike from all the cells composing the activity pattern should arrive within the integration time window of τ_m. This superposition is more reliable if each cell in the activity pattern produces a short burst instead of a single spike.

The simulations presented in this chapter show that the results from spatial memory experiments can be reproduced assuming coding and processing in the form of activity patterns, and firing probability modulation. In addition, major spatial memory experimental results and observations can be explained with the model developed in chapter 7 for episodic memory.

Following the hippocampal model of chapter 7, place cells are cells that have been recruited to take part in an internal representation of the environmental configuration. In a new environment, emotional reactions such as fear or novelty, trigger the formation of hippocampal internal representations which represent particular environmental configurations. Which features of the new environmental configuration are combined to form and store an (episodic) internal representations, and how many internal representations are stored, depend on the consistency of the environmental features. The short-term dynamics of mossy fibres may play an active role in selecting, among the multi-modal information, which properties are relevant. Most probably, other mechanism may also intervene in this process.

Visual, auditive, proprioceptive, somesthetic, and other sensory domains, are a function of space, as they are also a function of time. For example, if several different odour sources are distributed in an environment, the information that arrives at the olfactory system is a function of the concentration of each odour molecule which, in turn, is a function of the distance to each source. In the static arenas of the spatial memory experiments, a particular combination of perceptual (and cognitive) states is therefore linked to a particular position in space.

If the hypothesis that a place cell is simply a cell belonging to an internal representation of a specific environmental configuration, then cells in the model for episodic memory (chapter 7) should verify all fundamental properties of place cells. The simu-

lation results that will be now presented address this point. This analysis was grounded in the following assumptions:

- in a new environment, one or more internal representations are formed (and stored) combining the most relevant and consistent elements or features;

- the storage process requires exploration and the perceptual signal provided by particular features is stronger in specific regions;

- in a static arena/environment, the activity patterns in the entorhinal cortex are modulated mostly by space (position);

- the segregation in the EC means that changes in individual perceptual channels produce small changes (high degree of overlap) in the activity patterns in EC.

9.2.1 Simulations

A population of 200 neuron units representing CA1 pyramidal neurons was added to the hippocampal network shown in figure 8.6. The CA1 population received inputs from the EC ($C_d = 44$, $S_{peak} = 6.3nS$) and from CA3 ($C_d = 100$, $S_{max} = 11nS$). Instead of producing a population representing ECIII, the EC inputs to the CA1 were taken directly from ECII. The connections from EC to CA1 were parametrised so that an activity level of 10% in EC would produce an activity level of 5% in CA1. The Schaffer collaterals had $S_{peak} = 0nS$ initially and were subject to associative learning ($\Delta LTP = 5.5nS$ and $\Delta LTD = 0.5nS$). A population of 20 inhibitory interneurons, with properties identical to CA3i (see figure 8.6), was created for the CA1 population.

Internal representations were stored in the complete memory system for 5 different EC activity patterns, following the storage process described in section 8.4.

It is important to notice that, in this hippocampal network, any EC activity pattern produces a CA1 activity pattern. Nevertheless, after learning, the CA1 activity patterns that are produced by EC patterns that resemble one of the 5 pure patterns, are mostly driven by the Schaffer collaterals. That is, for noisy versions of a pure pattern, the Schaffer collaterals override the effect of the perforant path and rebuild the CA1 representation that was originally produced in the storage process. The competition provided by inhibitory interneurons is fundamental for this mechanism.
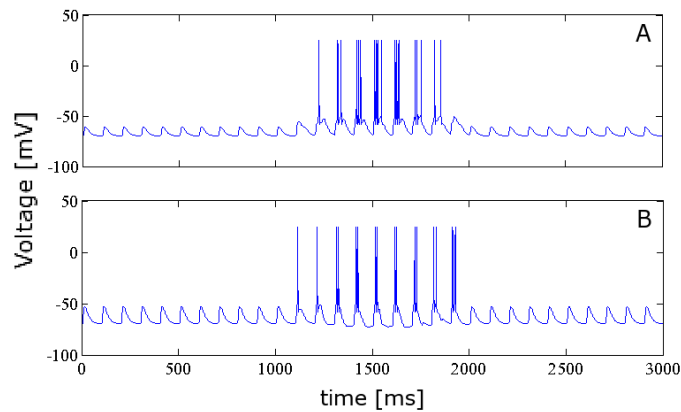
Figure 9.2: Firing profile of a CA3 neuron units as a function of the EC activity pattern. In **A**, the features associated with a stored representation are concentrated in a small area of the environment. In **B**, the same features are spread over a larger area of the environment, as a consequence of a geometrical expansion. In this case, the recall process produces a larger area of activation of the internal representation. Parameters from figure 8.6. See text for more details.

Figure 9.2 shows the firing characteristics of a CA3 neuron unit as function of noisy versions of one (stored) pure EC activity pattern. Following the concept of segregation in the EC, the sequence of noisy versions presented consisted of corrupting a portion (60%) of the pattern. The level of corruption in the portion progressed from 60% to 0% and returning to 60%. The total sequence had a length of 3 seconds, with each EC spatial pattern being presented at a rate of 10 Hz (theta frequency). This modulation in time of the perceptual input is equivalent to the space-modulated perceptual input that a rat is subject to when moving in a static environment. The EC activity pattern matches the pure pattern at 1500 ms.

The CA3 cell, as expected from the hippocampal theory presented in the previous chapter, starts to fire when the EC pattern is similar to the pure pattern, which occurs in the middle of the time sequence. At this stage, the pattern completion mechanisms take place and the complete internal representation is activated, including the measured unit. A few more spikes are fired when the environmental configuration (expressed in the input pattern) matches very closely the pure pattern. In other words, although spikes start to fire at similar activity patterns, the spiking activity is highest when all
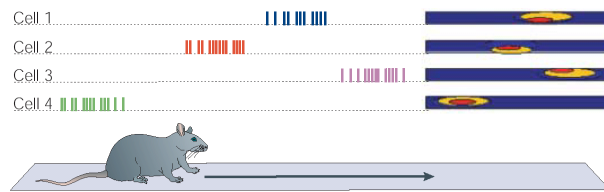
Figure 9.3: Place field in a linear path. Adapted from Nakazawa *et al* (2004).

the original features, which were present in the pure pattern, are combined again. This results from the fact that more reliable information leads to more input current to the representation neuron, which, in turn, originates more spikes in the neuron unit. Again, this is not related to rate coding. This is seen, in this hippocampal model, as an increase in the fidelity of the information encoded (see section 8.4).

The firing profile of this neuron can be compared with the firing of place cells as the animal enters its place fields. Recordings from place cells of a rat moving on a linear path are shown in figure 9.3. Conceptually, this experiment is identical to the simulation.

If the spiking profiles shown are converted into a firing-rate profiles, one obtains the structure of place fields (see figure 9.4).

If instead of a crude filter such as the running time window, a Gaussian kernel is used to convert the spike counts to firing-rate, one obtains firing-rate profile with the structure of the one shown in figure 9.5.

It can be seen that some firing fields are asymmetric with the firing peak closer to the end of the field. This means that the activity grows steadily as the field is entered but drops fast closer to the end of the field. This is a consequence of the dynamics of the continuous spiking networks that model each subfield of the hippocampus. The state of the system is continuously changing and affecting the following states. In the case of the CA3 network, a noisy version of a pure activity pattern may not be sufficient to invoke the activation of an internal representation, but it will put the constellation of intervening neurons closer to threshold. A following noisy pattern will require less effort to recall the representation. At the end of the field, the state of the CA3 network is still influenced by the low noise input patterns. When this influence is lost and the noise of the input pattern is high, the activity of the neurons belonging to the representations drops very fast.

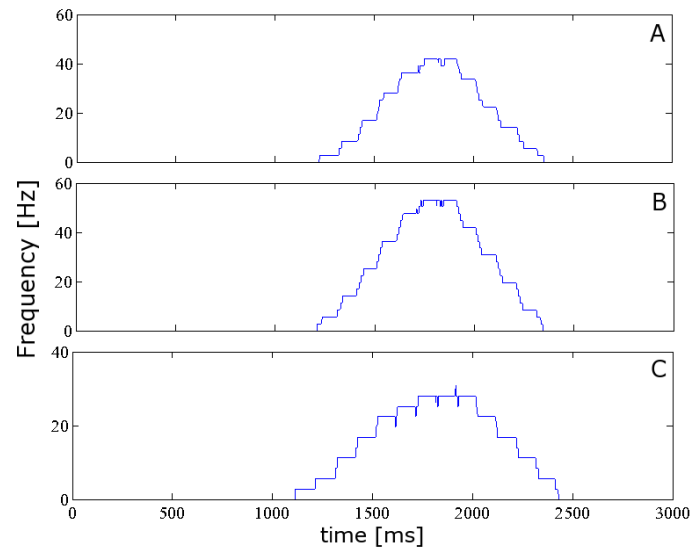Figure 9.4: A running-average filter was used to convert the spike counts from figure 9.2 to firing-rates. The amplitude of the moving time window was 500 ms. This simple filtering was used to avoid the production of any artefacts in the data.
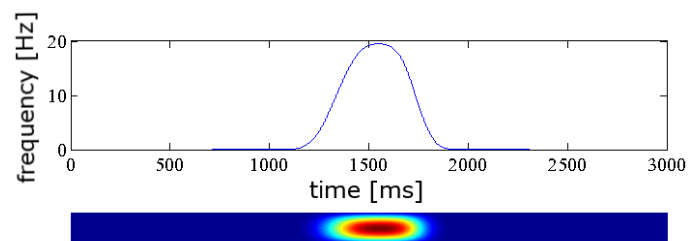

Figure 9.5: Smooth firing-rate profiles. A Gaussian kernel ($\sigma = 200ms$) was used to convert the spike counts of figure 9.4-**C** to firing rates.

