STDP, Rate-coded Hebbian Learning and Auto-associative Network Models of the Hippocampus

Daniel Bush

Submitted for the degree of DPhil
University of Sussex
September 2008

STDP, Rate-coded Hebbian Learning and Auto-associative Network Models of the Hippocampus

Daniel Bush

ABSTRACT

Auto-associative network models have proven extremely useful in modelling the hypothesised function of the CA3 region of the hippocampus in declarative memory. To date, the majority of these models have made use of Hebbian plasticity rules mediated by correlations between mean firing rates. However, recent neurobiological evidence suggests that synaptic plasticity in the hippocampus, and many other cortical regions, also depends explicitly on the temporal relationship between afferent action potentials and efferent spiking – a phenomena known as spike-timing dependent plasticity (STDP). Few attempts have been made to reconcile previous rate-coded Hebbian learning rules or auto-associative network function with this novel plasticity formulation. Further complications arise from the fact that there are many computational interpretations of the empirical data regarding STDP, each of which can have a unique effect on emergent network properties. The aims of this research are therefore three-fold: firstly, to provide a comprehensive description of the emergent dynamics generated by a wide range of STDP implementations within a biologically inspired spiking recurrent neural network; secondly, to reconcile these emergent dynamics with those generated by previous rate-coded Hebbian learning rules, as characterised by the BCM formulation; and finally, to employ these STDP implementations within a simple rate-coded auto-associative network model. The results presented demonstrate that several incarnations of the STDP rule can mediate rate-coded Hebbian learning and replicate numerous other features of synaptic and neural dynamics that are realistic of the CA3 region. These forms of STDP are consequently demonstrated to mediate efficient and robust auto-associative network function, although several issues which might affect the long-term stability and performance of the plasticity rule are identified and discussed. The results should therefore allow the successes of previous auto-associative network models of the hippocampus to be replicated, whilst increasing their versatility and computational power and providing them with a firmer basis in modern neurobiology.

Submitted for the degree of DPhil
University of Sussex
September 2008

ACKNOWLEDGEMENTS

I owe an enormous debt of gratitude to numerous people for providing help and support during the course of my DPhil research. First and foremost, I would like to thank Laura, for being an invaluable source of inspiration, comfort and encouragement, and for tolerating my caprice and absence during the course of writing this thesis. Similarly, I would like to thank my family, for their guidance and good humour; and David and Theresa, for their hospitality and tolerance. The constant provision of tea from all of those mentioned above was also a critical contribution to the success of this work. Moreover, I would like to thank Andy, who spent (and continues to spend) a great deal of time assisting me with this research; Phil and Mick, for supervising my DPhil studies and tolerating my unique working style; and finally, to the BBSRC, for the generous funding which has allowed me to indulge my thirst for knowledge in many far flung corners of the globe.

“Man is a reasonable being; and as such, receives from science his proper food and nourishment: But so narrow are the bounds of human understanding, that little satisfaction can be hoped for in this particular, either from the extent or security of his acquisitions. Man is a sociable, no less than a reasonable being: But neither can he always enjoy company agreeable and amusing, or preserve the proper relish for them. Man is also an active being; and from that disposition, as well as from the various necessities of human life, must submit to business and occupation: But the mind requires some relaxation, and cannot always support its bent to care and industry. It seems, then, that nature has pointed out a mixed kind of life as most suitable to human race, and secretly admonished them to allow none of their biases to *draw* too much, so as to incapacitate them for other occupations and entertainments. Indulge your passion for science, says he, but let your science be human, and such as may have a direct reference to action and society. Abstruse thought and profound researches I prohibit, and will severely punish, by the pensive melancholy which they introduce, by the endless uncertainty in which they involve you, and by the cold reception which your pretended discoveries shall meet with, when communicated. Be a philosopher; but, amidst all your philosophy, be still a man"

Hume (1777)

CHAPTER ONE: INTRODUCTION AND OVERVIEW OF THESIS

1.1 Introduction

The phenomenon of synaptic plasticity is widely believed to mediate memory and the activity dependent development of neural circuits in a diverse range of species. The majority of empirical observations of synaptic plasticity have been made in the hippocampus – a brain region which has long been implicated in human declarative memory function and represents an ideal testing ground for the synaptic plasticity and memory hypothesis. Observations of long-term potentiation and depression (LTP and LTD) have most commonly been induced in this region by the tetanic stimulation of pre- and post- synaptic neurons. Similarly, computational models of synaptic plasticity have conventionally been mediated by correlations between pre- and post- synaptic firing rates. However, these models face a significant problem with stability, due to the positive feedback system that consequently arises between potentiation and post-synaptic firing rates.

Recent neurobiological investigation in the hippocampus has delineated a theory of spike-timing dependent plasticity (STDP) which dictates that changes in the strength of synapses are primarily determined by the relative timing of pre- and post- synaptic spikes. STDP has proven popular in computational modelling studies because it appears to address the issue of positive feedback by stimulating competition between synapses for the control of post-synaptic firing. This competition can provide the unsupervised regulation of neural activity by reducing synaptic weights in response to increases in pre-synaptic firing rate. Although this inherent homeostasis is advantageous for the stable operation of artificial neural networks (ANNs), it directly contradicts previous computational models and observations of synaptic plasticity. Furthermore, due to the difficulty of making experimental observations at the synaptic level, there is little empirical data to inform the modelling choices involved in developing a biologically realistic STDP rule.

There is a pressing need, therefore, to reconcile STDP with both rate-coded models of synaptic plasticity and the role of the hippocampus in declarative memory. The suspected function of the hippocampus has been most successfully investigated using auto-associative network models – recurrent ANNs which implement synaptic plasticity in order to store applied activity patterns and recall them at a later date from partial or noisy cues. These models are attractive because they offer a unified description of encoding, storage and recall processes within a single framework. Surprisingly, few studies have examined the emergent dynamics of STDP within a recurrent neural architecture, and fewer still have attempted to reconcile this asymmetric plasticity rule with auto-associative network function. Furthermore, common auto-associative network models of the hippocampus rely explicitly on a Hebbian learning mechanism, and therefore represent an ideal context in which to attempt a reconciliation of STDP with previous rate-coded plasticity models.

Consequently, the aims of this research are three-fold: firstly, to provide an extensive characterisation of the emergent synaptic and neural dynamics generated by a variety of STDP implementations within a spiking recurrent neural network; secondly, to more specifically assess how STDP might be

reconciled with the properties of rate-coded Hebbian learning and observations of synaptic plasticity made in the hippocampus; and finally, to employ those STDP implementations that can mediate rate-coded processing in a simple auto-associative network model, and thereby assess the functional performance of the plasticity model within this context. It is demonstrated that STDP can mediate rate-coded Hebbian learning and reproduce several features of synaptic plasticity observed in the CA3 region, providing that certain constraints are placed on the profile of the asymmetric learning window and spike pairing scheme employed. This allows STDP to be directly reconciled with the BCM formulation of metaplasticity. Consequently, it is demonstrated that STDP can also mediate efficient and robust auto-associative network function, and some of the issues relating to the long-term stability of these models are discussed.

The research in this thesis therefore provides the first comprehensive description of STDP in a spiking recurrent network model that is consistently inspired by the neurobiology of the hippocampus, the cortical region in which synaptic plasticity is most frequently observed *in vivo*. The fact that multiple STDP implementations are examined within a single framework facilitates comparison and allows the specific effects of each parameter that defines the plasticity rule to be extricated. Several forms of STDP that can mediate and integrate both rate and temporally coded data are identified, and these allow the complexity, computational power and biological realism of previous Hebbian learning rules to be improved without compromising their function. Finally, the first known demonstration of asymmetric STDP successfully mediating auto-associative network function is provided, and some of the issues regarding the role of the plasticity rule within this context are discussed.

1.2 Thesis Overview

Chapter Two provides a detailed review of the existing literature regarding the biological and computational properties of Hebbian learning rules and STDP. A particular emphasis is placed on the multiple STDP implementations which have arisen in the absence of sufficient biological data, and how these can affect the emergent dynamics of the plasticity rule. The common disparities between the properties of STDP and rate-coded Hebbian learning rules are also highlighted, and some possible approaches to reconciling these plasticity models are described.

Chapter Three provides a detailed review of the neuroscience of the hippocampus, the cortical region in which empirical observations of synaptic plasticity are most often made, and describes the postulated role of this brain region in declarative memory. The theory that this mnemonic function is mediated, at least in part, by auto-associative network dynamics is introduced, and the fact that synaptic plasticity has a critical role in this process is emphasised. The successes of auto-associative network models of the hippocampus are then described, and their biological foundations examined in detail.

Chapter Four identifies a shortfall in previous computational neuroscience research regarding the integration of STDP, rate-coded Hebbian learning and auto-associative network models of the

hippocampus. A more detailed examination of the known properties of synaptic and neural dynamics within the CA3 region and how these will constrain the artificial neural network model examined in this research is then made.

Chapter Five defines the experimental methods and details of the neural spiking, synaptic plasticity and network models being implemented in more detail.

Chapter Six describes the results of the first set of simulations examined in this research, which aim to provide a detailed description of the neural and synaptic dynamics that are generated by STDP within a spiking recurrent network. Several key features of the STDP rule are manipulated in order to assess their effects on emergent properties, including the parameters which define the profile of the asymmetric learning window, the use of weight-dependent potentiation (i.e. multiplicative STDP) and the spike pairing scheme employed.

Chapter Seven describes the results of the second set of simulations examined in this research, which aim to more specifically reconcile STDP with previous rate-coded Hebbian learning rules and observations of synaptic plasticity made *in vivo*. It is consequently demonstrated that several distinct STDP implementations can generate the emergent synaptic dynamics that are required for successful auto-associative network function.

Chapter Eight describes the results of the final set of simulations examined in this research, which assess the functional performance of STDP within a simple auto-associative network model. Those forms of STDP that can mediate rate-coded learning are demonstrated to rapidly re-arrange synaptic weights in order to reflect discrete binary patterns of rate-coded activity, and produce significant and specific recall in unstimulated parts of these patterns when partial cues are subsequently applied. This represents the first demonstration of an asymmetric STDP rule successfully replicating the properties of rate-coded Hebbian learning and thereby mediating auto-associative network function.

Chapter Nine describes the conclusions of this research and outlines probable directions of future research. These will most likely address some outstanding issues with biological realism and long-term stability, and also expand and develop the auto-associative network model to integrate temporal coding, neuromodulation and more complex learning scenarios.

1.3 Previous Publications

Several sections of the results presented in this thesis have appeared elsewhere in peer-reviewed publications. These are listed below.

Chapter Six:

Bush D, Philippides A, Husbands P and O'Shea M. "Embodied Learning : Investigating Synaptic Weight Distributions in a Spiking Neural Network" in Kovacs T and Marshall JAR (eds) "The Proceedings of AISB'06: Adaptation in Artificial and Biological Systems". AISB (2006)

Bush D, Philippides A, Husbands P and O'Shea M. "Investigating STDP and LTP in a Spiking Neural Network" in Nolfi S, Baldassarre G, Calabretta R, Hallam J, Marocco D, Miglino O, Meyer J-A and Parisi D (eds) "From Animals to Animats 9: Proceedings of the Ninth International Conference on Simulation of Adaptive Behaviour". Springer Verlag (2006)

Chapters Seven and Eight:

Bush D, Philippides A, Husbands P and O'Shea M. "STDP and Auto-associative Network Function" in Mayor J, Ruh N and Plunkett K (eds) "Connectionist Models of Behavior and Cognition II: Proceedings of the Eleventh Neural Computation and Psychology Workshop". World Scientific (2009)

Chapter Eight and Future Directions:

Bush D, Philippides A, Husbands P and O'Shea M. "Theta Phase Coding and Acetylcholine Modulation" in Asada M, Hallam JCT, Meyer J-A and Tani J (eds) "From Animals to Animats 10: Proceedings of the Tenth International Conference on Simulation of Adaptive Behaviour". Springer Verlag (2008)

CHAPTER TWO: SYNAPTIC PLASTICITY

"The cerebral cortex is like a garden planted with innumerable trees ... which, thanks to intelligent cultivation, can multiply their branches and sink their roots deeper, producing fruits and flowers of ever greater variety and quality."

Ramon y Cajal, 1894

Aims

- To introduce and define the concept of Hebbian plasticity, which underlies the vast majority of computational models of learning and activity dependent development, and the phenomena of long-term potentiation and depression, which provide biological correlates for this process
- To introduce the synaptic plasticity and memory (SPM) hypothesis, which postulates that synaptic plasticity is necessary and sufficient to mediate memory formation and storage in the brain
- To highlight the fact that Hebbian learning is inherently unstable, and requires additional mechanisms that can introduce competition between synapses in order to regulate network activity and dynamically maintain the appropriate conditions for learning
- To describe some common mechanisms that are employed to achieve this competition, and the biological data upon which they are founded
- To introduce spike-timing dependent plasticity (STDP), a novel plasticity rule that operates using temporal coding, and can generate the homeostatic regulation of neural activity
- To describe the various common STDP implementations that have been examined in previous modelling studies, and the emergent neural and synaptic dynamics which they generate
- To describe the apparent inconsistencies between STDP and traditional observations of LTP / Hebbian learning rules, and how these might be reconciled

2.1 Hebbian Learning

In the late nineteenth century, William James formulated an elementary principle of psychology, "…[that] when two brain processes are active together or in immediate succession, one of them, on reoccurring, tends to propagate its excitement into the other" (James, 1890). A year earlier, Ramon y Cajal had postulated that the human brain is not a continuous web, but consists of distinct neurons which communicate at specialised junctions (Ramon y Cajal, 1889). A generalisation of James' concept; that activity-dependent changes in the strength of these junctions, which were later named 'synapses', might underlie the phenomena of learning and memory, followed shortly (Sherrington, 1897). This general hypothesis of synaptic plasticity, still based solely on theoretical considerations, was more precisely delineated by Donald Hebb almost sixty years later. He proposed "… [that] when an axon of cell A is near enough to excite cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased" (Hebb, 1949). Hebbian learning rules, which have come to include any plasticity formulation that is driven by causal correlations between pre- and post-synaptic activity, have been at the centre of computational models of memory and the activity dependent development of neural circuits ever since. Equation 2.1 describes a standard form of Hebbian learning that has been extensively utilised in computational modelling studies (Kempter, Gerstner and van Hemmen, 1999).

$$\Delta w_{ij} = \varepsilon r_i r_j$$

Equation 2.1: The simplest form of a Hebbian learning rule. Changes in the weight of a synapse from neuron j to neuron i (Δw_{ij}) are equal to the relative rates (r) or those neurons, multiplied by a learning rate (ε)

2.2 Long-term Potentiation

Hebb's hypothesis laid the foundations for a theoretical consideration of the neurological process by which transient experience is translated into enduring memory. Empirical evidence to support the theory of long-term changes in synaptic efficacy was not obtained until many years later, however, from the hippocampal formation of anaesthetised rabbits by Lomo and Bliss in 1973. They demonstrated that the pre-synaptic delivery of tetanic, high frequency stimulation (HFS) would cause an increase in the size of the evoked post-synaptic potential (EPSP) that could persist for many days. Since this seminal finding, a wealth of neurobiological investigation has elaborated on the mechanisms of this long-term potentiation (LTP) of synapses. Similarly, empirical support has allowed the computational examination of Hebbian learning rules within artificial neural networks (ANNs) to flourish. Despite the massive research interest, the problems associated with examining changes at the synaptic level mean that the underlying biological mechanisms have proven difficult to elucidate, and contrasting research findings have provoked significant debate. What has been made apparent, however, is that long-term plasticity is an almost universal property of excitatory synapses throughout the mammalian brain.

2.3 Synaptic Biology and LTP Induction

Despite the protean nature of activity-dependent synaptic change, research has shown that distinct characteristics are exhibited at varying locations within the brain. It has therefore been suggested that LTP may be more usefully viewed as a class of cellular and synaptic phenomena, expressed differently depending on cortical region, stage of development, and means of induction (Malenka and Bear 2004). These disparities may reflect the contrasting functions which long-term plasticity mediates in various processes of learning and development. NMDAr-dependent LTP in the hippocampus remains the most extensively studied, and can be observed throughout the mammalian brain (Malenka and Nicoll, 1999; Malenka and Bear, 2004; see also Sheng and Kim, 2002). In the hippocampus, the majority of synaptic transmission is mediated by the neurotransmitter glutamate. Upon stimulation, this amino acid is released from the pre-synaptic bouton, flows across the synaptic cleft and binds to receptors on the dendrites of the post-synaptic neuron. These receptors generally fall into two categories - N-methyl-D-aspartate (NMDA) and α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic (AMPA) - which have distinct characteristics, responses and distributions among synapses.

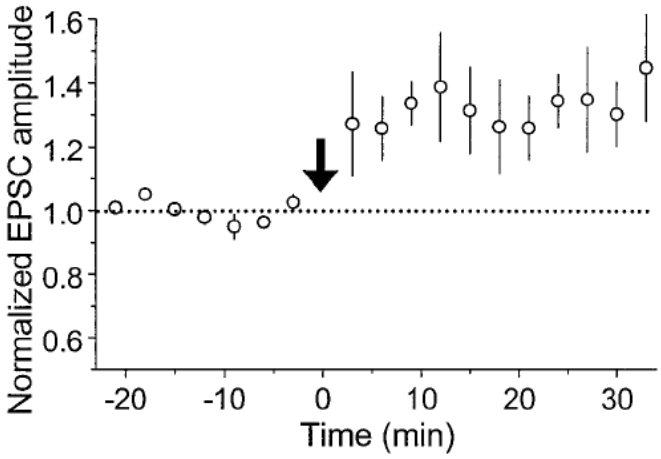

Figure 2.1: A typical increase in the amplitude of evoked post-synaptic currents (EPSCs) after an LTP induction protocol (at the black arrow) from Bi and Poo (1998)

The AMPA receptor (AMPAr) is permeable to monovalent cations (Na^+ and K^+), and provides the majority of inward current for generating post-synaptic responses when a cell is close to its resting membrane potential. Conversely, the NMDA receptor (NMDAr) demonstrates an extreme voltage dependence, as its channel is normally blocked by extra-cellular Magnesium ions. When the post-synaptic cell is hyperpolarised, therefore, this channel contributes little, but with sufficient depolarisation, Mg^{2+} dissociates from the NMDAr channel and drifts into the synaptic cleft, freeing the channel to bind with glutamate and allow an influx of Calcium ions (Bliss and Collingridge, 1993; Malenka and Nicoll, 1999; Bourtchouladze, 2002; Malenka and Bear, 2004). This rise in intracellular Ca^{2+} concentration signals the beginning of a cascade of neural events which mark the induction of long-term plasticity (for reviews see Bliss and Collingridge, 1993; Bear and Malenka, 1994; Malenka and Nicoll, 1999; Malenka and Bear, 2004; Cooke and Bliss, 2006). Hence, two events are essential to induce this prototypical form of LTP – the post-synaptic membrane must be sufficiently depolarised, in order to expel Mg^{2+} from the NMDA channels; and glutamate must simultaneously be present in order

to bind to these open channels and allow the influx of Ca^{2+}. The NMDA receptor is therefore often considered to be a molecular coincidence detector, which only operates under conditions of significant coincident pre- and post- synaptic activity.

2.4 The Expression of Long-term Plasticity

The molecular mechanisms which mediate the expression of LTP proceed through several stages following an influx of Ca^{2+} into the post-synaptic dendrite. Each of these stages is complex and none has been fully characterised by biological investigation. It is particularly difficult to differentiate between mediators - those molecules which are essential for LTP expression - and modulators - which merely regulate or adjust this expression. Well over a hundred different molecules have been implicated in one or another of these roles (Sanes and Lichtman, 1999; Lisman 2003; Malenka and Bear, 2004). For example, there are a wide variety of modulatory neurotransmitters (such as Acetylcholine and Dopamine) and circulating hormones (particularly those related to stress) which can have various effects on the induction, expression and maintenance of synaptic plasticity (e.g. Hasselmo, 2006; O'Carroll et al., 2006; Mockett et al., 2007). It is also clear that, during development, these mechanisms can change (Esteban et al., 2003; Jensen et al., 2003; Yasuda et al. 2003). It is interesting to note, however, that post-synaptic Ca^{2+} influx is consistently essential for the induction of LTP, even at synapses where the process is not NMDAr dependent (Johnston et al., 1992; Yeckel, Kapur and Johnston, 1999).

Initially, debate existed over whether long-term changes occur primarily at the pre- or post- synaptic terminal (Malenka and Nicoll, 1999; Malenka, 2003). Some evidence suggests that a rapid increase in the probability of neurotransmitter release (or the number of release sites) contributes, at least partly, to the expression of LTP, although findings are often contradictory (i.e. Larkman et al., 1992; see Crozier et al. 2004 and Lauri et al., 2007 for reviews). The existence of pre-synaptic changes, provoked by post-synaptic signals, would necessitate the existence of a retrograde messenger – a signal which moves backwards across the synaptic cleft from the dendrite to the bouton in order to convey the information that LTP has been induced (Malenka, 2003). There is some evidence – although it is not conclusive – that the gas Nitric Oxide (NO) operates in this role (Arancio et al., 1996; Holscher, 1997; Hawkins, Son and Arancio, 1998). However, a more significant body of research now suggests that the main (and perhaps sole) locus of LTP expression is post-synaptic.

The first stage of LTP expression lasts approximately 30 – 45 minutes and is generally described as short-term potentiation (or STP). Little is known about the mechanisms of STP, other than the fact that they are independent of protein kinase activity (Sweatt, 1999). This is demonstrated by the fact that broad spectrum protein kinase inhibitors limit the duration of synaptic potentiation to under an hour, but do not prevent it altogether (Soderling and Derkach, 2000; Lisman, Schulman and Cline, 2002; Nguyen and Woo, 2003; Bliss, Collingridge and Morris, 2007). Much more is known about the following stage – early LTP (E-LTP). The influx of Ca^{2+} into a dendrite activates several kinases which reside within the post-synaptic density (PSD) - a concentration of AMPA and NMDA receptors,

associated signalling proteins, scaffolding and adhesion molecules situated at the post-synaptic membrane. Among these are several critical mediators of E-LTP, including the Ca^{2+} / calmodulin dependent protein kinase (CaMKII), calcium / phospholipid-dependent protein kinase C (PKC), cAMP-dependent protein kinase (PKA), the family of tyrosine kinases and mitogen-activated protein kinase (MAPK).

It is hypothesised that these Calcium activated protein kinases phosphorylate certain proteins, which in turn regulate various properties of the post-synaptic neurotransmitter receptors in order to increase their efficacy (Benke et al., 1998; Malenka and Nicoll, 1999; Soderling and Derkach, 2000; see Kelly, 1991; Schulman, 1993; Sheng and Kim, 2002; Malenka and Bear, 2004 and Crozier et al. 2004 for reviews). The AMPA and NMDA receptors contain many target sites for phosphorylation, with the effect, among other things, of regulating their open channel probability, channel conductance and rate of desensitisation. Recent research, for example, has identified a CaMKII / PKC phosphorylation site on the AMPA receptor (Ser^{831}) which is likely modified during LTP, and which increases the single channel conductance of the receptor (see Sheng and Kim, 2002 for a review). Transgenic mice which lack Ser^{831} phosphorylation show impaired LTP (Lee et al., 2003).

In addition to modulating the properties of existing ion channels, it seems that these kinases also regulate the rate of insertion of AMPA receptors into the PSD from a latent sub-synaptic population – a process known as AMPAr trafficking (Lynch and Baudry, 1984). Research has demonstrated that significant changes in the concentration of AMPA receptors accompany long-term plastic changes in the strength of hippocampal synapses *in vivo* and *in vitro*, that these changes are NMDAr and CaMKII dependent, and that they are reversible (Hayashi et al., 2000; Heynen et al., 2000; Lu et al., 2001; see Malinow and Malenka, 2002; Song and Huganir, 2002; Crozier et al. 2004 and Derkach et al., 2007 for reviews). An increase in the number and efficacy of post-synaptic AMPA receptors provides "…an elegant means of producing stable, long-term biochemical changes in synaptic efficacy" (Crozier et al. 2004).

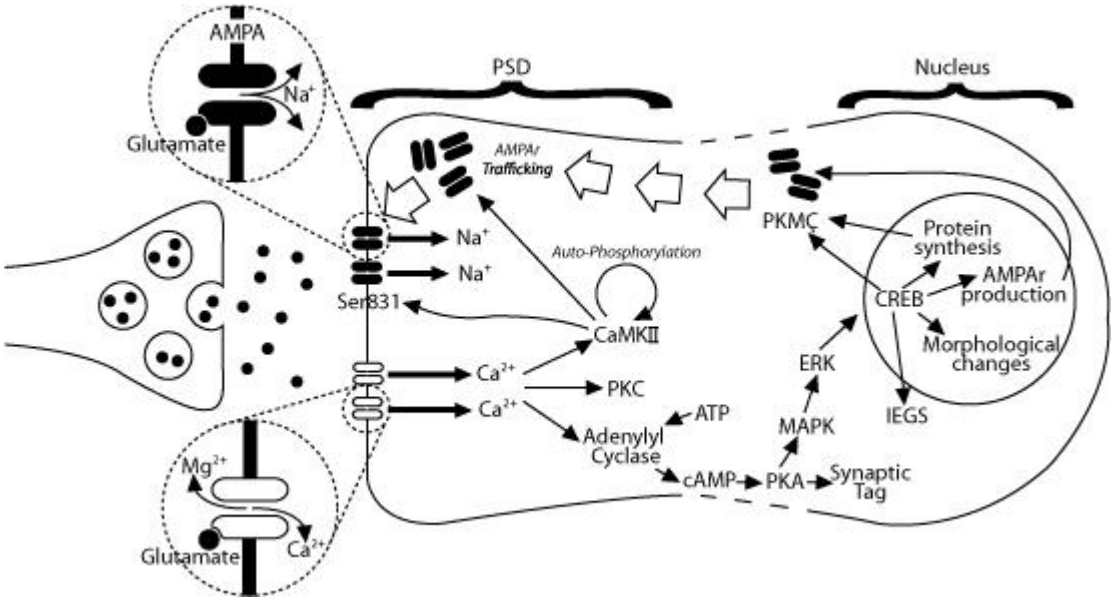

Figure 2.2: The biological mechanisms of LTP

The protein kinase CaMKII is now considered to be the primary mediator of the early stages of LTP. It is particularly prominent in the PSD, making up ~3% of the total protein therein, and contains a dozen or so catalytic subunits that each have several phosphorylation sites, thus providing multiple targets for regulation. It is also capable of auto-phosphorylation, which means that once it has been sufficiently activated, the activity of CaMKII can proceed independently of an elevated Ca^{2+} concentration, thus allowing the mechanisms of LTP to outlast the signal which originally provoked them (Miller and Kennedy, 1986; Lisman and Goldring, 1988; Lisman, Schulman and Cline, 2002; Cooke and Bliss 2006). It has been observed that the levels of auto-phosphorylated CaMKII enzyme and CaMKII mRNA increase in the PSD of cell bodies and dendrites in the hippocampus following LTP induction. Similarly, the pharmalogical inhibition of endogenous CaMKII prevents LTP induction, and transgenic mice which lack the alpha-isoform of this kinase exhibit deficits in LTP expression (Silva et al., 1992a, b; Yasuda et al., 2003). However, several other protein kinases are implicated in E-LTP expression, and some of these (such as PKC) also demonstrate the property of auto-phosphorylation. Furthermore, it is important to reiterate that the kinase cascades involved in E-LTP change profoundly during development (Bliss, Collingridge and Morris, 2007).

During this early stage of LTP expression, an increase in synaptic efficacy can easily be reversed (i.e. 'de-potentiated') by low-frequency stimulation (LFS). This form of input activates protein phosphatases, which initiate several signalling cascades that reverse the effects of protein kinases, and thus reduce or extinguish the magnitude of potentiation induced (Fujii et al., 1991; O'Dell and Kandel, 1994; Zhuo et al., 1999; see Section 2.7). For an increase in synaptic efficacy to become consolidated, protein synthesis and gene transcription are required (Krug, Loessner and Ott, 1984; Frey et al., 1988; Huang et al., 1996; Pittenger and Kandel, 2003; Abraham and Williams, 2003; for reviews see West, Griffith and Greenberg, 2002; Miyashita et al., 2008). This is demonstrated by studies which employ protein synthesis or transcription inhibitors to limit the duration of LTP to 5 or 6 hours (Frey et al.,

1988; Huang and Kandel, 1994; Nguyen, Abel and Kandel, 1994). Although this critical stage of late (or L-) LTP is only expressed after several hours, the molecular cascades which mediate it are induced shortly after the initial influx of Ca^{2+}. This is illustrated by the fact that pharmacological inhibitors are only effective if delivered shortly after the initial induction of E-LTP (see Kelleher, Govindarajan and Tonegawa, 2004 for a review).

The protein synthesis cascades involved in L-LTP are most likely triggered by the persistent activation of several Calcium activated kinases, such as MAPK, PKA and the adenylyl cyclases (Abraham and Tate, 1997). In particular, the extra-cellular signal related kinase (ERK) subfamily of MAPKs may represent the molecular link between early and late phases of LTP, since many of the signalling cascades involved in E-LTP converge on ERK (see Kelleher, Govindarajan and Tonegawa, 2004 for a review). Analogously to the expression mechanisms of E-LTP, the function of ERK may be to phosphorylate several cytoplasmic and nuclear molecules including transcription factors such as the cAMP response element binding protein (CREB). The effect of this phosphorylation is to stimulate de novo protein synthesis, including an increased production of AMPA receptors, and morphological changes, such as an increase in dendritic spine number and synaptic surface area.

Few of the proteins whose synthesis is modulated by these processes have been identified, although the protein kinase Mζ (PKMζ) has been demonstrated to increase its presence following LTP induction, and also seems to be required for the expression of L-LTP (Hrabetova and Sacktor, 1996). The function of this kinase is to direct the trafficking and reorganisation of proteins in the PSD, and inhibiting PKMζ blocks the expression of already established L-LTP (Serrano, Yao and Sacktor, 2005). Similarly, several immediate early genes (IEGs), such as activity-regulated cytoskeletal-associated protein (Arc) and Homer 1A (H1A), and products of IEGs, such as zif268, have been identified as playing some role in the long-term expression of plasticity processes (Knapska and Kaczmarek, 2004; Vazdarjanova et al., 2006). It has been demonstrated that disrupting the expression of these proteins can impair LTP and memory consolidation (Guzowski et al., 2000).

The input specificity of plasticity processes requires that *de novo* protein synthesis and gene expression processes are directed to the appropriate synapse, rather than distributed throughout the cell. It was originally suggested that protein synthesis and targeted mRNAs may be activated locally, at the particular synapse that is undergoing LTP (Steward and Levy, 1982; Steward and Fass, 1983; see Kelleher, Govindajaran and Tonegawa, 2004 for a review). However, it is now generally accepted that these processes occur more broadly across the dendrites and soma of the neuron, and so some mechanism must exist which can specifically direct newly synthesised mRNAs and proteins to recently active synapses, or enable them to capture these products (i.e. Steward et al., 1998; Moga et al., 2004). It has been theorised that this may be achieved by a process of synaptic tagging, and there is now a broad body of evidence to support this hypothesis (Frey and Morris, 1997, 1998). Some of the plasticity related proteins which are captured by tagged synapses (such as PKMζ) have been identified, and there is a growing body of evidence that PKA, and its consequent cAMP mediated signalling

cascade, may be the kinase responsible for the tagging process (Sajikumar et al., 2005; Young et al., 2006). This research has also made it clear that consequent LFS can interfere with this PKA/cAMP signalling mechanism, and thus impair the expression of L-LTP by disrupting the synaptic tag (Kauderer and Kandel, 2000; Young and Nguyen, 2005; Young et al., 2006).

2.5 The Role of LTP in Learning and Memory

The biological mechanisms outlined above endow LTP with several attractive properties that have led to it becoming the dominant neuronal model of learning and memory processes or "…a molecular metaphor for Hebbian plasticity" (Bear, 2003). Firstly, and most obviously, is the length of time which these changes in synaptic strength can endure for – several days *in vitro*, and as long as a year *in vivo* (Bliss and Collingridge, 1993; Abraham et al. 2002). Secondly, the process is pathway specific, in that only those particular synapses which have both pre-synaptic glutamate release and post-synaptic depolarisation are potentiated. This vastly increases the theoretical storage capacity of any neural network (Malenka and Nicoll, 1999). Thirdly, it is associative, in that LTP provoked by a strong stimulus on one synapse can facilitate induction at neighbouring synapses in response to weak stimulation (Harvey and Svoboda, 2007). This has often been viewed as a cellular analogue of classical conditioning (Malenka and Nicoll, 1999). The property of associativity may follow from the critical role of post-synaptic depolarisation, which can spread from afferent fibres in order to facilitate the unblocking of NMDA channels at less excited synapses (Bliss and Collingridge, 1993; Malenka, 2003). The process of synaptic tagging may also explain the co-operativity of LTP, as stimulation of an input which is too weak to provoke L-LTP may, nevertheless, produce a synaptic tag and thus capture the proteins that have been generated by stronger stimulation elsewhere on the dendritic tree (Abraham, 2008). Finally, the fact that LTP can be most reliably generated in brain regions such as the hippocampus, which are known to be involved in learning and memory, has also been cited as evidence for its functional relevance (Malenka and Nicoll, 1999; Sanes and Lichtman, 1999; Malenka, 2003).

These similarities between the properties of LTP and memory, as well as its ubiquitous nature in the brain, have led to the formulation of a synaptic plasticity and memory (or SPM) hypothesis. Generally speaking, this theory states that "…activity-dependent synaptic plasticity is induced at appropriate synapses during memory formation, and is both necessary and sufficient for the information storage underlying the type of memory mediated by the brain area in which that plasticity is observed" (Martin, Grimwood and Morris, 2000). A large amount of evidence explicitly highlights the link between LTP and memory in mammals, but it is difficult to ascertain whether synaptic changes alone are sufficient to mediate the disparate forms of human memory. A variety of experimental methods have been applied to test the SPM theory. For example, studies of aging rats have demonstrated that deteriorating performance in spatial tasks is correlated with the reduced persistence of synaptic potentiation (Gage, Dunnett and Bjorkland, 1989; Fischer et al., 1992; Barnes et al., 1997). Researchers have also attempted to saturate synaptic strengths *in vivo*, in the hope that an inability to generate any further potentiation will correlate with an inability to learn. A positive result was originally obtained using this methodology (Castro et al., 1989) but further studies have been unable to reproduce these findings

(Jeffery and Morris, 1993; Korol et al., 1993). More recent work has again demonstrated that the occlusion of plasticity in the hippocampus does impair spatial learning – but scepticism of both the method and result remains (Moser et al., 1998).

