UNIVERSITY OF SOUTHAMPTON

A Synaptic and Temporal Ensemble Interpretation of Spike-Timing-Dependent Plasticity

by

Peter A. Appleby

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the
Faculty of Engineering, Science and Mathematics
School of Electronics and Computer Science

March 2006

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS
SCHOOL OF ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by Peter A. Appleby

In this thesis we are concerned with activity-dependent neuronal plasticity in the nervous system, in particular the phenomenon of spike-timing-dependent-plasticity or STDP. We find that the experimental evidence for STDP may be interpreted in a variety of ways. Inspired by this observation, we propose a new model of spike-timing plasticity in the form of a synaptic switch rule. The switch rule governs changes at individual synapses, and only when the rule is averaged over multiple synapses and multiple spike-pairs does an STDP-like rule emerge. The STDP-like rule is therefore an ensemble property of our model, one that is nowhere instantiated at any individual synapse. We find that our switch rule explains a variety of spike- and rate-based plasticity results as a result of its intrinsic structure. We also find that stable, competitive dynamics emerge naturally due to multi-spike interactions. At no stage are we required to introduce additional modifications to accommodate particular experimental results or avoid otherwise undesirable learning behaviours. Indeed, ensuring consistency with various experimental results serves to neatly constrain the parameters of our model in a concise manner. This is in contrast to many other models of STDP, which are often required to introduce additional modifications and non-linearities to explain experimental results on a case-by-case basis. Furthermore, our synaptic switch rule is considerably simpler than many competing models of STDP and places a much lower computational burden on individual synapses. We are therefore freed from the need to postulate precise coincidence detections mechanisms and, as a result, out synaptic switch rule is broadly consistent with a range of possible biological implementations.

Contents

List of Figures

List of Tables

Nomenclature

LTP	Long-term potentiation	3
LTD	Long-term depression	3
NMJ	Neuro-muscular junction	4
LGN	Lateral geniculate nucleus	5
BCM	Bienenstock Cooper Munroe	7
STDP	Spike-timing-dependent plasticity	7
ODC	Ocular dominance column	9
PTP	Post-tetanic potentiation	11
NMDA	N-methyl-D-aspartate	12
APV	Aminophosphovalerate acid	12
PSP	Postsynaptic potential	12
EPSP/C	Excitatory postsynaptic potential/current	12
IPSP/C	Inhibitory postsynaptic potential/current	12
HFS	High-frequency stimulation	16
LFS	Low-frequency stimulation	16
D-AP-V	D-aminophosphovalerate acid	17
mGluR	Metabotropic NMDA receptor	17
MCPG	Alpha-methyl-4-carboxyphenylglycine	17
GABA	G-amino butyric acid	17
CA1/3	Cornu Ammonis 1/3	19
MPP	Medial perforent path	19
LPP	Lateral perforent path	19
ADS	Activity-dependent scaling	56
AMPA	Alpha-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid	60
CaMKII	Calcium/Calmodulin-dependent protein kinase II	61
BPAP	Back-propagating action potential	69
VGCC	Voltage-gated calcium channel	72
ISI	Inter-spike interval	81

Acknowledgements

Thanks to my supervisor Terry Elliott, and to Nigel Shadbolt.

Dedicated to my parents.

Chapter 1

Development and Plasticity of the Nervous System

1.1 Neuronal Plasticity

The structured connectivity of neurons in the brain underlies its distributed approach to information storage and processing. The pattern and weighting of these connections are not static, but may change dynamically, both in development and over the lifetime of an animal. The set of processes by which the pattern and weights of these connections change are collectively referred to as neuronal plasticity. Neuronal plasticity is an apparently ubiquitous feature of the nervous system, and has been observed across a variety of species, brain regions, and at different stages of development.