The same simulation experiments with CA1 neurons produce similar profiles but with an interesting difference: the activation of individual CA1 cells requires that only specific subsets of the EC input pattern are correct. This point will be discussed in more detail in section 9.3

Environmental manipulations

In this theory, environmental manipulations are simply geometric operations in the perceptual signals. In rats, the perceptual signals are invariant to the rotations of the whole environment. Animals with internal orientation mechanisms (internal compasses) could nevertheless provides interesting experimental results.

In terms of scaling manipulations, the individual perceptual fields are dilated or compressed, depending on the case. In the simulation shown in 9.2 this corresponds to an increase or decrease in the gradient of the noise in the sequence. In sequences where the transition from different to similar input pattern is faster, corresponding to a decrease of the environment, the activation field is reduced. When the transition is slower, meaning expansion of the environment, the region where pattern completion can take place is equally expanded leading to a bigger activation field.

Environmental manipulations involving transformations of subsets of the environment only test the power of pattern completion in CA3 and to which features a CA1 neuron responds to.

Due to the results of this section, and for the rest of this thesis, the term "place field" will be avoided and substituted by "activation field" which is more general.

A study of phase precession in this hippocampal model is left for a future work. It may, nevertheless, be related to the drifting mechanisms in the firing initiation time that have been shown in section 8.4 (third stage in the storage process).

9.3 Different types of activation fields

In this section I present my personal view of the true nature of "place-fields". This discussion, although speculative, is motivated by simulation results and by the hippocampal model functional properties. It remains, however, a mere hypothesis.

As mentioned before, not only CA1 and CA3 neurons exhibit space selectivity. Neurons with similar spatial correlates have also been described in the medial entorhi-

nal (Fyhn *et al*, 2004), perirhinal and postrhinal cortices (Burwell and Hafeman, 2003). There is nevertheless a strong difference in the activation fields of these cortical inputs to the hippocampus: most cells exhibit activation fields with split or multiple subfields (Burwell and Hafeman, 2003; Fyhn *et al*, 2004). In the case of postrhinal firing fields, the majority of the cells (84%) adopted new spatial correlates when experimental cues were rotated, but did so neither predictably nor concordantly. I therefore propose the hypothesis that, all these different cells that have been homogeneously called "place cells" in the literature, are in fact cells responding to signals of different nature.

The representations in CA3 are chosen through a process that projects into a high dimensional space the EC activity patterns. The multi-model information from the EC is completely mixed and combined to form an internal representation. On the other hand, the representations in CA1 are formed directly from the EC, although only the representations that are associated with a storage process in CA3 are stabilised by the Schaffer collaterals. While the CA3 representations are totally distributed, the CA1 representations tend to overlap more due to the absence of a separation mechanism. And this can be seen at the single neuron level: for two overlapping EC input patterns, a neuron from a CA1 representation will have a much higher probability of firing in both input patterns than a CA3 neuron. This is in accordance with experimental results (Leutgeb *et al*, 2004).

The very important point to make is that while CA3 neurons respond to the complete combination of all the original features in the EC multi-modal input, CA1 neurons can still respond to individual features. Therefore the interpretation given to the activation fields of this populations should be slightly different. The situation with the spatial selectivity in medial entorhinal, postrhinal and perirhinal is, however, completely different.

I propose here that the different properties of the activation fields of the cortical input areas reflect the coding of individual perceptual features. That is, assuming segregation in these regions (which is in fact true, at least, for the medial entorhinal cortex), each subregion is conveying information from different sensory associational areas in the neocortex. Therefore the firing of each cell refers to a single dimension of the whole perceptual (and cognitive) state. If the coding mechanism in these regions is also based in constellations of active neurons (activity patterns), a multi-peaked active field of a neuron represents its presence in more than one representation. In an abusive

example, one may think of a neuron which is shared in the internal, high level representation of two different odours. When recording from this neuron, its activation field will be peaked in two regions of the space where the distinct odours appear. Again, it is important to emphasise that it is not "space" but the modulation that it provides: its not a location that is being detected, but the presence of a feature instead.

Therefore I expect that the "place fields" measured in pre-hippocampal regions reflect the detection of individual features of specific sensory or cognitive channels, while CA3 "place fields" are the only fields representing a combination of features which describe an event, episode or environmental (perceptual) configuration. The CA1 activation fields are an intermediate situation between the entorhinal cortex and CA3.

It may be possible that, due to functional constraints, the output of the hippocampal memory system has to be the concatenation of individual features instead of a representation of the combination of the features. Nevertheless, the recognition and recall of an event, episode or perceptual configuration has to be performed at the whole level. This way it is plausible to think that when an event, episode or perceptual configuration is recalled through the activation of a CA3 internal representation, the original CA1 representation, sensible to individual features, is reconstructed. The CA1 connections back to the entorhinal cortex deep layers may then reconstruct, or complete with the missing parts, the EC activity patterns. This decoding mechanism makes use of the information stored in the CA3 internal representation to recreate the original EC activity pattern.

9.4   Final Remarks

With very few exceptions, spatial memory experiments have explored the influences of sensory channels in the activity of CA1 and CA3 cells. In these few exceptional cases (e.g. Wood *et al*, 1999), the activity in hippocampus has been found to be consistently related to perceptual, behavioural or cognitive events, regardless of the location where these events occurred.

If this hippocampal model is correct, all dimensions in the perceptual space affect the activation fields of CA1 and CA3 cells. More experiments are required to test this hypothesis.

A final remark goes for the coding problem. If the assumption of this thesis - that most information is expressed, transferred and processed in the form of constellations of active neurons (activity patterns) - is correct, then analysing the activation fields of pre-hippocampal structures is analysing the neural code in which perceptual features are expressed. These features should not be seen as "distance to a corner" or "head orientation". They are high-level representations of perceptual components which may, or may not, be merged together.

Chapter 10

Conclusions and Final Remarks

The objective of this chapter is to briefly summarise the major contributions of this thesis and present a critical analysis of its weak points.

10.1 Contributions of this thesis

This thesis has introduced a detailed functional model of the hippocampus which clarifies the roles of this memory system. The hypothesis of the hippocampus as a spatial map (O'Keefe and Dostrovsky, 1971) is rejected in favour of the view of the hippocampus as a general memory space (Eichenbaum *et al*, 1999). In this general, but not permanent memory system, multi-modal information, coming from different associational areas of the neocortex, is combined and stored together. By storing an internal representation of a particular combination, in time, of several perceptual and cognitive features, the hippocampus is storing the description of a specific event, episode or perceptual (environmental) configuration (see figure 10.1). While in humans one can speak in terms of events and episodes, it is more plausible to speak about perceptual or environmental configurations for rats.

With the aid of a simulation environment, intimately related to the hippocampal computational model, all the novel ideas of the hippocampal theory were shown to effectively work. The simulations not just reproduced experimental results but also added important information on specific components of this memory system. The detailed level in which the hippocampus was represented allowed many manipulations (such as the removal of the mossy fibres - Lassalle *et al*, 2000) and production of new
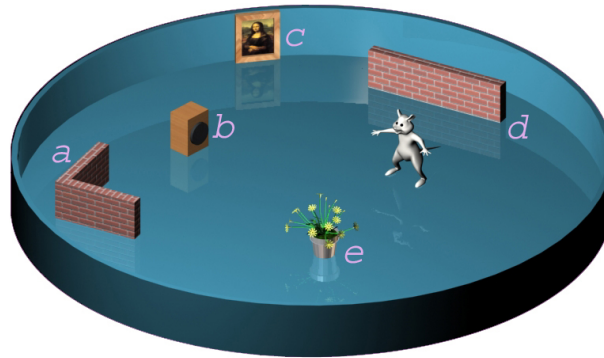
Figure 10.1: The hippocampus works as a general memory system. Triggered by the medial septum, the hippocampus stores internal representations of important perceptual configurations. The relevant features of the perceptual/environmental configuration may consist of geometric (**a** and **d**), auditive (**b**), visual (**c**) or olfactive (**e**) information.

hypothesis (such as the true nature of place fields).

Although based on models such as Marr (1971), Treves and Rolls (1992; 1994), O'Reilly and McClelland (1994) and Hasselmo *et al* (1995; 1996), this model introduced many novel issues not addressed in the previous models. Significant contributions are:

- A different coding mechanism, based on activity patterns (or constellations), is used instead of the typical firing-rate coding. This constellation coding mechanism explains, in much more detail, the storage, recall and signal processing dynamics in the short time scale of tenths of milliseconds. While the coding mechanism itself is an assumption, the biologically plausible strategies introduced in this model to deal with this equally plausible code mechanism (Vaadia *et al*, 1995) are a significant contribution of this thesis.

- A detailed quantitative description on how interneurons control the activity levels in the principal populations is provided. The importance of the activity level in the dentate gyrus during the process of selection of the CA3 representation suggests a precise role for the mossy cells. In addition, the mechanisms behind the proposed activity control suggest a reason for the variety of inhibitory interneurons in a single subfield: since each inhibitory population can only ef-

fectively control the activity of one excitatory source, principal populations with different sources require different populations of inhibitory cells.