Pharmacological agents have also been employed to block the molecular mechanisms of LTP, in order to assess the effects that this may have on learning and memory function. One of the first to research this link was Richard Morris at the University of Edinburgh, who set out to test the impact of NMDAr-inhibitors on navigational ability. The eponymous water maze used in these experiments has come to be one of the best known tests of spatial learning and memory. It consists of a small platform, submerged in a pool of opaque liquid and set within an arena with various cues on the walls. Rats swim in the pool until they locate the platform by trial and error. Subsequently, they can navigate more directly to the platform using their memory of its location. Morris demonstrated that rats whose hippocampii were perfused with an NMDAr inhibitor showed severe deficiencies in their ability to complete this task, and in their induction of hippocampal LTP (Morris et al., 1986). Several subsequent studies have employed this method to establish a link between LTP and other forms of spatial learning, olfactory learning, and contextual fear conditioning (see Neves, Cooke and Bliss, 2008 for a review).

Further studies have specifically aimed to block the expression of L-LTP, and demonstrated how the administration of a PKMζ inhibitor into the hippocampi of rats results in retrograde amnesia (Pastalkova et al., 2006). This not only provides support for the role of LTP in memory consolidation, but also for the role of this protein in maintaining L-LTP expression. The SPM theory suggests that drugs which enhance LTP might consequently enhance learning. Ampakines, which take their name from the AMPA receptor upon which they act, are a family of compounds which enhance and prolong glutamergic neurotransmission and have been demonstrated to facilitate the induction of hippocampal LTP (Staubli et al., 1994). There is a considerable body of evidence to suggest that these compounds can enhance performance on memory tasks which depend on different brain structures across diverse species and a range of timescales (Hampson et al., 1998a, b; Lynch, 2004).

These results offer strong support for the SPM hypothesis, but there are problems associated with drug diffusion, sensorimotor side effects, and the fact that NMDAr antagonists may affect neuronal processes other than LTP. For example, the behaviour of animals that have received pharmacological infusions are often, although not always, qualitatively different than control groups –they can demonstrate loss of balance, apparent loss of sight, thigmotaxis (an adherence to the walls of an arena) and even schizophrenic behaviour (Cain, 1998). Clearly, learning cannot proceed (and therefore the experimental method is unsound) if the animal cannot see, move, or feel properly. Elsewhere, it has been demonstrated that the adverse effects of NMDAr inhibitors on spatial performance in the Morris water maze can be almost completely eliminated if an animal is trained on a different maze immediately beforehand, without the administration of any drug (Bannerman et al., 1995). This result suggests that a variety of processes interact in a complex manner during spatial learning and recall – and that the role of NMDA receptors and LTP in this behaviour may have been oversimplified.

In more recent studies, transgenic mice have been used to examine behavioural differences which arise when the induction or expression mechanisms of LTP are disabled or manipulated (see Martin, Grimwood and Morris, 2000; Tsien, 2000 and Neves, Cooke and Bliss, 2008 for reviews). For example, a strain of mice has been engineered to have an increased concentration of NMDA receptors in their forebrain, which enhances the expression of LTP. These 'Doogie' mice, named after the precocious TV doctor Doogie Howser, demonstrate superior ability in various learning and memory tasks (Tang et al. 1999). Similarly, mice that are genetically engineered to lack NMDAr in certain regions of their hippocampus demonstrate reduced learning and recall abilities (Tsien, Huerta and Tonegawa, 1996; Nakazawa et al., 2002, 2003). Other strains with a deleted isoform of CaMKII have been bred, and these creatures demonstrate impaired hippocampal LTP and specific deficits in spatial learning tasks (Silva et al., 1992a, b). Experimenters have also made targeted mutations to the alpha-CaMKII protein, which removes the ability of this kinase to auto-phosphorylate, and thus to perpetuate LTP expression in the absence of elevated Calcium concentration. Mice which carry this mutation are unable to express LTP in the hippocampus under standard stimulation protocols, and demonstrate severe deficiencies in spatial learning and memory (Giese et al., 1998; Cooke et al., 2004). These results also support the hypothesis that CaMKII autonomously sustains LTP expression after the Calcium concentration subsides to a basal level (Cooke and Bliss, 2006).

One of the most compelling genetic knock out studies generated an inducible and reversible mutation of NR1 receptors, which are functional subunits of NMDAr, in a specific region of the hippocampus. This allowed the development of the mice to proceed as normal until such a time as the experimenter chose to make a targeted suppression of LTP in this brain region, which had the effect of preventing the consolidation of new episodic memories (Shimizu et al., 2000). Other genetic manipulations have been explicitly targeted at preventing L-LTP, and thus the consolidation of synaptic potentiation, demonstrating a consequent impairment in forming long-term memories (Abel et al., 1997; Wong et al., 1999). Again, it should be noted that transgenic mice may be impaired in many other abilities which could contribute to their failure to display normal learning behaviour. For example, it has been reported that mice generated to express a vastly reduced proportion of the NR1 subunit can exhibit several behavioural abnormalities which bear a striking resemblance to schizophrenia (Mohn et al., 1999).

Despite the body of evidence that delineates a specific link between NMDAr-dependent LTP and learning behaviour, the SPM theory is not considered to be conclusive. In particular, few studies have managed to address the assertion that LTP is sufficient for memory storage because of the obvious difficulties with experimental design. Some examples of mutant mice with deficient learning and memory, but normal LTP expression (or vice versa) have also been presented, although the methodology of these studies means that the findings may be questioned (Migaud et al., 1998; Zamanillo et al., 1999; Fragkouli et al., 2005; see Tsien, 2000 for a review of other earlier studies). These results suggest that mnemonic processes are likely to be far more complex than simple LTP, and

may involve many more cellular and systems level processes. It may be that synaptic plasticity represents the preferable mechanism, when available, but not the only one that the mammalian brain can employ in order to encode and retrieve memory traces (Neves, Cooke and Bliss, 2008).

2.6 Computational Modelling and the Need for Competition

Alongside the neurobiological investigation of LTP, computational neuroscientists have aimed to explore the emergent properties of Hebbian learning rules in studies of artificial neural networks (ANNs). This modelling, consistently inspired by biological data, has helped to shed light on the emergent functions and possible weaknesses of LTP and Hebb's postulate. These processes are synapse specific, and thus provide great scope and depth for memory storage. However, this property also makes it intrinsically difficult to co-ordinate and stabilise neural and synaptic dynamics across an entire network. Thus, it is clear that 'pure' Hebbian learning is insufficient to fully describe the plasticity processes which must occur *in vivo* – or to allow the stable operation of ANNs (Song, Miller and Abbott, 2000; van Rossum, Bi and Turrigiano, 2000; Roberts and Bell, 2002). In order to account for the processes of activity-dependent development and forgetting, and to maximize the capacity for information storage, a successful plasticity rule must stimulate competition between inputs to a neuron and generate a stable distribution of synaptic weights (Sejnowski, 1977; Bienenstock, Cooper and Munro 1982). Pure Hebbian learning cannot achieve this, not least because it fails to make any mention of synaptic weakening processes, but also because those inputs which correlate with post-synaptic firing are repeatedly strengthened and thus grow to infinitely high values. This tends to destabilise post-synaptic firing rates – driving them either towards zero or infinitely high, as synaptic weights similarly evolve towards zero or the maximum possible value (Miller and MacKay, 1994). In order for learning to proceed in a homeostatic manner, a successful plasticity model must allow Hebbian changes in the efficacy of synapses while also describing conditions under which synapses will be weakened, or provide a general rule for the decay of all synaptic strengths so that only those that are persistently reinforced will remain potentiated.

2.7 Long-term Depression

In light of the need for bi-directional plasticity in neural networks, it is reassuring that long-term depression (LTD), like LTP, appears to be a common feature of synapses throughout the brain (Malenka and Bear, 2004). LTD is an enduring, activity-dependent decrease in synaptic efficacy that was first discovered in the hippocampus *in vitro* by Lynch, Dunwiddie and Gribkoff (1977). Since that time, a number of different forms of homo-synaptic (i.e. specific to single junctions) and hetero-synaptic LTD (i.e. affecting all the afferent or efferent synapses of a single neuron) have been identified throughout the cortex (for reviews, see Malinow and Malenka, 2002; Crozier et al., 2004; Malenka and Bear, 2004).

Research has focussed on two main forms of homosynaptic LTD which are suspected to mediate distinct functions for memory processing and activity-dependent development (Delgado and O'Dell, 2005). The first form has many features in common with the standard model of LTP - it is NMDAr

dependent, triggered by a post-synaptic influx of Ca^{2+}, reversible, and sensitive to a process of 'de-depression' for a short period immediately following its induction. Although it is not yet entirely clear, the precise spatio-temporal dynamics of post-synaptic Ca^{2+} concentration seem to determine which signalling pathways are activated, and consequently whether a synapse is strengthened or weakened. Experimental evidence suggests that LTD is induced by a modest but enduring elevation in the intracellular concentration of Ca^{2+} (~0-7μM for ~60s), while LTP is evoked by a large but brief elevation (~10 μM for ~3s) (Lisman, 1989; Yang, Tang and Zucker, 1999; Bi, 2002). The direction and degree of NMDAr-dependent synaptic plasticity, then, depends directly on the level of post-synaptic depolarisation and the spatiotemporal profile of Ca^{2+} concentration, two variables which are intimately linked by the function of the NMDA receptor.

NMDAr-dependent LTD is expressed via a serine-threonine phosphotase cascade which activates molecules such as calcineurin and protein phosphotase-1 (Dudek and Bear, 1992; Kirkwood and Bear, 1994; Malenka and Siegelbaum, 2001; Shouval, Bear and Cooper, 2002; Malenka and Bear, 2004; Graupner and Brunel, 2007). In effect, while LTP mediates the phosphorylation of target protein kinases (such as CaMKII) in order to increase synaptic efficacy, LTD proceeds by de-phosphorylating similar protein kinases (PKA and PKC, for example, but not CaMKII) in order to decrease synaptic efficacy. This has similar, but contrasting, effects to the expression mechanisms of LTP – de-phosphorylation of the Ser^{845} substrate of PKA, for example, decreases the open channel probability of AMPA receptors, and transgenic mice who lack Ser^{845} cannot express hippocampal LTD (Banke et al., 2000; Lee et al., 2003). LTD also affects the trafficking of AMPA receptors and requires protein synthesis (but not gene transcription) in order to be consolidated. Hence, it seems likely that the signalling cascades mediated by LTP and LTD converge on the same target processes and molecules, and that the overall change in the strength of a synapse is determined by the relative activation of these kinase and phosphatase cascades (see Malenka and Bear, 2004 for a review). Similarly, late phase LTD has also been shown to rely on a process of synaptic tagging, and the delivery of plasticity related proteins to depressed and potentiated synapses on the same dendritic tree has been demonstrated to incur a variety of modulatory and non-linear integration effects (Sajikumar and Frey, 2004).

A second common form of LTD also exists, which is dependent on metabotropic glutamate receptors (mGluR). It is known that mGluR activation can release internal stores of Calcium in the dendrite – providing a slow and sustained rise in Ca^{2+} concentration which may then trigger the same signalling cascades as NMDAr-dependent LTD. However, despite the study of mGluR dependent LTD being in an early stage, there is significant evidence to suggest that it has a significant pre-synaptic expression component, as well as affecting post-synaptic AMPAr trafficking (Malenka and Bear, 2004; Delgado and O'Dell, 2005). Similarly, while it has been demonstrated that NMDAr dependent LTD can effectively erase recently induced (i.e. within ~30 mins) LTP, mGluR-dependent LTD can exist concomitantly with LTP, and merely suppresses its effects (Solger et al., 2004; Delgado and O'Dell, 2005). This allows the effective de-potentiation of synapses after a much greater time difference, and

these two forms of LTD have consequently been hypothesised to mediate separate functions for mnemonic processing.

2.8 Mechanisms of Homeostasis

Ideally, neural circuits should remain both "highly adaptive and sturdily robust" (Abbott, 2003) so that they can change significantly during learning or development whilst retaining their stability and flexibility. Computational neuroscientists have therefore turned their efforts to examining how the processes of LTP and LTD might be homeostatically regulated – an important consideration in light of the positive feedback nature of Hebbian learning rules. Several mechanisms have been proposed which control global activity levels in one of three manners: by adjusting the ability of synapses to undergo further plasticity, by placing limits on the resources available to synaptic inputs, or by adjusting the intrinsic excitability of a neuron. Each of these mechanisms can introduce competition between synapses, and this is essential for memory function and the development of neural circuits to proceed efficiently. The first of these processes is broadly referred to as metaplasticity – or the plasticity of synaptic plasticity – as it hetero-synaptically adjusts the propensity for future depression and potentiation of neural connections based on previous activity (Abraham and Tate, 1997; Abraham and Bear, 2004; see Abraham, 2008 for a review).

2.9 Mechanisms of Regulation: The BCM Formulation of Metaplasticity

The most well known theoretical model of metaplasticity is the BCM model, which achieves stability through the use of a modification threshold (θ_m in Figure 2.3), a crossover point between depression and potentiation which is slowly modulated by pre- or post- synaptic activity (Bienenstock, Cooper and Munro, 1982). This process makes the strengthening of a synapse more likely when average activity is low, and vice versa, thus generating competition between inputs and regulating firing rates. The BCM formulation was originally based on experimental data from the developing visual cortex, and computational modelling has demonstrated that it can replicate many known features of that brain region, such as orientation selectivity and ocular dominance (Shouval, Intrator and Cooper, 1997; Blais, Cooper and Shouval, 2000; Cooper et al., 2004). A number of studies have extracted BCM-type curves from neurobiological data taken from the visual cortex, illustrating the dependence of synaptic plasticity on post-synaptic firing rate (Dudek and Bear, 1992; Wang and Wagner, 1999; Tang et al., 1999; Quinlan, Olstein and Bear, 1999; Mellentin, Moller and Jahnsen, 2006).

The shape of the BCM curve is widely believed to represent the shift in the direction and degree of plasticity dictated by increasing intracellular Calcium concentration. There is also a significant amount of evidence from the sensory cortices for a shift in the BCM threshold following chronic levels of activity (Crair and Malenka, 1995; Kirkwood et al., 1995; Kirkwood, Rioult and Bear, 1996; Rioult-Pedotti et al., 2000; Quinlan et al., 2004; see Desai, 2003 for a review). Interestingly, there is little consensus among empirical or computational neuroscientists as to whether the BCM curve (illustrated in Figure 2.3) and corresponding theoretical modification threshold (θ_m) depend primarily on pre- or post- synaptic activity. Elevated pre- or post- synaptic firing rates can each induce LTP experimentally, but the coincident release of glutamate and unblocking of NMDA channels by depolarisation most

reliably produce the strengthening of a synaptic connection (see Section 2.3). It might be suggested, therefore, that the BCM curve and theoretical modification threshold represent a dependence on the level of both pre- and post- synaptic activity.

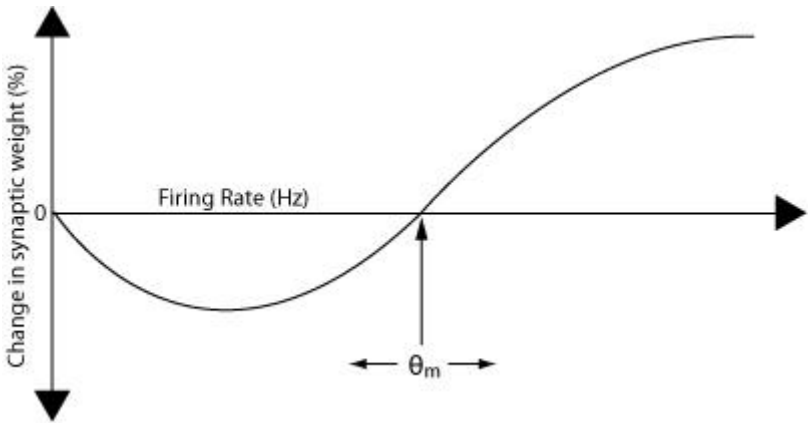

Figure 2.3: A typical BCM curve. The position of the sliding theoretical modification threshold (θ_m) is modulated by the recent history of pre- or post- synaptic activity. High firing rates produce an increase in the value of θ_m, such that potentiation becomes more difficult to induce, while low firing rates produce a decrease in the value of θ_m, such that depression becomes more difficult to induce. This has the effect of regulating neural activity and introducing synaptic competition into the Hebbian framework

Further studies have suggested that a BCM-type sliding modification threshold may also exist in the hippocampus. Several experiments have demonstrated that the induction of LTP by a strong tetanus is inhibited if a weak ('priming') tetanus is previously delivered to the same pathway (Huang et al., 1992; O'Dell and Kandel, 1994; Abraham and Tate, 1997; Wang and Wagner, 1999; Abraham et al., 2001; Solger et al., 2004; Mellentin, Moller and Jahnsen, 2005; Zhang et al., 2005; see Abraham, 2008 for a review). The delivery of a strong priming stimulus also facilitates depotentiation and the induction of LTD at neighbouring synapses for a period of around two hours (Holland and Wagner, 1998; Wang and Wagner, 1999). This priming effect has been demonstrated to be NMDA-receptor dependent, although it can be expressed without pre-synaptic stimulation if post-synaptic firing occurs, and does not require the induction of plasticity (Huang et al., 1992; Zhang et al., 2005; Abraham, 2008). The activation of NMDA receptors provokes a shift in the modification threshold which lasts around one hour, and further potentiation is not blocked, but merely requires a more significant level of stimulation, while the induction of depression is facilitated (Huang et al., 1992; Fujii et al., 1996; Youssef, Addae and Stone, 2006).

The exact mechanisms of this NMDAr dependent metaplasticity are not clear, but seem to involve similar signalling pathways to the expression of LTD (i.e. protein phosphotases and calcineurin), and may also entail the depression of NMDA receptor mediated currents (Fujii et al., 1996; Woo and Nguyen, 2002; see Abraham, 2008 for a review). The result of these changes is that a greater level of post-synaptic depolarisation is required to induce LTP, and a lower level is required to induce LTD (Ngezahayo, Schachner and Artola, 2000). Conversely, the activation of group 1 metabotropic glutamate receptors (mGlur1) has been demonstrated to facilitate the induction and persistence of LTP. This effect seems to be generated by enhancing levels of post-synaptic depolarisation (Bortolotto et al.,

1994; Manahan-Vaughan and Reymann, 1996; Cohen et al., 1999). The mechanisms of LTP enhancement are not entirely clear, but may be produced by an increased level of AMPA and NMDA mediated currents provoked by an increased rate of receptor trafficking (Oh et al., 2006). The majority of these studies have confirmed that the metaplastic changes observed are hetero-synaptic, in that they affect LTP induction at all inputs to a single post-synaptic neuron – a key assumption of the BCM model (Abraham et al., 2001; Abraham, 2008). Interestingly, changes in the propensity of a synapse to undergo LTP or LTD can also be affected by the behavioural state of an animal (i.e. Xu et al., 1998).

Several other possible mechanisms of metaplasticity have been proposed, but as yet there is insufficient evidence to conclusively support any of them. Some researchers have highlighted the propensity of the protein kinase CaMKII (or more specifically, one of its α- or β- holoenzymes) for bi-directional modulation via auto-phosphorylation. It is well documented that auto-phosphorylation (at the Threonine 286 site) allows the function of CaMKII to proceed in the absence of elevated Calcium, and thus outlast the induction signal for LTP. It has also been demonstrated that low level stimulation of CaMKII allows auto-phosphoryation at a different site, Threonine 305 / 306, which decreases the affinity of the kinase for the PSD, reducing its presence at the synaptic membrane and thus raising the threshold for LTP (Tompa and Friedrich, 1998; Elgersma et al., 2002; Krucker et al., 2002). Transgenic mice that do not express Thr^{305} / Thr^{306} show no changes in the position of the modification threshold with priming, and also demonstrate a significantly higher value of θ_m (Mayford et al., 1995).

Elsewhere in the hippocampus, chronic levels of activity have been shown to bi-directionally regulate the ratio of α- to β- forms of CaMKII. With activity blockade, for example, the relative concentration of βCaMKII increases. This form of the kinase is more sensitive to Ca^{2+}, and thus allows a neural network to continue to function in a quieter state – a perfect example of homeostasis (see Thiagarajan, 2007 for a review). However, this β- form of the kinase also acts to further diminish the relative concentration of αCaMKII, which is more specifically implicated in LTP expression cascades, and thus the ability of synapses to undergo potentiation is reduced. The implications of these observations have yet to be explored, and conflicting results (such as cases in which priming stimulation facilitates the induction of LTP) have been reported (Abraham and Tate, 1997; Abbott and Nelson, 2000).

An alternative theory of the mechanisms which may underlie BCM metaplasticity implicates a process of NMDA receptor trafficking and modulation in adjusting the position of the modification threshold. Altered NMDA receptor concentration or function will directly affect intracellular Calcium dynamics without modulating the basal response of a synapse, and thus represents an attractive mechanism for implementing metaplasticity. NMDA receptors consist of a single NR1 subunit, and one of four NR2 subunits (NR2A-D) which have different effects on the decay kinetics of Calcium signals, and therefore the direction and degree of plasticity (Hrabetova et al., 2000; Barria and Malinow, 2002). In the sensory cortices, it is known that the concentration of NMDA receptors is upregulated some time after the changes in AMPAr expression which are induced by LTP (Watt et al., 2004). In fact, the ratio of AMPA to NMDA receptors in these brain regions remains remarkably consistent, despite wide

changes brought about by LTP and LTD. This might suggest that LTP would be more easily induced following a history of potentiation, further destabilising synapses and directly contradicting metaplasticity theory.

However, it has also been demonstrated that visual experience and olfactory learning increases the concentration of NR2A-containing receptors at cortical synapses, and that visual deprivation has the opposite effect (Quinlan et al., 1999a; Quinlan et al., 1999b). As the NR2B subunit mediates longer and larger Calcium currents, the threshold of activation needed to induce LTP is effectively lowered by activity blockade, so that, despite the decreased concentration of NMDA receptors, the ease with which LTP can be induced actually increases. The observation that NR2A knock-out mice exhibit no change in the position of the modification threshold with visual deprivation or experience suggests that it is this subunit whose expression is modulated in order to control the value of θ_m (Philpot, Cho and Bear, 2007). The mechanisms of regulation of these specific subunits are most likely controlled by G-coupled protein receptors, which can activate various kinases and phosphotases in order to affect the properties and trafficking of NMDAr (see MacDonald, Jackson and Beazely, 2007 for a review).

A similar process has been identified in the hippocampus, with the ratio of NR2A to NR2B subunits of the NMDA receptor being bi-directionally manipulated by levels of activity (Vissel et al., 2001; Grosshans et al., 2002; Montogomery et al, 2005; Morishita, Marie and Malenka, 2005; Hellier et al., 2007). It has been demonstrated that the NR2A subunit is essential for the expression of potentiation, while the NR2B subunit is critical for depression, as selective pharmacological inhibition of these receptors will eliminate one or the other of these processes (Gerkin et al., 2007). The authors suggest that Calcium influx via the NR2A and NR2B channels may preferentially activate CaMKII and Calcineurin respectively, in order to mediate these changes. However, some conflicting observations have also been made. For example, it has been demonstrated that NR2B over-expressing transgenic mice exhibit a shift to the left of the BCM curve (i.e. facilitated potentiation), attributed to the longer opening time of this receptor (Tang et al., 1999). Elsewhere, it has been observed that the NR2B subunit is phosphorylated during LTP, an effect that peaks several hours after induction. This enhances the opening time of the NMDA channel, and thus affects subsequent attempts to induce plasticity, with little effect on AMPA receptor mediated currents (and therefore on the synaptic response). Conversely, LFS produces a depression of NMDAr mediated synaptic currents, and a particular reduction in Calcium flux, by activating a phosphotase and PKA dependent pathway (Sobcyzk and Svoboda, 2007). This raises the threshold for LTP – which is analogous to a priming effect, but has the effect of making potentiation more difficult for depressed synapses, contrary to classical interpretations of metaplasticity. There is also evidence for NMDAr insertion in the synaptic membrane following LTP, which may make further potentiation more likely (Williams et al., 2007). It has been suggested that these changes might serve to stabilise the state of synapses following potentiation or depression events, so that stored memories are not destroyed by ongoing activity (Hellier et al., 2007). Clearly, this contradicts the theory of NMDAr trafficking in the hippocampus as a homeostatic mechanism. Although the data is not conclusive, and often contradictory, the concept of metaplasticity through

NMDAr trafficking is attractive, and analogous to the role of AMPAr trafficking in expressing long-term plasticity.

Elsewhere, examinations of the entorhinal cortex have demonstrated that the activation of NMDA receptors favours LTP, as antagonists of this receptor shift the modification threshold to the left, while the activation of voltage-dependent Calcium channels (VDCCs) favours LTD, as antagonists of these channels shift the modification threshold to the right (Solger et al., 2004). In the hippocampus, mGluRs have been demonstrated to modulate NMDAr mediated synaptic currents, thus affecting the induction of long-term plasticity, and to modulate the process of AMPAr trafficking, which represents the main expression of long-term plasticity. Furthermore, the trafficking of mGluRs has been observed, which can directly affect the process of mGluR-dependent plasticity, and therefore directly the magnitude and direction of NMDAr-dependent LTP and LTD (Cheyne and Montgomery, 2008). Elsewhere, a growing field of research has highlighted the importance of inhibitory interneurons, their plasticity, and the consequent effects on depolarisation and the propensity for synaptic change at excitatory neurons (Freund and Buszaki, 1996; Kullman and Lamsa, 2007). Finally, theoretical examinations have suggested that the intrinsic plasticity of neurons, which affects their excitability, and thus the level of firing they produce in response to current injection, can interact with standard plasticity implementations to produce a BCM like curve (Triesch, 2007). These changes are most likely mediated by the modulation of several voltage dependent ion channels (Zhang and Linden, 2003).

The mechanisms may not be clear, but whatever the case, it seems that chronic and acute changes in activity levels can significantly affect the conditions required for the induction of LTP and LTD. It should be noted that the expression of robust plasticity does not vary; rather, it is the level of post-synaptic depolarisation or Calcium influx needed to induce synaptic changes that is affected by a short-term history of activity. These metaplastic changes (towards favouring depression) may also be able to explain the deficits in learning and memory that are exhibited as a result of the aging process, diabetes mellitus and behavioural stress (see Artola, 2008 for a review). The BCM model currently represents the best theoretical approach to modelling these changes, and while the majority of experimental support for this theory has been obtained from the visual cortex, there is now substantial evidence for a BCM-type curve and shifting threshold for plasticity in the hippocampus and other cortical regions.

2.10 Mechanisms of Regulation: Synaptic Scaling

An alternative method of regulating activity across a neural network is to make direct, global changes to synaptic weights, rather than indirectly modulating the capacity for synaptic plasticity. Synaptic scaling is a rapid, hetero-synaptic process that can balance the homo-synaptic changes caused by LTP and LTD. It involves the adjustment of all input strengths to a given post-synaptic neuron based on its level of activity – thus introducing competition into the Hebbian framework (Miller and MacKay, 1994; Abbott and Nelson, 2000; Turrigiano and Nelson, 2004). It has been demonstrated that this modification can allow neurons to operate as principal component analysers, whereby they extract the statistically most significant factor from their input (Oja, 1982; Abbott and Nelson 2000). There is

experimental evidence for synaptic scaling acting to dynamically maintain firing rates within a certain range in the visual cortex, hippocampus and spinal cord (Turrigiano et al., 1998; Lissin et al., 1998; Burrone, O'Byrne and Murthy, 2002; Desai et al., 2002). Adjustments to synaptic strengths can be either multiplicative, whereby each synaptic strength is adjusted by the same factor, or subtractive, whereby each synaptic strength is adjusted by the same absolute value (Miller and MacKay, 1994). Computational studies demonstrate that either of these approaches will stabilise Hebbian plasticity, but direct synaptic development in markedly different ways. Multiplicative scaling retains the relative weight differences that evolve (via Hebbian plasticity) to reflect input correlations, generating a graded synaptic weight distribution. Subtractive scaling, on the other hand, generates much more competition between synapses, and thus a bimodal weight distribution develops, whereby those inputs with the greatest correlation are maximally potentiated, at the expense of any others (Miller and MacKay, 1994).

In principle, these changes in total synaptic strength could be mediated by changes in postsynaptic receptor clustering (in a manner analogous to AMPAr trafficking), pre-synaptic release or re-uptake of neurotransmitter, or the number or size of functional synapses. There is empirical evidence for each of these mechanisms playing some role in synaptic scaling processes. In the visual cortex, AMPAr trafficking is altered and changes in inhibitory GABAergic synapses proceed in order to compensate for global changes in activity levels (see Desai, 2003 for a review). These changes are global, multiplicative and based on a slow integration of activity levels. As we have seen, the ratio of AMPA to NMDA receptors is roughly constant in the neocortex, and so a change in the modification threshold for LTP and LTD is also likely to occur (Abbot and Nelson, 2000; Watt et al., 2000; Turrigiano and Nelson, 2004). Changes in the size of synapses, the probability of transmitter release and the quantal amplitude which take a period of hours, or even days, to develop, have also been reported (see Turrigiano, 1999; Abbott and Nelson, 2000; Desai, 2003; Zhang and Linden, 2003 for reviews).

In the hippocampus, there is evidence that a reduction in neural activity increases the number of pre-synaptic vesicle stores, while increased activity decreases the rate of AMPAr trafficking (see Turrigiano and Nelson, 2004 for a review). Several other studies have implicated changes in the rate of receptor insertion and removal at the post-synaptic membrane as a mechanism by which synaptic scaling may be implemented (Turrigiano et al., 1998; Lissin et al., 1998; Desai, 2003). These changes have been demonstrated to be cell autonomous, as blocking activity in a single neuron results in a scaling up of pre-synaptic transmitter release to that neuron over a period of several days (Burrone, O'Byrne and Murthy, 2002). Recent experimental data has suggested that changes in AMPAr trafficking – which must interact with the processes of LTP and LTD, as they are mediated in part by the same mechanisms – are directed by a single molecule, the pro-inflammatory cytokine tumour-necrosis factor-α (TNF-α). The presence of TNF-α directly increases the surface expression of AMPA receptors in hippocampal cultures and slices, but does not appear to have any involvement in the processes of LTP or LTD. The presence of a TNF-α scavenger molecule also prevents the changes in AMPAr trafficking which normally result from chronic activity blockade. It has been hypothesised that

this molecule is produced by the glia (which is ideally situated to monitor activity levels, via glutamate spillover for example), accumulates extra-cellularly and affects both an increase in AMPAr expression and a decrease in inhibitory synaptic strengths (Stellwagen and Malenka, 2006). Brain-derived neurotrophic factor (or BDNF) is widely believed to mediate the opposing process of scaling down strengths during periods of consistently elevated activity (Rutherford, Nelson and Turrigiano, 1998). As with the BCM model, several other biological correlates have been suggested, including a process of synaptic scaling by altering the intrinsic excitability of neurons. There is some evidence from the neocortex and cerebellum which supports such a theory, as the firing rate produced in response to a set level of current input is dynamically altered in order to maintain stability (Desai, Rutherford and Turrigiano, 1999; Brickley et al., 2001).

2.11 Mechanisms of Regulation: Synaptic Redistribution

The theory of synaptic redistribution stems from observations of short-term depression at a number of cortical synapses, thought to arise from the depletion of neurotransmitter vesicles which are ready to release (Markram and Tsodyks, 1996; Abbott et al., 1997; Tsodyks and Markram, 1997; Abbott and Nelson, 2000). This causes a decrease in the likelihood of release with each transmission event, followed by an exponential recovery to the baseline release probability. If there is a pre-synaptic element to LTP expression, then potentiated synapses may have a larger number of pre-synaptic transmitter vesicles, and / or an increased probability of transmitter release. The process of synaptic redistribution reflects the equilibrium which arises between enhanced amplitude of synaptic transmission and the subsequently enhanced magnitude and duration of short-term depression at the pre-synaptic terminal (O'Donovan and Rinzel, 1997). Effectively, although individual synaptic currents are larger, the time taken to recover is longer, and therefore excessive levels of synaptic transmission are avoided. This process prevents runaway increases in post-synaptic firing rates, even in recurrently connected networks. There is some experimental evidence for a process of synaptic redistribution at neocortical synapses, but not those in the hippocampus (Buonomano, 1999; Finnerty, Roberts and Connors, 1999; Selig, Nicoll and Malenka, 1999).

2.12 Mechanisms of Regulation: Functional Inhibition

Within the human brain there are a number of neural networks, such as parts of the hippocampal formation, which exhibit extensive recurrent connectivity. In these systems, it seems unlikely that processes such as synaptic scaling are sufficient to regulate activity, as excitation is rapidly recycled and can therefore escalate on a much shorter timescale than the regulatory processes described above (Turrigiano and Nelson, 2004). In these cases, a balance between feedback inhibition and excitation is critical to prevent runaway activity (Changnac-Amitai and Connors, 1989). This functional inhibition can also have an effect on plasticity processes, by consistently hyperpolarising a neuron and thus preventing firing and the unblocking of NMDA receptors, for example (Hensch et al., 1998; see Paulsen and Moser, 1998 for a review of other findings). Additional mechanisms that dynamically adjust the levels of excitatory and inhibitory feedback have also been identified (Marty et al., 1996; Rutherford, Nelson and Turrigiano, 1998; Kilman, van Rossum and Turrigiano, 2002). For example, it

has been demonstrated that the blockade of activity in cultured hippocampal networks produces a reversible decrease in functional inhibition, mediated both by receptor trafficking and changes in the number of synapses (Rutherford et al., 1997; Kilman, van Rossum and Turrigiano, 2002). Hence, several systems can interact to maintain activity levels within a reasonable range in recurrent networks: low activity levels increase excitation and reduce inhibitory feedback, while high activity levels reduce excitation (except to inhibitory interneurons) and increase inhibitory feedback (see Turrigiano and Nelson, 2004 for a review).