Neuronal plasticity plays an important role in both the initial development of the nervous system and in its ongoing function. Examples of developmental plasticity are often dramatic. For example, the organisation of the vertebrate visual cortex into columns of alternating ocular dominance may not be completely preprogrammed, but may arise, in part, from the effect of postnatal activity patterns (Hubel and Wiesel, 1962, 1965, 1969, 1977). An impoverished environment during a critical period in early life can exert a profound influence on the development of normal patterns of connectivity (Wiesel and Hubel, 1963a,b, 1965; Shatz and Stryker, 1978). Plasticity in the adult, while perhaps not as dramatic as that during earlier stages of development, continues to play an important role. Ongoing plasticity allows limited recovery from damage arising from insults to the nervous system, and anatomical and physiological changes in neuronal connectivity are also often cited as the possible mechanisms for learning and memory (Martin et al., 2000). Both learning and memory are important in even the simplest of multi-celled organisms; associative learning in the Aplysia can lead

to short-term sensitisation to a stimulus, which is a basic survival mechanism (Castellucci et al., 1970; Pinsker et al., 1970; Castellucci and Kandel, 1976). In higher-level organisms, the increasing complexity of the nervous system further elevates the role of learning and memory, which, it may be argued, underlie almost all higher-level cognitive functions. A full understanding of synaptic plasticity will therefore be one of the key elements in a complete description of how the nervous system develops and operates.

1.2 Activity-Dependent Neuronal Plasticity

A variety of forms of neuronal plasticity are expressed throughout the nervous system. A common distinction is made between activity-dependent and activity-independent plasticity.

Activity-dependent plasticity is broadly defined to be that which depends in some way upon the electrical activity of the participating neurons. Activity-dependent changes in the connectivity of neuronal circuits apparently underlie several important developmental processes, such as the development of orientation selectivity and ocular dominance columns in the visual system. A great deal of experimental work has been directed towards exploring exactly how different patterns of activity affect neuronal connectivity. Patterns of activity may either be instructive or permissive. Instructive refers to the case where the precise pattern of activity play a direct role in shaping the changes in neuronal architecture. Permissive refers to the case where the presence of activity is all that is required, with another process actually directing the changes.

Consider, for example, the refinement of orientation selectivity in the primary visual cortex. Under normal developmental conditions an initially broad tuning of orientation selectivity is refined postnatally. This refinement is driven by the presence of activity along the visual pathway and does not occur if the subject is dark-reared. If both eyes are covered with diffusive patches, however, so that there is a loss of structure in the visual input without a significant reduction in the overall level of activity, refinement of orientation selectivity does not occur. There is therefore a requirement for patterned activity, with the presence of activity alone being insufficient to drive refinement, and this process is an example instructive activity.

Activity-dependent plasticity can persist for varying amounts of time, and a distinction is often made between short- and long-term plasticity. Short-term plasticity is generally considered to be that which decays over the course of a few minutes. Long-term plasticity is generally considered to be that which is stable over the course of perhaps thirty minutes or more. Although convenient,

in the absence of a complete understanding of the mechanisms of induction and expression these distinctions remain somewhat arbitrary. A well studied form of long-term plasticity is long-term potentiation and depression (LTP/LTD) (Bliss and Lømo, 1973; Bliss and Gardner-Medwin, 1973; Malenka and Nicoll, 1999). This form of plasticity apparently operates in a broad range of brain areas and species, notably the hippocampus, neocortex and cerebellum. In theoretical work, a variety of modelling studies have attempted to extract learning rules from these results and explore the implications for a wide range of tasks. Recently, experimental work has offered the possibility of uncovering the molecular basis of LTP and LTD. A growing body of work suggests is that it is the time-course and amplitude of intracellular calcium levels that drive LTP and LTD. This has formed the basis of some recent attempts to derive a biophysical model of LTP according to the "calcium-control hypothesis" (Senn et al., 2000; Shouval et al., 2002).

The various forms of activity-dependent neuronal plasticity observed in the nervous system may be influenced by a variety of factors. It has been suggested that dynamic adjustments to an existing plasticity rule may strongly influence the outcome of an activity-dependent process. This phenomenon, known as meta-plasticity, may subserve a variety of functions. For example, changes in the long-term average activity levels can modify the intrinsic excitability of neurons, altering their firing characteristics. Such a process may contribute to homeostasis, maintaining postsynaptic firing rates in some dynamic range despite large changes in input strengths.

Activity-independent plasticity is not dependent on the activity of the neurons involved. An example would be the chemical guidance of axonal growth cones towards appropriate targets in the initial stages of neuro-muscular innervation. While important for the development of the nervous system, this kind of plasticity does not display the kind of dramatic dependence on the environment that activity-dependent processes do.