- Synaptic short-term dynamics are introduced and shown to play important roles in the hippocampal functional behaviour. In particular, synaptic short-term in mossy fibres are shown to have strong implications on the dynamics of the system. This hippocampal model refined the responsibilities of these fibres: besides contributing to the creation of sparse CA3 representations, mossy fibres select which features in the input space are consistent in time. This produces internal representations which are more correlated with the relevant features of a particular environmental/perceptual configuration and less dependent on noise and variable input channels.

- This hippocampal model includes representations of the dentate gyrus, CA3 and CA1. Each component works in accordance with experimental data and, most importantly, the system as a whole, mimics and explains the memory system created by the hippocampus. The processes based in the coding mechanism and activity control explain how information is transferred from subfield to subfield maintaining the stability of the system.

- Taking advantage of the use of spiking neurons, this model opens a window to the short-time scale which is averaged out in firing-rate models.

The last point is in fact very important and is now discussed in more detail since it is behind relevant contributions of this model. The use of spiking neurons is crucial for this hippocampal model. They open a window to the milliseconds time-scale of the system's dynamics, and thanks to this temporal precision, it is possible to obtain information which is unreachable using firing-rate neurons. For example, with the hippocampal model based in spiking neurons, it was possible in this thesis to:

- calculate the time it takes to recall an internal representation, that is, the time between the presentation of a known activity pattern in EC2 and the full activation of the associated CA3 constellation;

- analyse the impact of short-term dynamics of the mossy fibres in the process of selecting features in the EC2 input space;

- verify the extreme importance of coherence and synchrony in neuronal activations in the propagation and processing of signals in the hippocampal subfields;

- calculate the time that it takes to build and complete an internal representation in the CA3 subfield;

- propose efficient and biologically realistic inhibitory control mechanisms which are able to respond even to transient excesses of excitation.

The above results could not be achieved using firing-rate models. The detailed analysis of the implications of spike-timing dependent plasticity in the perforant path and Schaffer collaterals, which was left to future-work, is also only possible with the use of spiking models.

Spiking models and firing-rate models only provide equivalent results in the cases where the firing of neurons in a network is uncorrelated, with little synchronous firing, and where precise patterns of spike timing are unimportant (Dayan and Abbott, 2001). This is exactly not the case here since the hippocampal model processes and transmits information using synchronised firing in constellations of neurons.

One significant achievement of this thesis is the creation of an hippocampal model which integrates spatial memory and episodic memory in a single framework. In the hippocampus literature, few are the models that try to combine episodic with spatial memory in a single computational framework. One such combined model was proposed by Rolls *et al* (2002) and is based in two connected attractor networks, one discrete and the other continuous. While the discrete attractor network creates a system capable of storing and recalling binary patterns (working as an associative network), the continuous attractor network recreates some of the properties of place fields, the fundamental building block of spatial memory. This model is nevertheless not biologically realistic and presents several limitations.

Another contribution of the thesis refers to the presented methods to integrate several spiking networks in a system that works harmoniously. The use of spiking neurons introduces several problems regarding stability. This thesis has presented ways of producing biologically plausible activity control mechanisms which secure the proper functional behaviour of spiking networks. Furthermore, it has shown how information can be expressed, transferred and processed in the form of activity patterns. This the

coding mechanism seems to me more sensible than the rate coding used in most hippocampal models. The implications of constellation-coding are vast and interesting. Some of them were addressed in this thesis for the case of the hippocampal system (for example, the production of a constellation in CA3 by the dentate gyrus).

10.2 Critical Analysis of the thesis

A weak point in this hippocampal theory is the absence of all the diversity of hippocampal interneurons. Although some specific roles have been given for particular interneurons (such as the mossy cells in the dentate gyrus), all the richness of the possible computations executed by the interneurons was left untouched. Related to this fact is the absence of a precise description for the mechanisms that segregate individual activity patterns. In the hippocampal model, the EC input activity patterns were presented at theta frequency. Although the hippocampal computational model is capable of functioning which activity patterns being continuously presented (as long as an external homogeneous input at theta frequency is applied) the performance of storage and recall is degraded. The cause for the reduction in performance is derived from the absence of coherence in the activation of the neurons belonging to the internal constellation (figure 10.2). Another cause for the reduction in performance is the spurious firings originated by the accumulation of continuous excitation.

Another point deserving some criticism is the absence of a more detailed analysis of the implications of the use of spike-timing dependent plasticity. One of many interesting consequences of using the STDP rule is that the activation order of internal representations becomes important. For example, if the representations composed by the constellations of neurons A, B and C were stored with an activation order of BCA, then the recall process will always proceed in this order. The use of STDP can therefore be relevant if sequences of representations are to be stored or if temporal associations of internal representations are important. Information regarding the transition between internal representations can be stored using STDP, and this is important to understand the role of the hippocampus in sequence memory/learning and navigation.
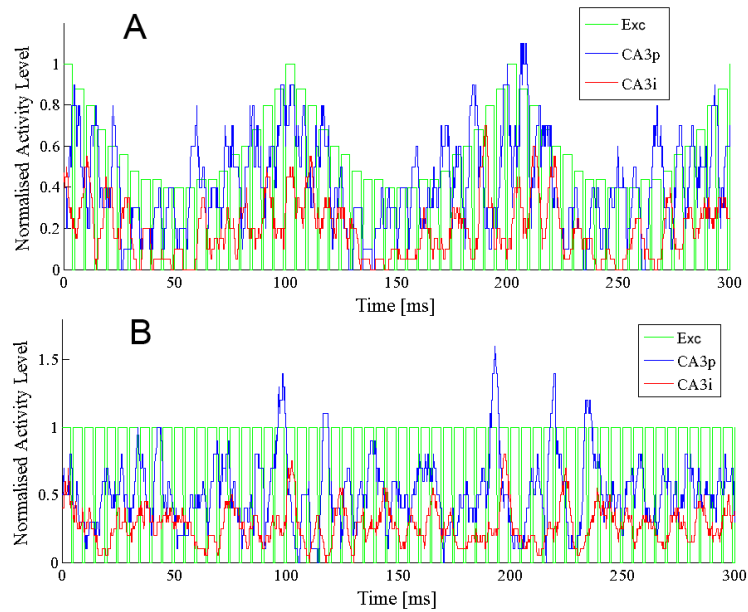
Figure 10.2: Importance of theta oscillation for the segregation of the activity patterns. In **A**, the excitatory inputs to CA3p are under theta oscillations. Although the amplitude of the oscillations is not very big, it is enough to secure that the neurons belonging to an internal representation are activated coherently. In the situation depicted in **B**, where the excitatory level is constant, the activation is not synchronised and there is no guarantee that all neurons belonging to the internal representation will be active within a time window of τ_m.

Chapter 11

Future Work

The hippocampal model of this thesis can be explored in many fronts. In particular, many of its assumptions (such as coding in constellations) and hypothesis (such as the different nature of place cells, and the notion of activation fields to substitute place fields), can be further analysed in order to produce more predictions which can be tested against experimental results.

There are many interesting points, initially part of the thesis objectives, which were prevented from being addressed here due to lack of time. In this chapter, I describe the most relevant and promising ones.

Topics for future work

1. Undoubtedly, the major future objective is to take the preliminary results shown in chapter 9, and investigate in detail the properties of the hippocampal model regarding memory for spatial configurations. Much can be done in showing that paradigmatic results from spatial memory experiments, which see the hippocampus as a specialised structure for spatial memory, can be recreated *and explained* using the hippocampal model of this thesis, which, in turn, sees the hippocampus as a general memory space for environmental/perceptual configurations.

2. Following the previous point, the issue of phase-precession should be addressed in the hippocampal model. This property should be verified in neurons taking part in ordinary internal representations. In order to achieve a conclusive result, a detailed model of the interactions between the septum and the hippocampus will probably be required.

3. Another interesting future objective is to explore the effects that short-term dynamics in Schaffer collaterals may produce (Dobrunz and Stevens, 1997) in the hippocampal memory system. In the present hippocampal model, the response of a subfield to an incoming activity pattern depends on the present state of the subfield. With synapses exhibiting short-term dynamics, temporal correlations between patterns can be detected and temporal features can be decoded. This is the door to sequence learning and navigation. The simulation environment satisfies all the requirements to conduct experiments and analysis in this topic.

4. Slowly the computational importance of interneurons is being acknowledged. This theory introduced specific roles for interneurons but much more can be done and it would be extremely interesting to populate the hippocampal model with more interneuron types. Fortunately there is a considerable amount of morphological and physiological experimental data regarding hippocampal interneurons which can be used to produce hypothesis about their specific functional roles.

5. The simulation environment that resulted from this thesis is a very rich and general tool. It uses a unique body of assumptions, approximations and algorithms that can be of use in many simulation experiments, not only regarding hippocampal functional behaviour. In the context of heterogeneous spiking networks it is much more efficient than NEURON. It can, nevertheless, be substantially improved in terms user interface, output functions, among other things.

6. The discussion about the storage capacity of the spiking network model of CA3 presented in this thesis was quite empirical. This subject is nevertheless of great relevance and it would be very interesting to analyse it theoretically. Particularly, I would like to see if the expression for the storage capacity holds in much larger networks.