2.13 Spike-timing Dependent Plasticity

In recent years, a novel Hebbian plasticity model has generated considerable interest, because of the manner in which it can inherently stabilise network activity by stimulating competition between synapses (Song, Miller and Abbott, 2000; Abbott, 2003). It is known as spike timing dependent plasticity (STDP) because it dictates that the direction and degree of changes in synaptic efficacy are determined by the temporal order of neural firing. STDP is directly inspired by neurobiology (Markram et al., 1997; Magee and Johnston, 1997; Bi and Poo, 1998; Debanne, Gahwiler and Thompson, 1998; Feldman, 2000; see Roberts and Bell, 2002 and Dan and Poo, 2004 for reviews). As with previous models of long-term plasticity, further investigation has revealed a myriad of different forms of STDP throughout the cortex. However, in its most common incarnation, only those pre-synaptic spikes which provoke post-synaptic firing within a short temporal window (~20ms) potentiate a synapse, while those which arrive after post-synaptic firing cause depression. STDP implicitly generates competition between synapses, favouring reliable inputs with shorter latencies or strong mutual correlations at the expense of others (Kistler and van Hemmen, 2000; Song, Miller and Abbott, 2000; Kepecs et al., 2002; Roberts and Bell, 2002). If a hard limit is placed on the achievable strength of a synapse, then this can automatically lead to a normalization of the total input to a postsynaptic neuron in a competitive self-organized process, removing the need for any synaptic scaling mechanism and generating intrinsically stable network operation (Kempter, Gerstner and van Hemmen, 2001; Abbott, 2003; Sakai, Nakano and Yoshizawa, 2004). Under certain conditions, STDP can also act to adjust synaptic weights in response to changes in input rate, allowing an unsupervised, homeostatic regulation of output firing rates. These properties have kindled great interest in this plasticity model.

2.14 From Rate Coding to Spike Timing

The discovery of STDP and its use in ANNs marks an important paradigm shift in computational neuroscience. Previous models of learning have consistently used correlations between firing rates in order to direct synaptic plasticity – reflecting the tetanic stimulation protocols commonly used to induce LTP and LTD in the lab (Kempter, Gerstner and van Hemmen, 1999). These rate-coded learning rules have now given way to plasticity formulations which make an explicit consideration of spike-timing in determining the direction and degree of weight change. STDP rules are still consistent with Hebb's postulate, as the asymmetrical time window (see Figure 2.4) stresses the importance of causality as opposed to coincidence in neuronal firing (Abbott and Nelson 2000; Kepecs et al., 2002). A number of computational studies have demonstrated how STDP can selectively potentiate temporal

correlations embedded within a background of uncorrelated firing activity (i.e. van Rossum, Bi and Turrigiano, 2000; Song, Miller and Abbott, 2000). Furthermore, spiking models are considered to be third-generation neural networks, superseding earlier rate based systems in their computational power and abilities (Hopfield, 1995; Maass 1997; Maass and Markram, 2004). The amount of information which can be stored by precise and reliable spike times in a single neuron is much greater than that contained in a mean firing rate, and is available immediately, rather than after the delay needed for an averaging period (van Rullen and Thorpe, 2002). There are also an increasingly large number of studies which demonstrate that spike timing patterns of sub-millisecond precision can be reliably generated *in vivo* (for a review, see Izhikevich, 2006).

2.15 The Mechanisms of STDP

Although STDP has become popular in computational modelling studies, its biological mechanisms have not received such significant research attention. It is generally assumed that they adhere to the NMDAr-dependent model of LTP and LTD, but emphasise the importance of the temporal sequence of synaptic events in determining the level of post-synaptic intracellular Calcium concentration, and thus the direction and magnitude of synaptic change. In essence, when pre-synaptic spiking occurs first, glutamate will be present when NMDA receptors are unblocked by post-synaptic depolarisation, thus allowing a large and rapid influx of Ca^{2+}. This depolarisation is widely believed to be mediated by backpropagating action potentials (bAPs) returning from the axon to the dendrites (Magee and Johnston, 1997). The fact that simultaneous pre- and post- synaptic spiking generally causes maximal depression is often taken as evidence for the importance of the bAP. Conversely, when post-synaptic activity comes first, the rise in intracellular post-synaptic Calcium concentration is slow but sustained, via voltage-dependent calcium channels (VDCCs) and possibly partial NMDAr activation (Koester and Sakmann, 1998; Malenka and Siegelbaum, 2001; Dan and Poo, 2004; Dan and Poo, 2006). The direction and degree of plasticity is thus dictated by the relative timing of the bAP and the incoming pre-synaptic spike. This again raises the possibility that feed-forward and feed-back inhibition may play a critical role in plasticity processes, as GABAergic interneurons which target post-synaptic dendrites can control both the strength and extent of the bAP (see Paulsen and Moser, 1998 for a review).

It has been demonstrated that the LTP component of STDP is NMDAr-dependent, but the timescale (~20ms) is generally much shorter than that of glutamate dissociation from NMDA receptors. Hence, the temporal profile of the potentiation window is most likely explained by the kinetics of Mg^{2+} in unblocking NMDA receptors, as well as other post-synaptic cytoplasmic processes (Lin et al., 2003; Kampa et al., 2004; Dan and Poo, 2006; Kampa, Letzkus and Stuart, 2007). However, some research has suggested that the bAP is not necessary or sufficient to allow potentiation, as it is too brief to unblock NMDA channels (Golding, Staff and Spruston, 2002; Lisman and Spruston, 2005). The presence of sufficient post-synaptic depolarisation may instead be caused by the co-operative activity of several inputs, sufficiently high firing rates from a single input (such as the tetanic stimulation protocols traditionally used to induce LTP), or the presence of dendritic spikes (those which are

generated in the dendrite and propagate to the axon, but may or may not generate an axonal action potential on arrival). It is also possible that, due to electric attenuation, EPSPs which may appear modest in size at the axon are actually much larger in the dendrites – and therefore sufficient to allow maximal Ca^{2+} influx through NMDA receptors (Golding, Staff and Spruston, 2002; Lisman and Spruston, 2005). The recent finding that STDP is also dependent on the exact location of a synapse on the dendritic tree further complicates theoretical modelling (see Froemke, Poo and Dan, 2005; Dan and Poo, 2006 or Kampa, Letzkus and Stuart, 2007 for reviews). However, computational models which incorporate NMDA channel characteristics and post-synaptic depolarisation profiles can replicate this dependence of plasticity on dendritic location, as well as predict changes in the shape of the STDP curve that are dictated by the profile of post-synaptic depolarisation (Saudargiene, Porr and Worgotter, 2005).

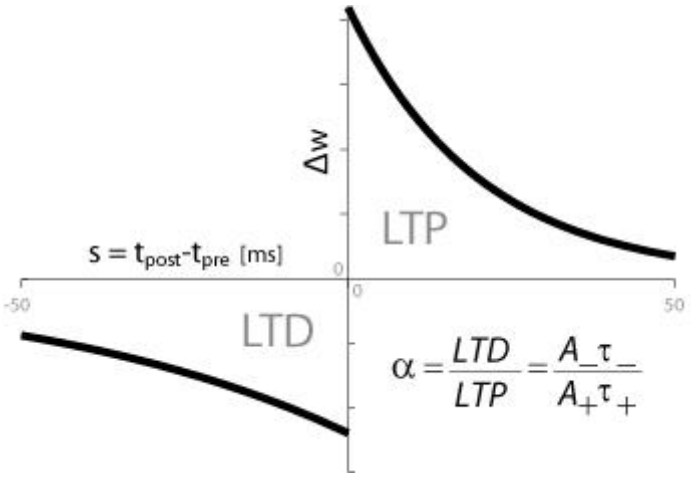

Figure 2.4: The asymmetric learning window of STDP. The absolute change in weight of a synapse (Δw) is a function of the latency between pre- and post- synaptic firing (s) and the parameters which define the shape of the asymmetric learning window ($A_{\pm}$ and $\tau_{\pm}$). The ratio of integrals of the depression and potentiation windows is defined as α, and thus $\alpha>1$ indicates an overall dominance of depression and $\alpha<1$ indicates an overall dominance of potentiation

The simplest explanation offered for the profile of the depression window is that an after-depolarization caused by post-synaptic spiking partially opens NMDA receptors, allowing a prolonged and low-level influx of Ca^{2+}. However, this model also predicts a second LTD window, caused by pre- and then post-synaptic spiking at intervals longer than those which cause LTP, as the post-synaptic concentration of Ca^{2+} would have to pass through the range in which LTD is induced before returning to baseline levels. This effect has been observed in hippocampal slices, but these findings have not been replicated (Nishiyama et al., 2000). Other possible mechanisms for the LTD window have been offered, based on evidence that depression can be dependent on metabotropic glutamate receptors (mGluRs) or post-synaptically released retrograde messengers (such as endocannabinoids). The activation of VDCCs and the release of Ca^{2+} from internal post-synaptic stores have also been shown to be necessary for LTD, and this further complicated the dynamics of post-synaptic Ca^{2+} concentration, and thus the determinants of the temporal window for LTD (Bi and Poo, 1998; Nishiyama et al., 2000; Dan and Poo, 2006; Kampa, Letzkus and Stuart, 2007).

2.16 STDP in Models of Memory and Learning

Models of STDP have proved extremely effective in modelling simple learning and memory processes. Particular successes have been made in areas where biophysical evidence for STDP is rife, such as the visual cortex. The critical emphasis placed on the temporal order of spiking by STDP is ideally suited for predictive sequence learning - a key feature of memory processes (Abbott and Blum, 1996). Many models have therefore employed this plasticity rule to great effect in learning and recalling specific sequences of neural spiking (Nowotny, Rabinovich and Abarbanel, 2003; Lengyel et al., 2005; Yamaguchi et al., 2007). Elsewhere, feed-forward networks which implement STDP have been demonstrated to allow the unsupervised learning of selectivity to direction and intermediate-complexity visual features in natural images (Buchs and Senn, 2002; Wenisch, Noll and van Hemmen, 2005; Masquelier and Thorpe, 2007). Spike-timing plasticity rules have also been used to model various features of the auditory system, such as the development of temporal feature maps and sound localisation based on sub-millisecond coding (Leibold, Kempter and van Hemmen, 2002; Burkitt and van Hemmen, 2003). Elsewhere, STDP has been reconciled with earlier models of reinforcement (or temporal-difference, TD) learning, in order to allow the selective enhancement of synaptic weights based on millisecond scale mechanisms, despite the reward which dictates this change being mediated on a more realistic behavioural timescale (Florian, 2005; Izhikevich, 2007). STDP has also been demonstrated to allow flexible single-trial learning in embodied robots (Di Paolo, 2003).

2.17 Key Properties of STDP: Size of Potentiation and Depression Windows

Computational modelling studies have offered significant insight into the functional implications and emergent behaviour of STDP. These have often focussed on several key features that can have a profound effect on the neural and synaptic dynamics of ANNs that implement this plasticity rule. The first of these regards the relative size of the 'learning' windows for potentiation and depression (see Figure 2.4). When the size of the window for depression is greater than that for potentiation (i.e. $\alpha>1$) STDP prevents the development of excitatory feedback loops, as two interconnected neurons cannot both consistently spike at short time intervals behind the other. This is considered to be essential for the stability of neural networks with extensive connectivity, and the primary reason for the instability of traditional models of Hebbian learning (Song and Abbott, 2001). It has also been demonstrated that, under these constraints, synapses are driven into a state of balanced excitation, whereby the membrane potential of the post-synaptic neuron remains at a level near the threshold for spiking. This forces neurons into a noisy but temporally sensitive mode of operation, which closely resembles observations of firing activity made *in vivo* (Abbott and Nelson, 2000; Feldman, 2000; Dan and Poo, 2006).

For example, one computational modelling study describes how STDP can regulate the variability of post-synaptic firing, as described by the co-efficient of variation, and consistently maintain this parameter within a biologically realistic range (Song, Miller and Abbott, 2000). In the majority of cases, the size of the time windows for LTP and LTD that have been observed empirically also have a value of $\alpha > 1$ (Dan and Poo, 2006). In particular, the original characterisation of STDP in the hippocampus by Bi and Poo (1998) measured values of 17 ± 9 ms for the time constant of LTP, and 34

± 13 ms for the time constant of LTD. It is more difficult to experimentally determine the maximum amount of potentiation and depression which can be induced by a single spike pair (i.e. $A_{\pm}$ in Figure 2.4), but it is clear that these values are equally important in determining the value of the integral over the STDP window, and thus the emergent behaviour of the plasticity rule.

One of the most attractive features of STDP is its ability, under certain conditions, to homeostatically regulate network activity. This property has been frequently replicated with a value of $\alpha>1$ (and additive STDP, see Section 2.18), as this incurs an increase in depression when input firing rates are raised, thus reducing the consequent increase in post-synaptic activity (van Rossum, Bi and Turrigiano, 2000; Abbott and Nelson 2000; Song, Miller and Abbott 2000; Kempter, Gerstner and van Hemmen, 2001; Tegner and Kepecs, 2002; Abbott, 2003). It has also been reported that the post-synaptic firing rate can be controlled by modulating the value of α, within the constraint that it must retain a value greater than unity (Song and Abbott, 2001). This negative rate correlation is anti-Hebbian in a time-averaged sense, but Hebbian correlations between input and output spikes are still preferentially potentiated. However, the problem remains that, while these stabilising properties are very attractive, they directly contradict earlier observations of LTP, according to which increased pre- and post-synaptic rates produce a greater degree of potentiation. This suggests that some non-linearity must be introduced, to allow the STDP rule to selectively reward temporal correlations at low frequencies, and rate coded input at higher frequencies (Shadlen and Newsome, 1998; Senn, 2002).

Conversely, when $\alpha=1$, previous research suggests that correlated inputs are not preferentially rewarded above uncorrelated spiking activity, and post-synaptic activity is more sensitive to input fluctuations, which detracts from some of the beneficial features of STDP (i.e. Song and Abbott, 2001). When $\alpha<1$, the behaviour of the plasticity model changes again, and bi-directional associations, which are often seen in real cortices, can arise (Kepecs et al., 2002). However, intrinsic rate stabilisation properties are lost completely, and increased firing rates generate greater net potentiation, which is in line with traditional models of LTP but requires additional mechanisms to regulate activity levels (Kempter, Gerstner and van Hemmen, 2001; Song and Abbott, 2001).

2.18 Key Properties of STDP: Weight Dependency

The second feature of computational STDP implementations that has a profound effect on network dynamics relates to the empirical observation that the degree of plasticity which is provoked by a spike pair can depend on the initial strength of a synapse (Bi and Poo 1998, Debanne, Gahwiler and Thompson, 1999; Wang et al., 2005). Effectively, this means that the values of A_{+} and / or A_{-} (see Figure 2.4) depend on the current weight of a synapse, rendering the empirical measurement of these values more difficult. Computational STDP implementations which ignore this supplementary feature and use an additive update rule (see equation 2) tend to generate bi-modal weight distributions, similar to those produced by pure Hebbian learning. Synaptic weights under this regime are inherently unstable, due to the positive feedback process which arises between potentiation and the decreased latency of post-synaptic firing. However, the inclusion of a hard limit on the maximum achievable

strength can prevent a runaway increase in synaptic currents. Additive STDP is characterised by strong competition between inputs for the control of post-synaptic firing – and the symmetry breaking which this allows can be a useful feature for the development of neural circuits (i.e. Wenisch, Noll and van Hemmen, 2005). However, the strength of this competition – and the bi-modal weight distribution that it incurs – means that graded representations are impossible, and synapses are effectively reduced to a binary state. The stability of weights, once evolved, also means that it is difficult for synapses to re-organise to reflect changes in input, or to erase incorrect associations. This implies that weight distributions which do not reflect input correlations may develop (Song, Miller and Abbot, 2000; Kempter, Gerstner and van Hemmen, 2001; Rubin, Lee and Sompolinsky, 2001; Tegner and Kepecs, 2002).

$$\Delta w = A_{\pm} e^{-(s/\tau_{\pm})}$$
$$\Delta w = e^{-fw} A_{\pm} e^{-(s/\tau_{\pm})}$$

Equation 2.2: Typical Additive (top) and Multiplicative (bottom) STDP rules. The change in synaptic weight (Δw) is determined by the time difference between pre- and post- synaptic firing (s), the maximum change in synaptic weight per spike pair ($A_{\pm}$) and the time constant of decay of this value ($\tau_{\pm}$) – see Figure 2.4. For multiplicative STDP rules, the change in strength of a synapse is also scaled by the initial weight value.

Multiplicative update rules, on the other hand, tend to produce stable, unimodal distributions of synaptic weights which closely resemble those observed *in vitro* (Bekkers, Richerson & Stevens, 1990; van Rossum, Bi and Turrigiano, 2000; Kistler and van Hemmen, 2000; Rubin, Lee and Sompolinsky, 2001). Under this regime, synaptic dynamics are much less sensitive to perturbations in input or network parameters, and graded representations can be produced. However, little competition exists between synapses, and this effectively means that all synaptic strengths are directed to a similar equilibrium value which is determined by network parameters. Previous research suggests that the pre-synaptic firing rate has no effect on the weight of synapses in this context, which is at odds with traditional models of LTP and Hebbian learning, whereby inputs which are firing at a higher than background rate are potentiated (Burkitt, Meffin and Grayden, 2004). Although they do not require the inclusion of hard limits on synaptic strength, as additive STDP implementations do, multiplicative STDP rules are apparently unable to generate the unsupervised homeostasis described by previous research for additive STDP models, presumably because of the absence of synaptic competition. Temporal input correlations can be selectively potentiated, but in the absence of competition there is no mechanism to ensure that these connections endure once the input stimulus is removed or to simultaneously depress uncorrelated activity. Furthermore, the rate normalisation properties outlined earlier (for $\alpha>1$) have not been demonstrated with this form of the plasticity rule. Hence, multiplicative STDP implementations seem to face the same problem as earlier Hebbian learning rules, and require additional mechanisms (such as synaptic scaling) in order to introduce competition between synapses (van Rossum, Bi and Turrigiano, 2000; Gutig et al., 2003).

The neurobiological data on STDP is very much open to interpretation, however, and a number of different methods of regulating the amount of potentiation and depression in proportion to synaptic strength have been examined (van Rossum, Bi and Turrigiano, 2000; Gutig et al., 2003; Morrison, Aertsen and Diesmann, 2007; Standage, Jalil and Trappenberg, 2007). Other models have attempted to interpolate between additive and multiplicative STDP implementations, in order to make use of the beneficial properties of each (Gutig et al., 2003; Meffin et al., 2006). It is also unclear if the competition seen in networks with additive STDP arises directly from the form of the plasticity rule or from the differences in the balance of excitation and inhibition that emerge from this implementation (Sakai, Nakano and Yoshizawa, 2004). For example, it has been demonstrated that, if the weight dependency of both depression and potentiation processes are set equal, then the bi-modal weight distribution and innate competition that characterise the additive STDP rule can be generated by a multiplicative STDP implementation (Kepecs et al., 2002; Sakai, Nakano and Yoshizawa, 2004).

2.19 Key Properties of STDP: Reconciling Models of STDP and LTP

The analysis of STDP has been primarily based on isolated pairs of pre- and post- synaptic action potentials, while observations of LTP are commonly mediated by the application of prolonged spike trains. Few studies have attempted to reconcile these separate stimulation protocols within a single theoretical framework, and hence it is unclear how STDP causes synaptic weights to develop with more complex input patterns that are characteristic of normal brain activity, and which involve many possible spike pairings. It is commonly assumed that STDP and classic models of LTP are mediated by the same biophysical processes, and essentially represent two induction protocols of the same cellular mechanisms. However, the homeostatic properties of STDP that have been demonstrated in previous simulations (for additive update and $\alpha>1$) directly contradict results obtained by the tetanic stimulation protocols used to invoke LTP. Critically, observations of STDP are typically made from quiescence, when membrane depolarisation and intracellular Calcium concentration are at their resting values. Clearly, during normal brain activity, the short term history of neural firing will have a direct effect on these parameters, and thus on the manner in which individual spike pairs contribute to the overall modification of a synapse. Making adaptations to common models of STDP so that they can be reconciled with earlier models of LTP, and can accurately model synaptic plasticity processes provoked by complex activity, is therefore an active and important area of research.

For example, one of the first studies to highlight the importance of spike timing in directing synaptic plasticity also demonstrated that spike-timing dependent changes are heavily influenced by overall firing rates. A threshold of around 5Hz and a saturation level of around 30-40Hz were identified for LTP in the neocortex (Markram et al., 1997). More recently, research has confirmed that the degree and direction of plasticity in the rat visual cortex depends conjointly on spike timing and firing rate, with a greater degree of potentiation provoked by individual spike pairs when delivered at a higher rate (Sjostrom, Turrigiano and Nelson, 2001). It has been theorised that this may result from the residual depolarisation which is produced by high frequency firing. This hypothesis is supported by the observation of a novel form of co-operativity in STDP at low frequencies. Single pairs of pre- and post-

synaptic spikes are often ineffective at potentiating a synapse, while a modest degree of post-synaptic depolarisation preceding the pairing allowed robust LTP to be induced. Similarly, in the hippocampus, post-synaptic bursting (which also results in a residual depolarisation) has been demonstrated to allow more robust and pronounced LTP expression (Debanne, Gahwiler and Thompson, 1996; Kobayashi and Poo, 2004). However, it has not yet been determined whether depolarisation requirements alone can explain the difference in the magnitude of potentiation invoked by spike pairs at very low (~0.1Hz) and high (~50Hz) frequencies in the visual cortex. There may be some further downstream stage in the signalling cascade that is sensitive to pairing frequency or the rate and magnitude of calcium influx, and which also contributes to differences in the expression of weight change.

In contrast, the degree of depression incurred by STDP has been shown to be independent of firing rates at low frequencies in the visual cortex, but entirely absent above ~40Hz pairing frequency (Sjostrom, Turrigiano and Nelson, 2001). Hence, the effective window for STDP changes dramatically as firing rates increase – at very low frequencies (~0.1 Hz), depression results from pre- then post- or post- then pre- synaptic spike pairs; at moderate frequencies (~5 to ~40Hz), the curve is temporally asymmetric (see Figure 4); and at high frequencies (>40Hz), potentiation results from all forms of spike pairing. There is some analogy here with the BCM formulation, thus reconciling STDP with rate-coded models of LTP, and each of these properties can be explained in terms of post-synaptic depolarisation and Calcium dynamics. However, this result raises further issues regarding stability, as the absence of potentiation at low frequencies and depression at high frequencies suggests that extreme instabilities in network activity could develop. The use of feedback inhibition, weight-dependent potentiation (i.e. multiplicative STDP) or activity-dependent synaptic scaling may all play a part in stabilising neural and synaptic dynamics.

Computational modelling studies have also attempted to replicate these additional rate dependent features. Plasticity rules which explicitly model intracellular Calcium dynamics or membrane depolarisation have been developed, and had some success in replicating the dependency on firing rate described above (i.e. a BCM-type curve) as well the dependency on dendritic location and the asymmetric temporal profile of plasticity (Castellani et al., 2001; Shouval, Bear and Cooper, 2002; Abarbanel et al., 2003; Yeung et al., 2004; Saudargiene, Porr and Wortgotter, 2005; Cai et al., 2007). One common problem faced by these models is the appearance of a second window for depression at longer pre- then post- synaptic spike pairings, as mentioned previously (see Section 2.15). This is widely considered to be biologically unrealistic, although it has been observed *in vitro* (Nishiyama et al., 2000). By making plasticity processes explicitly dependent on the time course of intracellular Calcium elevation, rather than the peak concentration achieved, the second LTD window can be eliminated, and this protocol may also be more biophysically realistic (Rubin et al., 2005). This model also offers some explanation for the notable heterogeneity in observations of synaptic plasticity, as membrane, ion channel and synaptic dynamics each contribute to the temporal profile of Calcium elevation, allowing much greater scope for variation. However, each of these models is computationally expensive and requires a considerable amount of parameter fitting. Elsewhere,

theoretical models which model the dynamics of transmitter vesicles (in a manner analogous to synaptic redistribution), receptor kinetics and secondary messenger pathways can also replicate the non-linear dependence of STDP on input firing rate, as well as some aspects of the BCM model (Senn, Markram and Tsodyks, 2001). Other research suggests that it may also be possible to reconcile models of STDP and LTP with more abstract modifications of the plasticity rule, without resort to the modelling of more complex dynamics.

Some studies have also attempted to integrate STDP with other forms of metaplasticity in order to dynamically regulate network activity, rather than replicate the firing rate dependency described above. In the most abstract form, this can be achieved by placing the value of α under the control of some post-synaptic variable, such as firing rate or Ca^{2+} concentration, so that its value decreases with increased activity (Tegner and Kepecs, 2002). The value of the BCM modification threshold can then be effectively modulated by adjusting the parameters which describe the asymmetric learning window. For example, by altering the values of A_+ and A_- based on a short term history of post-synaptic spiking, the metaplastic regulation of network activity can be achieved with STDP (Benuskova and Abraham, 2007; Benuskova and Kasabov, 2007).

2.20 Key Properties of STDP: Spike Pair Restrictions

One approach to modulating the behaviour of STDP with more complex input is to place restrictions on the interactions of multiple spike pairs falling within the effective time windows of potentiation and depression. Although there is no *a priori* reason for choosing a particular spike pair restriction model and no direct empirical evidence on which to base the decision (but see Froemke and Dan, 2002; Froemke et al., 2006), modelling studies have most commonly assumed either that each pre-and post-synaptic spike pair will sum linearly and independently with all others within the learning window (the all-to-all model) or that only temporally adjacent pre- and post- synaptic spikes contribute to plasticity, as the internal state of neurons is effectively reset by spiking activity (the nearest neighbour model). Analytical studies have suggested that the choice of spike pair restrictions in models of STDP can have a significant effect on synaptic weight dynamics (Sjostrom, Turrigiano and Nelson, 2001; Izhikevich and Desai, 2003; Burkitt, Meffin and Grayden, 2004). For example, it has been demonstrated that (with certain parameter values) STDP can be directly reconciled with a BCM-type curve, provided that the nearest neighbour implementation is employed (Sjostrom, Turrigiano and Nelson, 2001; Izhikevich and Desai, 2003). Further studies have suggested that the nearest neighbour restriction can be relaxed to include numerous neighbouring spike pairs, while still allowing the integration of STDP and the BCM formulation, particularly at lower firing rates, where the number of spike pairs falling within the effective window for plasticity decreases, and synaptic weights do not converge to their equilibrium values (Standage, Jalil and Trappenberg, 2007).

While a value of $\alpha>1$ is generally used to generate stable network activity, these studies have highlighted the fact that subtle differences in the values of individual parameters which define the profile of the asymmetric STDP window (i.e. A_+, A_-, τ_+ and τ_-) can have a profound effect on network

behaviour. For example, when the values of A_+ and A_- are set equal, and the time constant for depression is larger than that for potentiation, additional mechanisms are required to ensure that potentiation overrides depression at higher firing rates with nearest neighbour STDP (Sjostrom, Turrigiano and Nelson, 2001). However, if $A_+ > A_-$ and the integral over the learning window is still negative (i.e. $\alpha>1$) then no such mechanism is required, and STDP can be directly reconciled with the BCM rule (Izhikevich and Desai, 2003). Similarly, other studies have noted that selective potentiation of higher rate inputs can be achieved with $A_+ = A_-$ and $\tau_- > \tau_+$ if an 'input-restricted' spike pairing scheme is employed, such that each output EPSP interacts only with the most temporally proximate pre-synaptic spikes to direct plasticity (Burkitt, Meffin and Grayden, 2004; Baras and Meir, 2007). However, it is important to note that this result was obtained with a multiplicative STDP model, and therefore in the absence of competition between synapses.

2.21 Key Properties of STDP: Inter-Spike Suppression

An alternative approach to replicating the non-linear features of STDP in computational modelling studies is to employ some mechanism which modulates plasticity based on recent synaptic events (a form of short-term metaplasticity). Experiments in the visual cortex, for example, have identified a suppressive inter-spike interaction which allows the first spike pair in a train of action potentials to play a dominant role, while the effects of those pairs which follow within a certain time window are suppressed (Froemke and Dan, 2002; Froemke et al., 2006). Interestingly, this can help to stabilise firing rates in the face of fluctuations in input activity, much like a conventional competition or homeostasis inducing mechanism (Dan and Poo, 2004). It has been hypothesised that this suppression of plasticity caused by spike pairs which come later in a burst or train of activity is the result of a synaptic redistribution process. However, analytical modelling has suggested that this form of inter-spike suppression will reduce the selective potentiation of higher rate inputs which can be achieved with input restricted STDP (Burkitt, Meffin and Grayden, 2004).

Conversely, studies in cultured hippocampal neurons have demonstrated that depression will cancel previously activated potentiation, while potentiation tends to override previously activated depression if it is activated within a time window of ~70ms (Wang et al. 2005; Rubin et al., 2005; Gerkin et al., 2007). This is consistent with the absence of depression at high frequencies that was discussed earlier (Sjostrom, Turrigiano and Nelson, 2001). This suppression mechanism again raises issues of stability, as the dominance of LTP at higher firing rates might drive synaptic weights and post-synaptic rates to saturation. However, the authors were careful to note the existence of size-dependent potentiation acting as a stabilising mechanism (Wang et al. 2005). This still leaves an absence of competition between synapses, however, and the implications of this finding for the operation of neural networks are unclear. The differences in inter-spike suppression mechanisms identified in the VC and hippocampus are surprising, given that the properties of synaptic plasticity in these two regions are otherwise very similar. Computational modelling has suggested that these differences might be attributed to differences in synaptic and dendritic dynamics of these two systems. For example, neurons in the hippocampus have a lower probability of vesicle release, exhibit significant paired-pulse

facilitation, and have a much longer time constant of recovery from paired-pulse depression. When these observations are incorporated into synaptic models, the differences in inter-spike suppression can be replicated (Cai et al., 2007).

In summary, there have not been enough studies to really characterise the development of synapses under STDP with more complex spike trains, but it seems clear that a linear summation is insufficient to replicate the processes which are observed *in vivo*. Furthermore, STDP models in which each spike pair can dictate either depression or potentiation make synaptic weights very susceptible to change – and thus information stored in a network prone to corruption or erasure. If there are several requirements for plasticity – such as post-synaptic depolarisation, repeated pairings to trigger secondary messenger pathways or biochemical integrators, and some form of modulatory input – then plasticity events become more rare, and synaptic weights less fragile (Lisman and Spruston, 2005). This complexity of induction can therefore act as an additional mechanism to promote network stability.

2.22 STDP and Spike Triggering Effects

The majority of previous computational investigations into STDP have been analytical or employed independent stochastic spike trains, and in both cases spike triggering effects are effectively neglected (e.g. Burkitt, Meffin and Grayden, 2004; Pfister and Gerstner, 2006). However, in real neural networks, some form of correlation between input and output action potentials must be assumed, as synaptic currents inherently affect the state of their target neurons. Furthermore, the asymmetric nature of STDP rules specifically rewards temporal correlations between pre- and post- synaptic spiking, and might therefore be expected to result in a strongly correlated unidirectional propagation of activity through a network in a manner analogous to synfire chains (Tsodyks, 2002). It has been demonstrated that repeated temporal correlations can cause drastic changes in the dynamics of networks which implement STDP, reversing the BCM type curve described in previous studies (Standage, Jalil and Trappenberg, 2007). In order to retain a dialogue with learning and memory processes in real biological systems, therefore, the inclusion of spike triggering effects in computational studies of STDP is critical – particularly in the modelling of brain regions such as the hippocampus, where highly correlated, rhythmic firing patterns (such as gamma and theta oscillations) are common.

2.23 Conclusions

Synaptic plasticity is a ubiquitous phenomenon that is observed throughout the cortices of a wide variety of species, and has been historically linked with the processes of learning, memory and the development of neural circuits. Spike-timing dependent plasticity represents the most contemporary biological model of synaptic change, and offers the possibility of vastly increasing the computational power and complexity of artificial neural networks as well as providing further insight into the operation of real brains. However, STDP has yet to be fully characterised, and the difficulties involved in making empirical observations at the synaptic level mean that a wide variety of computational interpretations exist. These different forms of STDP can generate distinctly different synaptic and

neural dynamics, and very few studies have attempted to draw comparisons between the possible function of these emergent properties.

Computational research has thus far placed an emphasis on forms of STDP that can provide unsupervised homeostasis and stability, as these represent the key issues facing previous Hebbian learning implementations and are attractive properties from a modelling standpoint. Although there is evidence that the value of $\alpha>1$ *in vivo*, it is equally clear that this form of the plasticity rule cannot be reconciled with previous rate-coded models of Hebbian learning and LTP (such as the BCM formulation), as these dictate that an increase in synaptic weights should be produced by increased firing rates. In order to replicate all of the observed properties of synaptic plasticity, mediated by both rate and temporal coded stimulation protocols, it is critical that the STDP and LTP models be reconciled within a single framework. This may require that the attractive homeostatic properties of STDP be abandoned in order to more accurately reflect data from real biological systems.

Because the properties of synaptic plasticity vary so widely throughout the cortex, it is useful to constrain the modelling decisions involved in a computational examination of STDP by focussing on biological data from a single region. This provides a more limited number of properties that need to be replicated by the plasticity rule, allows the emergent dynamics of modelling studies to be reconciled with systems level function, and may even permit the results of computational simulations to re-inform the direction of further biological investigation. A review of the empirical studies described in this chapter demonstrates that the vast majority of observations of synaptic plasticity are made in a single cortical region – the hippocampus. Hence, a detailed examination of the structure, dynamics and suspected function of this region will be made in the following chapter, in order to delineate the specific role that synaptic plasticity may perform in this context.

CHAPTER THREE: THE HIPPOCAMPUS AND AUTO-ASSOCIATIVE MEMORY MODELS

"What matters in life is not what happens to you but what you remember and how you remember it"

Gabriel Garcia Marquez

Aims

- To highlight the fact that the majority of empirical studies of synaptic plasticity are made in the hippocampus
- To describe the theory regarding the role of the hippocampus in declarative memory and, more specifically, the role it may play in mediating the processing and short term storage of semantic memory, episodic and/or spatial memory
- To introduce the concept of auto-associative networks, describe their properties, and outline how these correspond to the postulated function of the hippocampus
- To provide a detailed review of the theory that the CA3 region acts as an auto-associative network which stores sparse, orthogonalised, activity patterns corresponding to a conjunctive encoding of current sensory input, and recalls these patterns from partial cues for recoding and redistribution
- To emphasise that the vast majority of computational models of the hippocampus make use of auto-associative networks, that these models have had great success in replicating emergent function, and that they depend critically on synaptic plasticity

3.1 Introduction

The vast majority of neurobiological research into synaptic plasticity has focussed on a single brain region – the hippocampus, which is a part of the medial temporal lobe (MTL). This "neural Rosetta Stone" (Andersen et al., 2007a) has several features which have consistently made it a popular target for neuroscience researchers. Firstly, its straightforward organisation – with all principal cells and synaptic inputs arranged in well defined layers – facilitates recording studies. Thus, it was in this region that many basic principles of contemporary neuroscience such as unidirectional excitatory transmission and the properties of excitatory and inhibitory synapses were first established. Secondly, it is possible to make recordings both in awake, freely moving animals and in the hippocampal slice, which can be prepared with relative ease and survive for long periods in the lab (Andersen et al., 2007b). Finally, the basic anatomy of the hippocampal formation is similar in all mammals – making comparison and generalisations straightforward. These features have allowed the pyramidal cells which constitute around 90% of the hippocampus to become "the most extensively studies neurons in the brain" (Andersen et al., 2007a; see also Bird and Burgess, 2008).