1.3 Plasticity and Development of the Nervous System

Anatomical and physiological changes in neuronal connectivity play important roles during early development of the nervous system. After neuronal cells have differentiated, a period of migration is followed by the outgrowth of axons. These axons are guided towards appropriate targets by various chemical cues, eventually forming functional synapses there. In several systems, for example at the

neuromuscular junction (Sanes and Lichtman, 1999) and along the retinogeniculocortical pathway (LeVay et al., 1978, 1980), the precise architecture of connectivity then undergoes refinement based on a less refined initial innervation (Purves and Lichtman, 1985). Such processes have been shown to occur both pre- and postnatally, and to depend on neuronal activity. It is often the case that the role of activity-dependence is still unclear. In particular the question of whether activity is instructive, where the activity patterns govern directly the expression of plasticity, or permissive, where only the presence of activity is needed, is often difficult to resolve.

1.3.1 The Neuromuscular Junction

The vertebrate neuromuscular junction (NMJ) is one of the more commonly cited example of activity-dependent developmental plasticity (for review, see Sanes and Lichtman (1999)). The term NMJ refers to the point of synaptic contact between motor-neurons and muscle fibres. The NMJ forms when initially distant motor-neurons extend axonal processes to the muscle fibres. Both the motor-neurons and muscle fibres are predifferentiated before contact, and possess characteristic pre- and postsynaptic elements such as vesicles and neurotransmitter receptors. Establishment of an initial synaptic connection, albeit of low efficacy, is therefore rapidly achieved on arrival of the motor-neuron growth-cone at the muscle fibre surface. At birth, each muscle fibre is innervated by several different motor neurons. This stage of polyneuronal innervation does not persist long, typically only up to postnatal week one in mice, and a process of refinement known as synapse elimination soon begins. During synapse elimination, inputs compete for control of the target muscle fibre and, gradually, one input gains complete control. The other inputs atrophy, and eventually withdrawing from the fibre completely.

Synapse elimination is activity-dependent, in the sense that the more active motor neuron is more likely to win the competition. Lowering the overall level of activity has also been shown to slow the rate of synapse elimination. The precise mechanisms that govern this competition remain unknown, but it is probable that a range of factors influence the outcome of the competition. By postnatal week two in mice each muscle fibre is left innervated by a single motor-neuron, although some exceptions invariably exist and there is a great deal of variation between species. For example, in amphibians up to one third of muscles fibres remain polyinnervated into adulthood. There is little cell death during synapse elimination, with the total number of motor units (defined to be a motor neuron and the muscle fibres it controls) remaining unchanged.

The refinement of the NMJ appears to be a flexible process for regulating the precise pattern of muscle innervation. The refinement also apparently ensures that the overall level of innervation is appropriate. In the developing chick embryo, for example, the number of neurons innervating a particular limb decreases by about half before hatching.

1.3.2 The Primary Visual Cortex

Another well studied example of activity-dependent developmental plasticity is the formation of topographic maps in the immature vertebrate visual system. Topographic maps are a characteristic feature of the vertebrate brain, where projections from one area of the brain maintain an ordered innervation in a target region. This preserves the spatial relationships between them, with neighbouring projecting cells innervating neighbouring targets. The retino-geniculo-cortical pathway, for example, shows striking organisation on several levels. Photoreceptors trigger, via intermediate bipolar cells, spiking of retinal ganglion cells, whose axons form the optic nerve. In higher vertebrates, such as carnivores and primates, axons from the two eyes converge onto the lateral geniculate nucleus (LGN), but occupy distinct regions. Information is therefore carried to the lateral geniculate nucleus in a topographically ordered manner, preserving a point-to-point correspondence between the retina and LGN. Stimulating a small area of the retina will therefore give a localised response in a small area of the LGN. The axonal arbors of LGN neurons are themselves segregated within layer IV of the primary visual cortex into a pattern of alternating ocular dominance patches (Hubel and Wiesel, 1962, 1965, 1969, 1977). This precise arrangement arises through activity-dependent refinement of an initially much coarser map established by chemical cues (Purves and Lichtman, 1985). In primates, cats, and ferrets, the axonal terminals of ganglion cells of the two eyes initially share common territories within the LGN, but through a process that eliminates inappropriately placed branches, projections from the two eyes become restricted to their appropriate layer. Likewise, ocular dominance columns emerge from an initial pattern of overlapping arbors representing the two eyes. The experiments of Hubel and Wiesel first gave insight into how visual experience affects the formation and patterning of ocular dominance columns. Monocular deprivation (closure of one eye by lid-suture) lead to a shift in responsiveness toward the open eye (Wiesel and Hubel, 1963a,b, 1965; Shatz and Stryker, 1978), demonstrating that such changes are indeed driven by neuronal activity. The tendency for segregation of projections from the two eyes is apparent even in lower vertebrates; when a third eye is implanted in frogs, eye-specific stripes, normally absent, appear in the optic tectum (Constantine-Paton and Law, 1978).