Appendix A

Mossy fibre butons as dynamical synapses

Abstract: In this work the consequences on the activity of CA3 pyramidal neurons of the dynamic properties (facilitation and depression) associated with hippocampal mossy fibre boutons are explored; namely, a study is carried on the dependence of the pyramidal cell activity with the input stimulus frequency to the mossy fibre boutons as well as the influence of the single vesicle post-synaptic response magnitude on this dependence. We calculate tuning curves for the mossy fibre boutons excitatory post-synaptic potentials as a function of pre-synaptic frequency and we show how this property, together with channels that contribute to spike frequency adaptation, makes the activity of CA3 pyramidal neurons strongly dependent on the temporal patterns of mossy fibres pre-synaptic signals.

A.1   Introduction

Hippocampal mossy fibres have their origin in dentate granule cells and terminate on the proximal dendrites of CA3 pyramidal cells where their activity produces fast glutamergic excitatory post-synapticpotentials (EPSPs). They have a strong influence on CA3 pyramidal neurons activity not only due to the proximal location but also because of the multiple release sites, up to 37 (Chicurel *et al*, 1992), which generate large EPSPs. Due to this strong efficiency and efficacy, it is natural to assume that the activity dynamics of the boutons will be strongly transfered to the soma.

Viewing the boutons as dynamical synapses with short-term memory mechanisms like facilitation and depression, with the associated filtering properties, and considering

the fact that CA3 pyramidal neurons have spike frequency adaptation ionic currents present in their proximal dendrites and soma, brings interesting consequences in terms of the neuron's activity dependence with the input temporal patterns of mossy fibre boutons.

A.2   Methods

All simulations were performed with *NEURON* (Hines and Carnevale, 1997) using a realistic model of a CA3 pyramidal neuron (Migliore *et al*, 1995) which includes *Na,t*, *L* high threshold, *N*, *T* low threshold, *A*, *K*, *M*, *K,Ca* and *Ca* ionic currents. Some of them (*M*) contribute to spike frequency adaptation or (*A* and *K,Ca*) affect interspike time interval. All the parameters used were based in experimental results. The temperature was set to 30 Celsius.

Using *NMODL* (*NEURON* Model Description Language) a model was created for the mossy fibre synapse based on the probabilistic model for dynamic synapses suggested in Fuhrmann *et al* (2002), which accounts for depression processes. According to that, P_r, the probability of release for every release site at time of a spike, t_{sp}, is given by $P_r(t_{sp}) = P_{vr} \cdot P_{va}$ where P_{vr} is the probability of vesicle release for a release site with an available vesicle (affected directly by facilitation) and P_{va} is the probability of a vesicle to be available for release at time t (directly affected by depression). The dynamics are:

$$\frac{dP_{va}}{dt} = \frac{1 - P_{va}}{\tau_{rec}} - P_{va} \cdot P_{vr} \cdot \delta(t - t_{sp}) \tag{A.1}$$

where δ is the delta function, t_{sp} is the time of arrival of a spike and τ_{rec} the relaxation time constant of depression.

For the facilitation process, P_{vr} was made proportional to the calcium concentration inside the bouton with $P_{vr} = 1$ when $[Ca^{2+}] \geq 0.2\mu M$.

The post-synaptic response is then calculated according to the classical quantal model of release: $PSR = q \cdot n \cdot P_r$, where q is the the post-synaptic response to the release of each vesicle and n is the number of vesicle release sites. p and n are related with pre-synaptic processes/mechanisms whereas q is post-synaptic related. The conductance q is assumed to be affected by learning (i.e. long term potentiation/depression).

The value used for τ_{rec} was 300 ms which is in close agreement with biological measured values (Markram *et al*, 1998). The post-synaptic conductance was modelled as a two state kinetic scheme synapse described by rise time of 0.5 ms and decay time constant of 5 ms (Jonas *et al*, 1993).

On this model a small number of mossy fibre synapses (5) with 30 release sites were connected to the CA3 pyramidal neuron with a associated unitary, i.e. per vesicle, conductance ranging from 1 to 20 nS. The synapses were placed in the proximal region (proximal apical dendrites) and all received the same spike train characterised by a specific constant frequency (no variability was added to the interspike time intervals).

We here assume that among the approximately 50 contacts that each pyramidal cell receives from granule cells (Johnston and Amaral, 1998), only a small percentage were active in a limited time window and with the same, or very similar, temporal pattern.

The dependence of the activity of the pyramidal neuron with the frequency of mossy fibre inputs and the post-synaptic conductance was measured as the number of spikes generated in a specific time window.

The effect of the dynamic properties of the mossy fibre boutons on the activity of the neuron was also studied when the soma is bombarded by small EPSPs generated by several poissonian point processes randomly distributed in the dendritic tree. The mean frequency of the Poisson spike train for each point process representing an non-dynamic dual exponential synapse, was randomly chosen in the 0 to 60 Hz spectrum. In this more realistic context, the activity of the pyramidal neuron was measured again in terms of its average firing rate as a function of the mossy fibre spike train frequency.

A.3 Results

The tuning curves for all the analysed unitary conductances (q) were plotted and the optimal input frequency was calculated fitting the data points with a dual exponential expression (with typical $R^2 \geq 0.9$). An example for $q = 4nS$ data points and the fitted tuning curves for other three different conductances is shown in Fig.1.

Changing the post-synaptic response through the variable q, has the effect of moving the peak frequency. The calculated fitting surface, as a function of q and the input frequency is shown in the left side of Fig.2. The dark line represent the optimal input frequency (associated with the highest soma activity) for each unitary conductance.
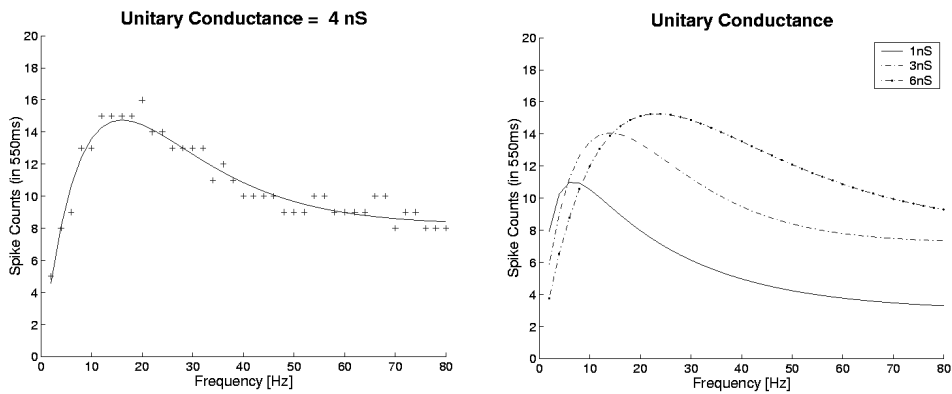
Figure A.1: Tuning curves for neuron's activity.
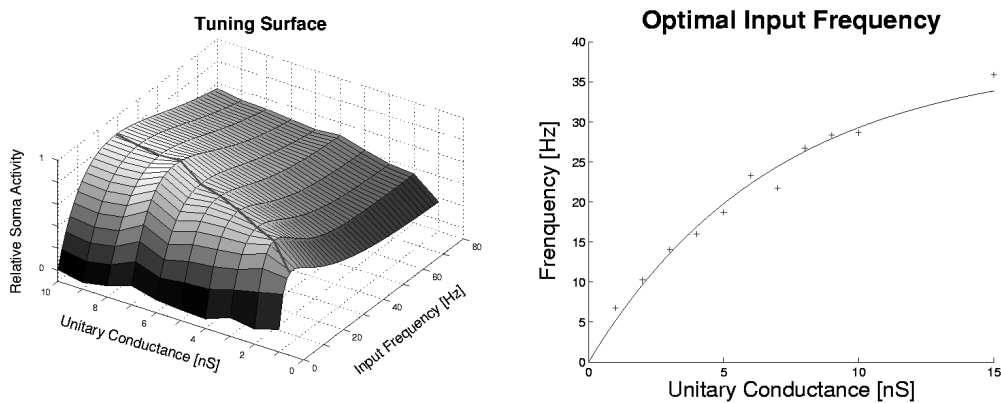
Figure A.2: Tuning surface and optimal average rate for the mossy fibre input spike train.

This optimal frequency curve is presented in more detail in the right side of Fig.2.

For the more realistic situation where the soma is being bombarded by small EP-SPs, the properties of the dependence of the pyramidal neuron average firing rate with the mossy fibre input frequency, and unitary conductance, were not affected:

For $q = 4$nS, the calculated corresponding optimal input frequency is $f_{opt} \approx 62$Hz. With the mossy fibres receiving spike trains with this frequency together with the fixed distribution of poissonian point processes with fixed random average rate and small fixed random conductances, the obtained average firing rate of the pyramidal neuron was 24 Hz. This maximal value was significantly higher (50%) than the neuron's firing rate values obtained for input spike train frequencies equal to half the optimal

frequency and twice the optimal frequency (14 and 16 Hz respectively). The rate activity of the neuron with silent mossy fibres was 4 Hz.

A.4 Conclusions

The bandpass filter behaviour of the mossy fibre boutons and the existence of an optimal frequency are mainly consequences of the short-term memory dynamics (facilitation and depression). If instead of efficiency we also analyse the efficacy of the boutons, i.e. their contribution for firing the neuron, we observe that the associated conductances also have interesting effects in the non-static, temporal domain, behaviour due to the existence of ionic currents that affect interspike time interval or contribute to spike frequency adaptation.