3.2 The Biology of the Hippocampus

The MTL consists of the perirhinal cortex, entorinhal cortex (EC, which is divided into lateral, medial and deep layers), dentate gyrus (DG), Cornu Ammonis fields (CA, named for their resemblance to the horns of a ram), presubiculum, parasubiculum, and subiculum (or subicular complex; Amaral and Witter 1995; Amaral and Lavenex, 2007). The DG, CA1 and CA3 fields are often referred to as the hippocampus proper, and the main source of neocortical input to this region comes via the superficial layers (II and III) of the EC, which are in turn fed by projections from each of the rhinal cortices. All the sensory modalities project into these rhinal cortices – and thus the hippocampus receives multisensory, abstracted, highly processed information from the neocortex. Other major sources of input to this region are from the brain stem, which might provide idiothetic feedback to the MTL, and amygdala, which might provide emotive or motivational input (Churchland and Sejnowski, 1992; Lever et al., 2003; Rolls, 2008). There are three major fibre systems associated with the hippocampal formation – the fimbria-fornix pathway, which connects the hippocampal formation to the basal forebrain, hypothalamus and brain stem regions; the angular bundle, also known as the perforant path, which (uni-directionally) connects the EC to the other parts of the hippocampus; and the dorsal and ventral commissures, which form a connection between the hippocampal formations of each cerebral hemisphere (Amaral and Lavenex, 2007). Backprojections from the hippocampus to the neocortex and the anterior nucleus of the thalamus are the two major outputs for neural activity in this region. Thus the hippocampus is ideally placed to integrate and inform activity in a wide variety of cortical regions (McLelland, McNaughton and O'Reilly, 1995; Rolls, 1996; Rolls and Treves, 1998).

The most widely known circuit of connections within the hippocampus proper is the tri-synaptic loop, which is comprised of unidirectional connections between the EC and DG (the perforant path); DG and CA3 (the mossy fibres); CA3 and CA1 (the Schaffer collaterals). Neurons in CA1 project back to the subiculum and deep layers of the EC. This uni-directionality is rare among the reciprocally connected

circuits which characterise the majority of the cortex (Amaral and Lavenex, 2007). However, direct projections from the EC to CA1 and CA3 fields exist, are also part of the perforant path, and may be functionally significant. The DG is principally comprised of granule cells, while the CA fields are mostly comprised of pyramidal cells, which exhibit different neural and synaptic dynamics. Aside from this, it is important to note two aspects of hippocampal biology that have been of great importance in modelling studies. The first is the existence of dense, recurrent collaterals in CA3; and the second is the large number of neurons in the DG – an order of magnitude more than either it's main afferent (the EC) or efferent (the CA3). The DG is also one of the rare structures within the human brain that increases in size (via neurogenesis) throughout life (Andersen et al., 2007b).

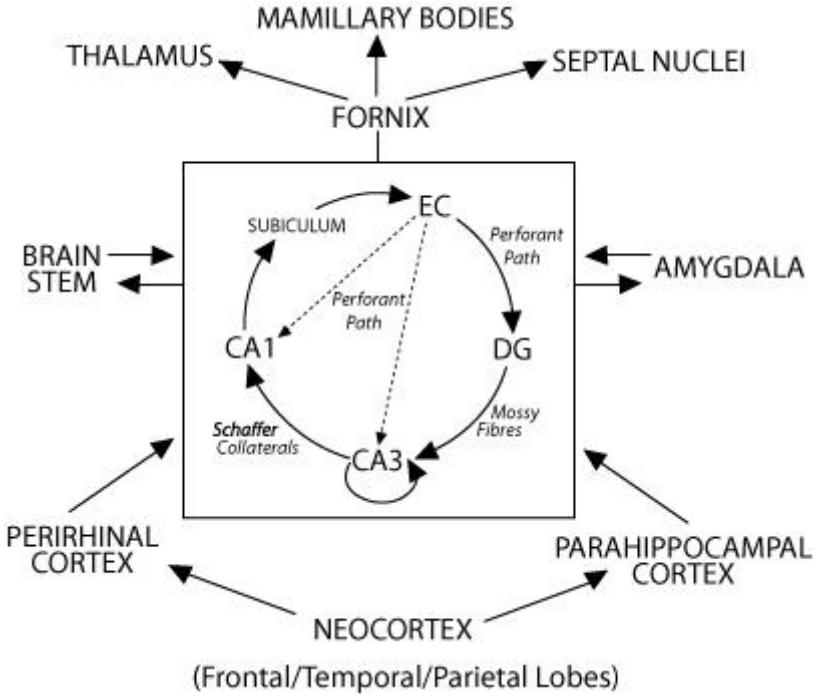

Figure 3.1: The functional connections of the medial temporal lobe

3.3 The Hippocampus, Synaptic Plasticity and Memory

As described in the previous chapter, the hippocampus has proved to be one of the brain regions where synaptic plasticity can be most easily and reliably induced and observed. It was in this cortical region that LTP was first identified and the majority of consequent studies which have helped to elucidate the mechanisms of synaptic change have been carried out. At a functional level, there is also a wealth of evidence which suggests that the hippocampus has a specific role in human memory, making it an ideal testing ground for the SPM theory (Rolls and Treves, 1998; Murray, 2000; Andersen et al., 2007b). Human memory is commonly divided into two categories: declarative memory, which comprises of the knowledge of facts and events; and procedural memory, which comprises of the knowledge of physical skills (Tulving, 1972; Cohen and Squire, 1980; Tulving, 1990). These divisions, to some extent, reflect the different loci of neural memory processes within the brain.

A role for the medial temporal lobe (MTL) in declarative memory function was postulated as far back as the late nineteenth century (Brown and Schafer, 1888). During surgery on epileptic patients in the

mid-20th century, Wilder Penfield observed that stimulation of the hippocampus could lead to the vivid recall of past experiences (Penfield, 1958). This link was later cemented by the study of patients with damage to the hippocampal formation (such as the famous case of 'HM') who exhibit severe anterograde amnesia (see Rempel-Clower et al., 1996; Spiers, Maguire and Burgess, 2001 and Corkin, 2002 for reviews). In an attempt to alleviate severe epilepsy, which had resisted other forms of treatment, the patient HM had large parts of his MTL removed. Although this reduced the severity and frequency of seizures, it also left the patient with profound memory loss. In particular, HM is completely unable to form new memories regarding facts or events, while his general intellect, perceptual ability, short-term ('working') memory, and procedural skills are unaffected (see Stark, 2007 for an in-depth review). These findings led Scoville and Milner (1957) to propose that the hippocampus is "critically concerned in the retention of current experience" and formed the basis for several theories regarding the role of the medial temporal lobe in declarative memory (e.g. Cohen and Squire, 1980; Squire and Cohen, 1984; Squire, 1986; Squire and Zola-Morgan, 1991; Squire, 1992; Schacter and Tulving, 1994; Rolls and Treves, 1998; Eichenbaum and Cohen, 2001; Squire, Stark and Clark, 2004; Moscovitch et al., 2005; Rolls, 2008).

Subsequent research has provided significant support for the theory that the general function of the hippocampus and surrounding structures of the MTL is to mediate declarative memory. Numerous studies of human patients have correlated damage to this region with the impairment of various aspects of declarative memory function, while patients usually retain their procedural (i.e. learning skills and habits) and working (i.e. short-term) memory as well as a normal vocabulary and perceptual abilities (e.g. Smith and Milner, 1981; Gaffan et al., 1984; Parkinson, Murray and Mishkin, 1988; Gaffan, 1994; Teng and Squire, 1999; Manns, Hopkins and Squire, 2003; Manns et al., 2003; Crane and Milner, 2005; Martin, de Hoz and Morris, 2005; Hunsaker, Lee and Kesner, 2008; for reviews see Rolls, 1996; Spiers, Maguire and Burgess, 2001; Wood and Dudchenko, 2003; Squire, Stark and Clark, 2004; Morris, 2007; Rolls, 2008). More recently, electrophysiological and molecular imaging have provided further correlative evidence of a role for the hippocampus in declarative memory (e.g. Eldridge et al., 2000; Bartha et al., 2003; Rosenbaum et al., 2004; Kirwan and Stark, 2004; see Suzuki and Eichenbaum, 2000 for a review of earlier studies). Interestingly, it has also been noted that premature babies with significantly smaller than average sized (~10%) hippocampi exhibit marginally impaired declarative memory function (Isaacs et al, 2000).

3.4 Retrograde Amnesia and the Multiple Trace Theory

The fact that the patient HM also demonstrates some temporally graded retrograde amnesia, such that he cannot recall events from a limited period immediately preceding the operation, suggests that the hippocampus also has a short-term role to play in declarative memory storage. Initially it was suggested that, once a memory trace has been encoded, it is slowly transferred from the MTL to the neocortex for long-term storage, possibly during sleep (Squire, Cohen and Nadel, 1984; Squire and Zola-Morgan, 1991; Squire, 1992; O'Reilly and Rudy, 2000). There is a considerable amount of evidence that recently acquired memories are replayed during slow wave sleep, and that this activity

plays some part in both their consolidation and transfer (e.g. Sutherland and McNaughton, 2000; Stickgold and Walker, 2005; Ji and Wilson, 2007). This configuration of a pairing between networks with fast and slow learning rates allows the rapid acquisition of new information and its integration with an established hierarchy without interference (Marr, 1971; McLelland, McNaughton and O'Reilly, 1995). It has also been observed that the cortex is too sparsely connected to support the encoding and storage of arbitrary associations (Rolls and Treves, 1998). It is important to note, however, that the hippocampus may need to retain some index of consolidated memories, in order that they can be integrated with consequently acquired data.

Human patients with damage to, or lesions of, the MTL demonstrate a wide range of retrograde amnesia, some suffering no loss of memories formed prior to the damage while others are unable to recall any declarative memories whatsoever (Spiers, Maguire and Burgess, 2001). This has raised some controversy regarding the extent to which the medial temporal lobe remains involved with older mnemonic traces. The ability of many patients to freely recall distant events suggests that they may eventually become completely independent of this brain region (Stark, 2007). However, it has also been observed that remote declarative recall is much less detailed in amnesic patients, compared to those with healthy brain function (Moscovitch et al, 2005; Rosenbaum et al., 2008). Functional imaging studies of healthy adults also demonstrate that the recall of recent and remote declarative memories activates the hippocampus equally (Steinvorth, Levine and Corkin, 2005; Cabeza and St. Jacques, 2007). It is possible, of course, that this activity is epiphenomenal, and does not actually indicate the dependence of remote declarative memory on this brain structure.

The multiple trace theory has been developed in an attempt to integrate these various observations of retrograde amnesia into a single framework. The theory posits that each time a declarative memory is retrieved, it is re-encoded as a new (or reinforced) set of connections between the hippocampus and neocortex. Thus, although older (or more frequently recalled) memories are still reliant on the hippocampus, the multiple traces that underlie these memories render them more resistant to MTL damage (Nadel and Moscovitch, 1997; Moscovitch et al., 2005, Nadel, Campbell and Ryan, 2007). This model can also account for the common observation that declarative memory is frequently modified during consolidation, making recall at least partially reconstructive (Bartlett, 1932; Hupbach et al., 2007). Recent functional imaging studies have provided evidence that the MTL involvement in declarative memory is independent of the age of the memory, and that activity in the neocortex during recall (as well as the speed and quality of recall) increases with the frequency of recollection, suggesting a continuous process of consolidation (Nadel, Campbell and Ryan, 2007).

The multiple trace theory is not universally accepted, however, and while it is supported by some reviews of human lesion studies (i.e. Fujii, Moscovitch and Nadel, 2000) others suggest that a more straightforward transfer from MTL to neocortex would better describe the temporal gradient of retrograde amnesia that is often observed (i.e. Spiers, Maguire and Burgess, 2001). One study has demonstrated that AMPAr (and therefore, activity) blockade in the hippocampus impedes the recall of

recently formed paired associations, but if this task is repeated over a period of weeks, then blocking AMPA receptors in the hippocampus no longer blocks recall – suggesting that some form of memory consolidation had taken place outside of this region (Day, Langston and Morris, 2003). The observation that impairment of hippocampal function disrupts the consolidation of memories only if applied within a few days or weeks after the event is also supported by numerous animal studies (i.e. Kim and Fanselow, 1992; Shimizu et al., 2000). Elsewhere, studies that use the imaging of immediate early genes (IEGs, see Section 2.4) to assess network activity and plasticity induction have demonstrated that these proteins are expressed only in response to the recall of recent, as opposed to remote, memories (see Kubik, Miyahsita and Guzowski, 2007 for a review). Whether these results support a time-limited involvement of the hippocampus in memory recall, or simply a time-limited involvement of synaptic plasticity in consolidating multiple traces between the hippocampus and neocortex is a matter of debate, however. Hence, the timescale and nature of MTL involvement in declarative memory consolidation remains a contentious issue, but it is clear that synaptic plasticity plays some role in the initial formation and storage of memories in this brain region.

3.5 Semantic Memory

Declarative memory is commonly divided into several sub-categories: episodic memory (the recollection of autobiographical events), semantic memory (regarding the knowledge of abstract facts or propositions) and recognition memory (regarding the ability to recognise some stimulus as familiar, and thereby recall details about it) (Tulving, 1990; Burgess, 2002). Studies of patients with hippocampal lesions have highlighted the fact that these separate forms of memory are differentially impaired by damage, which suggests that some form of functional dissociation might exist. For example, the retrograde amnesia exhibited by human patients with damage to the hippocampus is generally restricted to episodic memory, while semantic memory is unaffected (for reviews, see Rolls and Treves, 1998; Redish, 1999; Fujii, Moscovitch and Nadel, 2000; Squire, Clark and Knowlton, 2001; Squire, Clark and Bayley, 2004; Moscovitch et al., 2005; Steinvorth, Levine and Corkin, 2005; Bird and Burgess, 2008). This is particularly well illustrated by a study of patients who suffered severe hippocampal damage shortly after birth, and consequently demonstrated the expected impairments in episodic memory (such as an inability to remember what they had done that day, their way around, the time or date). Surprisingly, each of these patients coped remarkably well in mainstream education, as their capacity for semantic processing appears to be relatively intact. The nature of these patient's lesions has led to the hypothesis that semantic memory, which is perhaps less complex in nature, may be mediated by the entorhinal and perirhinal cortices while the hippocampus proper is concerned with mediating context-rich episodic memories (Vharga-Khardem et al., 1997).

Subsequent investigation has suggested that this might not be the case, however, and it is more likely that the slower rate of learning which is generally attributed to the neocortex allows facts, which are more often presented than a unique experience, to be consolidated there in the absence of the rapid acquisition system which the hippocampus usually provides (Rolls and Treves, 1998; Squire, Stark and Clark, 2004; Bayley et al., 2008). Semantic memories may also be inherently 'sparser' and therefore

less affected by damage to the MTL as a whole (Manns, Hopkins and Squire, 2003, Morris, 2007). It should be noted, however, that patients with damage that is more widespread in the MTL suffer a similar level of retrograde amnesia for semantic and episodic memory (Fujii, Moscovitch and Nadel, 2000; Spiers, Maguire and Burgess, 2001; Manns, Hopkins and Squire, 2003). Whether this implies that semantic memory is stored in the parahippocampal structure or that the MTL is essential for the retrieval of semantic memory from the neocortex has yet to be determined, but these findings indicate that some structural dissociation between semantic and episodic memory processing may exist (Davies et al., 2008). Whatever the case, the vast majority of studies emphasise the fact that the whole MTL must be relatively intact to mediate normal declarative memory of any kind, as some deficit is usually seen after lesions to any part (Morris, 2007). It is also possible that those who suffer MTL damage shortly after birth are able to develop alternative neural mechanisms for mediating semantic memory, and that this can account for the relatively small impairment in the learning of facts exhibited by these patients.

3.6 Recognition Memory and Novelty Detection

One aspect of declarative memory for which research has more clearly implied a structural dissociation of function within the MTL is recognition memory. The process of recognition is commonly divided into two elements – familiarity (i.e. knowing that you have seen someone before, but not the details of that encounter) and recollection (i.e. consequently recalling the episode in which you saw them). It has been suggested that these two facets may be unified, as familiarity merely represents a very weak form of recollection, but other groups contend that there are distinctly different processes involved in each stage, and that these may have different loci within the brain (Sauvage et al., 2008). The concept of familiarity is closely aligned with that of novelty detection – another property that has a historical association with the hippocampus (Knight, 1996; Parkin, 1997; Kesner, 1998; Eichenbaum, 1999; Morris, 2007). A commonly cited example of this function is the behavioural differences observed in rats in the Morris water maze following hippocampal lesion. These animals are less systematic in their exploration, slower to become habituated to the arena and less likely to explore novel objects placed therein (O'Keefe and Nadel, 1978; Morris, 2007).

More specifically, a number of studies of animals with MTL lesions have noted that impairments in visual, olfactory and tactile recognition memory are greater if the perirhinal cortex is damaged (for a review, see Brown and Aggleton, 2001; Aggleton and Brown, 2006; Viskontas et al., 2006; Kumaran and Maguire, 2007). Interestingly, electrophysiological recordings in rats and monkeys have revealed that neurons in the perirhinal cortex (as well as other areas of the anterior, inferior temporal cortex, but not the hippocampus) exhibit a reduction in response to visual stimuli that have previously been encountered. This behaviour is symptomatic of a familiarity system, being rapidly acquired (after a single exposure), automatic (occurring even when many stimuli are observed, without attention, even under anaesthesia), and lasting for many days (e.g. Rolls et al., 1993; Xiang and Brown, 1998; see Brown and Aggleton, 2001; Aggleton and Brown, 2006; Rolls and Kesner, 2006; Kumaran and Maguire, 2007 for reviews). IEG imaging studies also demonstrate an increased neural response in the

perirhinal cortex in response to novel (as oppose to familiar) stimuli (Zhu et al., 1995). It has been postulated that, by repressing activity in response to familiar stimuli, the rhinal cortices preferentially direct the downstream processing power of the hippocampus towards encoding novel stimuli (Fernandez and Tendolkar, 2006; Kumaran and Maguire 2007).

It seems, therefore, that the perirhinal cortex mediates familiarity judgements independently of the hippocampus. However, this proposal requires further qualification, as several groups have supplied data which indicate that the whole medial temporal lobe must be intact in order to perform this function (e.g. Spiers, Maguire and Burgess, 2001; Manns et al., 2003). A review of the literature suggests that the role of the perirhinal cortex is restricted to the familiarity of individual stimuli. Lesions of the hippocampus provoke relatively little impairment in performance on tests of single stimulus familiarity – such as spontaneous object exploration, the delayed non-match to sample (DNMS) task, and item recognition (Murray and Mishkin, 1998; Mumby, 2001; Kumaran and Maguire, 2007). Conversely, tests which require familiarity with the configuration or context of stimuli seem to depend more explicitly on the hippocampus (Brown and Aggleton, 2001; Aggleton and Brown, 2006; Kumaran and Maguire, 2007; Sauvage et al., 2008). However, the neural correlates of this process are more difficult to isolate, as studies have demonstrated both increases and decreases in activity in response to associative familiarity (Suzuki and Eichenbaum, 2000; Fyhn et al., 2002; Kirwan and Stark, 2004; Kohler et al., 2005; Viskontas et al, 2006; see Kumaran and Maguire, 2007 for a review). This may reflect the complementary processes of increased activity (passed on from the perirhinal cortex) reflecting the preferential encoding of novel stimuli, and increased activity due to the recollection of previously encountered associative stimuli (Kumaran and Maguire, 2007). It is also important to note that successful associative recognition memory by the hippocampus requires attention (i.e. Clark and Squire, 1998; Uncapher and Rugg, 2008).

3.7 Encoding, Retrieval and Conjunctive Coding

Electro-physiological examinations using humans and animals have also identified further differences between the rhinal cortices and hippocampus. For example, studies of epileptic patients with electrode implants have described how neural activity in the hippocampus during the learning of word pairs was correlated with successful later recall, while activity in the entorhinal cortex during recall was indicative of successful earlier learning (Cameron et al., 2001). The hypothesis that encoding is mediated by the hippocampus, while recall is mediated by the subiculum and rhinal cortices, has been supported by a number of other neuroimaging studies (e.g. Kirwan and Stark, 2004; See Stark, 2007 for a review). It has also been observed that neurons in the entorhinal and perirhinal cortices can fire selectively in response to various sensory stimuli (such as odour or colour), while firing in the hippocampus proper corresponds to conjunctions of these sensory stimuli, to more complex sensory features (such as facial expressions) or to abstracted representations which incorporate several sensory elements (Fried, MacDonald and Wilson, 1997; Wood, Dudchenko and Eichenbaum, 1999; see Suzuki and Eichenbaum, 2000 for a review).

For example, in a study in which rats were trained to visit several cups of sand with various odours, a significant proportion of conjunctive coding neurons were identified in the hippocampus – some of which fired only during the approach to a cup, some when the odour encountered was unexpected, and some in response to highly specific combinations of stimuli, behaviour and location (Wood, Dudchenko and Eichenbaum, 1999). Of particular interest is a study in which human subjects were shown a variety of images of individuals, animals, objects and landmark buildings on a computer screen. The activity of single neurons which responded preferentially to some of these stimuli, regardless of their context, was identified. For example, a single neuron in the left posterior hippocampus was demonstrated to fire at a significantly higher rate whenever images of the actress Jennifer Aniston were presented, compared to a baseline rate in response to pictures of other celebrities, animals, objects or landmarks. Even more surprisingly, the same neuron would fire when the actress' name was presented as text on the screen. Single cell responses were also reported for Halle Berry, the Sidney Opera House, The Beatles and Bill Clinton (Quiroga et al., 2005). Similarly, a more recent study has identified individual neurons in the mouse brain which respond to any object that may allow the affordance of a bed or nest (Lin et al., 2007). In summary, these results suggest a general increase in the complexity of neural representations, along with the flow of information, from basic sensory representations, such as simple and complex cells, in the sensory cortices, to single stimulus representations in the rhinal cortices, to the integration of multiple stimuli with spatial, contextual, and emotional information in the tri-synaptic loop of the hippocampus proper (e.g. Gabrieli et al., 1997, see Squire, Stark and Clark, 2004; Stark, 2007 for a review).

3.8 Critique of the Declarative Memory Theory

All of the evidence discussed thus far provides general support for the notion that declarative memory function is mediated by the hippocampal formation, and this remains the most cited theory of memory in neuroscience (Morris, 2007). While there is a large body of neuropsychological evidence from humans and other mammals to support the hypothesis, it is by no means conclusive, and has been criticised on several grounds (e.g. Gaffan, 2002; Murray and Wise, 2004). It is difficult to explicitly examine episodic memory in animals, and the use of human subjects is often impractical. Therefore, the majority of support for the theory is obtained from lesion studies, but the extent and loci of damage which exists following a lesion is difficult to judge – and the hippocampus is an integral part of several neural processing circuits, so impairments do not explicitly imply this cortical region as a functional mediator. Similarly, those who have suffered MTL damage from birth may have developed alternative neural strategies for various forms of learning and memory. Results obtained from epileptic patients can also be questioned on the grounds that these subjects already exhibit abnormal hippocampal function (see Redish, 1999; Bird and Burgess, 2008 for a review). The observations of single neuron activity in response to sensory stimuli are relatively rare, and recent, and so the theory has also often been criticised on the grounds that it lacks the specificity that is necessary to make novel predictions, or to identify the computations that are performed in the hippocampus (O'Keefe, 2007; Bird and Burgess, 2008). A different class of observations regarding single neuron activity have also given birth to a new,

or at least partially modified, theory of hippocampal function, which has largely superseded the declarative memory theory in popularity.

3.9 The Hippocampus and Spatial Memory

Spatial memory is a sub-division of declarative memory which relates to the knowledge of ones environment and has come to be closely associated with the postulated function of the hippocampus. The discovery of distinct sets of pyramidal cells whose activity corresponds directly to various spatial features has propelled this theory into the foreground. While examining the hippocampal EEG of rats, O'Keefe and Dostrovsky noted that the activity of certain cells correlated directly with the position of the animal in its environment (O'Keefe and Dostrovsky, 1971). These pyramidal neurons were named place cells, and the corresponding areas of the environment which preferentially activated them dubbed their place field. Since that time, two other classes of hippocampal neuron with strong spatial activity correlates have also been discovered: the head direction cell (in various parts of the hippocampal formation, most notably the anterior thalamus and dorsal pre-subiculum; Ranck, 1984; Taube, Muller and Ranck, 1990a, b) and the grid cell (in the medial entorhinal cortex, Fyhn et al., 2004; Hafting et al., 2005). Evidence has also been obtained for neural activity which correlates with idiothetic information (i.e. proprioceptive, vestibular and reafferent signals from intended movements), such as velocity cells in the pre-subiculum and hippocampus (Sharp, 1996; O'Keefe et al., 1998) and whole body motion cells (O'Mara et al., 1994). Furthermore, it has been observed that the firing rate of place cells can be modulated by speed (Huxter, Burgess and O'Keefe, 2003; see Bird and Burgess, 2008 for a review). The demonstration that place cell activity can be updated in the absence of sensory input suggests that these idiothetic inputs may serve to allow the representation of an animal's location to be maintained during motion in the dark (see O'Keefe, 2007 for a review).

The discovery of place cells has led to the development of a cognitive map theory of the hippocampus that now competes strongly with the declarative memory theory (Tolman, 1948; O'Keefe and Nadel, 1978). Broadly speaking, this states that the hippocampus maintains a map of known locations together with the distance and direction between them in order to aid efficient navigation. A place representation (i.e. collection of place cells) can be activated by direct sensory input, or by the prior activity of place cells together with idiothetic information and a prior knowledge of the distance and direction between that place and the current location. Head direction and grid cell activity might provide the distance and direction metrics needed to form a cognitive map of the environment that consists of overlapping place fields, and idiothetic input a means of navigating this map in the absence of sensory input (O'Keefe and Nadel, 1978; O'Keefe, 2007). A comparison can be made with the single neuron activity regarding non-spatial memory that has since been reported (and was discussed above) – as discrete (head direction and distance) representations are conjunctively coded by the hippocampus to create an abstract representation – in this case, location. The original cognitive map theory also purports that, while the function of the hippocampus may be purely spatial in animals, in humans it also subserves episodic and linguistic memory. It is important to note that place cell firing does not simply reflect sensory input, as it does not change with rotating view in a single location –

rather, it appears to be an abstracted, conjunctively coded sense of location, mediated by internal bearings and distal landmarks (Bird and Burgess, 2008). Later research has demonstrated that, in humans and primates, hippocampal cells encode spatial view, while parahippocampal cells encode location, allowing both egocentric and allocentric spatial representations to be maintained (Rolls, Robertson and Georges-Francois, 1997; Matsumura et al., 1999; Ekstrom, 2003). These spatial view cells can also be updated in the dark by idiothetic signals, based on previous knowledge of an environment and internal feedback regarding head direction or eye position (Rolls and Kesner, 2006).
While it is well documented that human patients with damage to the hippocampal formation are generally unable to become acquainted with new environments, further evidence in support of the cognitive map theory has been provided by numerous experiments on rats with similar lesions who show deficits in spatial learning on the Morris water maze and other tasks (e.g. Morris et al., 1982; see Section 2.5; Martin, Grimwood and Morris, 2000; and see Rolls, 2008 for a review). Studies of macaques and humans with hippocampal damage have also shown that they are impaired in tasks where the object seen and the location in which it was seen must be remembered (e.g. Burgess, Maguire and O'Keefe, 2002; Crane and Milner, 2005; see Rolls, 2008 for a review of earlier studies). In humans, fMRI studies have demonstrated that hippocampal activity is correlated with active navigation tasks (Maguire et al., 1998; Thomas et al., 2001; Hartley et al 2003; Parslow et al., 2004). Similarly, studies have shown that the hippocampi of London taxi drivers are significantly larger and undergo navigation related structural changes (Maguire et al., 2000). More recently, it has been demonstrated that injections of stem cells into the hipocampi of mice allows them to regain previously abolished spatial capacity (Yamasaki et al., 2007). Interestingly, it has also been observed that damage to the hippocampus impairs the ability of patients to imagine complex scenes, particularly in terms of spatial detail (Hassabis et al., 2007).

Furthermore, there have been a substantial number of observations of place cell activity which have helped to elaborate on the specific manner in which the hippocampus might encode information about the topology of an environment (see O'Keefe, 2007 for a review). This neural activity has some clear correlates with more general concepts of declarative memory. For example, spatial memory is also affected by temporally graded retrograde amnesia following hippocampal damage, as demonstrated by a study of the patient EP (see Squire, Clark and Bayley, 2004 for a review). Place cells retain the activity patterns associated with an environment for a period of months, and their activity can be correlated with performance on tests of spatial memory (e.g. O'Keefe and Speakman, 1987; Thompson and Best, 1990). After this period, cognitive maps seem to be stored outside of the hippocampus, as research demonstrates that the greater the delay between learning a route and undergoing a hippocampal lesion, the better the recall of that route (Redish and Toretzky, 1998; Teng and Squire, 1999; Squire, Clark and Bayley, 2004). Similarly, the retrieval of spatial maps that were formed in the remote past does not activate the hippocampus (Moscovitch et al., 2007). The process of spatial memory transfer is, again, associated with sleep – as it has been shown that recently learned place cell sequences tend to fire in the same order during slow wave sleep (Wilson and McNaughton, 1994; Redish and Touretzky, 1998).

3.10 Critique of the Cognitive Map Theory

Debate has raged over whether spatial processing represents the primary function of the hippocampus, or whether it is merely a subset of declarative memory that is more coherently experimentally observable. For example, it is easy to establish where a rat is located within a maze, but harder to assess what conscious recollections of the past it may be making. Whether animals possess the capability for episodic memory at all has been a matter of some controversy, although several recent research papers suggest that many species are capable of episodic recall (Griffiths, Dickinson and Clayton, 1999; Eacott, Easton and Zinkivskay, 2005; Kart-Teke et al., 2006). Some proponents of the cognitive map theory have suggested that the impaired performace on non-spatial tasks which results from hippocampal lesions could be a result of spatial strategies that are employed in these tasks (Stark, 2007). However, such criticism cannot rule out the growing number of observations of non-spatial correlates of hippocampal firing. For example, recent studies have demonstrated that place cell activity in the rodent hippocampus simultaneously encodes goal-directed actions (Gothard et al., 1996a, 1996b, 2001; Ekstrom, 2003; for a review, see Wood and Dudchenko, 2003). This is illustrated by studies of rats running in a T-shaped maze, whose place cell activity is distinctly different (for overlapping sections of the route) depending on whether they plan to turn right or left at the junction (Wood et al., 2000; see Redish, 1999 for a comprehensive review of earlier findings). Furthermore, several studies have demonstrated that changes in task within a single environment can stimulate changes in the place fields of pyramidal neurons (Wiener, Paul and Eichenbaum, 1989; Markus et al., 1995).

The manipulation of non-spatial landmarks in an environment can also drastically alter place cell firing (see Eichenbaum, 1996 for a review), and the rate of place cell firing has also been demonstrated to code for non-spatial aspects of behaviour, such as objects or other stimuli encountered (Wood, Dudchenko and Eichenbaum, 1999; Wood and Dudchenko, 2003; O'Keefe, 2007). Similarly, place cells are often context dependent - their activity being modulated by reward and punishment, for example (Wood and Dudchenko, 2003). Studies of macaque monkeys support these findings – the occurrence and firing rate of spatial view cells being modulated by expected reward or encountered stimuli, and the existence of cells which respond to the presence of objects or conjunctive presence of objects and spatial view being documented (Rolls and Xiang, 2005; Rolls, Xiang and Franco, 2005; Rolls and Xiang, 2006). As described earlier, a growing number of animal studies have also identified the selective firing of hippocampal cells in response to non-spatial cues, such as olfactory stimuli (Wood, Dudchenko and Eichenbaum, 1999), specific objects (Tamura et al, 1992; Rolls, Xiang and Franco, 2005; Rolls and Xiang, 2006), arousing stimuli such as food or a known face (Vidyasagar, Salzmann and Creutzfeldt, 1991; Salzmann, Vidyasagar and Creutzfeldt, 1993; Fried, McDonald and Wilson, 1997) and in humans, in response to various images (Kreiman, Koch and Fried, 2000; Fried et al., 2002; see Redish, 1999; Eichenbaum, 2004 and O'Keefe, 2007 for a review of all of these).

It is important to remember that the cognitive map theory does not deny the involvement of the hippocampus in episodic memory function, but merely postulates that spatial memory is the primary

function of this brain region. It has therefore been suggested by proponents of the cognitive map theory that each of the responses described above can be gated by varying the location in which these stimuli are encountered (O'Keefe, 2007). For example, the majority of place fields in the hippocampus retain their firing location when non-spatial cues are altered, provided that there is no ambiguity regarding the animals current location – although the firing rates do change significantly (Leutgeb et al. 2005a, b). In studies of epileptic patients with hippocampal electrode implants, it was noted that a significantly higher number of cells would preferentially respond to the presentation of spatial images, such as vistas, houses or scenes, as oppose to non-spatial images, such as faces or objects (Kreiman, Koch and Fried, 2000). Conversely, it should be noted that tests of spatial memory often employ sparse environments with only distal visual landmarks, and therefore an absence of local, non-spatial cues that could be encoded (Eichenbaum, 1996). In one study where a standard maze task was furnished with an abundance of local visual and tactile cues, an approximately equal number of neurons were demonstrated to independently encode spatial and non-spatial cues, while the majority of cells encoded conjunctions of non-spatial and spatial information (Young, Fox and Eichenbaum, 1994). Similar experiments have also identified a greater number of cells with non-spatial, rather than spatial, correlates (Wood, Dudchenko and Eichenbaum, 1999).