Thus, both activity-dependent and independent neuronal plasticity potentially subserve a range of functions in the brain, both in developing and mature animals. Although a great deal remains to be uncovered about their precise operation, including the actual extent to which activity takes an instructive role in many apparently activity-dependent processes, activity-dependent mechanisms make attractive candidates for underpinning a range of phenomena due to their intrinsic flexibility.

1.4 Outline of the Thesis

We now give an overview of the content of this thesis.

Chapter 2 presents the main experimental results on neuronal plasticity. We first establish some useful nomenclature, and then go on to briefly discuss the distinctions between short- and long-term plasticity, and between anatomical and physiological plasticity. Some general experimental issues are then covered, such as recording methods and stimulation protocols. We also outline the three main methods used for plasticity experiments (neuronal cultures, slices and in vivo recordings) and discuss the influence of subject age on the observed plasticity. Finally, we briefly discuss some of the pharmacological manipulations commonly used in plasticity experiments, such as the introduction of calcium buffers or neurotransmitter receptor antagonists.

We then go on to discuss the experimental basis of neuronal plasticity in some depth. The literature is extensive, and we select papers based either on their importance or for the illustrative nature of the experiments performed. We divide results into four groups which roughly follow on from each other. Within each group experiments are discussed chronologically. Early LTP/LTD experiments refers to those which relied upon a tetanic stimulation protocol. Later LTP/LTD experiments refer to those using bursts of pre- and postsynaptic action potentials, and it is here that we first see evidence for timing-dependence in neuronal plasticity. Timing experiments are those which probe this timing-dependence in more detail, and typically use single pre- and postsynaptic spikes rather than trains or bursts. Finally we present a section on heterogeneity in neuronal plasticity. This section covers a variety of results, from metaplasticity to the recently observed all–or–none potentiation of putative single connections in the hippocampus. We also discuss the issue of whether different neuronal plasticity rules possibly reflect functional differences in the brain areas under examination.

Chapter 3 discusses the various theoretical attempts to model neuronal plasticity from a rate-based perspective. After establishing some notation, we first

discuss Hebbian rules where changes are made based on the correlation between pre- and postsynaptic firing. We formulate a general Hebbian rule in the form of an expansion in the pre- and postsynaptic firing rates, then extract particular examples from it. We discuss the basic Hebb rule and the issues arising from it, such as runaway learning, then discuss various modifications, such as hard bounds on synaptic strengths. We then extend our discussion to pre- and postsynaptically gated rules. Following on from this, we examine the covariance rule, and the Bienenstock-Cooper-Munro (BCM) rules. We discuss the success of the BCM rule in explaining developmental phenomena in the visual system, such as development of orientation selectivity and ocular dominance. We then discuss competition and stability in rate based rules, and examine subtractive and multiplicative normalisation schemes.

In Chapter 4 we consider spike-based models of neuronal plasticity. These models are built around a description of individual spikes rather than describing pre- and postsynaptic activity in purely rate-based terms. We examine two classes of spike-based model: phenomenological and biophysical. Phenomenological models are those which take the experimentally observed spike-timing results over directly to form a plasticity rule. We examine both additive and multiplicative phenomenological models, and discuss the learning dynamics that two particular implementations exhibit. We also examine the competitive nature of these models and their equilibrium distributions of afferent weights. We then discuss biophysical models, outlining a commonly cited biophysical mechanism in the form of the calcium control hypothesis. We examine some particular implementations of this hypothesis, and the success these models have had in explaining a variety of experimental results.