The fact that in the time domain and in terms of efficacy, the filtering properties of the dynamical synapses change with their strength introduces a new complexity in the understanding of learning mechanisms and it is a very important result if we accept that information may be present in the spike trains time domain.

Furthermore, we can also speculate that the existence of facilitation and depression dynamics in the mossy fibre boutons may create a *transfer switch* in the sense that for the same spatio/temporal input pattern presented to the CA3 neuron through all its synapses, the activity of the neuron will be considerably different when the inputs to the mossy fibre boutons are close to the their respective optimal input frequencies. Through temporal patterns, dentate gyrus neurons would be able to select which populations of CA3 neurons to activate.

Important Note *The study described here was done before the theory for the roles of each pathway in the hippocampus was concluded. Although this study presents valid results, I have introduced some modifications on the views for the functional contributions given by the mossy fibres' short-term dynamics. For a detailed discussion see chapter 7.*

Bibliography

Abbott, L. F., Varela, J. A., Sen, K., and Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science*, 275(5297):220–224.

Acsdy, L., Kamondi, A., Sk, A., Freund, T., and Buzski, G. (1998). GABAergic cells are the major postsynaptic targets of mossy fibers in the rat hippocampus. *J Neurosci*, 18(9):3386–3403.

Agmon-Snir, H., Carr, C. E., and Rinzel, J. (1998). The role of dendrites in auditory coincidence detection. *Nature*, 393(6682):268–272.

Aguiar, P. and Willshaw, D. (2004). Hippocampal mossy fibre boutons as dynamical synapses. *Neurocomputing*, 58-60:699–703.

Amaral, D. G., Ishizuka, N., and Claiborne, B. (1990). Neurons, numbers and the hippocampal network. *Prog Brain Res*, 83:1–11.

Amaral, D. G. and Witter, M. P. (1989). The three-dimensional organization of the hippocampal formation: a review of anatomical data. *Neuroscience*, 31(3):571–591.

Amit, D. (1989). *Modelling Brain Function*. Cambridge University Press, New York.

Amit, D. J. and Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb Cortex*, 7(3):237–252.

Arleo, A. and Gerstner, W. (2000). Spatial cognition and neuro-mimetic navigation: a model of hippocampal place cell activity. *Biol Cybern*, 83(3):287–299.

Artola, A. and Singer, W. (1993). Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends Neurosci*, 16(11):480–487.

Aviel, Y., Horn, D., and Abeles, M. (2005). Memory capacity of balanced networks. *Neural Comput*, 17(3):691–713.

Aviel, Y., Mehring, C., Abeles, M., and Horn, D. (2003). On embedding synfire chains in a balanced network. *Neural Comput*, 15(6):1321–1340.

Baddeley, R., Abbott, L. F., Booth, M. C., Sengpiel, F., Freeman, T., Wakeman, E. A., and Rolls, E. T. (1997). Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proc Biol Sci*, 264(1389):1775–1783.

Barnes, C. A., McNaughton, B. L., Mizumori, S. J., Leonard, B. W., and Lin, L. H. (1990). Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Prog Brain Res*, 83:287–300.

Bennett, M. R., Gibson, W. G., and Robinson, J. (1994). Dynamics of the CA3 pyramidal neuron autoassociative memory network in the hippocampus. *Philos Trans R Soc Lond B Biol Sci*, 343(1304):167–187.

Bi, G. Q. and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci*, 18(24):10464–10472.

Bienenstock, E. L., Cooper, L. N., and Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci*, 2(1):32–48.

Bland, B. H. (1986). The physiology and pharmacology of hippocampal formation theta rhythms. *Prog Neurobiol*, 26(1):1–54.

Bliss, T. V. and Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J Physiol*, 232(2):331–356.

Boss, B. D., Turlejski, K., Stanfield, B. B., and Cowan, W. M. (1987). On the numbers of neurons in fields CA1 and CA3 of the hippocampus of Sprague-Dawley and Wistar rats. *Brain Res*, 406(1-2):280–287.

Bouron, A. (2001). Modulation of spontaneous quantal release of neurotransmitters in the hippocampus. *Prog Neurobiol*, 63(6):613–635.

Bragin, A., Jandó, G., Nádasdy, Z., Hetke, J., Wise, K., and Buzsáki, G. (1995). Gamma (40-100 Hz) oscillation in the hippocampus of the behaving rat. *J Neurosci*, 15(1 Pt 1):47–60.

Brown, T. H., Fricke, R. A., and Perkel, D. H. (1981). Passive electrical constants in three classes of hippocampal neurons. *J Neurophysiol*, 46(4):812–827.

Brun, V. H., Otnæss, M. K., Molden, S., Steffenach, H.-A., Witter, M. P., Moser, M.-B., and Moser, E. I. (2002). Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. *Science*, 296(5576):2243–2246.

Brunel, N. and Trullier, O. (1998). Plasticity of directional place fields in a model of rodent CA3. *Hippocampus*, 8(6):651–665.

Buonomano, D. V. (2000). Decoding temporal information: A model based on short-term synaptic plasticity. *J Neurosci*, 20(3):1129–1141.

Buonomano, D. V., Hickmott, P. W., and Merzenich, M. M. (1997). Context-sensitive synaptic plasticity and temporal-to-spatial transformations in hippocampal slices. *Proc Natl Acad Sci U S A*, 94(19):10403–10408.

Buonomano, D. V. and Merzenich, M. M. (1995). Temporal information transformed into a spatial code by a neural network with realistic properties. *Science*, 267(5200):1028–1030.

Buonomano, D. V. and Merzenich, M. M. (1998). Net interaction between different forms of short-term synaptic plasticity and slow-IPSPs in the hippocampus and auditory cortex. *J Neurophysiol*, 80(4):1765–1774.

Burgess, N., Jeffery, K., and O'Keefe, J. (1999). *The hippocampal and parietal foundations of spatial cognition*. Oxford University Press.

Burgess, N., Maguire, E. A., and O'Keefe, J. (2002). The human hippocampus and spatial and episodic memory. *Neuron*, 35(4):625–641.

Burwell, R. D. and Hafeman, D. M. (2003). Positional firing properties of postrhinal cortex neurons. *Neuroscience*, 119(2):577–588.

Buzsáki, G. (1989). Two-stage model of memory trace formation: a role for "noisy" brain states. *Neuroscience*, 31(3):551–570.

Buzsáki, G. (2002). Theta oscillations in the hippocampus. *Neuron*, 33(3):325–340.

Buzsáki, G., Leung, L. W., and Vanderwolf, C. H. (1983). Cellular bases of hippocampal EEG in the behaving rat. *Brain Res*, 287(2):139–171.

Cannon, R. C., Hasselmo, M. E., and Koene, R. A. (2003). From biophysics to behavior: Catacomb2 and the design of biologically-plausible models for spatial navigation. *Neuroinformatics*, 1(1):3–42.

Chicurel, M. E. and Harris, K. M. (1992). Three-dimensional analysis of the structure and composition of CA3 branched dendritic spines and their synaptic relationships with mossy fiber boutons in the rat hippocampus. *J Comp Neurol*, 325(2):169–182.

Chrobak, J. J. and Buzski, G. (1998). Gamma oscillations in the entorhinal cortex of the freely behaving rat. *J Neurosci*, 18(1):388–398.

Chun, M. M. and Phelps, E. A. (1999). Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. *Nat Neurosci*, 2(9):844–847.

Colicos, M. A., Collins, B. E., Sailor, M. J., and Goda, Y. (2001). Remodeling of synaptic actin induced by photoconductive stimulation. *Cell*, 107:605–616.

Cowan, W. M., Südhof, T. C., and Stevens, C. F., editors (2001). *Synapses*. The Johns Hopkins University Press. ISBN 0-8018-6498-4.

Dayan, P. and Abbott, L. (2001). *Theoretical Neuroscience*. The MIT Press. ISBN 0-262-04199-5.

Dayan, P. and Willshaw, D. J. (1991). Optimising synaptic learning rules in linear associative memories. *Biol Cybern*, 65(4):253–265.

Delaney, K. R. and Tank, D. W. (1994). A quantitative measurement of the dependence of short-term synaptic enhancement on presynaptic residual calcium. *J Neurosci*, 14(10):5885–5902.

Dittman, J. S., Kreitzer, A. C., and Regehr, W. G. (2000). Interplay between facilitation, depression, and residual calcium at three presynaptic terminals. *J Neurosci*, 20(4):1374–1385.

Dobrunz, L. E. and Stevens, C. F. (1997). Heterogeneity of release probability, facilitation, and depletion at central synapses. *Neuron*, 18(6):995–1008.

Doerksen, K. (2000). Pulsed neural networks: Temporal signal processing using atificial neural networks with dynamic synapses. Technical report, CYSF.

Dolorfo, C. L. and Amaral, D. G. (1998). Entorhinal cortex of the rat: topographic organization of the cells of origin of the perforant path projection to the dentate gyrus. *J Comp Neurol*, 398(1):25–48.

Dudchenko, P. A., Wood, E. R., and Eichenbaum, H. (2000). Neurotoxic hippocampal lesions have no effect on odor span and little effect on odor recognition memory but produce significant impairments on spatial span, recognition, and alternation. *J Neurosci*, 20(8):2964–2977.