Elsewhere, it has been noted that place cells demonstrate some properties which suggest that their function is not directed towards cognitive mapping. For example, it is not clear how their activity might aid navigation, and the overall activity of place cells does not suggest a cohesive, continuous mapping system (Morris, 1990; see Eichenbaum et al., 1999 for a review). Several studies have suggested that – rather than forming a homogenous map of known environments – place cell activity reflects an independent representation of spatial cues. Place fields are often clustered around areas with local stimuli, such as the walls of an arena, and are bound to these cues, so that they will move with the cue, rather than remaining anchored to the spatial representation of the rest of the environment (Gothard et al., 1996a, b; O'Keefe and Burgess, 1996; Hetherington and Shapiro, 1997; Shapiro, Tanila and Eichenbaum, 1997; Tanila et al., 1997a, b, c). It has also been demonstrated that rats can still accomplish path integration after lesions to the MTL, and some human patients demonstrate intact spatial processing despite major hippocampal damage (Alyan and McNaughton, 1999; Teng and Squire, 1999; see Redish, 1999 and Morris, 2007 for reviews). However, contradictory results have also been presented, in which damage to the hippocampal formation impairs path integration, although the exact contribution which it makes to this process has not been clarified (Maaswinkel, Jarrard and Whishaw, 1999; Save, Guazzelli and Poucet, 2001). These conflicting findings might suggest that, while the hippocampus mediates the acquisition and storage of new spatial knowledge, the processing of this knowledge, as well as more basic navigational function, occurs outside of this region. For example, it has been noted that damage limited to the parahippocampal and parietal cortices can specifically impair spatial computation in humans and monkeys (see Rolls and Kesner, 2006 for a review).

In essence, this debate can be reduced to a matter of opposing viewpoints – proponents of the cognitive map theory argue that there are considerably more significant neural responses to spatial information, which suggests that this is the primary function of the hippocampus, although it may have evolved to also play a role in the wider processing of declarative memory by embedding the memory of non-spatial stimuli within the cognitive map (O'Keefe, 2007). Proponents of the declarative memory theory, on the other hand, suggest that space is merely a mental construct, no different from any other relational form of associative memory, and that while neural correlates of this more tangible sensory stimulus are easier to come by, it remains a subset of declarative memory (Morris, 2007). Supporters of the cognitive map theory suggest that the kind of navigation tasks which are impaired by hippocampal damage involve knowledge and computation beyond what is directly perceived, so that it cannot be considered merely a subset of episodic memory. However, it has been demonstrated that spatial processing can be accomplished in the absence of an intact hippocampus, as described previously.

Neither theory is perceived as cogent or able to explain all experimental findings, although the cognitive map theory has been able to supply more directly testable predictions and cohesive observations, and thus allowed the development of a significant number of more precise models of hippocampal function. It should also be noted that there are several other theories regarding the function of the hippocampus which take a more abstract view of its function – suggesting that it may merely allow the recall of context upon entering a new environment, or that it mediates associations and relations between all manner of stimuli, for example (Redish, 1999; Eichenbaum, 2004; see Morris, 2007 for a review). Regardless of which of these theories may be the most accurate or able to explain the wealth of experimental data obtained from the hippocampus, they generally concur broadly when it comes to computational modelling of the MTL, and thereby suggesting how this brain structure may perform the different functions attributed to it. Broadly speaking, all groups can agree that rich, multi-modal sensory information enters the rhinal cortices, where it is processed into more abstract, conjunctional representations which are then stored in the hippocampus for some period of time using auto-associative network dynamics.

3.11 Auto-associative Networks

Auto-associative memory models comprise of a neural network with recurrent connectivity which implements some form of Hebbian learning rule. Stable patterns of input activity can be stored in the synaptic weights of the recurrent connections, which then drive the dynamics of the network during later activity so that entire patterns can be recalled from incomplete or noisy external cues (Marr, 1971; Hopfield, 1982; Amit, 1995**;** Rolls and Treves, 1998; Redish, 1999; Tsodyks, 2005; Rolls, 2008). Non-linear activation functions or global inhibitory input are often employed in order to stabilise this positive feedback system (Rolls and Treves, 1998). These networks represent a powerful tool in computational neuroscience, offering a unified implementation of memory encoding, storage and retrieval within a robust framework. They also exhibit many features that are analogous with human associative memory (Rolls and Treves, 1998; Tsodyks, 2005; Rolls, 2008). Multiple patterns can be stored after a single presentation, and these models have the ability to generalise between similar input

patterns, allowing unsupervised prototype formation and a mechanism of noise reduction (Rolls and Treves 1998, Kesner and Rolls 2001). Furthermore, due to the distributed nature of storage, auto-associative networks exhibit graceful degradation – such that successful recall is only incrementally impaired by the loss of constituent synapses and/or neurons. One further property of auto-associative networks that has proved critical to their use in mnemonic modelling is the ability to sustain recall activity in the absence of input via recurrent self-excitation. Effectively, activity in the network can settle into an attractor state in which it will remain until a new input cue is presented that shifts activity to a new basin of attraction. This feature has often been likened to a short-term memory mechanism, and these attractor networks have been hypothesised to mediate short term memory function in many brain regions, including the frontal eye fields, visual cortex, prefrontal cortex and rhinal cortices (Miyashita and Chang, 1988; Rolls and Tovee, 1994; Amit, 1995; Goldman-Rakic, 1996; see Rolls and Treves, 1998 and Rolls, 2008 for reviews).

3.12 Auto-associative Network Models of the Hippocampus

The vast majority of computational theories regarding the function of the hippocampus in human memory make use of the properties of auto-associative networks, since Marr's classic model was published in 1971. These modelling studies have successfully replicated various aspects of navigation, spatial and episodic memory – and related these models to known properties of various hippocampal regions (see Tsodyks, 1999; Redish, 1999; Lever et al., 2003; Rolls, 2008 for reviews). Although the basic architecture of these models is identical, they differ significantly in their properties, often according to the postulated function which they assign to the hippocampus. Those who emphasise the role of the hippocampus in episodic memory generally suggest that input from a variety of sensory areas (as well as the brain stem and amygdala) is integrated by the rhinal cortices and then encoded as abstract, discrete patterns of activity in the hippocampus proper (i.e. Rolls and Treves, 1998; Hasselmo and Eichenbaum, 2005; Rolls, 2008). The encoded memories are stored as discrete attractors by the plastic synapses of the CA3 network – which exhibits an extensively recurrent architecture, and is one of the first networks on the tri-synaptic loop where information is integrated rather than segregated (Suzuki and Amaral, 1994a, b; Rolls and Treves, 1998). The later input of partial cues, representing some portion of the original memory, can then drive the dynamics of the CA3 network back to this attractor state, so that output activity corresponding to the original pattern of sensory stimulation in the rhinal cortices (and other neocortical areas) can be reinstated – allowing recall of the memory, and possible transfer for long-term storage in the cortex. Another of the main outputs of the hippocampus, via the fornix to the anterior nucleus of the thalamus, may also allow this output to influence direct action (Rolls and Treves, 1998). In light of this model, semantic memory may also be considered as prototype extraction from episodic memory (Eichenbaum, 2004).

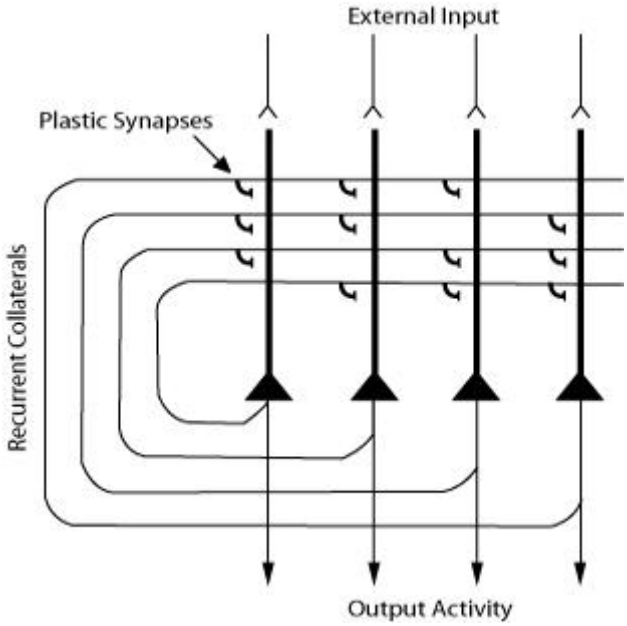

Figure 3.2: The architecture of a typical auto-associative network

In auto-associative network models, interference between the recall of learned patterns is minimised if the overlap of representations is kept to a minimum, and it has been theorised that the size of the DG indicates that its function is to generate sparse and distinct patterns of input based on the information it receives from the EC - a process known as pattern separation or orthogonalisation (Marr, 1971; Rolls, 2008). The capacity of attractor networks (and therefore, perhaps, the CA3 region) to maintain recall firing patterns may also act as a putative short term memory mechanism. Discrete auto-associative networks, by nature, demand symmetric weight matrices, so that a complete pattern can be recalled from any set of cues (Rolls, 2008). However, they can be easily extended to allow the storage of sequences, which represents another key feature of declarative memory (Levy, 1996; Wallenstein, Eichenbaum and Hasselmo, 1998; Lisman, 1999). In this extension of the model, asymmetric connections are necessary to ensure that the temporal order of a pattern is preserved during recall, and these connections must operate on a slower time scale than the process of pattern completion, so that each section of the dynamic pattern can be fully recalled before activity shifts to the subsequent section (Rolls and Treves, 1998). Most models have yet to identify a biologically plausible manner of implementing this constraint, however.

Conversely, those who emphasis the role of the hippocampus in spatial memory generally make use of continuous attractor network models, which employ a flat energy landscape and global inhibition to maintain a single localised 'bubble' of input activity via recurrent excitation (i.e. Samsonovich and McNaughton, 1997; Tsodyks, 1999; Fuhs and Toretzky, 2006; McNaughton et al., 2006; Stringer and Rolls, 2006). This activity packet can then be shifted by external input to move smoothly across the network – analogously to the encoded position of an animal within a cognitive map being updated by idiothetic cues. In this context, a one dimensional (ring) network can be used to represent head direction, while a two-dimensional (toroidal) network can be used to represent location, both being updated by external sensory input or feedback from the brain stem. These models can be more directly

based on, and compared with, empirical data, due to the wealth of observations regarding the activity of place, head direction and other spatially correlated cells. There is also some more subtle computation involved in spatial attractor network models – as the output should not just be the recall of the appropriate cognitive map, but also some form of processing regarding goals, actions or short cuts within that locale. Continuous attractor network models also rely on symmetric weight matrices, as associations between adjacent head directions or known places must be identical, regardless of the direction of approach. However, like discrete attractor networks, they can easily be extended to allow the learning of sequences – which could, in this case, correspond to well known routes. Because they utilise the same architecture, it is also possible that the CA3 network can integrate both spatial and episodic memory function, and thus recall discrete objects or experiences from locational input and vice versa (Rolls, Stringer and Trappenberg, 2002).

3.13 Biological Correlates of Auto-associative Function in the Hippocampus

Auto-associative network models of declarative memory function offer several testable predictions, based on the manner in which they divide processing tasks among different regions of the MTL (see Eichenbaum, 2004; Kesner, Lee and Gilbert, 2004; Rolls and Kesner, 2006; Kesner, 2007a, b; Leutgeb and Leutgeb, 2007 and Rolls, 2008 for reviews). These models theorise that the dentate gyrus acts to produce sparse non-redundant representations of arbitrary sets of concurrent neural activity in the pyramidal neurons of the CA3 region, as such patterns are required for successful auto-association. The entorhinal cortex demonstrates preferential activation in response to novel stimuli, so that particular emphasis is placed on the representation of these stimuli in downstream processing stages. These patterns are rapidly encoded in the synaptic weights of the CA3 network using Hebbian plasticity, and the synapses of the perforant path are also adjusted to allow direct connections from the entorhinal cortex to stimulate future recall from noisy or incomplete cues. The function of the CA1 network is to recode the abstract, conjunctive code which is stored by CA3 into the form in which it originally arrived from the sensory cortices.

There is a substantial amount of evidence to suggest that the DG is involved in a process of pattern separation and the encoding of novel memories. For example, studies have demonstrated that mossy fibre input (from the DG to CA3) is essential for the process of learning in the Morris water maze, but not for the recall of routes which had been learned earlier (Lasalle, Bataille and Halley, 2000; Lee and Kesner, 2004; Jerman, Kesner and Hunsaker, 2006). It has also been demonstrated that the perforant path is consequently essential for memory retrieval (Lee and Kesner, 2004). Functional imaging studies have correlated activity in the DG during the early stages of a trial with successful spatial learning, and shown that this activity dictates that in the CA3 field. During later training, when recall processes may begin to dominate, it is activity in the CA fields which correlates with successful performance, and the influence of DG activity is vastly reduced (Poirier, Amin and Aggleton, 2008). Elsewhere, it has been demonstrated that enriched environments provoke a larger DG, with more granule cells, in rats (Kempermann, Kuhn and Gage, 1997).

It is also well known that mossy fibre synapses are very large, and may thus be capable of driving the firing of CA3 neurons. It is essential that activity in an auto-associative network approximates external input during learning, otherwise a pattern which combines input and recall activity will be stored by the recurrent synapses. Studies of synaptic plasticity in the CA3 network have demonstrated that the depolarisation caused by input from the mossy fibres is critical for robust potentiation to be induced in the CA3 network (Kobayashi and Poo, 2004). The evidence obtained in this study also suggests that the mossy fibres release a mGluR mediated signal to further facilitate LTP in the recurrent synapses, as well as providing sufficient input to provoke activity in the post-synaptic pyramidal neuron. Hence, input from the dentate gyrus is essential for the induction of LTP in the CA3 network. Neural activity in the DG is also sparse, with the place fields of granule cells being small and highly reliable. The hetero-synaptic, non-associative form of plasticity observed at mossy fibre synapses may also act to enhance the signal to noise ratio of these neural representations (see Rolls, 1996; Rolls and Treves, 1998; Kesner, Lee and Gilbert, 2004 and Rolls, 2008 for a review of these findings).

Selective lesions of the dentate gyrus have been observed to significantly reduce the reliability of CA3 neural firing fields, and when DG lesioned rats are exposed to similar spatial locations, their inability to discriminate between them is correlated with the extent of their lesions, and with the similarity of the environments – implying a deficit in pattern separation (McNaughton et al., 1989; Gilbert, Kesner and DeCoteau, 1998; Gilbert, Kesner and Lee, 2001). Rats with lesions of the DG also exhibit little exploratory interest after the reconfiguration of object locations, when compared to control / CA1 lesioned rats, suggesting that they are unable to distinguish between the changed environments due to an absence of pattern separation (Goodrich-Hunsaker, Hunsaker and Kesner, 2005). Similarly, DG-lesioned rats have great difficulty learning a water maze when their start position is changed for each trial (Lee and Kesner, 2004). The need for a pattern separation process, mediated by the DG, is further illustrated by studies which show that a deficit in learning object-place paired associations is only apparent in lesioned mice if the stimuli are close together (Gilbert and Kesner, 2003).

Of course, it is the conjunction of the dentate gyrus and CA3 region which is theorised to mediate successful auto-associative learning of arbitrary activity patterns. We have seen that the DG is specifically implicated in encoding, as oppose to retrieval, processes. Similar learning deficits are observed in animals with CA3 lesions, and in animals with contralateral lesions (in which the CA3 region is lesioned in one hemisphere, and the DG in the other) (Jerman, Kesner and Hunsaker, 2006). These results clearly imply that these regions of the hippocampus operate in unison to encode and store sparse activity patterns relating to sensory experience. More recent functional imaging studies support the theory that a process of pattern separation takes place in the dentate gyrus and CA3 field, while the effects of pattern completion and recoding can be observed in the CA1 field and parahippocampal structures (Bakker et al., 2008). Further studies have demonstrated that neurotoxic lesions of CA3 impair the learning of a novel environment, while similar lesions to CA1 have no such effect (Lee and Kesner, 2003). It has also been demonstrated that place fields within a novel environment are acquired and altered much more rapidly in CA3 than in CA1 (Lee, Rao and Knierem, 2004). This varying time

course of CA1 and CA3 activity is further illustrated by studies which demonstrate that, when learning a maze, rats with CA3 lesions are impaired on same-day performance, while those with CA1 lesions are only impaired in retrieval across days (Jerman, Kesner and Hunsaker, 2006; Vago et al., 2007). It has also become clear that the CA3 region is important for tasks which require multiple trials to learn, such as object-place and odour-place paired association learning (Gilbert and Kesner, 2003; Lee and Kesner, 2003). For example, CA3 lesioned mice are often unable to learn the water maze, which requires multiple trials (Brun et al., 2002).

Elsewhere, studies have demonstrated that distinct, sparse activity patterns are generated in the CA3 region in response to virtually identical environments – a clear example of a pattern separation process, probably inherited from the dentate gyrus (Tanila, 1999; Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004). The sparse patterns that exist in CA3 have also been demonstrated to change discontinuously in response to changes in the external environment, suggesting a process of orthogonalisation, while activity in CA1 is altered more systematically in response to environmental manipulations (Vazdarjanova and Guzowksi, 2004; Leutgeb et al., 2005c). Much fewer CA3 (~18%) than CA1 (~35%), and even fewer DG (~2%) neurons are active during each stage of these tests, which further attests to the sparseness of representations and the progression of the pattern separation and completion process along the tri-synaptic loop (Vazdarjanova and Guzowski, 2004). Similarly, there are distinct differences in the overlap of activity patterns within each region of the hippocampus when animals are placed in very similar environments. In CA3, distinct subsets of pyramidal cells which overlap with no greater than chance probability are activated, regardless of the similarity of the testing enclosures, while in CA1 the activated populations overlap, and the degree of this overlap increases with the similarity of the enclosures (Vazdarjanova and Guzowksi, 2004; Leutgeb et al., 2004).

The CA3 region is also hypothesised to play a role in pattern completion during the recall of stored episodic memories from partial cues, and this property has been supported by various empirical findings. It has been demonstrated that both CA3 and CA1 are required for the retrieval of spatial knowledge, although larger deficits are produced by CA1 lesions (Hunsaker, Lee and Kesner, 2008). For example, rats with CA3 lesions performing a task in which a number of cues are available during learning exhibit a linear degradation in recall which correlates directly with the removal of cues. Thus, the rats seem unable to generalise from the recall environment, with a reduced number of cues, to the original learning environment, via a process of pattern completion (Gold and Kesner, 2005). Similarly, transgenic mice that specifically lack NMDA receptors in the CA3 field are unable to identify known environments when a subset of the local cues are removed (Nakazawa et al., 2002). Another study examined the behaviour of place cells as an arena was slowly morphed between different shapes, and demonstrated that activity in the CA3 region abruptly switched between two distinct patterns. This implies the use of discrete attractor dynamics, as activity is resistant to small changes in environment, but will then shift drastically to a new basin of attraction once the differences become significant (Wills et al., 2005).

It is important to note that the conflicting nature of the pattern separation and completion processes occurring within CA3 can lead to confusing results – for example, we have seen how perturbations in an environment will lead to the activation of a distinct set of pyramidal neurons in CA3, as evidence of pattern separation. It has been noted elsewhere that CA3 activity can remain constant during mild changes in an environment – which may seem contradictory, but has been taken as further evidence for generalisation and attractor dynamics operating therein (Vazdarjanova and Guzowksi, 2004). The conflicting processes of pattern separation and completion can be considered as a sigmoidal process between input and output similarity. If input patterns are only marginally dissimilar, then a process of pattern completion will maintain activity within a single basin of attraction, but once some threshold of similarity is crossed, pattern separation will occur and activity within the CA3 region will shift significantly to some new basin of attraction (O'Reilly and McClelland, 1994; McClelland and Goddard, 1996). This process may also explain the commonly observed phenomena of remapping in place cells (i.e. Fyhn et al., 2007).

Other research studies have identified persistent activity in the CA3 region during the delay period of short-term memory tasks, and it has been also been demonstrated that spatial view neurons maintain their activity patterns in the absence of visual input (Rolls et al., 1989; Wirth et al., 2003). These findings imply that the CA3 network may maintain firing patterns in an attractor state until a new cue is presented, acting as a short-term memory buffer. This role has been postulated for the region, due to its recurrent connectivity, but is not essential to its function, or directly proven by these results, as partial cues obtained from sensory experience could be maintained (or 'clamped') by some other, upstream processing unit such as the dentate gyrus or entorhinal cortex (Kesner and Rolls, 2001; Rolls and Kesner, 2006).

3.14 Auto-associative Network Models and Synaptic Plasticity

The fact that synaptic plasticity is so easily and reliably observed in the hippocampus provides further impetus to the theory that this region may operate as an auto-associative network. A process of synaptic plasticity is integral to the performance of these models, and many studies have provided explicit evidence for the role of LTP in hippocampal function (see Section 2.5). Indeed, the principles of associative learning are closely linked with those of Hebbian plasticity – that connections between concurrent events (behavioural or neural) are strengthened, so that in the future, the stimulation of one will provoke stimulation of the other. In particular, it has been repeatedly demonstrated that obstructing plasticity in the CA3 region prevents the rapid acquisition of new information (see Rolls and Kesner, 2006 for a review). It has also been shown that blocking plasticity (using an NMDAr antagonist) in CA3 impairs the learning of novel environments, whereas impairing LTP in CA1 or DG had no such effect (Lee and Kesner, 2002). Further work has elaborated on this result, confirming that NMDA antagonists impair declarative encoding and learning processes, with no effect on recall, and also demonstrating that AMPA receptor blockade impaired both learning and recall (Day, Langston and Morris, 2003). Similar results were obtained with transgenic mice, engineered to lack NMDA receptors specifically in the CA3 field. A number of studies using these mice have demonstrated that plasticity in

CA3 in essential for the rapid and automatic formation of episodic memories (Nakazawa et al., 2002; Nakazawa et al., 2003; Cravens et al., 2006). Some studies have gone further, in demonstrating that synaptic plasticity processes are specifically activated when animals encounter novel configurations of sensory stimuli for the first time (Lee and Kesner, 2002). There is also evidence that enriched environments lead to the regulation of several properties of neural and synaptic dynamics in the hippocampus, such as the increased decay of LTP, which are indicative of increased information turnover (Irvine et al., 2005). The use of immediate early gene (IEG) imaging has also provided support for the critical role of synaptic plasticity in hippocampal function. For example, various plasticity related IEGs are expressed in the hippocampus following spatial exploration, particularly in novel environments or after the re-arrangment of local cues (see Kubik, Miyashita and Guzowski, 2007 for a review).

3.15 Conclusions

In this chapter, a substantial body of experimental data which provides support for the theory that the hippocampus plays a significant role in declarative memory processing has been presented. This brain region is particularly implicated in the initial formation and short-term storage of episodic and/or spatial memories, and a substantial body of evidence suggests that this function is mediated (at least, in part) by auto-associative network dynamics. Broadly speaking, it is hypothesised that the dentate gyrus creates a sparse, abstract and conjunctive code from the sensory representations present in the rhinal cortices (with stronger activation for novel stimuli), and that this activity pattern is then rapidly stored in the recurrent collaterals of the CA3 region by synaptic plasticity. This general model can be considered an inclusive, neurobiological theory of MTL function which integrates observations of synaptic plasticity, neural activity and emergent dynamics in this region (Morris, 2007).

It can only be presumed that the process of natural selection has developed this system as a robust and efficient means of mediating declarative memory in the mammalian brain, and that by attempting to replicate the neural and synaptic dynamics of the hippocampus (and surrounding parts of the MTL) as closely as possible, a similarly successful computational model of declarative mnemonic processing might also be developed. Such a model will have clear benefits for the study of artificial intelligence and various mobile robotics applications. Although the complexity of this task is beyond the scope of contemporary computational neuroscience, developing a biologically realistic auto-associative network model appears to represent the critical first step towards this ultimate goal. These networks underpin wider models of both episodic and spatial memory function, and are worthy of study in their own right as minimal network architectures that can perform useful computational processes.

The discussion presented in the preceding chapters makes it clear that synaptic plasticity plays a critical role in mediating both the function of auto-associative networks and mnemonic processing in the hippocampus. In the following chapter, the learning rules that have been commonly implemented in auto-associative network models will be assessed with regard to the emergent properties of synaptic plasticity in the hippocampus which were described in the previous chapter. Simultaneously, the neural

dynamics and wider details of common auto-associative network models will be assessed in light of their relation to the observed neurobiology of the CA3 region. This should provide an insight into which features of these models might have their biological realism improved, and how this might contribute to developing a more complex and computationally useful auto-associative network model of the CA3 region. The resulting auto-associative network could then be integrated into various wider models of the postulated function of the hippocampus in both episodic and spatial memory, in order to provide them with a firmer basis in modern neurobiology.

CHAPTER FOUR: EXPERIMENTAL AIMS

"Neither one should hesitate about dedicating oneself to philosophy when young, nor should get tired of doing it when one is old, because no one is ever too young or too old to reach one's souls healthy"

Epicurus

Aims

- To highlight the fact that few previous modelling studies have investigated the emergent dynamics of STDP specifically within a recurrent architecture, or attempted to employ this plasticity rule within an auto-associative network
- To emphasise the fact that auto-associative networks represent an ideal framework in which to attempt a reconciliation of temporally coded and rate coded plasticity models
- To identify the key properties of Hebbian learning rules that contribute to efficient auto-associative network function, and which must therefore be replicated by STDP within this context
- To identify other shortfalls in biological realism that commonly exist in auto-associative network models of the hippocampus
- To discuss how the properties of STDP which have been identified by previous modelling studies might address these shortfalls, and describe any features of the plasticity rule that may be detrimental to auto-associative network function
- To define the aims of this research and describe the general features of the network model that will be examined in the following chapters

4.1 Introduction

In the preceding chapters, a wealth of evidence linking the process of long-term potentiation with associative memory function in the hippocampus has been presented, although it is also clear that synaptic plasticity alone may not be sufficient to fully explain the complexities of this behaviour. STDP represents the most contemporary plasticity rule formulated from electrophysiological recordings of this brain region and offers several advantages for the operation of neural networks such as increased information processing power and inherent homeostasis. However, many different computational interpretations of STDP exist, and these have yet to be unified or reconciled with previous synaptic plasticity models and the suspected function of the hippocampus. Auto-associative networks represent the most successful approach to modelling this postulated function, and can directly reconcile observations of synaptic plasticity and neural dynamics with systems level properties. These models are attractive because they offer a unified description of the encoding, storage and retrieval processes that constitute declarative memory, and can replicate features of both episodic and spatial memory function. Because auto-associative function relies critically on rate-coded Hebbian learning, these networks also represent an ideal framework in which to attempt a reconciliation of STDP with previous rate-coded plasticity rules.

To date, no computational studies have specifically aimed to characterise the emergent neural and synaptic dynamics of STDP within a recurrent neural architecture, and only one has attempted to reconcile this plasticity rule with the successes of previous auto-associative models of the hippocampus (Samura and Hattori, 2005). The vast majority of previous computational modelling studies that have sought to characterise the emergent properties of STDP have made use of feed-forward networks or multiple afferent synapses of a single neuron. It is surprising that examinations of STDP have so often proceeded without reference to its role in hippocampal information processing, as it is in this region that the majority of empirical observations of synaptic plasticity and mammalian memory function are made. Where STDP has been implemented within a recurrent network, the focus of research has most frequently been on modelling the function of the visual cortex, the role of sub-threshold membrane potential oscillations which occur during slow wave sleep, or the self-organisation of networks through spontaneous activity (i.e. Wenisch, Noll and van Hemmen, 2005; Yasui et al., 2007; Kang, Kitano and Fukai, 2008).

More recently, several research projects have employed STDP within an auto-associative framework in order to learn and recall dynamic patterns (i.e. temporal sequences of neural spiking) – a form of learning for which this plasticity rule is ideally suited (i.e. Lengyel et al. 2005; Yamaguchi et al., 2007; see also Section 2.16). However, these projects have made no attempt to reconcile STDP with the process of learning and recalling static rate-coded activity patterns (be they discrete or continuous) which characterises the majority of previous auto-associative modelling of the hippocampus. In the one case where static, bi-directional auto-associative memory function has been examined, a symmetric form of STDP that may not be realistic of the CA3 region was utilised (Samura and Hattori, 2005). A substantial amount of analytical work has provided invaluable insight into the limits and nature of

emergent synaptic dynamics, but these results cannot be directly applied to, or inform the nature of, more complex network behaviour, which might incorporate spike triggering effects and non-Poissonian firing patterns, for example (i.e. Izhikevich and Desai, 2003; Burkitt, Meffin and Grayden, 2004; Standage, Jalil and Trappenberg, 2007).

In summary, synaptic plasticity is critical to the function of auto-associative networks, and is most frequently, easily and reliably observed in the hippocampus. Contemporary neurobiological investigation of this region has delineated a spike-timing dependent plasticity model that differs significantly from the rate-coded Hebbian learning rule most commonly implemented. Hence, it seems that there is a pressing need to reconcile the properties of STDP with the dynamics of auto-associative networks, as these processes are both likely to be critical to the function of the hippocampal formation. These models are also worth studying for their own sake, as minimal neural network architectures which can perform effective and useful computational processes.

The primary aims of this research, therefore, are three-fold. Firstly, the emergent synaptic dynamics of numerous models of STDP will be examined within a recurrent neural network that is consistently inspired by the neurobiology of the CA3 region, in order to characterise the properties of the plasticity rule within this framework. Secondly, a more specific attempt to reconcile STDP with the emergent dynamics of previous rate-coded models of Hebbian learning and properties of synaptic plasticity in the CA3 region will be made. Finally, any forms of STDP that are capable of mediating rate-coded processing will be implemented within a simple auto-associative network model of the CA3 region, in order to further illuminate any factors which influence the efficient operation of the plasticity rule within this context. This should allow an assessment of the functional implications of STDP within the CA3 region of the hippocampus to proceed, and provide previous auto-associative network models with a firmer basis in modern neurobiology. Several issues that constrain and inform the modelling decisions involved in constructing a framework for this research will now be discussed in more detail.

4.2 Biological Realism in Previous Auto-associative Network Models of the Hippocampus

There are several issues that arise when attempting to reconcile the operation of common auto-associative network models with empirical observations of activity in the hippocampus. These problems are not sufficient to bring the general framework into doubt, but rather represent areas where biological realism might be improved. For example, the firing patterns generated in recurrent neural networks often differ significantly from those observed *in vivo*. In the majority of previous models, active neurons fire close to saturation, at a level much higher than that observed in the CA3 region, while background neurons are effectively silent (i.e. Marr, 1971; Hopfield, 1982; Roudi and Latham, 2007). Intrinsic plasticity mechanisms, inhibitory input and sigmoidal activation curves have each been employed to address these issues (i.e. Rolls, Stringer and Trappenberg, 2002; Triesch, 2007). Furthermore, typical network models generally operate using either time averaged firing rates or extremely regular spiking dynamics, while activity patterns observed *in vivo* are generally much more irregular and can exhibit several distinct examples of oscillatory behaviour (i.e Samsonovich and

McNaughton, 1997; but see Mongillo et al., 2005). This problem is exacerbated by the fact that active neurons in working memory network studies often fire more regularly than background neurons, which directly contradicts empirical data (see Roudi and Latham, 2007 for a review).

Frequently, learning and recall processes have also been arbitrarily separated in previous modelling studies. In auto-associative networks, it is essential that external input dominates the dynamics of the network during learning, in order to avoid interference between this process and the recall of previously established patterns (Rolls and Treves, 1998). In the hippocampus, however, it seems likely that plasticity processes continue at all times, although they may be modulated by a number of different factors (e.g. Hasselmo, 2006; see Section 2.4). In particular, research discussed in the preceding chapters has highlighted the separate roles of input from the dentate gyrus for learning, and from the entorhinal cortex for recall, as well as the manner in which input from the DG can facilitate synaptic plasticity (see Section 3). In the absence of modulation, however, it is critical that both learning and recall processes can be accounted for while synaptic plasticity remains continually active.

The ability to maintain firing patterns as a stable attractor state of the system once external stimulation has been removed has often been incorporated into auto-associative models of the hippocampus as a putative short term or working memory mechanism. This feature is supported by observations of persistent activity during working memory tasks *in vivo*. Although this function is attractive from a computational modelling standpoint, and a subject of research into working memory models in its own right, it is not considered to be an essential feature of auto-associative memory models of the hippocampus. It is possible that external cues are maintained, even after the sensory stimuli that created them have been removed, by a structure that is upstream of the CA3 region, such as the entorhinal cortex (Rolls, 2008).

4.3 Synaptic Plasticity in Previous Auto-associative Network Models of the Hippocampus

In previous computational studies, auto-associative network models of episodic memory function have frequently made use of synaptic plasticity in order to develop weight matrices which reflect patterns of external input. Conversely, continuous attractor network models of spatial memory function have often omitted plasticity altogether, choosing instead to arbitrarily implement the desired 'Mexican Hat' connectivity in recurrent synaptic weights and make no allowance for changes in this weight matrix during simulations (i.e. Samsonovich and McNaughton, 1997; Menschik, Yen and Finkel, 1999; Conklin and Eliasmith, 2003; Fuhs and Touretzky, 2006). It has been established that the performance of such models is extremely reliant on an appropriate synaptic weight matrix being developed and maintained, with a change of as little as 1% in synaptic strengths causing a significant (and unrealistic) drift of activity (see Brody, Romo and Kepecs, 2003 for a review). Other research has successfully employed synaptic plasticity to develop self-organising continuous attractor network models of various spatial memory processes, however (i.e. Stringer, Rolls and Trappenberg, 2005; Stringer and Rolls, 2006).

In the vast majority of studies that do utilise plasticity, Hebbian learning rules which strengthen connections based on correlations between mean firing rates are employed (i.e. Rolls, Stringer and Trappenberg, 2002; Stringer and Rolls, 2006; but see also Mongillo et al., 2005). In some cases, unconstrained plasticity is allowed, with global feedback inhibition being used to allow efficient recall from perpetually escalating synaptic strengths (i.e. Rolls, Stringer and Trappenberg, 2002). However, some form of heterosynaptic LTD has been identified as generally essential for maintaining efficient operation in auto-associative networks (Rolls, 1996; Rolls and Treves, 1998; Rolls and Kesner, 2006). This is commonly achieved through the use of a synaptic scaling mechanism, and there is a substantial body of empirical evidence to support the existence of such a process in the hippocampus (i.e. Stringer and Rolls, 2006 and see Section 2.10).