In Chapter 5 we present a fuller analysis of an existing spike-based model of STDP in the form of the Song model (Song et al., 2000). We give a more detailed derivation of the model, drawing out the various issues that are involved in its formulation. Using our own simulations, we reproduce the main findings of the original work. We then extend these simulations to explore the learning dynamics of the Song model more fully. In particular, we discuss the issue of stability which was only partially addressed in the original work. We then compare our findings to those of the original work. Finally, we discuss the various issues surrounding this particular formulation of STDP in more detail.

In Chapter 6 we present a new model of synaptic plasticity. This model, originally presented in the paper "A Synaptic and Temporal Ensemble Interpretation of Neuronal Plasticity" (Appleby and Elliott, 2005), is in the form a three-state synaptic switch. We first discuss the motivation of the rule, drawing upon observations made in Chapter 2, and then formulate the switch rule in its simplest

possible form. After establishing notation which we use throughout the discussion of the model, we proceed with a limited analysis based on the possible interactions of two spikes. We show that an STDP-like modification curve is a robust prediction of our model, arising from the ensemble averaging of our synaptic switch rule across multiple synapses and multiple spike-pairings. We also show that a natural explanation of spike-triplets emerges from our model as a consequence of its intrinsic structure, with no need to introduce additional modifications. Furthermore, we show that the average, long-term behaviour of the switch rule is consistent with experimental results on rate-based plasticity. We derive constraints under which this long-term behaviours is qualitatively BCM-like. We then present numerical simulations, and show that the analytically observed behaviours are robust even under variable pre- and postsynaptic spike timings. We show that permitting multi-spike interactions does not destroy these results, which were derived for 2-spike interactions only. Finally, we discuss various issues surrounding this initial exploration of the switch rule, such as stability and convergence of the 2- and multi-spike rules to the infinite-spike limit (where interactions are not limited in any way).

Chapter 7 extends the analysis of the switch rule to multiple spike interactions. That is, we now consider three, four and higher spike trains instead of limiting interactions to two spikes only. This analysis is rather lengthy, and was originally presented in the paper "Multi-Spike Interactions in a Stochastic Model of Spike Timing Dependent Plasticity: Derivation of Learning Rules" (Appleby and Elliott, 2005) (submitted). We proceed by first establishing a number of preliminary results which are of use throughout the derivation. We then extend the model slightly by allowing, in certain circumstances, the stochastic process underlying the synaptic switch to be reset. We make use of the resetting form of the model to derive the expectation value for the change in synaptic strength for any multi-spike train. We consider the effect of adding additional spikes in front of a multi-spike train, and examine how the resulting expected change in synaptic strength induced by that train changes. In this manner, we may produce a set of recurrence relations, solving these relations produces the required expressions. We refer to these expressions as multi-spike learning rules, as they consider the interaction of more than two spikes.

We then show that the expressions obtained by solving the recurrence relations may be obtained in a much simpler manner by considering the underlying Markovian nature of our switch model. We go on to derive a generating function for the expected change for any multi-spike train. Finally, we present a comparison of our analytical results and the corresponding numerical simulations. We also investigate how rapidly the multi-spike rules converge to the result where spike interactions are not limited in any way. We also show that the 2-spike learning

rule is of a qualitatively different character compared to all other, multi-spike rules.

In Chapter 8, which is based on the paper "Stable Competitive Dynamics Emerge from Multi-Spike Interactions in a Stochastic Model of Spike Timing Dependent Plasticity" Appleby and Elliott, (2005) (submitted), we return to this issue of the difference between the 2-spike and multi-spike interaction functions, and explore the learning dynamics of our switch rule in much greater detail. In particular, we look for the presence or otherwise of stable, competitive dynamics. Examining the 2-spike rule, we see that, without further modification, it always leads to pathological learning behaviours. The irredeemably pathological nature of the 2-spike rule is illustrated by considering synaptic strengths in various limiting cases.