Eichenbaum, H. (1999). Conscious awareness, memory and the hippocampus. *Nat Neurosci*, 2(9):775–776.

Eichenbaum, H. (2001). The hippocampus and declarative memory: cognitive mechanisms and neural codes. *Behav Brain Res*, 127(1-2):199–207.

Eichenbaum, H. and Cohen, N. J. (1988). Representation in the hippocampus: what do hippocampal neurons code? *Trends Neurosci*, 11(6):244–248.

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(2):209–226.

Eichenbaum, H., Wiener, S. I., Shapiro, M. L., and Cohen, N. J. (1989). The organization of spatial coding in the hippocampus: a study of neural ensemble activity. *J Neurosci*, 9(8):2764–2775.

Eldridge, L. L., Knowlton, B. J., Furmanski, C. S., Bookheimer, S. Y., and Engel, S. A. (2000). Remembering episodes: a selective role for the hippocampus during retrieval. *Nat Neurosci*, 3(11):1149–1152.

Feder, R. and Ranck, J. B. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. II. Hippocampal slow waves and theta cell firing during bar pressing and other behaviors. *Exp Neurol*, 41(2):532–555.

Foster, D. J., Morris, R. G., and Dayan, P. (2000). A model of hippocampally dependent navigation, using the temporal difference learning rule. *Hippocampus*, 10(1):1–16.

Fourcaud, N. and Brunel, N. (2002). Dynamics of the firing probability of noisy integrate-and-fire neurons. *Neural Comput*, 14(9):2057–2110.

Fox, S. E. and Ranck, J. B. (1975). Localization and anatomical identification of theta and complex spike cells in dorsal hippocampal formation of rats. *Exp Neurol*, 49(1 Pt 1):299–313.

Fox, S. E. and Ranck, J. B. (1981). Electrophysiological characteristics of hippocampal complex-spike cells and theta cells. *Exp Brain Res*, 41(3-4):399–410.

Frank, L. M., Brown, E. N., and Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron*, 27(1):169–178.

Freund, T. and Buzsáki, G. (1996). Interneurons of the hippocampus. *Hippocampus*, 6(4):347–470.

Fuhrmann, G., Segev, I., Markram, H., and Tsodyks, M. (2002). Coding of temporal information by activity-dependent synapses. *J Neurophysiol*, 87(1):140–148.

Fusi, S., Drew, P. J., and Abbott, L. F. (2005). Cascade models of synaptically stored memories. *Neuron*, 45(4):599–611.

Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., and Moser, M.-B. (2004). Spatial representation in the entorhinal cortex. *Science*, 305(5688):1258–1264.

Gerstner, W. and Kistler, W. (2002). *Spiking neural models*. Cambridge University Press. ISBN 0-521-89079-9.

Gibson, W. G., Robinson, J., and Bennett, M. R. (1991). Probabilistic secretion of quanta in the central nervous system: granule cell synaptic control of pattern sepa-

ration and activity regulation. *Philos Trans R Soc Lond B Biol Sci*, 332(1264):199–220.

Giovanello, K. S. and Verfaellie, M. (2001). Memory systems of the brain: a cognitive neuropsychological analysis. *Semin Speech Lang*, 22(2):107–116.

Graham, B. P. (2001). Pattern recognition in a compartmental model of a CA1 pyramidal neuron. *Network*, 12(4):473–492.

Greicius, M. D., Krasnow, B., Boyett-Anderson, J. M., Eliez, S., Schatzberg, A. F., Reiss, A. L., and Menon, V. (2003). Regional analysis of hippocampal activation during memory encoding and retrieval: fMRI study. *Hippocampus*, 13(1):164–174.

Hasselmo, M. E. (2005). What is the function of hippocampal theta rhythm?-Linking behavioral data to phasic properties of field potential and unit recording data. *Hippocampus*, 15(7):936–949.

Hasselmo, M. E., Fransen, E., Dickson, C., and Alonso, A. A. (2000). Computational modeling of entorhinal cortex. *Ann N Y Acad Sci*, 911:418–446.

Hasselmo, M. E., Schnell, E., and Barkai, E. (1995). Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region CA3. *J Neurosci*, 15(7 Pt 2):5249–5262.

Hasselmo, M. E., Wyble, B. P., and Wallenstein, G. V. (1996). Encoding and retrieval of episodic memories: role of cholinergic and GABAergic modulation in the hippocampus. *Hippocampus*, 6(6):693–708.

Haykin, S. (1999). *Neural Networks*. Prentice Hall International, Inc., second edition. ISBN 0-13-908385-5.

Hebb, D. (1949). *The Organization of Behavior: A Neuropsychological Theory*. New York: Wiley.

Hertz, J., Krogh, A., and Palmer, R. (1991). *Introduction to the theory of neural computation*. Westview Press. ISBN 0-201-51560-1.

Hines, M. L. and Carnevale, N. T. (1997). The NEURON simulation environment. *Neural Comput*, 9(6):1179–1209.

Holland, J. (1992). *Adaptation in Natural and Artificial Systems*. The MIT Press. ISBN: 0262581116.

Hooper, S. L. (1998). Transduction of temporal patterns by single neurons. *Nat Neurosci*, 1(8):720–726.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci U S A*, 79(8):2554–2558.

Hopfield, J. J. (1995). Pattern recognition computation using action potential timing for stimulus representation. *Nature*, 376(6535):33–36.

Howard, M. W., Fotedar, M. S., Datey, A. V., and Hasselmo, M. E. (2005). The temporal context model in spatial navigation and relational learning: toward a common explanation of medial temporal lobe function across domains. *Psychol Rev*, 112(1):75–116.

Huang, Y. Y., Kandel, E. R., Varshavsky, L., Brandon, E. P., Qi, M., Idzerda, R. L., McKnight, G. S., and Bourtchouladze, R. (1995). A genetic test of the effects of mutations in PKA on mossy fiber LTP and its relation to spatial and contextual learning. *Cell*, 83(7):1211–1222.

Huerta, P. T. and Lisman, J. E. (1993). Heightened synaptic plasticity of hippocampal CA1 neurons during a cholinergically induced rhythmic state. *Nature*, 364(6439):723–725.

Izhikevich, E. M., Desai, N. S., Walcott, E. C., and Hoppensteadt, F. C. (2003). Bursts as a unit of neural information: selective communication via resonance. *Trends Neurosci*, 26(3):161–167.

Jaeger, H. (2001a). The "echo state" approach to analysing and training recurrent neural networks. Technical report, GMD Report 148, German National Research Center for Information Technology.

Jaeger, H. (2001b). Short term memory in echo state networks. Technical report, GMD Report 152, German National Research Center for Information Technology.

Jensen, O. and Lisman, J. E. (2000). Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase coding. *J Neurophysiol*, 83(5):2602–2609.

Johnston, D. and Amaral, D. (1998). Hippocampus. In Shepherd, G., editor, *The Synaptic Organization of the Brain*. Oxford University Press, fourth edition edition.

Jonas, P., Major, G., and Sakmann, B. (1993). Quantal components of unitary EPSCs at the mossy fibre synapse on CA3 pyramidal cells of rat hippocampus. *J Physiol*, 472:615–663.

Jung, M. W. and McNaughton, B. L. (1993). Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus*, 3(2):165–182.

Káli, S. and Dayan, P. (2000). The involvement of recurrent connections in area CA3 in establishing the properties of place fields: a model. *J Neurosci*, 20(19):7463–7477.

Káli, S. and Dayan, P. (2004). Off-line replay maintains declarative memories in a model of hippocampal-neocortical interactions. *Nat Neurosci*, 7(3):286–294.

Kesner, R. P., Lee, I., and Gilbert, P. (2004). A behavioral assessment of hippocampal function based on a subregional analysis. *Rev Neurosci*, 15(5):333–351.

Koch, C. (1999). *Biophysics of Computation*. Oxford University Press. ISBN 0-19-510491-9.

Koch, C. and Segev, I., editors (1998). *Methods in Neuronal Modeling*. The MIT Press, second edition edition. ISBN 0-262-11231-0.

Koch, C. and Segev, I. (2000). The role of single neurons in information processing. *Nat Neurosci*, 3 Suppl:1171–1177.

Korn, H. and Farber, D. (1993). Transmitter release mechanism: Relevance for neural network functions. *in Exploring Brain Functions: Models in Neuroscience, John Wiley*, 5:1–7.

Land, L. and Zucker, R. S. (1994). ca^{2+} cooperativity in neurosecretion measured using photolabile Ca^{2+} chelators. *J Neurophysiol*, 72(2):825–830.

Lapicque, L. (1907). Recherches quantitatifs sur l'excitation electrique des nerfs traitée comme une polarization. *J. Physiol. Paris*, 9:620–635.

Larson, J., Wong, D., and Lynch, G. (1986). Patterned stimulation at the theta frequency is optimal for the induction of hippocampal long-term potentiation. *Brain Res*, 368(2):347–350.

Lassalle, J. M., Bataille, T., and Halley, H. (2000). Reversible inactivation of the hippocampal mossy fiber synapses in mice impairs spatial learning, but neither consolidation nor memory retrieval, in the Morris navigation task. *Neurobiol Learn Mem*, 73(3):243–257.

Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M.-B., and Moser, E. I. (2004). Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science*, 305(5688):1295–1298.

Levy, W. B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus*, 6(6):579–590.