STDP has several properties which may contribute to improving the biological realism of auto-associative network models, but some that may also be detrimental to its performance in this role. For example, previous research has suggested that STDP specifically promotes irregular firing activity by maintaining membrane potentials at a level close to that required to provoke spiking (Song, Miller and Abbott, 2000; see Section 2.17). This is manifested as a value for the co-efficient of variation of neural activity that is relatively constant in the face of increases in pre-synaptic firing rate, and which occupies a range comparable to that observed *in vivo* (see Roudi and Latham, 2007 for a review). Conversely, the asymmetric learning window which characterises STDP implicitly promotes asymmetric connections, while bi-directional connections are essential in auto-associative networks to allow the recall of a pattern from any partial cue. The asymmetric nature of STDP does make it particularly suitable for the storage and retrieval of sequences, which is a key feature of associative memory. It might be suggested that this plasticity rule is therefore more suitable for processing temporally coded input, which has much greater computational power than rate coded data (see Section 2.14). However, the majority of auto-associative network models of the hippocampus that have been developed thus far rely on rate-coded data, and these models have successfully replicated a range of functions that are ascribed to this brain region (see Section 3.12). It seems plausible that a dual code may exist; allowing the CA3 region to process both temporal and rate coded input (i.e. dynamic sequences and static patterns) in order to create mnemonic associations. As all forms of STDP are known to selectively potentiate temporal correlations (see Sections 2.16 to 2.21), identifying which forms can also selectively potentiate firing rate correlations will allow the development of a single plasticity model that can process this dual code.

Hence, if STDP is to operate successfully within a discrete auto-associative network model of the hippocampus, then it must replicate the properties of previous, rate-coded models of Hebbian learning and LTP. This requires the potentiation of synapses which connect neurons that are concurrently firing at a higher than background rate. This has been identified as an area where many previous models of STDP might fail, as they produce either a decrease or no change in mean synaptic weights with increased pre-synaptic firing rates (see Sections 2.17 and 2.18). Simultaneously, it is preferable that STDP can generate hetero-synaptic depression at lower firing rates, so that synaptic weights are

maintained within a stable and responsive range, and redundant connections can be erased. The plasticity model must also operate stably, and create a stable weight distribution, so that network operation can proceed unsupervised, although some external mechanism (such as feedback inhibition) may contribute to this process. Each of these processes are effectively characterised by the BCM model (see Section 2.9), which allows the homeostatic control of potentiation and depression processes on a long term basis by modulation of the modification threshold.

4.4 Experimental Aims: The Network Model

As this research is directly related to the postulated function of CA3, the artificial neural network employed will be consistently inspired by neurobiological data from this brain region. Neurons in the CA3 region are known to fire both single spikes and short, high frequency bursts of action potentials *in vivo* (McNaughton, Barnes and O'Keefe, 1983; Debanne, Gahwiler and Thompson, 1998; Staba et al., 2002; Kobayashi and Poo, 2004). The spiking model used in this research is capable of replicating both firing patterns, and this will allow the impact that each has on network dynamics to be explored (Izhikevich, 2004). It is also known that firing rates in the CA3 region are generally lower than elsewhere in the cortex – varying between a background firing rate of ~1Hz and an active firing rate of 10 to 20Hz. This might help to stabilise the network by reducing recurrent excitation, and assist the process of pattern separation by restricting the capability of generating recall activity to very strong synapses (Rolls and Treves, 1998). Hence, attention will be focussed on relatively low firing rates during the simulations carried out in this research. The level of recurrent connectivity in CA3 is estimated to be between 2 and 20% *in vivo* (Rolls and Treves, 1998; Eichenbaum, 2004; Amaral and Lavenex, 2007; Rolls, 2008) and up to 33% in cultures of hippocampal neurons (Debanne, Gahwiler and Thompson, 1998). In the simulations presented here, however, a wider variety of connectivity will be examined, both to assess any impact that this may have on network dynamics, and because fully connected networks have most frequently been employed in previous auto-associative modelling studies.

Each neuron in the network model is assigned an axonal delay, randomly chosen from a range which is representative of the CA3 region. This seems a prescient approach to examining a plasticity model which places a particular emphasis on temporal coding. The presence of strong inhibitory input is a characteristic of the CA3 region which may also play a critical role in synaptic plasticity. A subset of inhibitory neurons will therefore be incorporated into these network simulations, comprising 20% of the total number of neurons therein. This figure is in line with both previous modelling studies and estimates of the ratio of GABA-ergic to pyramidal neurons in the CA3 region (Izhikevich, 2006). The inhibitory neurons in the network model examined here are governed by fast spiking dynamics, which are believed to correspond most closely to those observed in biological studies (Izhikevich, 2004; Izhikevich and Edelman, 2008). In these simulations, only constant mean firing rates of inhibitory neurons will be examined, although feedback inhibition is more likely to operate in the hippocampus *in vivo*, and might provide a more efficient method of maintaining stable neural and synaptic dynamics in a recurrent neural network. Feedback inhibition has also been commonly employed in previous auto-

associative modelling of the hippocampus (i.e. Rolls, Stringer and Trappenberg. 2002). Similarly, the plasticity of inhibitory synapses will not be examined, although such a process is known to occur in the hippocampus (see Kullman and Lamsa, 2007 for a review).

4.5 Experimental Aims: The Plasticity Model

While the most often studied synapses in the hippocampus have been those of the Schaffer collaterals which terminate on CA1 pyramidal cells, there has been sufficient research into the properties of the associative/commissural (A/C) recurrent synapses of the CA3 region to largely characterise their plasticity. Where there is an absence of conclusive biological data, each plausible possibility will be examined, with the hope that emergent dynamics may be matched with those observed empirically. Synaptic plasticity of the A/C synapses is known to be NMDAr-dependent, requiring both post-synaptic depolarisation and pre-synaptic glutamate release (Debanne, Gahwiler and Thompson, 1998; Kobayashi and Poo, 2004; see Sections 2.4 and 2.15). It is well known that this form of plasticity is effectively dictated by elevations in the level of intracellular Calcium concentration. A common approach to replicating the various non-linearities which relate to STDP in real networks with realistic input is to directly simulate intracellular Calcium dynamics (see Section 2.19). Conversely, the aim of this research is to identify which, if any, of the more abstract STDP implementations can be reconciled with these properties, without resort to a more complex process of modelling glutamate receptor kinetics, intracellular Calcium stores and/or neurotransmitter release. This phenomenological plasticity model can be implemented without significant computational cost, while being able to replicate the key emergent properties of synaptic plasticity observed in the CA3 region. Similarly, rather than focusing this investigation on the more complex mechanisms of STDP, such as its dependence on dendritic location (see Section 2.15), this research will concentrate on how the plasticity rule might mediate the functional role required of it within an auto-associative network model.

4.5.1 Weight Dependency

Both additive and multiplicative STDP implementations will be examined, to allow a comparison of the network dynamics produced with and without weight dependent plasticity. Although there is now a substantial body of evidence to suggest that the degree of synaptic potentiation *in vivo* is dependent on initial strength, it is equally clear that the synaptic competition which arises from an additive rule is attractive from a functional standpoint (see Section 2.18). Indeed, the observation that NMDA receptor trafficking stabilises synaptic weights in the hippocampus (see Section 2.9) might suggest that additive STDP creates a more realistic simulation of long-term synaptic dynamics in this region. While multiplicative STDP generally produces a unimodal distribution of synaptic weights that closely resembles that found *in vivo*, a combination of additive STDP and a Gaussian distribution of weight limits could also account for this observation, and the saturation of synaptic potentiation processes could explain the reduced potentiation of strong synapses. Although a number of multiplicative STDP rules have previously been examined (see Section 2.18), it appears that little functional difference exists between them, and hence only a single implementation, with an inverse exponential dependence of potentiation on synaptic weight, will be employed in this research. Previous studies have suggested

that some form of interpolation between additive and multiplicative STDP rules might offer attractive functionality, and while this remains an intriguing possibility, it will not be examined in this research.

4.5.2 Size of Potentiation and Depression Windows

It is well documented that the exact profile of the asymmetric learning window which characterises STDP can have a profound effect on emergent behaviour (see Section 2.17). Empirical data suggests that the time window of depression is significantly longer than that of potentiation in the CA3 region - 75ms and 25ms respectively, according to the investigations of Debanne, Gahwiler and Thompson (1998). However, several different ratios of STDP time constants will be tested in this research, in order to more fully characterise the emergent behaviour of the model. Similarly, several different ratios of the maximum amount of potentiation or depression which can be induced by a single spike pair (i.e. A_+ and A_- in Figure 2.4) will be examined. It is more difficult to estimate these figures from empirical data, particularly as values are usually expressed as relative changes following repeated pairings and, as discussed previously, initial weight may have some effect on the degree of change. However, the majority of evidence suggests that both $A_+ > A_-$ and $\tau_- > \tau_+$ in the hippocampus *in vivo* (Bi and Poo, 1998; Debanne, Gahwiler and Thompson, 1998; see Dan and Poo, 2006 for a review of other data). For the data of Bi and Poo (1998), which was obtained from dissociated hippocampal cultures, the parameters which define the STDP window were estimated as $A_+=0.86$; $A_-=-0.25$; $\tau_+=19$ms and $\tau_-=34$ms – these representing relative changes following 60 pairings (Bi, 2002). Hence, it is clear that the absolute magnitude of change in synaptic strength for a single spike pair is very small - in fact, it has often been speculated that single spike pairs cannot induce plasticity, as statistically significant weight change cannot be measured. It is possible that only the repeated stimulation of Calcium triggered pathways is sufficient to induce significant L-LTP. However, previous computational modelling of STDP has demonstrated that the relative values of the parameters which define the asymmetric learning window affect qualitative behaviour, while the scale of the values only affects the speed with which stable or asymptotic regimes are achieved (i.e. Burkitt, Meffin and Grayden, 2004).

What is also known, from some of the studies that originally characterised the spike-timing dependency of synaptic plasticity in the hippocampus, is that synchronous pre- and post- synaptic spiking should lead to almost maximal levels of depression (Debanne, Gahwiler and Thompson, 1998; see Section 2.15). This observation is incorporated into the plasticity model employed in this research at all times. A separate, but related, consideration regards the maximum overall change that a synapse can undergo before it becomes saturated. Data from CA3 cultures and slices indicate that potentiation can increase the strength of a typical synapse by as much as ~1000% (Debanne, Gahwiler and Thompson, 1998). This figure is used to determine the absolute values of A_+ and A_- in relation to the maximum possible weight of a synapse (i.e. $A\pm$ should be approximately 10% of the maximum weight limit). Consequently, a modelling choice must also be made regarding the maximum value that synaptic weights can achieve. This is likely to have a direct effect on the level of recurrent excitation (and therefore spike triggering) that exists in the network. Again, it is difficult to estimate a value for this parameter from biological data, and so several will be examined. However, it is clear from empirical

studies that single recurrent synapses of the CA3 region are incapable of solely producing post-synaptic firing, and so this constraint will be observed in all simulations presented (Kobayashi and Poo, 2004).

It is also important to note that in all simulations with additive STDP, the value of the maximum weight limit is fixed and uniform across all excitatory synapses. While there is little direct biological data available to inform this modelling decision, it seems likely that differences in the size of pre-synaptic boutons or concentration of post-synaptic receptors might endow different synaptic contacts with a different range of possible strengths. This feature is omitted here to avoid unnecessary computational complexity, although it is likely to represent an important direction for future research.

More generally, in each simulation examined in this research, the form of STDP is assumed to be fixed, while some empirical data suggests that the profile of the asymmetric learning window is malleable and directly affected by spike pairing frequency, for example (e.g. Sjostrom, Turrigiano and Nelson, 2001). By examining fixed STDP implementations across a range of mean firing rates, it is hoped that a form of the plasticity rule which can mediate rate-coded Hebbian learning without resort to dynamic modulation can be identified. This will subsequently ease the computational burden of implementing this plasticity rule in future auto-associative network models. Furthermore, using fixed forms of STDP should allow those particular features of the plasticity rule which directly influence emergent dynamics to be more easily identified, thus providing a broader insight into how future models of STDP might be dynamically manipulated to produce desired behaviour.

4.5.3 Spike Pair Restrictions

The relative paucity of empirical data regarding STDP, particularly within the CA3 region, means that there are a number of features for which there is little or no biological evidence available to direct modelling decisions. The contribution of different spike pairings to overall plasticity at a synapse, for example, has yet to be characterised in the hippocampus *in vivo* or *in vitro* (see Section 2.20). Relevant data has been obtained from the visual cortex, but differences between the basic properties of synaptic plasticity in these two distinct cortical regions suggests that generalisation is inappropriate (Froemke and Dan, 2002; Wang et al., 2005; Froemke et al., 2006). There is some inferential evidence to support a temporally restricted spike pairing implementation (i.e. one in which the state of pre- and post-synaptic neurons is effectively reset by each action potential), as previous computational modelling studies have demonstrated that this form of STDP best replicates data obtained from biological studies of both the visual cortex and hippocampus (Sjostrom, Turrigiano and Nelson, 2001; Izhikevich and Desai, 2003; Burkitt, Meffin and Grayden, 2004). In order to provide a more comprehensive assessment of the dynamics incurred by STDP, however, a variety of spike pair restriction schemes will be examined in this research, as these represent biologically plausible adaptations of STDP which might have a drastic effect on emergent dynamics.

One further issue that requires consideration is the observation that prolonged periods of pre-synaptic activity, with no post-synaptic firing, produces robust LTD in cultures of CA3 pyramidal cells. In

particular, the delivery of pre-synaptic spiking at a rate of 3Hz for a period of 3 minutes has been shown to depress synapses to ~64% of their initial amplitude (Debanne, Gahwiler and Thompson, 1998; Unni et al., 2004). This form of plasticity is also generally incorporated into auto-associative network models, whereby synapses are weakened due to the negative Hebbian correlation between pre- and post- synaptic firing rates. Under standard models of STDP, however, this depression cannot possibly be produced in the absence of post-synaptic spiking, although it is known that pyramidal cells in CA3 display spontaneous activity at a rate of ~1Hz, which suggests that a period of three minutes without activity might be biologically unrealistic (Rolls and Treves, 1998; Rolls, 2008). While no explicit attempt is made to model this depression, generated by pre- but not post- synaptic activity, each STDP implementation examined will be assessed for its ability to replicate this feature.

4.5.4 Synaptic Scaling and Redistribution

There is no evidence for a process of synaptic redistribution taking place in the hippocampus, and so this homeostatic mechanism will not be examined (see Section 2.11). There is, however, evidence for a process of synaptic scaling, although this is most frequently obtained by inducing unrealistic situations such as total activity blockade or extremely high firing rates across a network (see Section 2.10). The adjustment of synaptic weights is a long-term process which is suspected to retain the weight distribution developed by Hebbian learning mechanisms. Because the focus of this research is on the rapid process of storing externally applied activity patterns via synaptic plasticity, slower processes that regulate the overall activity levels within the network in the long-term are of less concern. Previous studies have also demonstrated that a process of synaptic scaling does not affect the qualitative properties of STDP, only the time taken to reach an asymptotic state and the quantitative aspects of weight distributions (van Rossum, Bi and Turrigiano, 2000; Burkitt, Meffin and Grayden, 2004). It seems unlikely that synaptic scaling acts to regulate levels of neural activity in recurrent networks, and it has often been suggested that feedback inhibition, which is rapid, flexible and affects both neural activity and synaptic plasticity, represents the main stability mechanism operating in the CA3 region (see Section 2.12). While synaptic scaling may be an important feature to integrate into the model in the future, therefore, this research will aim to exploit the natural competition incurred by STDP and the effects of inhibitory input to maintain the homeostatic control of neural firing rates and synaptic weights in an operational regime.

4.5.5 Metaplasticity

Rather than adjusting synaptic weights directly, evidence also suggests that changes in the ability to induce synaptic plasticity are incurred by the history of neural activity. This process of metaplasticity is best characterised by the BCM formulation, and previous analytical studies have demonstrated that the parameters of STDP might be directly reconciled with the position of the sliding modification threshold (i.e. Izhikevich and Desai, 2003; see Section 2.19). Although this process, like synaptic scaling, is more directly related to the long-term development of synaptic weights, and unlikely to mediate short term network stability, the fact that it explicitly affects the induction of plasticity makes it more relevant to this research. Metaplasticity processes might also serve to stabilise or destabilise synaptic weight

distributions, and thereby play a more explicit role in declarative memory processing than merely restricting synaptic weight values to a certain range via slow modulation. Hence, while the modulation of a BCM-type modification threshold will not be implemented in this research, possible methods by which this form of metaplasticity might be introduced will be discussed.

4.5.6 The Timescale of Synaptic Dynamics

There are a number of key assumptions made in this research. One major simplification that is inherent in the network model employed here is that plasticity processes occur instantaneously – that is, changes in synaptic weight are incurred at the point of pre- or post- synaptic spiking. While it takes a few minutes for potentiation or depression to be manifested in real neural networks, this simplification is in line with the vast majority of previous computational modelling studies. In effect, only the early stages of plasticity (i.e. E-LTP, see Section 2.4) are being investigated. These changes are (relatively) rapidly induced, but unstable, as they can be easily overwritten by subsequent activity patterns. It is worth bearing in mind that, should the potentiation or depression of a synapse persist for a longer period, then it will be rendered less vulnerable to ongoing neural activity by the consolidation process of L-LTP. While these distinctions between different stages of long-term plasticity and their relative stability will not be explicitly examined, the transition from E-LTP to L-LTP represents a mechanism by which all mnemonic traces – no matter how unstable – might be consolidated only a relatively short period (~30-45 minutes) after the original experience which induced them.

A related point regards the synapse model employed in this research, which transfers current from pre- to post- synaptic terminals immediately at the point of afferent firing (see Section 5.2). While this method omits the stochastic and dynamic aspects of synaptic transmission observed in real neural networks, it is in line with previous computational modelling studies of numerous cortical regions (such as pulse coupled networks - see Johnson, Padgett and Ovidmar, 1999). Furthermore, the aim of this research is to create an abstract model of auto-associative learning using STDP, and in order to limit computational complexity and avoid unnecessary complications, this assumption seems justified.

4.6 Network Input

The form of external input applied to the network during simulations forms the distinction between the different stages of testing in this research. Initially, independent stochastic spike trains with uniform mean firing rates will be provoked in the network by the consistent application of random levels of excitatory current. This is in line with several previous studies of STDP (i.e. van Rossum, Bi and Turrigiano, 2000; Song, Miller and Abbott, 2000) and should allow the properties of the plasticity model within a recurrent network to be characterised. However, it is important to note that the presence of iterant excitation and inhibitory input means that spiking activity within the network may not always adhere to a Gaussian distribution of inter-spike intervals, although stochastic spike timings are generated in a single neuron by the form of input used (see Section 5.2.2). Critically, spike triggering effects can become prominent – although the weight of individual synapses is constrained such that they are unable to solely provoke post-synaptic firing (see Section 4.5.2) – and persistent correlations

between afferent and efferent spike timings may therefore arise. This methodology aims to compromise between a desire to characterise the dynamics incurred by STDP with independent stochastic pre- and post- synaptic spike trains, thus allowing comparisons with the BCM rule to be drawn, while also generating results that can inform auto-associative network modelling (which relies on spike triggering at recurrent synapses for the generation of recall activity) and be reconciled with biological data (where synaptic currents inherently affect the state of their target neurons).

Secondly, tests which aim to more specifically relate the emergent properties of each STDP implementation to rate-coded Hebbian learning are carried out. In order to do this, the firing rates in a subset of the excitatory neurons in the network are elevated to ~20Hz, which is in line with firing rates observed in the CA3 region, while the remainder of excitatory neurons fire at a realistic background rate (~1Hz). The dichotomy in mean firing rates that exists among neurons in these simulations forms the critical distinction with the initial set of simulations. By examining the synaptic dynamics of various connections within the network, it should be possible to establish whether processes of selective potentiation, hetero-synaptic depression (particularly of synapses with high pre- and low post-synaptic firing rates) and the generation of strong bi-directional connections can be replicated by STDP.

Finally, those forms of STDP which best provide the necessary emergent features will be implemented within a simple auto-associative network, inspired by previous modelling studies of the hippocampus and declarative memory function. Several sparse patterns of external input which elevate the mean firing rate in foreground neurons to ~20Hz will be repeatedly applied to the network, while background activity proceeds at a mean firing rate of ~1Hz in all other neurons. Following this learning period, partial cues based on these previously learned patterns will be used to assess whether concise and significant recall activity can be generated in the network via a process of pattern completion. At this stage, only discrete, binary and orthogonal activity patterns will be examined, because these are more easily stored and recalled, and because models which employ this form of input have most often implemented synaptic plasticity in previous auto-associative network modelling studies. The synaptic and neural dynamics present in the network during these simulations will also be examined, and a comparison made with the properties of both previous auto-associative network models and experimental data from the CA3 region that have been described in this and the preceding chapters.

Importantly, during each of these simulations, synaptic plasticity and recurrent excitation are continually active, while in the majority of previous auto-associative modelling studies they are arbitrarily separated. This assumption aims in inspired by biological realism, as it seems unlikely that processes of learning and recall are behaviourally distinct or mutually exclusive. However, it is important to note that there may be several biological realistic methods for modulating neural and synaptic dynamics in order to dynamically separate these processes (e.g. Hasselmo, 2006), and this may represent a critical direction for future research.

4.7 Conclusions

In the previous chapter, it was demonstrated that auto-associative networks can successfully replicate various aspects of the postulated function of the hippocampus in declarative memory. However, the discussion presented in this chapter also makes it clear that several issues regarding the biological realism of these models exist. Primary to the interests of this work is the fact that STDP has rarely been examined within a recurrent neural network and no known attempts have been made to reconcile this temporally asymmetric learning rule with emergent auto-associative function. Because discrete auto-associative network function relies so critically on the properties of rate-coded Hebbian learning rules, these models also represent an ideal framework in which to investigate how STDP might mediate synaptic plasticity based on correlations between pre- and post- synaptic firing rates. All forms of STDP are ideally suited to selectively process temporally correlated data, and so the identification of an STDP implementation which can also process rate-coded data should increase the functional versatility of this plasticity rule. The discussion presented in this chapter suggests that the STDP model might also improve the biological realism of neural firing within a recurrent ANN. By constraining modelling decisions based on an *a priori* knowledge of the neural and synaptic dynamics that exist in the CA3 region, an auto-associative network model of declarative memory function with a firmer basis in modern neurobiology might consequently be developed. Furthermore, this novel auto-associative network model should exhibit significantly improved flexibility and computational power, and might have many applications beyond those required for modelling declarative memory and postulated hippocampal function.

In the following chapter, a more specific description of the synaptic and neural dynamics implemented in these simulations will be presented. In the subsequent chapters, simulations using this model will proceed in three main stages. Firstly, in chapter six, numerous common STDP implementations will be examined using uniform levels of Poissonian activity in order to provide a comprehensive description of the emergent dynamics of the plasticity rule within a spiking, recurrent neural network. Secondly, in chapter seven, the nature of external input will be adjusted in order to more specifically assess the ability of these STDP implementations to process firing rate correlations in a manner analogous with previous Hebbian learning rules. Finally, in chapter eight, an examination of any STDP implementations which can mediate rate-coded processing and replicate other known features of synaptic plasticity in the CA3 region will be employed within a simple auto-associative network model, with the aim of providing further insight into the functional performance of STDP within this context. In chapter nine, a discussion of the main results of this research, the novel contributions to the study of computational neuroscience that it has presented, and the possible directions of future research will be described.

BIBLIOGRAPHY

Abarbanel HD, Gibb L, Huerta R, Rabinovich MI. Biophysical model of synaptic plasticity dynamics. *Biological Cybernetics* 89(3): 214-26 (2003)

Abbott LF, Nelson SB. Synaptic plasticity: taming the beast. *Nature Neuroscience* 3: 1178-83 (2000)

Abbott LF, Sen K, Varela JA & Nelson SB. Synaptic depression and cortical gain control. *Science* 275:220-222 (1997).

Abbott LF. Balancing homeostasis and learning in neural circuits. *Zoology* 106:365-371 (2003).

Abel T et al. Genetic demonstration of a role for PKA in the late phase of LTP and in hippocampus-based long-term memory. *Cell* 88:615-626 (1997).

Abraham WC & Williams JM. Properties and mechanisms of LTP maintenance. *Neuroscientist* 9:463-474 (2003).

Abraham WC, Bear MF. Metaplasticity: the plasticity of synaptic plasticity. *Trends in Neuroscience* 19(4): 126-30 (2004)

Abraham WC, Logan B, Greenwood JM, Dragunow M. Induction and experience-dependent consolidation of stable long-term potentiation lasting months in the hippocampus. *Journal of Neuroscience* 22(21): 9626-34 (2002)

Abraham WC, Mason-Parker SE, Bear MF, Webb S & Tate WP. Heterosynaptic metaplasticity in the hippocampus in vivo: A BCM-like modifiable threshold for LTP. *PNAS* 98:10924-10929 (2001).

Abraham WC, Tate WP. Metaplasticity: a new vista across the field of synaptic plasticity. *Progress in Neurobiology* 52(4): 303-23 (1997)

Abraham WC. Metaplasticity: tuning synapses and networks for plasticity. *Nature Reviews Neuroscience* 9(5): 387 – 399 (2008)

Aggleton JP, Brown MW. Interleaving brain systems for episodic and recognition memory. *Trends in Cognitive Science* 10(10): 455-63 (2006)

Alyan S & McNaughton BL. Hippocampectomized rats are capable of homing by path integration. *Behavioral Neuroscience* 113(1):19-31(1999).

Amaral DG & Witter MP. The Hippocampal formation. Paxinos G, editor, *The Rat Nervous System*. Academic Press, San Diego (1995).

Amit DJ. The Hebbian paradigm reintegrated: local reverberations as internal representations. *Behavioural and Brain Sciences* 18:617-657 (1995).

Andersen P, Morris R, Amaral D, Bliss T & O'Keefe J. Historical Perspective: Proposed functions, biological characteristics, and neurobiological models of the Hippocampus. Andersen P, Morris R, Amaral D, Bliss T & O'Keefe J, editors, *The Hippocampus Book.* Oxford University Press, Oxford (2007b).

Andersen P, Morris R, Amaral D, Bliss T & O'Keefe J. The Hippocampal Formation. Andersen P, Morris R, Amaral D, Bliss T & O'Keefe J, editors, *The Hippocampus Book.* Oxford University Press, Oxford (2007a).

Arancio O, Lev Ram V, Tsien RY, Kandel ER & Hawkins RD. Nitric oxide acts as a retrograde messenger during long-term potentiation in cultured hippocampal neurons. *Journal of Physiology (Paris)* 90(5-6): 321-2 (1996).

Artola A. Diabetes-, stress- and ageing-related changes in synaptic plasticity in hippocampus and neocortex--the same metaplastic process? *European Journal of Pharmacology* 585(1): 153-62 (2008)

Bakker A, Kirwan CB, Miller M & Stark CEL. Pattern Separation in the Human Hippocampal CA3 and Dentate Gyrus. *Science* 319:1640-1642 (2008).

Banke TG, Bowie D, Lee H, Huganir RL, Schousboe A, Traynelis SF. Control of GluR1 AMPA receptor function by cAMP-dependent protein kinase. *Journal of Neuroscience* 20(1): 89-102 (2000)

Bannerman DM, Good MA, Butcher SP, Ramsay M & Morris RG. Distinct components of spatial learning revealed by prior training and NMDA receptor blockade. *Nature* 378:182-186 (1995).

Baras D & Meir R. Reinforcement Learning, Spike-Time-Dependent Plasticity, and the BCM Rule. *Neural Computation* 19:2245-2279 (2007).

Barnes CA, Suster MS, Shen J & McNaughton BL. Multistability of cognitive maps in the hippocampus of old rats. *Nature* 388:272-275 (1997).

Barria A, Malinow R. Subunit-specific NMDA receptor trafficking to synapses. *Neuron* 35(2): 345-53 (2002)

Bartha L, Brenneis C, Schocke M, Trinka E, Koylu B, Trieb T, Kresmer C, Jaschke W, Bauer G, Poewe W & Benke T. Medial temporal lobe activation during semantic language processing: fMRI findings in healthy left- and right-handers. *Cognitive Brain Research* 17(2): 339–46 (2003).

Bartlett FC. *Remembering: A Study in Experimental and Social Psychology*. Cambridge University Press, Cambridge, UK (1932).

Bayley PJ, O'Reilly RC, Curran T & Squire LR. New semantic learning in patients with large medial temporal lobe lesions. *Hippocampus* 18(6):575-583 (2008).

Bear MF & Kirkwood A. Bidirectional plasticity of cortical synapses. Fazeli MS & Collingridge GL, editors, *In Cortical Plasticity: LTP and LTD*, pp. 191-205. Oxford Bios Scientific, Oxford (1996).

Bear MF and Malenka RC. Synaptic plasticity: LTP and LTD. *Current Opinion in Neurobiology* 4(3): 389-99 (1994)

Bear MF. Bidirectional synaptic plasticity: from theory to reality. *Philosophical Transactions of the Royal Society London B* 358(1432): 649-55 (2003)

Bekkers JM, Richerson GB, Stevens CF. Origin of variability in quantal size in cultured hippocampal neurons and hippocampal slices. *Proceedings of the National Academy of Sciences* 7(14): 5359-62 (1990)

Benke TA, Lüthi A, Isaac JT and Collingridge GL. Modulation of AMPA receptor unitary conductance by synaptic activity. *Nature* 393(6687): 793-7 (1998)

Benuskova L & Kasabov N. Modeling L-LTP based on changes in concentration of pCREB transcription factor. *Neurocomputing* 70(10-12):2035-2040 (2007).

Benuskova L, Abraham WC. STDP rule endowed with the BCM sliding threshold accounts for hippocampal heterosynaptic plasticity. *Journal of Computational Neuroscience* 22(2): 129-33 (2007)

Bi GQ, Poo MM. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience* 18(24): 10464-72 (1998)

Bienenstock EL, Cooper LN, Munro PW. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience* 2(1): 32-48 (1982)

Bird CM, Burgess N. The hippocampus and memory: insights from spatial processing. *Nature Reviews Neuroscience* 9(3): 182-94 (2008)

Blais B, Cooper LN & Shouval H. Formation of direction selectivity in natural scene environments. *Proceedings of the National Academy of Sciences* 12(5) (2000).

Bliss TVP & Collingridge GL. A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361:31-39 (1993).

Bliss TVP, Collingridge GL & Morris RM. Andersen P, Morris R, Amaral D, Bliss T & O'Keefe J, editors, *The Hippocampus Book* Oxford University Press, Oxford (2007).

Bourtchouladze R. *Memories are made of this – the biological building blocks of memory*. Phoenix, London (2002).

Brickley SG, Revilla V, Cull-Candy SG, Wisden W & Farrant M. Adaptive regulation of neuronal excitability by a voltage independent potassium conductance. *Nature* 409:88-92 (2001).

Brown MW, Aggleton JP. Recognition memory: what are the roles of the perirhinal cortex and hippocampus? *Nature Reviews Neuroscience* 2(1): 51-61 (2001)

Brown S & Schaffer EA. An investigation into the function of the occipital and temporal lobes of the monkey's brain. *Philosophical transactions of the Royal Society B* 179:303-327 (1888).

Buchs NJ & Senn W. Spike-based synaptic plasticity and the emergence of direction selective simple cells: simulation results. *Journal of Computational Neuroscience* 13(3):167-86 (2002).

Bunsey M & Eichenbaum H. Selective damage to the hippocampal region blocks long-term retention of a natural and non-spatial stimulus-stimulus association. *Hippocampus* 5:546-556 (1995).

Buonomano DV. Distinct functional types of associative long-term potentiation in neocortical and hippocampal pyramidal neurons. *Journal of Neuroscience* 19(16): 6748-54 (1999)

Burgess N, Becker S, King J & O'Keefe J. Memory for events and their spatial context: models and experiments. *Philosophical transactions of the Royal Society B* 356:1-11 (2001).

Burgess N, Maguire EA & O'Keefe J. The human hippocampus and spatial and episodic memory. *Neuron* 35:625-641 (2002).

Burkitt AN, Meffin H, Grayden DB. Spike-timing-dependent plasticity: the relationship to rate-based learning for models with weight dynamics determined by a stable fixed point. *Neural Computation* 16(5): 885-940 (2004)

Burkitt AN, van Hemmen JL. How synapses in the auditory system wax and wane: theoretical perspectives. *Biological Cybernetics* 89(5): 318-32 (2003)

Burrone J, O'Byrne M, Murthy VN. Multiple forms of synaptic plasticity triggered by selective suppression of activity in individual neurons. *Nature* 420(6914): 414-8 (2002)

Cabeza R & St Jacques P. Functional neuroimaging of autobiographical memory. *Trends in Cognitive Science* 11:219-227 (2007).

Cain DP. Testing the NMDA, long-term potentiation, and cholinergic hypotheses of spatial learning. *Neuroscience and Biobehavioural Reviews* 22(2):181–193 (1998).

Castellani GC, Quinlan EM, Cooper LN, Shouval HZ. A biophysical model of bidirectional synaptic plasticity: dependence on AMPA and NMDA receptors. *Proceedings of the National Academy of Sciences* 98(22): 12772-7 (2001)

Chagnac-Amitai Y & Connors BW. Horizontal spread of synchronized activity in neocortex and its control by GABA-mediated inhibition. *Journal of Neurophysiology* 61:747-758 (1989).

Cheyne JE, Montgomery JM. Plasticity-dependent changes in metabotropic glutamate receptor expression at excitatory hippocampal synapses. *Molecular and Cellular Neurosciences* 37(3): 432-9 (2008)

Chiba AA, Kesner RP & Reynolds AM. Memory for spatial location as a function of temporal lag in rats: Role of hippocampus and medial prefrontal cortex. *Behavioral and Neural Biology* 61(2):123-131 (1994).

Churchland P & Sejnowski T. *The Computational Brain*. MIT Press, Cambridge, Massachussetts (1992).

Clark RE & Squire LR. Classical Conditioning and Brain Systems: The Role of Awareness. *Science* 280:77-81 (1998).

Cohen AS, Coussens CM, Raymond CR & Abraham WC. Long-lasting increase in cellular excitability associated with the priming of LTP induction in rat hippocampus. *Journal of Neurophysiology* 82:3139-3148 (1999).

Cohen NJ & Squire LR. Preserved learning and retention of pattern-analysing skill in amnesia: dissociation of knowing how and knowing that. *Science* 210:207-210 (1980).

Cooke SF & Bliss TVP. Plasticity in the human central nervous system. *Brain* 129:1659-1673 (2006).

Cooke SF et al. Long-term potentiation at perforant path-granule cell synapses does not require phosphorylation of α-CaMKII at threonine 286. *Forum European Neuroscience Abstracts* 4.AO15.13. (2004).

Cooper LN, Intrator N, Blais BS & Shouval HZ. *Theory of cortical plasticity*. World Scientific, New Jersey (2004).