Examining the 3-spike rule, we see that the pathological behaviour is absent, and the rule displays stable, competitive dynamics under a wide range of parameters. We observe that this is also true for any multi-spike learning rule. We also show that the 2-spike rule may be rescued from its pathological behaviour by modifying it in such as a way as to introduce a BCM-like sliding of the threshold between potentiation and depression. We propose that the difference between the 2-spike and the multi-spike and modified 2-spike rules is that the latter two couple potentiation and depression together, whilst the former does not. This coupling takes place in the sense that prior pre- and postsynaptic spikes, and the resulting potentiation and depression, influences later plasticity. We support this claim by re-deriving the multi-spike rules for an uncoupled switch, as we see that something akin to the 2-spike rule is returned in all cases. We then perform a fixed-point analysis of a multi-spike learning rule, and show that the dynamics are governed in a well defined manner by the presence or absence of stable, segregated fixed-points. The existence of these stable, segregated fixed-points broadly constrains the parameters of the model to a certain regime, outside of which competition breaks down.

We then perform a numerical verification that the averaged, rate-based rules we consider can be respected by the real, spike-based system. We numerically obtain a condition on the magnitude of plasticity required for the rate-based mode of computation to dominate spike-to-spike interactions. We find that this threshold is in agreement with experimental data from spike-pair experiments. Finally, we perform a large-scale numerical simulation of ocular dominance column development. The stable, competitive dynamics present in our switch model are able to explain the formation of ODCs in a satisfactory manner. The observed learning behaviours therefore scale up successfully to larger scale systems.

In Chapter 9 we discuss the implications of our switch model of synaptic plasticity and the various issues surrounding it. We first discuss the interpretation of experimental data, and highlight some of the assumptions that are commonly made in theoretical studies. We then discuss emergent computation and the role of multi-spike interactions in our and other models of synaptic plasticity. We then discuss the issue of coincidence detection, and outline some possible molecular implementations of our synaptic switch rule. We then highlight some experimental tests that could distinguish the various forms of our synaptic switch rule as well as differentiate our model from competing models of STDP. We then place our model in context by comparing it to existing theories of STDP, and highlight the differences in our approach. Finally, we discuss various outstanding issues, and indicate the possible direction of future work.

Chapter 2

Experimental Basis of Neuronal Plasticity

2.1 Introduction

Experimental work has explored the induction and expression of activity-dependent neuronal plasticity in a range of species, brain areas, and at several stages of development. Different stimulation protocols can often give rise to changes that persist for different amounts of time. A distinction is therefore made between changes which rapidly decay (short-term plasticity) and those which appear to be relatively stable (long-term plasticity). Short-term plasticity is generally considered to be that which decays over the course of a few minutes. Long-term plasticity is generally considered to be that which is stable over the course of perhaps 30 minutes or more.

A distinction is also often made between early- and late-phase LTP. Early-phase LTP is generally considered to be protein synthesis-independent, and typically occurs on timescales of a few hours after the induction protocol. Late-phase LTP, on the other hand, requires protein synthesis, and can take place over several days following the induction of plasticity.

Although convenient, in the absence of a complete understanding of the mechanisms of induction and expression these distinctions remain somewhat arbitrary. Often, short and long term plasticity may be observed simultaneously in the same preparation. For example, sustained presynaptic stimulation giving rise to long-term depression is often preceded by an initial transient potentiation (an effect known as post-tetanic-potentiation, or PTP). In the plasticity experiments we describe, the focus is on long-term changes which typically appear on a time scale of a few seconds to a few minutes.

Due to its activity-dependence, stability, and wide-spread occurrence, long-term potentiation and depression (LTP and LTD) have often been linked to certain higher brain functions, such as learning and memory (for review, see Tsien (2000)). A variety of experimental evidence suggests such a link, such as the observation that infusion of the NMDA-receptor antagonist APV into the brain ventricles of mice (Morris et al., 1986), or the enhancement of LTP (and a concurrent decrease in LTD) by gene manipulation (Migaud et al., 1998), impairs performance in the Morris water maze test. However, given that the Morris water maze involves a variety of cognitive factors and that the experimental manipulation may have a wider influence throughout the brain, it is often difficult to interpret these results.