Levy, W. B., Desmond, N. L., and Zhang, D. X. (1998). Perforant path activation modulates the induction of long-term potentiation of the schaffer collateral–hippocampal CA1 response: theoretical and experimental analyses. *Learn Mem*, 4(6):510–518.

Liaw, J. S. and Berger, T. W. (1996). Dynamic synapse: a new concept of neural representation and computation. *Hippocampus*, 6(6):591–600.

Lisman, J., Schulman, H., and Cline, H. (2002). The molecular basis of CaMKII function in synaptic and behavioural memory. *Nat Rev Neurosci*, 3(3):175–190.

Lorente de Nó, R. (1934). Studies on the structure of the cerebral cortex ii. continuation of the study of the ammonic system. *J. Psychol. Neurol.*, 46:113–177.

Maass, W. and Markram, H. (2002). Synapses as dynamic memory buffers. *Neural Netw*, 15(2):155–161.

Maass, W., Natschläger, T., and Markram, H. (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput*, 14(11):2531–2560.

MacKay, D. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press. ISBN 0-521-64298-1.

Magee, J. C. and Cook, E. P. (2000). Somatic EPSP amplitude is independent of synapse location in hippocampal pyramidal neurons. *Nat Neurosci*, 3(9):895–903.

Maguire, E. A. (2001). The retrosplenial contribution to human navigation: a review of lesion and neuroimaging findings. *Scand J Psychol*, 42(3):225–238.

Maguire, E. A., Burgess, N., Donnett, J. G., Frackowiak, R. S., Frith, C. D., and O'Keefe, J. (1998). Knowing where and getting there: a human navigation network. *Science*, 280(5365):921–924.

Maguire, E. A., Spiers, H. J., Good, C. D., Hartley, T., Frackowiak, R. S. J., and Burgess, N. (2003). Navigation expertise and the human hippocampus: a structural brain imaging analysis. *Hippocampus*, 13(2):250–259.

Markram, H., Wang, Y., and Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *Proc Natl Acad Sci U S A*, 95(9):5323–5328.

Marr, D. (1971). Simple memory: a theory for archicortex. *Philos Trans R Soc Lond B Biol Sci*, 262(841):23–81.

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.

McNaughton, B. and Smolensky, P. (1991). Connectionist and neural modeling: Converging in the hippocampus. In Liister, R. and Weingartner, H., editors, *Perspectives on cognitive neurosciences*. New York Oxford Univ. Press.

McNaughton, B. L., Barnes, C. A., and O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp Brain Res*, 52(1):41–49.

Mechler, F., Victor, J. D., Purpura, K. P., and Shapley, R. (1998). Robust temporal coding of contrast by V1 neurons for transient but not for steady-state stimuli. *J Neurosci*, 18(16):6583–6598.

Mel, B. (1994). Information processing in dendritic trees. *Neural Computation*, 6:1031–1085.

Migliore, M., Cook, E. P., Jaffe, D. B., Turner, D. A., and Johnston, D. (1995). Computer simulations of morphologically reconstructed CA3 hippocampal neurons. *J Neurophysiol*, 73(3):1157–1168.

Migliore, M., Hoffman, D. A., Magee, J. C., and Johnston, D. (1999). Role of an A-type K+ conductance in the back-propagation of action potentials in the dendrites of hippocampal pyramidal neurons. *J Comput Neurosci*, 7(1):5–15.

Misják, F., Lengyel, M., and Érdi, P. (2001). Episodic memory and cognitive map in a rate model network of the rat hippocampus. *Lecture Notes in Computer Science*, 2130:1135–1140.

Morgane, P. J., Galler, J. R., and Mokler, D. J. (2005). A review of systems and networks of the limbic forebrain/limbic midbrain. *Prog Neurobiol*, 75(2):143–160.

Morris, R. G. (1989). Synaptic plasticity and learning: selective impairment of learning rats and blockade of long-term potentiation in vivo by the N-methyl-D-aspartate receptor antagonist AP5. *J Neurosci*, 9(9):3040–3057.

Muller, R. U., Kubie, J. L., and Ranck, J. B. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *J Neurosci*, 7(7):1935–1950.

Nakazawa, K., McHugh, T. J., Wilson, M. A., and Tonegawa, S. (2004). NMDA receptors, place cells and hippocampal spatial memory. *Nat Rev Neurosci*, 5(5):361–372.

Natschläger, T. and Maass, W. (2001). Computing the optimally fitted spike train for a synapse. *Neural Comput*, 13(11):2477–2494.

Natschläger, T., Maass, W., and Zador, A. (2001). Efficient temporal processing with biologically realistic dynamic synapses. *Network*, 12(1):75–87.

Nicholls, J., Martin, R., and Wallace, B. (1992). *From neuron to brain*. Sinauer, third edition edition. ISBN 0-87893-580-0.

Norman, K. A. and O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. *Psychol Rev*, 110(4):611–646.

Okabe, S., Kim, H. D., Miwa, A., Kuriu, T., and Okado, H. (1999). Continual remodeling of postsynaptic density and its regulation by synaptic activity. *Nat Neurosci*, 2(9):804–811.

O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Exp Neurol*, 51(1):78–109.

O'Keefe, J. and Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381(6581):425–428.

O'Keefe, J., Burgess, N., Donnett, J. G., Jeffery, K. J., and Maguire, E. A. (1998). Place cells, navigational accuracy, and the human hippocampus. *Philos Trans R Soc Lond B Biol Sci*, 353(1373):1333–1340.

O'Keefe, J. and Conway, D. H. (1978). Hippocampal place units in the freely moving rat: why they fire where they fire. *Exp Brain Res*, 31(4):573–590.

O'Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely moving rat. *Brain Research*, 31:573–590.

O'Keefe, J. and Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford University Press.

O'Keefe, J. and Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3(3):317–330.

O'Reilly, R. C. and McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall: avoiding a trade-off. *Hippocampus*, 4(6):661–682.

O'Reilly, R. C. and Munakata, Y. (2000). *Computational Explorations in Cognitive Neuroscience*. The MIT Press. ISBN 0-262-65054-1.

Otto, T., Eichenbaum, H., Wiener, S. I., and Wible, C. G. (1991). Learning-related patterns of CA1 spike trains parallel stimulation parameters optimal for inducing hippocampal long-term potentiation. *Hippocampus*, 1(2):181–192.

Patton, P. E. and McNaughton, B. (1995). Connection matrix of the hippocampal formation: I. The dentate gyrus. *Hippocampus*, 5(4):245–286.

Poirazi, P., Brannon, T., and Mel, B. W. (2003). Arithmetic of subthreshold synaptic summation in a model CA1 pyramidal cell. *Neuron*, 37(6):977–987.

Press, W., Teukolsky, S., Vetterling, W., and Flannery, B. (1994). *Numerical Recipes in C*. Cambridge University Press, second edition edition. ISBN 0-521-43108-5.

Rall, W. and Agmon-Snir, H. (1998). Cable theory for dendritic neurons. In Koch, C. and Segev, I., editors, *Methods in Neuronal Modeling*. The MIT Press, second edition edition.

Ranck, J. B. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. I. Behavioral correlates and firing repertoires. *Exp Neurol*, 41(2):461–531.

Rao, R. P. and Sejnowski, T. J. (2001). Spike-timing-dependent Hebbian plasticity as temporal difference learning. *Neural Comput*, 13(10):2221–2237.

Remondes, M. and Schuman, E. M. (2004). Role for a cortical input to hippocampal area CA1 in the consolidation of a long-term memory. *Nature*, 431(7009):699–703.

Richmond, B. J., Optican, L. M., and Spitzer, H. (1990). Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. I. Stimulus-response relations. *J Neurophysiol*, 64(2):351–369.

Richter-Levin, G. (2004). The amygdala, the hippocampus, and emotional modulation of memory. *Neuroscientist*, 10(1):31–39.

Rieke, F., Warland, D., van Steveninck, R., and Bialek, W. (1996). *Spikes: exploring the neural code*. The MIT Press. ISBN 0-262-18174-6.

Rizzoli, S. O. and Betz, W. J. (2005). Synaptic vesicle pools. *Nat Rev Neurosci*, 6(1):57–69.

Rolls, E. and Treves, A. (1998). *Neural Networks and Brain Function*. Oxford University Press. ISBN 0-19-852432-3.

Rolls, E. T., Stringer, S. M., and Trappenberg, T. P. (2002). A unified model of spatial and episodic memory. *Proc Biol Sci*, 269(1496):1087–1093.

Ross, S. (1996). *Stochastic Processes*. John Wiley & Sons, Inc., second edition. ISBN 0-471-12062-6.

Rotter, S. and Diesmann, M. (1999). Exact digital simulation of time-invariant linear systems with applications to neuronal modeling. *Biol Cybern*, 81(5-6):381–402.

Samsonovich, A. V. and Ascoli, G. A. (2005). A simple neural network model of the hippocampus suggesting its pathfinding role in episodic memory retrieval. *Learn Mem*, 12(2):193–208.

Schikorski, T. and Stevens, C. F. (1997). Quantitative ultrastructural analysis of hippocampal excitatory synapses. *J Neurosci*, 17(15):5858–5867.

Schneidman, E., Freedman, B., and Segev, I. (1998). Ion channel stochasticity may be critical in determining the reliability and precision of spike timing. *Neural Computation*, 10:1679–1793.

Scoville, W. B. and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *J Neurol Neurosurg Psychiatry*, 20(1):11–21.