Corkin S. What's new with the amnesic patient H.M.? *Nature Reviews Neuroscience* 3:153-160 (2002).

Crair MC & Malenka RC. A critical period for long-term potentiation at thalamocortical synapses. *Nature* 375:325-328 (1995).

Crane J & Milner B. What went where? Impaired object location in patients with right hippocampal lesions. *Hippocampus* 15:216-231 (2005).

Cravens CJ, Vargas-Pinto N, Christian KM & Nakazawa K. CA3 NMDA receptors are crucial for rapid and automatic representation of context memory. *European Journal of Neuroscience* 24(6):1771-1780(10) (2006).

Crozier RA, Philpot BD, Sawtell NB & Bear MF. Long-term plasticity of glutamatergic synaptic transmission in the cerebral cortex. Gazzaniga MS, editor, *The Cognitive Neurosciences III*, pp.109-127. MIT Press, Cambridge, Massachusetts (2004).

Dan Y & Poo M-M. Spike Timing-Dependent Plasticity: From Synapse to Perception. *Physiological Reviews* 86:1033-1048 (2006).

Dan Y, Poo MM. Spike timing-dependent plasticity of neural circuits. *Neuron* 44(1): 23-30 (2004)

Davies RR, Halliday GM, Xuereb JH, Kril JJ & Hodges JR. The neural basis of semantic memory: Evidence from semantic dementia. *Neurobiology of Aging* (2008, in press).

Day M, Langston R & Morris RGM. Glutamate-receptor-mediated encoding and retrieval of paired-associate learning. *Nature* 424:205-209 (2003).

Debanne D, Gähwiler BH & Thompson SM. Cooperative interactions in the induction of long-term potentiation and depression of synaptic excitation between hippocampal CA3-CA1 cell pairs in vitro. *PNAS USA* 93(20):11225-11230 (1996).

Debanne D, Gähwiler BH, Thompson SM. Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *Journal of Physiology* 507: 237-247 (1998)

Delgado JY, O'Dell TJ. Long-term potentiation persists in an occult state following mGluR-dependent depotentiation. *Neuropharmacology* 48(7): 936-48 (2005)

Derkach VA, Oh MC, Guire ES, Soderling TR. Regulatory mechanisms of AMPA receptors in synaptic plasticity. *Nature Reviews Neuroscience* 2007 Feb;8(2): 101-13

Desai NS, Cudmore RH, S.B. Nelson SB & Turrigiano GG. Critical periods for experience-dependent synaptic scaling in visual cortex. *Nature Neuroscience* 5:783-789 (2002).

Desai NS, Rutherford LC & Turrigiano GG. Plasticity in the intrinsic excitability of cortical pyramidal neurons, *Nature Neuroscience* 2:515-520 (1999).

Desai NS. Homeostatic plasticity in the CNS: synaptic and intrinsic forms. *Journal of Physiology (Paris)* 97:391-402 (2003).

Di Paolo EA. Evolving spike-timing-dependent plasticity for single-trial learning in robots. *Philosophical Transactions of the Royal Society London A* 361(1811): 2299-319 (2003)

Dudek SM & Bear MF. Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade. *PNAS* 89:4363-4367 (1992).

Eacott MJ, Easton A & Zinkivskay A. Recollection in an episodic-like memory task in the rat. *Learning and Memory* 12(3):221-3 (2005).

Eichenbaum H & Cohen NJ. *From conditioning to conscious recollection: Memory systems of the brain*. Oxford University Press, Oxford (2001).

Eichenbaum H, Dudchenko P, Wood E, Shapiro M, Tanila H. The hippocampus, memory, and place cells: is it spatial memory or a memory space? *Neuron* 23(2): 209-26 (1999)

Eichenbaum H. An Information Processing Framework for Memory Representation by the Hippocampus. Gazzaniga MS, editor, *The Cognitive Neurosciences.* MIT Press, Cambridge, Massachusetts (2004).

Eichenbaum H. Is the rodent hippocampus just for "place"? *Current Opinion in Neurobiology* 6:187-195 (1996).

Eichenbaum H. The hippocampus: The shock of the new. *Current Biology* 9(13):482-484 (1999).

Ekstrom AD, Kahana MJ, Caplan JB, Fields TA, Isham EA, Newman EL & Fried I. Cellular networks underlying human spatial navigation. *Nature* 425:184-188 (2003).

Ekstrom AD, Kahana MJ, Caplan JB, Fields TA, Isham EA, Newman EL, Fried I. Cellular networks underlying human spatial navigation. *Nature* 425(6954): 184-8 (2003)

Elgersma Y et al. Inhibitory autophosphorylation of CaMKII controls PSD association, plasticity, and learning. *Neuron* 36:493-505 (2002).

Esteban JA et al. PKA phosphorylation of AMPA receptor subunits controls synaptic trafficking underlying plasticity. *Nature Neuroscience* 6:136-143 (2003).

Feldman DE. Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron* 27:45-56 (2000).

Fernández G, Tendolkar I. The rhinal cortex: 'gatekeeper' of the declarative memory system. *Trends in Cognitive Science* 10(8): 358-62 (2006)

Finnerty GT, Roberts LS, Connors BW. Sensory experience modifies the short-term dynamics of neocortical synapses. *Nature* 400(6742): 367-371 (1999)

Fischer W, Chen KS, Gage FH & BjoÈrklund A. Progressive decline in spatial learning and integrity of forebrain cholinergic neurons in rats during aging. *Neurobiology of Aging* 13:9-23 (1992).

Florian RV. A reinforcement learning algorithm for spiking neural networks. Zaharie D, Petcu D, Negru V, Jebelean T, Ciobanu G, Cicortas A, Abraham A & Paprzycki M, editors, *Proceedings of the Seventh International Symposium on Symbolic and Numeric Algorithms for Scientific Computing* (SYNASC 2005), pp. 299–306. IEEE Computer Society (2005).

Fragkouli A et al. Loss of forebrain cholinergic neurons and impairment in spatial learning and memory in LHX7-deficient mice. *European Journal of Neuroscience* 21:2923-38 (2005).

Freund TF & Buzsaki G. Interneurons of the hippocampus. *Hippocampus* 6:347-470 (1996).

Frey U & Morris RG. Synaptic tagging: implications for late maintenance of hippocampal long-term potentiation. *Trends in Neurosciences* 21:181-188 (1998).

Frey U, Krug M, Reymann KG & Matthies H. Anisomycin, an inhibitor of protein synthesis, blocks late phases of LTP phenomena in the hippocampal CA region in vitro. *Brain Research* 452:57-65 (1988).

Frey U, Morris RG. Synaptic tagging and long-term potentiation. *Nature* 385(6616): 533-6 (1997)

Frick A, Magee J & Johnston D. LTP is accompanied by an enhanced local excitability of pyramidal neuron dendrites. *Nature Neuroscience* 7:126-135 (2004).

Fried I, Cameron KA, Yashar S, Fong R & Morrow JW. Inhibitory and Excitatory responses of single neurons in the human medial temporal lobe during recognition of faces and objects. *Cerebral Cortex* 12:575-584 (2002).

Fried I, MacDonald KA & Wilson CL. Single neuron activity in human hippocampus and amygdala during recognition of faces and objects. *Neuron* 18:753–765 (1997).

Froemke RC, Dan Y. Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature* 416(6879): 433-8 (2002)

Froemke RC, Poo MM & Dan Y. Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature* 434(7030):221-5 (2005).

Fuhs MC, Touretzky DS. A spin glass model of path integration in rat medial entorhinal cortex. *Journal of Neuroscience* 26(16): 4266-76 (2006)

Fujii S et al. The long-term suppressive effect of prior activation of synaptic inputs by low-frequency stimulation on induction of long-term potentiation in CA1 neurons of guinea pig hippocampal slices. *Experimental Brain Research* 111:305-312 (1996).

Fujii S, Saito K, Miyakawa H, Ito K-I & Kato H. Reversal of long-term potentiation (depotentiation) induced by tetanus stimulation of the input to CA1 neurons of guinea pig hippocampal slices. *Brain Research* 555:112-122 (1991).

Fyhn M, Hafting T, Treves A, Moser MB, Moser EI. Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* 446(7132): 190-4 (2007)

Fyhn M, Molden S, Hollup S, Moser MB & Moser E. Hippocampal neurons responding to first-time dislocation of a target object. *Neuron* 35(3):555-566 (2002).

Gabrieli JDE, Brewer JB, Desmond JE & Glover GH. Separate neural bases of two fundamental memory processes in the human medial temporal lobe. *Science* 276: 264–266 (1997).

Gaffan D, Saunders RC, Gaffan EX, Harrison S, Shields C & Owen MJ. Effects of fornix transection upon associative memory in monkeys: role of the hippocampus in learned action. *Quarterly Journal of Experimental Psychology* 26B:173-221 (1984).

Gaffan D. Against memory systems. *Philosophical Transactions of the Royal Society B: Biological Sciences* 357:1111-1121 (2002)

Gaffan D. Scene-specific memory for objects: a model of episodic memory impairment in monkeys with fornix transection. *Journal of Cognitive Neuroscience* 6:305-320 (1994).

Gage FH, Dunnett SB & Bjorklund A. Age-related impairments in spatial memory are independent of those in sensorimotor skills. *Neurobiology of Aging* 10:347-352 (1989).

Gerkin RC, Lau PM, Nauen DW, Wang YT, Bi GQ. Modular competition driven by NMDA receptor subtypes in spike-timing-dependent plasticity. *Journal of Neurophysiology* 97(4): 2851-62 (2007)

Giese KP, Fedorov NB, Filipkowski RK & Silva AJ. Autophosphorylation at Thr286 of the alpha calcium-calmodulin kinase II in LTP and learning. *Science* 279:870-3 (1998).

Gilbert PE, Kesner RP & DeCoteau WE. The role of the hippocampus in mediating spatial pattern separation. *Journal of Neuroscience* 18:804-810 (1998).

Gilbert PE, Kesner RP & Lee I. Dissociating hippocampal sub-regions: a double dissociation between the dentate gyrus and CA1. *Hippocampus* 11:626-636 (2001).

Gilbert PE, Kesner RP. Localization of function within the dorsal hippocampus: the role of the CA3 subregion in paired-associate learning. *Behavioural Neuroscience* 117(6): 1385-94 (2003)

Gold AE, Kesner RP. The role of the CA3 subregion of the dorsal hippocampus in spatial pattern completion in the rat. *Hippocampus* 15(6): 808-14 (2005)

Golding NL, Staff NP & Spruston N. Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature* 418:326-331 (2002).

Goldman-Rakic PS. The pre-frontal landscape: implications of functional architecture for understanding human mentation and the central executive. *Philosophical transactions of the Royal Society of London, Series B* 351:1445-1453 (1996).

Gothard KM, Hoffman KL, Battaglia FP, McNaughton BL. Dentate gyrus and CA1 ensemble activity during spatial reference frame shifts in the presence and absence of visual input. *Journal of Neuroscience* 21(18): 7284-92 (2001)

Gothard KM, Skaggs WE & McNaughton BL. Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. *Journal of Neuroscience*. 16:8027-8040 (1996b).

Gothard KM, Skaggs WE, Moore KM & McNaughton BL. (1996a). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *Journal of Neuroscience* 16:823-835 (1996a).

Graupner M & Brunel N. STDP in a bistable synapse model based on CaMKII and associated signaling pathways. *PLoS Computational Biology* 3(11): e221, 2299-2323 (2007).

Griffiths D, Dickinson A & Clayton N. Episodic memory: what can animals remember about their past? *Trends in Cognitive Sciences* 3(2):74-80 (1999).

Grosshans DR, Clayton DA, Coultrap SJ & Browning MD. LTP leads to rapid surface expression of NMDA but not AMPA receptors in adult rat CA1. *Nature Neuroscience* 5:27-33 (2002).

Guzowski JF et al. Inhibition of activity-dependent arc protein expression in the rat hippocampus impairs the maintenance of long-term potentiation and the consolidation of long-term memory. *Journal of Neuroscience* 20:3993-4001 (2000).

Hampson RE, Rogers G, Lynch G & Deadwyler SA. Facilitative Effects of the Ampakine CX516 on Short-Term Memory in Rats: Enhancement of Delayed-Nonmatch-to-Sample Performance. *The Journal of Neuroscience* 18(7):2740-2747 (1998a).

Hampson RE, Rogers G, Lynch G & Deadwyler SA. Facilitative effects of the ampakine CX516 on short-term memory in rats: correlations with hippocampal neuronal activity. *The Journal of Neuroscience* 18:2748-2763 (1998b).

Hardingham GE, Fukunaga Y & Bading H. Extrasynaptic NMDARs oppose synaptic NMDARs by triggering CREB shut-off and cell death pathways. *Nature Neuroscience* 5:405-414 (2002).

Hartley T, Maguire EA, Spiers HJ & Burgess N. The well worn route and the path less travelled: distinct neural bases of route following and wayfinding. *Neuron* 37:877-888 (2003).

Harvey CD & Svoboda K. Locally dynamic synaptic learning rules in pyramidal neuron dendrites. *Nature* 450:1195-1200 (2007).

Hassabis D, Kumaran D, Vann SD & Maguire EA. Patients with hippocampal amnesia cannot imagine new experiences. *Proceedings of the National Academy of Sciences of the United States of America* 104:1726-1731 (2007).

Hasselmo ME, Eichenbaum H. Hippocampal mechanisms for the context-dependent retrieval of episodes. *Neural Networks* 18(9): 1172-90 (2005)

Hasselmo ME.The role of acetylcholine in learning and memory. *Current Opinion in Neurobiology* 16(6): 710-5 (2006)

Hawkins RD, Son H & Arancio O. Nitric oxide as a retrograde messenger during long-term potentiation in hippocampus. *Progress in Brain Research* 118:155-72 (1998).

Hayashi Y, Shi SH, Esteban JA, Piccini A, Poncer JC and Malinow R. Driving AMPA receptors into synapses by LTP and CaMKII: requirement for GluR1 and PDZ domain interaction. *Science* 287(5461): 2262-7 (2000)

Hebb DO. The Organization of Behavior: A Neuropsychological Theory. *Wiley Press, New York* (1949)

Hellier JL, Grosshans DR, Coultrap SJ, Jones JP, Dobelis P, Browning MD, Staley KJ. NMDA receptor trafficking at recurrent synapses stabilizes the state of the CA3 network. *Journal of Neurophysiology* 98(5): 2818-26 (2007)

Hensch TK et al. Local GABA circuit control of experience-dependent plasticity in developing visual cortex. *Science* 282:1504-1508 (1998).

Hetherington PA & Shapiro ML. Hippocampal place fields are altered by the removal of single visual cues in a distance dependent manner. *Behavioral Neuroscience* 111:20-34 (1997).

Heynen AJ, Quinlan EM, Bae DC, Bear MF. Bidirectional, activity-dependent regulation of glutamate receptors in the adult hippocampus in vivo. *Neuron* 28(2): 527-36 (2000)

Holland LL & Wagner JJ. Primed facilitation of homosynaptic long-term depression and depotentiation in rat hippocampus. *Journal of Neuroscience* 18:887-894 (1998).

Hölscher C. Nitric Oxide, the Enigmatic Neuronal Messenger: Its Role in Synaptic Plasticity. *Trends in Neuroscience* 20(7):298-303 (1997).

Hopfield JJ. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences* 79(8): 2554-8 (1982)

Hopfield JJ. Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376(6535): 33-6 (1995)

Hrabetova S, Serrano P, Blace N, Tse HW, Skifter DA, Jane DE, Monaghan DT, Sacktor TC. Distinct NMDA receptor subpopulations contribute to long-term potentiation and long-term depression induction. *Journal of Neuroscience* 20(12): RC81 (2000)

Huang YY & Kandel ER. Recruitment of long-lasting and protein kinase A-dependent long-term potentiation in the CA1 region of hippocampus requires repeated tetanization. *Learning and Memory* (1):74-82 (1994).

Huang Y-Y, Colino A, Selig DK & Malenka RC. The influence of prior synaptic activity on the induction of long-term potentiation. *Science* 255:730-733 (1992).

Huang YY, Nguyen PV, Abel T & Kandel ER. Long-lasting forms of synaptic potentiation in the mammalian hippocampus. *Learning and Memory* 3:74-85 (1996).

Huerta PT, Sun LD, Wilson MA & Tonegawa S. Formation of Temporal Memory Requires NMDA Receptors within CA1 Pyramidal Neurons. *Neuron* 25(2):473-480 (2000).

Hunsaker MR, Lee B, Kesner RP. Evaluating the temporal context of episodic memory: the role of CA3 and CA1. *Behavioural Brain Research* 188(2): 310-5 (2008)

Hupbach A, Gomez R, Hardt O, Nadel L. Reconsolidation of episodic memories: a subtle reminder triggers integration of new information. *Learning and Memory* 14(1-2): 47-53 (2007)

Irvine GI, Logan B, Eckert M & Abraham WC. Enriched environment exposure regulates excitability, synaptic transmission, and LTP in the dentate gyrus of freely moving rats. *Hipppocampus* 16(2):149-160 (2005).

Isaacs EB, Lucas A, Chong WK, Wood SJ, Johnson CL, Marshall C, Vargha-Khadem F & Gadian DG. Hippocampal volume and everyday memory in children of very low birth weight. *Paediatric Research* 47(6):713-20 (2000).

Izhikevich EM, Desai NS. Relating STDP to BCM. *Neural Computation* 15(7): 1511-23 (2003)

Izhikevich EM. Polychronization: computation with spikes. *Neural Computation* 18(2): 245-82 (2006)

Izhikevich EM. Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cerebral Cortex* 17(10): 2443-52 (2007)

Izumi Y, Clifford DB & Zorumski CF. Inhibition of long-term potentiation by NMDA-mediated nitric oxide release. *Science* 257:1273-1276 (1992).

James W. *Psychology: Briefer Course*. Harvard University Press, Cambridge, MA (1890).

Jeffery KJ and Morris RGM. Cumulative long-term potentiation in the rat dentate gyrus correlates with, but does not modify, performance in the watermaze. *Hippocampus* 3(2): 133-140 (1993)

Jensen V et al. A juvenile form of postsynaptic hippocampal long-term potentiation in mice deficient for the AMPA receptor subunit GluR. *The Journal of Physiology* 553:843-856 (2003).

Jerman T, Kesner RP, Hunsaker MR. Disconnection analysis of CA3 and DG in mediating encoding but not retrieval in a spatial maze learning task. *Learning and Memory* 13(4): 458-64 (2006)

Ji D & Wilson MA. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience* 10:100-107 (2007).

Johnson JL, Padgett ML, Omidvar O. Guest Editorial Overview of the Special Issue on Pulse Coupled Neural Networks. *IEEE Transactions on Neural Networks* 10 (3), 461 – 463 (1999)

Johnston D, Williams S, Jaffe D & Gray R. NMDA-Receptor-Independent Long-Term Potentiation. *Annual review of Physiology* 54:489-505 (1992).

Kampa BM, Clements J, Jonas P & Stuart GJ. Kinetics of Mg2+ unblock of NMDA receptors: implications for spike-timing dependent synaptic plasticity. *The Journal of Physiology* 556:337-345 (2004).

Kampa BM, Letzkus JJ, Stuart GJ. Dendritic mechanisms controlling spike-timing-dependent synaptic plasticity. *Trends in Neuroscience* 30(9): 456-63 (2007)

Kart-Teke E, De Souza Silva MA, Huston JP & Dere E. Wistar rats show episodic-like memory for unique experiences. *Neurobiology of learning and memory* 85(2):173-82 (2006).

Katz B. *Nerve, Muscle and Synapse*. McGraw-Hill (1966).

Kauderer BS & Kandel ER. Capture of a protein synthesis-dependent component of long-term depression. *PNAS USA* 97:13342-13347 (2000).

Kelleher RJ 3rd, Govindarajan A, Tonegawa S. Translational regulatory mechanisms in persistent forms of synaptic plasticity. *Neuron* 44(1):59-73 (2004)

Kelly PT. Calmodulin-dependent protein kinase II. Multifunctional roles in neuronal differentiation and synaptic plasticity. *Molecular Neurobiology* 5(2-4): 153-77 (1991)

Kempermann G, Kuhn G & Gage FH. More hippocampal neurons in adult mice living in an enriched environment. *Nature* 386:493-495 (1997).

Kempter R, Gerstner W, van Hemmen JL. Intrinsic stabilization of output rates by spike-based Hebbian learning. *Neural Computation* 13(12): 2709-41 (2001)

Kemptner RP, Gerstner W & van Hemmem JL. Spike-based compared to Rate-based Hebbian learning. *Advances in Neural Information Processing Systems 11*, pp. 125-131. MIT Press, Cambridge, Massachusetts (1999).

Kepecs A, van Rossum MC, Song S, Tegner J. Spike-timing-dependent plasticity: common themes and divergent vistas. *Biological Cybernetics* 87(5-6): 446-58 (2002)

Kesner RP & Rolls ET. Role of long-term synaptic modification in short-term memory. *Hippocampus* 11:240-250 (2001).

Kesner RP, Lee I, Gilbert P. A behavioral assessment of hippocampal function based on a subregional analysis. *Reviews in the Neurosciences* 15(5): 333-51 (2004)

Kesner RP. A behavioral analysis of dentate gyrus function. *Progressive Brain Research* 163:567-76 (2007).

Kesner RP. A behavioral analysis of dentate gyrus function. *Progress in Brain Research* 163: 567-76 (2007)

Kesner RP. Behavioral functions of the CA3 subregion of the hippocampus. *Learning and Memory* 14(11): 771-81 (2007)

Kesner RP. Neurobiological views of memory. Martinez J & Kesner RP, editors, *The Neurobiology of Learning and Memory*, pp 361-416. Academic Press, New York (1998).

Kilman V, van Rossum MC & Turrigiano GG. Activity deprivation reduces miniature IPSC amplitude by decreasing the number of postsynaptic GABAA receptors clustered at neocortical synapses. *Journal of Neuroscience* 22:1328-1337 (2002).

Kim JJ & Fanselow MS. Modality-specific retrograde amnesia of fear. *Science* 256(5057):675-677 (1992).

Kirkwood A et al. Co-regulation of long-term potentiation and experience-dependent synaptic plasticity in visual cortex by age and experience. *Nature* 375:328-331 (1995).

Kirkwood A, Bear MF. Homosynaptic long-term depression in the visual cortex. *Journal of Neuroscience* 14(5): 3404-12 (1994)

Kirkwood A, Rioult MG & Bear MF. Experience-dependent modification of synaptic plasticity in rat visual cortex. *Nature* 381:526-528 (1996).

Kirwan CB & Stark CE. Medial temporal lobe activation during encoding and retrieval of novel face-name pairs. *Hippocampus* 14:919–930 (2004).

Kistler WM, van Hemmen JL. Modeling synaptic plasticity in conjuction with the timing of pre- and postsynaptic action potentials. *Neural Computation* 12(2): 385-405 (2000)

Knapska E & Kaczmarek L. A gene for neuronal plasticity in the mammalian brain: Zif268/Egr-1/NGFI-A/Krox-24/TIS8/ZENK? *Progress in Neurobiology* 74(4):183-211 (2004).

Knight R. Contribution of human hippocampal region to novelty detection. *Nature* 383(6597):256-259 (1996).

Kobayashi K, Poo MM. Spike train timing-dependent associative modification of hippocampal CA3 recurrent synapses by mossy fibers. *Neuron* 41(3): 445-54 (2004)

Kobayashi Y, Shimazaki H, Aihara T & Tsukada M. Spatial distributions of hippocampal LTP/LTD induced electrically from Schaffer collaterals and stratum oriens with relative timing. *The Brain and Neural Networks*, 8:57–64 (2001).

Koester HJ & Sakmann B. Calcium dynamics in single spines during coincident pre- and postsynaptic activity depend on relative timing of back-propagating action potentials and subthreshold excitatory postsynaptic potentials. *PNAS USA 95*:9596-9601 (1998).

Kohler S, Danckert S, Gati JS & Menon RS. Novelty responses to relational and non-relational information in the hippocampus and the parahippocampal region: A comparison based on eventrelated fMRI. *Hippocampus* 15(6):763-774 (2005).

Korol DL, Abel TW, Church LT, Barnes CA, McNaughton BL. Hippocampal synaptic enhancement and spatial learning in the Morris swim task. *Hippocampus* 3(2):127-32 (1993)

Kreiman G, Koch C & Fried I. Category specific visual responses of single neurons in the human medial temporal lobe. *Nature Neuroscience* 3:946-953 (2000).

Krucker T et al. Targeted disruption of RC3 reveals a calmodulin-based mechanism for regulating metaplasticity in the hippocampus. *Journal of Neuroscience* 22(13):5525-35 (2002).

Krug M, Loessner B & Ott T. Anisomycin blocks the late phase of long-term potentiation in the dentate gyrus of freely moving rats. *Brain Research Bulletin* 13:39-42 (1984).

Kubik S, Miyashita T & Guzowski JF. Using immediate-early genes to map hippocampal subregional functions. *Learning and Memory* 14(11):758-770 (2007).

Kullmann DM & Lamsa KP. Long-term synaptic plasticity in hippocampal interneurons. *Nature Reviews Neuroscience* 8:687-699 (2007).

Kumaran D, Maguire EA. Which computational mechanisms operate in the hippocampus during novelty detection? *Hippocampus* 17(9): 735-48 (2007)

Larkman A, Hannay T, Stratford K & Jack J. Presynaptic release probability influences the locus of long-term potentiation. *Nature* 360:70-73 (1992).

Lasalle JM, Bataille T & Halley H. Reversible inactivation of the hippocampal mossy fibre synapses in mice impairs spatial learning, but neither consolidation nor memory retrieval, in the Morris navigation task. *Neurobiology of Learning and Memory* 73:243-257 (2000).

Lauri SE, Palmer M, Segerstrale M, Vesikansa A, Taira T & Collingridge GL. Presynaptic mechanisms involved in the expression of STP and LTP at CA1 synapses in the hippocampus. *Neuropharmacology* 52:1-11 (2007).

Lee HK et al. Phosphorylation of the AMPA receptor GluR1 subunit is required for synaptic plasticity and retention of spatial memory. *Cell* 112(5):631-643 (2003).

Lee I & Kesner RP. Encoding versus retrieval of spatial memory: double dissociation between the dentate gyrus and the perforant path inputs into CA3 in the dorsal hippocampus. *Hippocampus* 14:66-76 (2004).

Lee I and Kesner RP. Differential contribution of NMDA receptors in hippocampal subregions to spatial working memory. *Nature Neuroscience* 5(2): 162-8 (2002)

Lee I and Kesner RP. Time-dependent relationship between the dorsal hippocampus and the prefrontal cortex in spatial memory. *Journal of Neuroscience* 23(4): 1517-23 (2003)

Lee I, Rao G & Knierem JJ. A double dissociation between hippocampal subfields: different time course of CA3 and CA1 place cells for processing changed environments. *Neuron* 42:803-815 (2004).

Lee I, Yaganarasimha D, Rao G & Knierem JJ. Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. *Nature* 430:456-459 (2004).

Leibold C, Kempter R, van Hemmen JL. How spiking neurons give rise to a temporal-feature map: from synaptic plasticity to axonal selection. *Physical Review E* 65(5): 051915 (2002)

Lengyel M, Kwag J, Paulsen O, Dayan P. Matching storage and recall: hippocampal spike timing-dependent plasticity and phase response curves. *Nature Neuroscience* 8(12): 1677-83 (2005)

Leutgeb JK, Leutgeb S, Treves A, Meyer R, Barnes CA, McNaughton BL, Moser MB, Moser EI. Progressive transformation of hippocampal neuronal representations in "morphed" environments. *Neuron* 48(2): 345-58 (2005)

Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB. Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* 309(5734): 619-23 (2005)

Leutgeb S, Leutgeb JK, Moser MB, Moser EI. Place cells, spatial maps and the population code for memory. Current Opinion in Neurobiology 15(6): 738-46 (2005)

Leutgeb S, Leutgeb JK. Pattern separation, pattern completion, and new neuronal codes within a continuous CA3 map. *Learning and Memory* 14(11): 745-57 (2007)

Lever C et al. Spatial coding in the hippocampal formation: input, information type, plasticity and behaviour. Jeffery KJ, editor, *The neurobiology of Spatial Behaviour*. Oxford University Press, Oxford (2003).

Lever C, Wills T, Cacucci F, Burgess N & O'Keefe J. Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature* 416:90-94 (2002).

Lin YW, Min MY, Chiu TH & Yang HW. Enhancement of associative long-term potentiation by activation of beta-adrenergic receptors at CA1 synapses in rat hippocampal slices. *Journal of Neuroscience* 23:4173-4181 (2003).

Lisman J, Schulman H and Cline H. The molecular basis of CaMKII function in synaptic and behavioural memory. *Nature Reviews Neuroscience* 3(3): 175-90 (2002)

Lisman J, Spruston N. Postsynaptic depolarization requirements for LTP and LTD: a critique of spike timing-dependent plasticity. *Nature Neuroscience* 8(7): 839-41 (2005)

Lisman J. Long-term potentiation: outstanding questions and attempted synthesis. *Philosophical Transactions of the Royal Society London B* 358(1432): 829-42 (2003)

Lisman J. Relating Hippocampal Circuitry to Function: Recall of Memory Sequences by Reciprocal Dentate/CA3 Interactions. *Neuron* 22:233-242 (1999).

Lisman JE, Goldring MA. Feasibility of long-term storage of graded information by the Ca2+/calmodulin-dependent protein kinase molecules of the postsynaptic density. *Proceedings of the National Academy of Sciences* 85(14): 5320-4 (1988)
Lisman JE. A mechanism for the Hebb and anti-Hebb processes underlying learning and memory. *PNAS* 86:9574-9578 (1989).

Lissin DV et al. Activity differentially regulates the surface expression of synaptic AMPA and NMDA glutamate receptors. *PNAS USA* 95:7097-7102 (1998).

Lomo T and Bliss TV. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of Physiology* 232 (2): 331-56 (1973)

Lu W-Y, Man H-Y, Ju W, Trimble WS, MacDonald JF, & Tian Y. Activation of Synaptic NMDA Receptors Induces Membrane Insertion of New AMPA Receptors and LTP in Cultured Hippocampal Neurons. *Neuron* 29:243-255 (2001).

Lynch GS & Baudry M. The biochemistry of memory: A new and specific hypothesis. *Science* 224:1057-1063 (1984).

Lynch MA. Long-term potentiation and memory. *Physiological Reviews* 84:87-136 (2004).

Maass W & Markram H. On the computational power of circuits of spiking neurons. *Journal of computer and system sciences* 69(4):593-616 (2004).

Maass W. Networks of spiking neurons: The third generation of neural network models. *Neural Networks* 10(9):1659-1671 (1997).

Maaswinkel H, Jarrard LE & Whishaw IQ. Hippocampectomized rats are impaired in homing by path integration. *Hippocampus* 9(5):553-561 (1999).

Magee JC, Johnston D. A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science* 275(5297): 209-13 (1997)

Maguire E et al. Knowing where and getting there: a human navigation network. *Science* 280:921-924 (1998).

Maguire EA et al. Navigation-related structural change in the hippocampi of taxi drivers. *Proceedings of the National Academy of Sciences of the United States of America 97* (8):4398-4403 (2000).

Malenka R. The long-term potential of LTP. *Nature reviews Neuroscience* 4:923-926 (2003).

Malenka RC & Bear MF. LTP and LTD: An embarrassment of riches. *Neuron* 44:5-21 (2004).

Malenka RC & Nicoll RA. Long-term Potentiation - A decade of progress? *Science* 285:1870-1874 (1999).

Malenka RC & Siegelbaum SA. Synaptic plasticity. Cowan WM, Sudhof TC & Stevens CF, editors, *Synapses*, pp. 393-453. Baltimore, MD: Johns Hopkins University Press (2001).

Malinow R, Malenka RC. AMPA receptor trafficking and synaptic plasticity. *Annual Reviews in the Neurosciences* 25: 103-26 (2002)

Manahan-Vaughan D, Reymann KG. Metabotropic glutamate receptor subtype agonists facilitate long-term potentiation within a distinct time window in the dentate gyrus in vivo. *Neuroscience* 74(3): 723-31 (1996)

Manns JR, Hopkins RO & Squire LR. Semantic memory and the human hippocampus. *Neuron* 38: 127–33 (2003a).

Manns JR, Hopkins RO, Reed JM, Kitchener EG, Squire LR. Recognition memory and the human hippocampus. *Neuron* 37:171-80 (2003b).

Manns JR, Stark CEL & Squire LR. The visual paired-comparison task as a measure of declarative memory. PNAS USA 97(22):12375-12379 (2000).

Marder E, Abbott LF, Turrigiano GG, Liu Z & Golowasch J. Memory from the dynamics of intrinsic membrane currents. *PNAS USA* 93:13481-13486 (1996).

Markram H, Lübke J, Frotscher M, Sakmann B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275(5297): 213-5 (1997)

Markram H, Tsodyks M. Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature* 382(6594): 807-10 (1996)

Markus EJ, Qin YL, Leonard B, Skaggs WE, McNaughton BL, Barnes CA. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *Journal of Neuroscience* 15(11): 7079-94 (1995)

Marr D. Simple Memory: A Theory for Archicortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 262(841):23-81 (1971).

Martin SJ, de Hoz L & Morris RG. Retrograde amnesia: neither partial nor complete hippocampal lesions in rats result in preferential sparing of remote spatial memory, even after reminding. *Neuropsychologia* 43:609–624 (2005).

Martin SJ, de Hoz L & Morris RG. Retrograde amnesia: neither partial nor complete hippocampal lesions in rats result in preferential sparing of remote spatial memory, even after reminding. *Neuropsychologia* 43:609-624 (2005).

Martin SJ, Greenwood PD & Morris RG. Synaptic plasticity and memory: an evaluation of the hypothesis. *Annual reviews in Neuroscience* 23:649-711 (2000).

Marty S, Berninger B, Carroll P & Thoenen H. GABAergic stimulation regulates the phenotype of hippocampal interneurons through the regulation of brain-derived neurotrophic factor. *Neuron* 16:565-570 (1996).

Matsumura N, Nishijo H, Tamura R, Eifuku S, Endo S & Ono T. Spatial and task dependent neuronal responses during real and virtual translocation in the monkey hippocampal formation. *Journal of Neuroscience* 19:2381-2393 (1999).

Mayford M, Wang J, Kandel ER & O'Dell TJ. CaMKII regulates the frequency-response function of hippocampal synapses for the production of both LTD and LTP. *Cell* 81:891-904 (1995).