Recently, a great deal of experimental work has been directed towards uncovering the molecular and cellular basis of activity-dependent neuronal plasticity. The morphological and physiological complexity of a typical neuron provide an abundance of possible sites for the expression of plasticity. For example, short-term changes in postsynaptic responsiveness may be supported by a variety of physiological mechanisms, such as the depletion of readily-releasable neurotransmitter vesicles docked at the presynaptic membrane. Longer-term changes may involve up- or downregulation of postsynaptic receptor numbers or function. A distinction is typically made between physiological and anatomical changes. Physiological plasticity refers to a change in postsynaptic responsiveness due to the adaptation of an existing population of synapses, for example an up- or downregulation in the number of postsynaptic neurotransmitter receptors. Anatomical plasticity refers to processes where new synapses may be added or, alternatively, withdrawn. This may involve an increase or decrease the number of synaptic contacts made by a neuron onto an existing target, or the formation of new connections where none previously existed. Anatomical changes allow the actual connectivity between neurons, rather than just the ability of one neuron to influence the firing of another, to change.

2.2 Experimental Methods

A typical experimental protocol involves recording from an identified target cell using an intra- or extracellular electrode. Intracellular recording is more difficult, but allows much more detailed information about postsynaptic responses to be gathered by measuring actual postsynaptic currents rather than external field potentials. With an intracellular recording electrode the postsynaptic cell may also, if desired, be depolarised above its resting membrane potential by applying direct current injection. Sustained postsynaptic depolarisation is useful when

stimulation of presynaptic inputs alone is insufficient to cause postsynaptic spiking. An intracellular electrode can also be used in a feedback system to clamp the postsynaptic cell at a desired voltage. Voltage-clamping is most often used in experiments where the same stimulation protocol is shown to have different effects depending on the voltage at which the postsynaptic cell is clamped. A depolarising current injection can also be used to deliberately drive the neuron past threshold and trigger postsynaptic spiking on demand.

Presynaptic stimulation is typically achieved through simple extracellular current injection. This usually activates a population of afferents, which individually may or may not form functional connections with the target cell. Stimulating a population of afferents increases the probability of recording a postsynaptic response. However, by the same token it also means that the measured postsynaptic response is due to the population of stimulated afferents rather than a single afferent. Even more importantly, the plasticity observed in response to a particular stimulation protocol will be in the form of a population-averaged change, not of individual afferents. The alternative is to stimulate single presynaptic neurons in turn, attempting to find one that is synaptically coupled to the target cell. Identifying an individual presynaptic neuron, and impaling it in the same way as the postsynaptic cell, is more difficult but allows a more detailed examination of plasticity across a single connection. As individual neurons may innervate the target cell with a relatively low probability, finding a synaptically coupled pair of neurons by this method is often laborious.

The connection strength of inputs may be assessed by stimulating presynaptic neurons at a low frequency and examining the initial slope of the evoked postsynaptic field potential (PSP). The initial slope, rather than the peak amplitude, of the PSP is often used as it is the component most likely to stem from direct connections to the target cell. This is especially important in preparations where there are a large number of secondary pathways from the inputs to the target, via intermediate interneurons. In the literature, direct connections are termed "monosynaptic", in the sense that there is only one set of synapses between the pre- and postsynaptic cell. Here, we prefer the term "direct" to avoid confusion later when we explicitly consider the number of functional synapses comprising a connection as part of our own modelling studies.

A commonly used technique for studying synaptically coupled cells is to harvest neurons from an animal which are then cultured on an appropriate supporting substrate. Neurons cultured in this way can often survive for periods of 4 weeks or more. Culturing is a popular technique due to the easy access to the neurons that it affords. This facilitates recording and stimulation of neurons. However, neuronal cultures suffer from certain drawbacks. Neurons often swell up, and form fewer, higher strength connections than would be the case in vivo.