Segev, I. and Burke, R. (1998). Compartmental models of complex neurons. In Koch, C. and Segev, I., editors, *Methods in neural modeling*. The MIT Press, second edition edition.

Senn, W., Markram, H., and Tsodyks, M. (2001). An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Comput*, 13(1):35–67.

Shadlen, M. N. and Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Curr Opin Neurobiol*, 4(4):569–579.

Shadlen, M. N. and Newsome, W. T. (1995). Is there a signal in the noise? *Curr Opin Neurobiol*, 5(2):248–250.

Shapiro, M. L., Tanila, H., and Eichenbaum, H. (1997). Cues that hippocampal place cells encode: dynamic and hierarchical representation of local and distal stimuli. *Hippocampus*, 7(6):624–642.

Shastri, L. (2002). Episodic memory and cortico-hippocampal interactions. *Trends Cogn Sci*, 6(4):162–168.

Shepherd, M. (1994). *Neurobiology*. Oxford University Press, third edition edition. ISBN 0-19-508842-5.

Shouval, H. Z., Bear, M. F., and Cooper, L. N. (2002). A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. *Proc Natl Acad Sci U S A*, 99(16):10831–10836.

Singer, W. (1994). Coherence as an organizing principle of cortical functions. *Int Rev Neurobiol*, 37:153–83; discussion 203–7.

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):149–172.

Smith, A. C., Wu, X. B., and Levy, W. B. (2000). Controlling activity fluctuations in large, sparsely connected random networks. *Network*, 11(1):63–81.

Softky, W. R. (1995). Simple codes versus efficient codes. *Curr Opin Neurobiol*, 5(2):239–247.

Softky, W. R. and Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci*, 13(1):334–350.

Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci*, 3(9):919–926.

Spruston, N. and Johnston, D. (1992). Perforated patch-clamp analysis of the passive membrane properties of three classes of hippocampal neurons. *J Neurophysiol*, 67(3):508–529.

Spruston, N., Lübke, J., and Frotscher, M. (1997). Interneurons in the stratum lucidum of the rat hippocampus: an anatomical and electrophysiological characterization. *J Comp Neurol*, 385(3):427–440.

Spruston, N., Schiller, Y., Stuart, G., and Sakmann, B. (1995). Activity-dependent action potential invasion and calcium influx into hippocampal CA1 dendrites. *Science*, 268(5208):297–300.

Squire, L. and Kandel, E. (1999). *Memory, from mind to molecules*. W. H. Freeman and Company. ISBN 0-7167-6037-1.

Squire, L., Shimamura, A., and Amaral, D. (1989). Memory and the hippocampus. In Byrne, J. and Berry, W., editors, *Neural models of plasticity: Experimental and theoretical approaches*, pages 208–239. Sab Diego, CA: Academic Press.

Squire, L. R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol Rev*, 99(2):195–231.

Squire, L. R. and Zola, S. M. (1996). Structure and function of declarative and non-declarative memory systems. *Proc Natl Acad Sci U S A*, 93(24):13515–13522.

Stern, C. E. and Hasselmo, M. E. (1999). Bridging the gap: integrating cellular and functional magnetic resonance imaging studies of the hippocampus. *Hippocampus*, 9(1):45–53.

Stewart, M. and Fox, S. E. (1990). Do septal neurons pace the hippocampal theta rhythm? *Trends Neurosci*, 13(5):163–168.

Sullivan, D. W. and Levy, W. B. (2004). Quantal synaptic failures enhance performance in a minimal hippocampal model. *Network*, 15(1):45–67.

Suzuki, W. A. and Amaral, D. G. (1994). Topographic organization of the reciprocal connections between the monkey entorhinal cortex and the perirhinal and parahippocampal cortices. *J Neurosci*, 14(3 Pt 2):1856–1877.

Swanson, L. W. and Cowan, W. M. (1977). An autoradiographic study of the organization of the efferent connections of the hippocampal formation in the rat. *J Comp Neurol*, 172(1):49–84.

Thompson, L. T. and Best, P. J. (1990). Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain Res*, 509(2):299–308.

Thomson, A. M. and Deuchars, J. (1994). Temporal and spatial properties of local circuits in neocortex. *Trends Neurosci*, 17(3):119–126.

Thurbon, D., Field, A., and Redman, S. (1994). Electrotonic profiles of interneurons in stratum pyramidale of the CA1 region of rat hippocampus. *J Neurophysiol*, 71(5):1948–1958.

Torioka, T. (1979). Pattern separability in a random neural net with inhibitory connections. *Biol Cybern*, 34(1):53–62.

Traub, R. D., Spruston, N., Soltesz, I., Konnerth, A., Whittington, M. A., and Jefferys, G. R. (1998). Gamma-frequency oscillations: a neuronal population phenomenon, regulated by synaptic and intrinsic cellular processes, and inducing synaptic plasticity. *Prog Neurobiol*, 55(6):563–575.

Treves, A. (2004). Computational constraints between retrieving the past and predicting the future, and the CA3-CA1 differentiation. *Hippocampus*, 14(5):539–556.

Treves, A. and Rolls, E. (1991). What determines the capacity of autoassociative memories in the brain? *Network*, 2:371–397.

Treves, A. and Rolls, E. T. (1992). Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus*, 2(2):189–199.

Treves, A. and Rolls, E. T. (1994). Computational analysis of the role of the hippocampus in memory. *Hippocampus*, 4(3):374–391.

Tsien, J. Z., Huerta, P. T., and Tonegawa, S. (1996). The essential role of hippocampal CA1 NMDA receptor-dependent synaptic plasticity in spatial memory. *Cell*, 87(7):1327–1338.

Tsodyks, M., Pawelzik, K., and Markram, H. (1998). Neural networks with dynamic synapses. *Neural Comput*, 10(4):821–835.

Tsodyks, M., Uziel, A., and Markram, H. (2000). Synchrony generation in recurrent networks with frequency-dependent synapses. *J Neurosci*, 20(1):RC50.

Tsodyks, M. V. and Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc Natl Acad Sci U S A*, 94(2):719–723.

Tulving, E. and Markowitsch, H. J. (1998). Episodic and declarative memory: role of the hippocampus. *Hippocampus*, 8(3):198–204.

Tulving, E. and Schacter, D. L. (1990). Priming and human memory systems. *Science*, 247(4940):301–306.

Urban, N. N., Henze, D. A., and Barrionuevo, G. (2001). Revisiting the role of the hippocampal mossy fiber synapse. *Hippocampus*, 11(4):408–417.

Vaadia, E., Haalman, I., Abeles, M., Bergman, H., Prut, Y., Slovin, H., and Aertsen, A. (1995). Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature*, 373(6514):515–518.

van Ooyen, A. (2001). Competition in the development of nerve connections: a review of models. *Network*, 12(1):R1–47.

van Rossum, M. C., Bi, G. Q., and Turrigiano, G. G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J Neurosci*, 20(23):8812–8821.

Vanderwolf, C. H. (1988). Cerebral activity and behavior: control by central cholinergic and serotonergic systems. *Int Rev Neurobiol*, 30:225–340.

Varela, J. A., Sen, K., Gibson, J., Fost, J., Abbott, L. F., and Nelson, S. B. (1997). A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex. *J Neurosci*, 17(20):7926–7940.

Vargha-Khadem, F., Gadian, D. G., Watkins, K. E., Connelly, A., Paesschen, W. V., and Mishkin, M. (1997). Differential effects of early hippocampal pathology on episodic and semantic memory. *Science*, 277(5324):376–380.

Vayssettes-Courchay, C. and Sessler, F. M. (1983). [Evidence for sensory convergences in rat entorhinal cortex]. *C R Seances Acad Sci III*, 296(18):877–879.

Vida, I. and Frotscher, M. (2000). A hippocampal interneuron associated with the mossy fiber system. *Proc Natl Acad Sci U S A*, 97(3):1275–1280.

White, J. A., Rubinstein, J. T., and Kay, A. R. (2000). Channel noise in neurons. *Trends Neurosci*, 23(3):131–137.

Willshaw, D. J., Buneman, O. P., and Longuet-Higgins, H. C. (1969). Non-holographic associative memory. *Nature*, 222(197):960–962.

Wilson, M. A. and McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, 261(5124):1055–1058.

Wood, E. R., Agster, K. M., and Eichenbaum, H. (2004). One-trial odor-reward association: a form of event memory not dependent on hippocampal function. *Behav Neurosci*, 118(3):526–539.

Wood, E. R., Dudchenko, P. A., and Eichenbaum, H. (1999). The global record of memory in hippocampal neuronal activity. *Nature*, 397(6720):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(3):623–633.

Wu, Z., Desmond, N. L., and Levy, W. B. (1998). Homosynaptic long-term depression of CA3-CA3 synapses in the in vivo hippocampus. *Brain Res*, 789(2):335–338.

Zador, A. (1998). Impact of synaptic unreliability on the information transmitted by spiking neurons. *J Neurophysiol*, 79(3):1219–1229.

Zador, A. M. and Dobrunz, L. E. (1997). Dynamic synapses in the cortex. *Neuron*, 19(1):1–4.

Zeineh, M. M., Engel, S. A., Thompson, P. M., and Bookheimer, S. Y. (2003). Dynamics of the hippocampus during encoding and retrieval of face-name pairs. *Science*, 299(5606):577–580.

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. *J Neurosci*, 6(10):2950–2967.