McClelland JL, Goddard NH. Considerations arising from a complementary learning systems perspective on hippocampus and neocortex. *Hippocampus* 6(6): 654-65 (1996)

McClelland JL, McNaughton BL, O'Reilly RC. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological Review* 102(3): 419-57 (1995)

McNaughton BL, Barnes C, Meltzer J & Sutherland RJ. Hippocampal granule cells are necessary for normal spatial learning but not for spatially-selective pyramidal cell discharge. *Experimental Brain Research* 76:485-496 (1989).

McNaughton BL, Barnes CA & O'Keefe J. The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research* 52:41-49 (1983).

Meffin H, Besson J, Burkitt AN, Grayden DB. Learning the structure of correlated synaptic subgroups using stable and competitive spike-timing-dependent plasticity. *Physical Review E* 73(4): 041911 (2006)

Migaud M et al. Enhanced long-term potentiation and impaired learning in mice with mutant postsynaptic density-95 protein. *Nature* 396:433-9 (1998).

Miller SG, Kennedy MB. Regulation of brain type II Ca2+/calmodulin-dependent protein kinase by autophosphorylation: a Ca2+-triggered molecular switch. *Cell* 44(6): 861-70 (1986)

Miyashita Y & Chang HS. Neuronal correlate of short-term pictorial short-term memory in the primate temporal cortex. *Nature* 331:68-70 (1988).

Mockett BG, Guévremont D, Williams JM & Abraham WC. Dopamine D1/D5 Receptor Activation Reverses NMDA Receptor-Dependent Long-Term Depression in Rat Hippocampus. *The Journal of Neuroscience* 27(11):2918-2926 (2007).

Moga DE et al. Activity-regulated cytoskeletal associated protein is localized to recently activated excitatory synapses. *Neuroscience* 125:7-11 (2004).

Mohn AR, Gainetdinov RR, Caron MG & Koller BH. Mice with reduced NMDA receptor expression display behavior related schizophrenia. *Cell* 98:427-436 (1999).

Montgomery JM, Selcher JC, Hanson JE & Madison DV. Dynamin-dependent NMDAR endocytosis during LTD and its dependence on synaptic state. *BMC Neuroscience* 6:48 (2005).

Morishita W, Marie H & Malenka RC. Distinct triggering and expression mechanisms underlie LTD of AMPA and NMDA synaptic responses. *Nature Neuroscience* 8:1043-1050 (2005).

Morris R. Theories of Hippocampal function. Andersen P, Morris R, Amaral D, Bliss T & O'Keefe J, editors, *The Hippocampus Book.* Oxford University Press, Oxford (2007).

Morris RG, Garrud P, Rawlins JN & O'Keefe J. Place navigation impaired in rats with hippocampal lesions. *Nature* 297:681-3 (1982).

Morris RGM, Anderson E, Lynch GS & Baudry M. Selective impairment of learning and blockade of long-term potentiation by an N-methyl-D-aspartate receptor antagonist, APV. *Nature* 319:774-776 (1986).

Morris RGM. Does the hippocampus play a disproportionate role in spatial memory? *Discussions in Neuroscience VI*:39-45 (1990).

Morrison A, Aertsen A, Diesmann M. Spike-timing-dependent plasticity in balanced random networks. *Neural Computation* 19(6): 1437-67 (2007)

Moscovitch M et al. Functional neuroanatomy of remote episodic, semantic and spatial memory: a unified account based on multiple trace theory. *Journal of Anatomy* 207(1):35-66 (2005).

Mumby DG. Perspectives on object-recognition memory following hippocampal damage: lessons from studies in rats. *Behavioural Brain Research* 127(1-2): 159-81 (2001)

Murray EA & Wise SP. What, if anything, is the medial temporal lobe, and how can the amygdale be part of it if there is no such thing? *Neurobiology of learning and memory* 82:178-198 (2004).

Murray EA, Mishkin M. Object recognition and location memory in monkeys with excitotoxic lesions of the amygdala and hippocampus. *Journal of Neuroscience* 18(16): 6568-82 (1998)

Murray EA. Memory for objects in non-human primates. Gazzaniga MS, editor, *The New Cognitive Neurosciences*, Second edition. MIT Press, Cambridge, Massachussetts (2000).

Nadel L, Campbell J, Ryan L. Autobiographical memory retrieval and hippocampal activation as a function of repetition and the passage of time. *Neural Plasticity* 90472 (2007)

Nadel L, Moscovitch M. Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology* 7(2): 217-27 (1997)

Nakazawa K et al. Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience. *Neuron* 38(2):305-315 (2003).

Nakazawa K et al. Requirement for hippocampal CA3 NMDA receptors in associative memory recall. *Science* 297:211-218 (2002).

Neves G, Cooke SF, Bliss TV. Synaptic plasticity, memory and the hippocampus: a neural network approach to causality. *Nature Reviews Neuroscience* 9(1): 65-75 (2008)

Ngezahayo A, Schachner M & Artola A. Synaptic activity modulates the induction of bidirectional synaptic changes in adult mouse hippocampus. *Journal of Neuroscience* 20:2451-2458 (2000).

Ngezahayo A, Schachner M, Artola A. Synaptic activity modulates the induction of bidirectional synaptic changes in adult mouse hippocampus. *Journal of Neuroscience* 20(7): 2451-8 (2000)

Nguyen PV and Woo NH. Regulation of hippocampal synaptic plasticity by cyclic AMP-dependent protein kinases. *Progress in Neurobiology* 71(6): 401-37 (2003)

Nguyen PV, Abel T & Kandel ER. Requirement of a critical period of transcription for induction of a late phase of LTP. *Science* 265:1104-1107 (1994).

Nishiyama M, Hong K, Mikoshiba K, Poo MM & Kato K. Calcium stores regulate the polarity and input specificity of synaptic modification. *Nature* 408:584-588 (2000).

Nowotny T, Rabinovich MI and Abarbanel HDI. Spatial representation of temporal information through spike timing dependent plasticity. *Physical Review E* 68: 011908 (2003)

O'Dell TJ & Kandel ER. Low-frequency stimulation erases LTP through an NMDA receptor-mediated activation of protein phosphatases. *Learning and Memory* 1:129-139 (1994).

O'Keefe J & Burgess N. Geometric determinants of the place fields of hippocampal neurons. *Nature* 381:425-428 (1996).

O'Keefe J. Hippocampal Neurophysiology in the behaving animal. Andersen P, Morris R, Amaral D, Bliss T & O'Keefe J, editors, *The Hippocampus Book.* Oxford University Press, Oxford (2007).

O'Keefe JA & Speakman A. Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research* 68:1-27 (1987).

O'Reilly RC & Rudy JW. Computational Principles of Learning in the Neocortex and Hippocampus. *Hippocampus* 10:389-397 (2000).

O'Carroll CM, Martin SJ, Sandin J, Frenguelli B and Morris RG. Dopaminergic modulation of the persistence of one-trial hippocampus-dependent memory. *Learning and Memory* 13(6): 760-9 (2006)

O'Donovan MJ, Rinzel J. Synaptic depression: a dynamic regulator of synaptic communication with varied functional roles. *Trends in the Neurosciences* 20(10): 431-433 (1997)

Oh MC, Derkach VA, Guire ES & Soderling TR. Extrasynaptic membrane trafficking regulated by GluR1 serine 845 phosphorylation primes AMPA receptors for long-term potentiation. *The Journal of Biological Chemistry* 281:752-758 (2006).

Oja E. A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology* 15:267-273 (1982).

O'Reilly RC, McClelland JL. Hippocampal conjunctive encoding, storage, and recall: avoiding a trade-off. *Hippocampus* 4(6): 661-82 (1994)

Parkin AJ. Human memory: novelty, association and the brain. *Current Biology* 7(12):768-769 (1997).

Parkinson JK, Murray EA & Mishkin M. A selective mnemonic role for the hippocampus in monkeys: memory for the location of objects. *Journal of Neuroscience* 8:4059-4167 (1988).

Parslow DM et al. Allocentric spatial memory activation of the hippocampal formation measure with fMRI. *Neuropsychology* 18(3):450-461 (2004).

Pastalkova E, Serrano P, Pinkhasova D, Wallace E, Fenton A & Sacktor T. Storage of spatial information by the maintenance mechanism of LTP. *Science* 313(5790):1141-1144 (2006).

Paulsen O, Moser EI. A model of hippocampal memory encoding and retrieval: GABAergic control of synaptic plasticity. *Trends in the Neurosciences* 21(7): 273-8 (1998)

Pfister JP, Gerstner W. Triplets of spikes in a model of spike timing-dependent plasticity. *Journal of Neuroscience* 26(38): 9673-82 (2006)

Philpot BD, Cho KK, Bear MF. Obligatory role of NR2A for metaplasticity in visual cortex. *Neuron* 53(4): 495-502 (2007)

Pittenger C & Kandel ER. In search of general mechanisms for long-lasting plasticity: Aplysia and the hippocampus. *Philosophical Transactions of the Royal Society B: Biological Sciences* 358:757-763 (2003).

Poirier GL, Amin E, Aggleton JP. Qualitatively different hippocampal subfield engagement emerges with mastery of a spatial memory task by rats. *Journal of Neuroscience* 28(5): 1034-45 (2008)

Quinlan EM et al. A molecular mechanism for stabilization of learning-induced synaptic modifications. *Neuron* 41:185-192 (2004).

Quinlan EM et al. Bidirectional, experience-dependent regulation of N-methyl-D-aspartate receptor subunit composition in the rat visual cortex during postnatal development. *PNAS USA* 96:12876-12880 (1999b).

Quinlan EM et al. Rapid, experience-dependent expression of synaptic NMDA receptors in visual cortex in vivo. *Nature Neuroscience* 2:352-357 (1999a).

Quinlan EM, Olstein DH & Bear MF. Bidirectional, experience-dependent regulation of NMDA receptor subunit composition in rat visual cortex during postnatal development. *PNAS* 96:12876 (1999).

Ramón y Cajal S. Conexion general de los elementos nerviosos. *La Medicina Práctica* 2:341-346 (1889).

Ramon y Cajal S. The Croonian lecture: la fine structure des centres nerveux. *Proceedings of the Royal Society* B65:444-467 (1894).

Redish AD & Touretzky DS. The role of the hippocampus in solving the Morris water maze. *Neural Computation* 10(1):73-111 (1998).

Redish AD. *Beyond the cognitive map: from place cells to episodic memory.* MIT Press, Cambridge, Massachusetts (1999).

Rempel-Clower NL, Zola SM, Squire LR & Amaral DG. Three Cases of Enduring Memory Impairment after Bilateral Damage Limited to the Hippocampal Formation. *The Journal of Neuroscience* 16(16):5233-5255 (1996).

Rioult-Pedotti MS et al. Learning-induced LTP in neocortex. *Science* 290:533-536 (2000).

Roberts PD, Bell CC. Spike timing dependent synaptic plasticity in biological systems. *Biological Cybernetics* 87(5-6): 392-403 (2002)

Rolls ET & Tovee MJ. Processing speed in the cerebral cortex and the neurophysiology of visual masking. *Proceedings of the Royal Society of London, Series B* 257:9-15 (1994).

Rolls ET & Treves A. *Neural Network and Brain Function*. Oxford University Press, Oxford (1998).

Rolls ET, Cahusac PM, Feigenbaum JD & Miyashita Y. Responses of single neurons in the hippocampus of the macaque related to recognition memory. *Experimental Brain Research* 93:299-306 (1993).

Rolls ET, Kesner RP. A computational theory of hippocampal function, and empirical tests of the theory. *Progress in Neurobiology* 79(1): 1-48 (2006)

Rolls ET, Miyashita Y, Cahusac PM, Kesner RP, Niki H, Feigenbaum JD, Bach L. Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. *Journal of Neuroscience* 9(6): 1835-45 (1989)

Rolls ET, Robertson RG & Georges-Francois P. Spatial view cells in the primate hippocampus. *European Journal of Neuroscience* 9:1789-1794 (1997).

Rolls ET, Stringer SM, Trappenberg TP. A unified model of spatial and episodic memory. *Proceedings of the Royal Society London B* 269(1496): 1087-93 (2002)

Rolls ET, Xiang JZ & Franco L. Object, space and object–space representations in the primate hippocampus. *Journal of Neurophysiology* 94:833-844 (2005).

Rolls ET, Xiang JZ. Reward-spatial view representations and learning in the primate hippocampus. *Journal of Neuroscience* 25(26): 6167-74 (2005)

Rolls ET, Xiang JZ. Spatial view cells in the primate hippocampus and memory recall. *Reviews in the Neurosciences* 17(1-2): 175-200 (2006)

Rolls ET. A Theory of Hippocampal Function in Memory. *Hippocampus* 6:601-620 (1996).

Rosenbaum RS, Moscovitch M, Foster JK, Schnyer DM, Gao F, Kovacevic N, Verfaellie M, Black SE & Levine B. Patterns of Autobiographical Memory Loss in Medial-Temporal Lobe Amnesic Patients. *Journal of Cognitive Neuroscience* 20(8):1490-1506 (2008).

Rosenbaum RS, Ziegler M, Winocur G, Grady CL & Moscovitch M. 'I have often walked down this street before': fMRI studies on the hippocampus and other structures during mental navigation of an old environment. *Hippocampus* 14:826–835 (2004).

Rutherford LC, DeWan A, Lauer HM & Turrigiano GG. Brain-derived neurotrophic factor mediates the activity dependent regulation of inhibition in neocortical cultures. *Journal of Neuroscience* 17:4527-4535 (1997).

Rutherford LC, Nelson SB & Turrigiano GG. BDNF has opposite effects on the quantal amplitude of pyramidal neuron and interneuron excitatory synapses. *Neuron* 21:521-530 (1998).

Sajikumar S & Frey JU. Late-associativity, synaptic tagging, and the role of dopamine during LTP and LTD. *The Neurobiology of Learing and Memory* 82:12-25 (2004).

Sajikumar S, Navakkode S, Sacktor TC, Frey JU. Synaptic tagging and cross-tagging: the role of protein kinase Mzeta in maintaining long-term potentiation but not long-term depression. *Journal of Neuroscience* 25(24): 5750-6 (2005)

Sakai Y, Nakano K & Yoshizawa S. Synaptic regulation on various STDP rules. *Neurocomputing* 58-60:351-357 (2004).

Salzmann E, Vidyasagar TR & Creutzfeldt OD. Fucntional comparison of neuronal properties in the primate posterior hippocampus and parahippocampus (area TF/TH) during different behavioural paradigms involving memory and selective attention. *Behavioural Brain Research* 53:133-149 (1993).

Samsonovich A, McNaughton BL. Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience* 17(15): 5900-20 (1997)

Sanes JR and Lichtman JW. Can molecules explain long-term potentiation? *Nature Neuroscience* 2(7): 597-604 (1999)

Saudargiene A, Porr B, Wörgötter F. Local learning rules: predicted influence of dendritic location on synaptic modification in spike-timing-dependent plasticity. *Biological Cybernetics* 92(2): 128-38 (2005)

Sauvage MM, Fortin NJ, Owens CB, Yonelinas AP & Eichenbaum H. Recognition memory: opposite effects of hippocampal damage on recollection and familiarity. *Nature Neuroscience* 11:16-18 (2008).

Schacter DL & Tulving E. *Memory Systems*. MIT Press, Cambridge, Massachusetts (1994).

Schulman H. The multifunctional Ca2+/calmodulin-dependent protein kinases. *Current Opinion in Cell Biology* 5(2): 247-53 (1993)

Sejnowski TJ. Statistical constraints on synaptic plasticity. *Journal of Theoretical Biology* 69(2): 385-9 (1977)

Selig DK, Nicoll RA, Malenka RC. Hippocampal long-term potentiation preserves the fidelity of postsynaptic responses to presynaptic bursts. *Journal of Neuroscience* 19(4): 1236-46 (1999)

Serrano P, Yao Y & Sacktor T. Persistent phosphorylation by protein kinase Mzeta maintains late-phase long-term potentiation. *Journal of Neuroscience* 25(8):1979-84 (2005)

Shadlen MN & Newsome WT. The variable discharge of cortical neurons: Implications for connectivity, computation and information coding. *Journal of Neuroscience*, 18:3870-3896 (1998).

Shapiro ML, Tanila H & Eichenbaum H. Cues that hippocampal place cells encode: dynamic and hierarchical representations of local and distal stimuli. *Hippocampus* 7:624-642 (1997).

Shen B & McNaughton BL. Modelling the spontaneous reactivation of experience specific hippocampal cell assemblies during sleep. *Hippocampus* 6(6):685-693 (1996).

Sheng M & Kim MJ. Postsynaptic Signaling and Plasticity Mechanisms. *Science* 298:776-780 (2002).

Shepherd GM. The neuron doctrine: a revision of functional concepts. *Yale Journal of Biological Medicine* 45:584-599 (1972).

Sherrington CS (1897). Foster M, editor, *Textbook of Physiology*. Macmillan, London (1891).

Shimizu E, Tang YP, Rampon C & Tsien JZ. NMDA receptor-dependent synaptic reinforcement as a crucial process for memory consolidation. *Science* 290:1170-1174 (2000).

Shouval H, Intrator N & Cooper LN. BCM network develops orientation selectivity and ocular dominance in natural scene environment. *Vision Research* 37:3339–3342 (1997).

Shouval HZ, Bear MF & Cooper LN. A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. *PNAS* 99:10831-10836 (2002).

Silva A, Paylor R, Wehner J & Tonegawa S. Impaired spatial learning in α-calcium-calmodulin-kinase II mutant mice. *Science* 257:206-211 (1992a).

Silva AJ, Stevens CF, Tonegawa S & Wang Y. Deficient hippocampal long-term potentiation in α-calcium-calmodulin-kinase II mutant mice. *Science* 257:201-206 (1992b).

Sjöström PJ, Turrigiano GG, Nelson SB. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32(6): 1149-64 (2001)

Smith ML & Milner B. The role of the right hippocampus in the recall of spatial location. *Neuropsychologia* 19:781-793 (1981).

Sobczyk A, Svoboda K. Activity-dependent plasticity of the NMDA-receptor fractional Ca2+ current. *Neuron* 53(1): 17-24 (2007)

Soderling TR and Derkach VA. Postsynaptic protein phosphorylation and LTP. *Trends in the Neurosciences* 23(2): 75-80 (2000)

Softky WR & Koch C. The Highly Irregular Firing of Cortical Cells Is Inconsistent Temporal Integration of Random EPSP's. *Journal of Neuroscience*, 13:334-530 (1993).

Solger J, Wozny C, Manahan-Vaughan D, Behr J. Distinct mechanisms of bidirectional activity-dependent synaptic plasticity in superficial and deep layers of rat entorhinal cortex. *European Journal of Neuroscience* 19(7): 2003-7 (2004)

Song I, Huganir RL. Regulation of AMPA receptors during synaptic plasticity. *Trends in the Neurosciences* 25(11): 578-88

Song S, Abbott LF. Cortical development and remapping through spike timing-dependent plasticity. *Neuron* 32(2): 339-50 (2001)

Song S, Miller KD, Abbott LF. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience* 3(9): 919-26 (2000)

Spiers H, Maguire E & Burgess N. Hippocampal Amnesia. *Neurocase* 7:357-382 (2001).

Squire LR & Cohen NJ. Human memory and Amnesia. Lynch G, McGaugh JL, Weinberger NM, editors, *Neurobiology of learning and memory*, p3-64. Gilford Press, New York (1984).

Squire LR & Zola-Morgan S. The medial temporal lobe memory system. *Science* 253:1380-1386 (1991).

Squire LR, Clark RE & Bayley PJ. Medial temporal lobe function. Gazzaniga MS, editor, *The Cognitive Neurosciences.* MIT Press, Cambridge, Massachusetts (2004).

Squire LR, Clark RE & Knowlton BJ. Retrograde Amnesia. *Hippocampus* 11(1):50-55 (2001).

Squire LR, Cohen NJ & Nadel L. The medial temporal region and memory consolidation: a new hypothesis. Weingartner H & Parker E, editors, *Memory Consolidation: Psychobiology of Cognition*, pp. 185-210. Hillsdale, NJ, USA, (1984).

Squire LR, Stark CE & Clark RE. The medial temporal lobe. *Annual reviews in the Neurosciences* 27:279-306 (2004).

Squire LR. Mechanisms of memory. *Science* 232:1612-1619 (1986).

Squire LR. Memory and the hippocampus: A synthesis from findings with rats, monkeys and humans. *Psychological Reviews* 99:195-231 (1992).

Staba RJ, Wilson CL, Fried I & Engel J Jr. Hippocampal neurons fire bursts during slow wave sleep: Single neuron burst firing in the human hippocampus during sleep. *Hippocampus* 12(6):724-34 (2002).

Standage D, Jalil S, Trappenberg T. Computational consequences of experimentally derived spike-time and weight dependent plasticity rules. *Biological Cybernetics* 96(6): 615-23 (2007)

Stark C. Functional role of the human hippocampus. Andersen P, Morris R, Amaral D, Bliss T & O'Keefe J, editors, *The Hippocampus Book.* Oxford University Press, Oxford (2007).

Stark CEL & Squire LR. Recognition memory and familiarity judgments in severe amnesia: No evidence for a contribution of repetition priming. *Behavioral Neuroscience* 114(3):459-467 (2000).

Staubli U et al. Centrally active modulators of glutamate receptors facilitate the induction of long-term potentiation in vivo. *PNAS USA* 91(23):11158-11162 (1994).

Steinvorth S, Levine B, Corkin S. Medial temporal lobe structures are needed to re-experience remote autobiographical memories: evidence from H.M. and W.R. *Neuropsychologia* 43(4): 479-96 (2005)

Stellwagen D, Malenka RC. Synaptic scaling mediated by glial TNF-alpha. *Nature* 440(7087): 1054-9 (2006)

Steward O & Fass B. Polyribosomes associated with dendritic spines in the denervated dentate gyrus: evidence for local regulation of protein synthesis during reinnervation. *Progress in Brain Research* 58:131-136 (1983)

Steward O & Levy WB. Preferential localization of polyribosomes under the base of dendritic spines in granule cells of the dentate gyrus. *Journal of Neuroscience* 2:284-291 (1982)

Steward O, Wallace CS, Lyford GL & Worley PF. Synaptic activation causes the mRNA for the IEG Arc to localize selectively near activated postsynaptic sites on dendrites. *Neuron* 21:741-751 (1998)

Stickgold R & Walker MP. Memory consolidation and reconsolidation: what is the role of sleep? *Trends in Neurosciences* 28(8):408-415 (2005).

Stringer SM, Rolls ET. Self-organizing path integration using a linked continuous attractor and competitive network: path integration of head direction. *Network* 17(4): 419-45 (2006)

Sutherland GR & McNaughton B. Memory trace reactivation in hippocampal and neocortical neuronal ensembles. *Current Opinion in Neurobiology* 10(2):180-186 (2000)

Suzuki WA & Amaral DG. Perirhinal and parahippocampal cortices of the macaque monkey: cortical afferents. *Journal of comparative neurology* 350:497-533 (1994a)

Suzuki WA & Amaral DG. Topographic organisation of the reciprocal connections between the monkey entorhinal cortex and the perirhinal and parahippocampal cortices. *Journal of Neuroscience* 14:1856-1877 (1994b)

Suzuki WA & Eichenbaum H. The neurophysiology of memory. *Annals of the New York Academy of Sciences* 911:175–191 (2000)

Sweatt JD. Toward a molecular explanation for long-term potentiation. *Learning and Memory* 6(5): 399-416 (1999)

Tamura R, Ono T, Fukuda M & Nakamura K. Spatial responsiveness of monkey hippocampal neurons to various visual and auditory stimuli. *Hippocampus* 2:307-322 (1992)

Tang Y-P et al. Genetic enhancement of learning and memory in mice. *Nature* 401:63 (1999).

Tanila H, Shapiro M, Gallagher M & Eichenbaum H. Brain aging: impaired coding of novel environmental cues. Journal of Neuroscience 17:5167-5174 (1997a)

Tanila H, Shapiro ML & Eichenbaum HE. Discordance of spatial representation in ensembles of hippocampal place cells. *Hippocampus* 7:613-623 (1997b)

Tanila H, Sipila P, Shapiro M & Eichenbaum H. Brain aging: changes in the nature of information coding by the hippocampus. *Journal of Neuroscience* 17:5155-5166 (1997c)

Tanila H. Hippocampal place cells can develop distinct representations of two visually identical environments. *Hippocampus* 9:235-246 (1999)

Teng E & Squire LR. Memory for places learned long ago is intact after hippocampal damage. *Nature* 400:675-7 (1999)

Thiagarajan TC, Lindskog M, Malgaroli A, Tsien RW. LTP and adaptation to inactivity: overlapping mechanisms and implications for metaplasticity. *Neuropharmacology* 52(1): 156-75 (2007)

Thomas KG, Hsu M, Laurance HE, Nadel L and Jacobs WJ. Place learning in virtual space III – Investigation of spatial navigation training procedures and their application to fMRI and clinical neuropsychology. *Behavior Research Methods, Instruments, & Computers* 33:21-37 (2001)

Thompson LT & Best PJ. Long-term stability of place field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain Research* 509:299-308 (1990).

Tompa P, Friedrich P. Synaptic metaplasticity and the local charge effect in postsynaptic densities. *Trends in the Neurosciences* 21(3): 97-102 (1998)

Triesch J. Synergies between intrinsic and synaptic plasticity mechanisms. *Neural Computation* 19(4): 885-909 (2007)

Tsien JZ, Huerta PT & Tonegawa S. The essential role of hippocampal CA1 NMDA receptor-dependent synaptic plasticity in spatial memory. *Cell* 87(7):1327-1338 (1996).

Tsodyks M. Attractor neural network models of spatial maps in hippocampus. *Hippocampus* 9(4): 481-489 (1999)

Tsodyks M. Attractor neural networks and spatial maps in hippocampus. *Neuron* 48(2): 168-9 (2005)

Tsodyks M. Spike-timing-dependent synaptic plasticity - the long road towards understanding neuronal mechanisms of learning and memory. *TINS* 25(12):599-600 (2002).

Tsodyks MV & Markram H. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *PNAS USA* 94:719-723 (1997).

Tulving E. Episodic and Semantic Memory. Tulving E & Donaldson W, editors, *Organisation and Memory*, p382-403. Academic Press, San Diego (1972).

Tulving E. Memory systems. Eysenck MW, editor, *The Blackwell Dictionary of Cognitive Psychology*, pp. 222-223. Oxford: Blackwell (1990).

Turrigiano GG & Nelson SB. Homeostatic plasticity in the developing nervous system. *Nature reviews Neuroscience* 5:97-107 (2004).

Turrigiano GG, Leslie KR, Desai NS, Rutherford LC & Nelson SB. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391:892-896 (1998).

Turrigiano GG. Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. *Trends in the Neurosciences* 22(5): 221-7 (1999)

Uncapher MR, Rugg MD. Fractionation of the component processes underlying successful episodic encoding: a combined fMRI and divided-attention study. *Journal of Cognitive Neuroscience* 20(2): 240-54 (2008)

Vago DR, Bevan A, Kesner RP. The role of the direct perforant path input to the CA1 subregion of the dorsal hippocampus in memory retention and retrieval. *Hippocampus* 17(10): 977-87 (2007)

van Rossum MC, Bi GQ, Turrigiano GG. Stable Hebbian learning from spike timing-dependent plasticity. *Journal of Neuroscience* 20(23): 8812-21 (2000)

Vazdarjanova A, Guzowski JF. Differences in hippocampal neuronal population responses to modifications of an environmental context: evidence for distinct, yet complementary, functions of CA3 and CA1 ensembles. *Journal of Neuroscience* 24(29): 6489-96 (2004)

Vazdarjanova A, Ramirez-Amaya V, Insel N, Plummer TK, Rosi S, Chowdhury S, Mikhael D, Worley PF, Guzowski JF, Barnes CA. Spatial exploration induces ARC, a plasticity-related immediate-early gene, only in calcium/calmodulin-dependent protein kinase II-positive principal excitatory and inhibitory neurons of the rat forebrain. *Journal of Comparative Neurology* 498(3): 317-29 (2006)

Vidyasagar TR, Salzmann E & Creutzfeldt OD. Unit activity in the hippocampus and parahippocampal temporobasal association cortex related to memory and complex behaviour in the awake monkey. *Brain Research* 544:269-278 (1991)

Viskontas IV, Knowlton BJ, Steinmetz PN & Fried I. Differences in mnemonic processing by neurons in the human hippocampus and parahippocampal regions. *Journal of Cognitive Neuroscience* 18(10):1654-62 (2006)

Vissel B, Krupp JJ, Heinemann SF & Westbrook GL. A use-dependent tyrosine dephosphorylation of NMDA receptors is independent of ion flux. *Nature Neuroscience* 4:587-596 (2001)

Vyazovskiy VV, Cirelli1 C, Pfister-Genskow M, Faraguna1 U & Tononil G. Molecular and electrophysiological evidence for net synaptic potentiation in wake and depression in sleep. *Nature Neuroscience* 11(2):200-208 (2008)

Wallenstein GV, Eichenbaum H, Hasselmo ME. The hippocampus as an associator of discontiguous events. *Trends in the Neurosciences* 21(8): 317-23 (1998)

Wang H & Wagner JJ. Priming-induced shift in synaptic plasticity in the rat hippocampus. *Journal of Neurophysiology* 82:2024-2028 (1999).

Wang HX, Gerkin RC, Nauen DW, Bi GQ. Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nature Neuroscience* 8(2): 187-93 (2005)

Watt AJ et al. A proportional but slower NMDA potentiation follows AMPA potentiation in LTP. *Nature Neuroscience* 7:518-524 (2004).

Watt AJ, van Rossum MCW, MacLeod KM, Nelson SB & Turrigiano GG. Activity Coregulates Quantal AMPA and NMDA Currents at Neocortical Synapses. *Neuron* 26(3):659-670 (2000).

Wenisch OG, Noll J & Hemmen JL. Spontaneously emerging direction selectivity maps in visual cortex through STDP. *Biological Cybernetics* 93(4):239-247 (2005).

West AE, Griffith EC, Greenberg ME. Regulation of transcription factors by neuronal activity. *Nature Reviews Neuroscience* 3(12): 921-31 (2002)

Wiener SI, Paul CA, Eichenbaum H. Spatial and behavioral correlates of hippocampal neuronal activity. *Journal of Neuroscience* 9(8): 2737-63 (1989)

Williams JM et al. Differential Trafficking of AMPA and NMDA Receptors during Long-Term Potentiation in Awake Adult Animals. *The Journal of Neuroscience* 27(51):14171-14178 (2007).

Wills T, Lever C, Cacucci F, Burgess N & O'Keefe J. Attractor dynamics in the hippocampal representation of the local environment. *Science* 308:873-876 (2005).

Wilson MA, McNaughton BL. Reactivation of hippocampal ensemble memories during sleep. *Science* 265(5172): 676-9 (1994)

Wirth S, Yanike M, Frank LM, Smith AC, Brown EN, Suzuki WA. Single neurons in the monkey hippocampus and learning of new associations. *Science* 300(5625): 1578-81 (2003)

Wong ST et al. Calcium-stimulated adenylyl cyclase activity is critical for hippocampus-dependent long-term memory and late phase LTP. *Neuron* 23:787-798 (1999).

Woo NH & Nguyen PV. "Silent" metaplasticity of the late phase of long-term potentiation requires protein phosphatases. *Learning and Memory* 9:202-213 (2002).

Wood ER, Dudchenko PA & Eichenbaum H. The global record of memory in hippocampal neuronal activity. *Nature* 397:561–563 (1999).

Wood ER, Dudchenko PA, Robitsek J & Eichenbaum H. Hippocampal Neurons Encode Information about Different Types of Memory Episodes Occurring in the Same Location. *Neuron* 27(3):623-633 (2000).

Wood ER, Dudchenko PA, Eichenbaum H The global record of memory in hippocampal neuronal activity. *Nature* 397: 613-616 (1999)

Xiang JZ & Brown MW. Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. *Neuropharmacology* (37)657-676 (1998).

Yamasaki TR, Blurton-Jones M, Morrissette DA, Kitazawa M, Oddo S & LaFerla FM. Neural stem cells improve memory in an inducible mouse model of neuronal loss. *Journal of Neuroscience* 27(44):11925-33 (2007).

Yang SN, Tang YG & Zucker RS. Selective induction of LTP and LTD by postsynaptic $[Ca^{2+}]i$ elevation. *Journal of Neurophysiology* 78:781-787 (1999).

Yasuda H, Barth AL, Stellwagen D & Malenka RC. A developmental switch in the signaling cascades for LTP induction. *Nature Neuroscience* 6:15-16 (2003).

Yeckel MF, Kapur A & Johnston D. Multiple forms of LTP in hippocampal CA3 neurons use a common postsynaptic mechanism. *Nature Neuroscience* 2:625-633 (1999).

Young BJ, Fox GD & Eichenbaum H. Correlates of hippocampal complex-spike cell activity in rats performing a nonspatial radial maze task. *Journal of Neuroscience* 14:6553-6563 (1994).

Young JZ, Isiegas C, Abel T, Nguyen PV. Metaplasticity of the late-phase of long-term potentiation: a critical role for protein kinase A in synaptic tagging. *European Journal of Neuroscience* 23(7): 1784-94 (2006)

Young JZ, Nguyen PV. Homosynaptic and heterosynaptic inhibition of synaptic tagging and capture of long-term potentiation by previous synaptic activity. *Journal of Neuroscience* 25(31): 7221-31 (2005)

Youssef FF, Addae JI & Stone TW. NMDA-induced preconditioning attenuates synaptic plasticity in the rat hippocampus. *Brain Research* 1073:183-189 (2006).

Zamanillo D et al. Importance of AMPA receptors for hippocampal synaptic plasticity but not for spatial learning. *Science* 284:1805-1811 (1999).

Zhang L et al. Hippocampal synaptic metaplasticity requires inhibitory autophosphorylation of Ca2+/calmodulin-dependent kinase II. *Journal of Neuroscience* 25(33):7697-707 (2005).

Zhang W & Linden DJ. The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nature Reviews Neuroscience* 4:885-900 (2003).

Zhu XO, Brown MW, McCabe BJ & Aggleton JP. Effects of novelty or familiarity of visual stimuli on the expression of the immediate early gene c-fos in rat brain. *Neuroscience* 69:821-829 (1995).

Zhuo M, Zhang W, Son H, Mansuy I, Sobel RA, Seidman J, Kandel ER. A selective role of calcineurin a-alpha in synaptic depotentiation in hippocampus. *Proceedings of the National Academy of Sciences* 96(8): 4650-5 (1999)