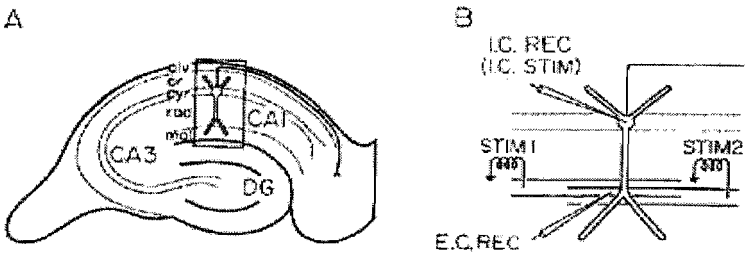

FIGURE 2.1: **A** A schematic diagram of a hippocampal slice. The CA3, CA1 (boxed), and dentate gyrus (DG) regions are shown, along with the characteristic layered structure of the hippocampus (alveus (alv), stratum oriens (or), stratum pyramidale (pyr), stratum radiatum (rad), stratum moleculare (mol)). **B** Diagram showing the arrangement of stimulating (STIM1 and STIM2), intracellular recording (I.C. REC) and extracellular (E.C. REC) electrodes. Such a set-up is typical of slice experiments, the extracellular stimulation of a population of presynaptic afferents increases the chances of finding one that is connected to the postsynaptic cell (from Gustafsson et al. (1987)).

Moreover, the connectivity that arises is random. Cultured networks of neurons are also prone to unnatural spontaneous activity patterns, such as epileptiform activity.

An alternative to culturing neurons is to prepare acute slices from the brain. These slices are typically a few hundred micrometres thick, and are stored in a specially prepared oxygenated salt solution with glucose before being transferred to a recording chamber. The recording chamber is constantly perfused with fresh maintenance solution. Acute slices allow the gross anatomical structure of the area under study to be quickly identified. Electrode placement and recording is then a relatively straight-forward process, and stable recordings may be made for several hours. A schematic diagram of a hippocampal slice, and a typical experimental set up, is shown in Fig. 2.1. The disadvantages of slice recordings are that the preparation is silent, with no background activity, and has undergone a massive loss of connectivity compared to its normal state in vivo. The absence of sustained, low-level activity can alter conductance levels and have important consequences for synaptic integration. The loss of connectivity means that each cell receives a fraction of the inputs it once did, and many of these inputs may have played modulatory roles. The extracellular environment is also very different to that in vivo, with the maintenance solution containing only the absolute minimum number of constituents. Finally, the slicing process itself may also traumatise the neurons under study.

A third alternative is to record directly from neurons in vivo. This is a much more difficult task than in culture or slice experiments. In deep recordings, individual neurons cannot be identified with certainty until after the experiment has

been concluded, the brain area stained, and the animal sacrificed and examined. Maintaining stable recordings can also be difficult, especially in an awake animal, with the electrodes having to be carefully mounted on the animal's skull. The natural expansion and contraction of the brain due to the circulation of blood can also cause recording electrodes to be unstable. This effect cannot be prevented without stopping the heart, killing the animal in the process. The main advantages of recording in vivo are that the neurons under study are relatively intact, with at least some level of background activity, and are supported by their normal extracellular environment. This is advantageous when performing plasticity experiments, as the neuronal circuits exist in a more natural environment. It is important, however, to remember that an applied stimulation protocols will be superimposed on the natural background activity and that there will be some interaction between the two. Indeed, the appearance of apparently non-input-specific LTD in adult rat hippocampal neurons may be attributable to the presence of natural activity patterns not present in slice or culture preparations (Doyere et al., 1997).

In the three methods described above, there are distinct advantages and disadvantages. The expression of plasticity may be influenced to some degree by the choice of experimental preparation. Results must therefore be carefully interpreted and placed within their experimental context.

2.3 Experimental Basis of Neuronal Plasticity

A great body of published work has documented the various forms of plasticity that may be expressed in a variety of experimental preparations. We briefly discuss some general experimental issues before moving on to consider individual experiments.

As briefly discussed in Chapter 1, there are periods during development when neuronal plasticity can be rather dramatic, such as during the refinement of ODCs (Hubel and Wiesel, 1962, 1965, 1969, 1977). The age of the animal, and its stage of development, may therefore have important consequences for the kinds of plasticity observed. Table 2.1 sets out approximate conventions for describing the age of a rat used in an experiment. Naturally, these definitions are slightly arbitrary but they remain useful for establishing the context within which a particular experimental result should be considered. Experiments are often performed on the hippocampus. This is partly due to the relative ease with which the hippocampus may be accessed, and partly due to the observation that plasticity may be readily induced across hippocampal synapses. The hippocampal formation is shown in Fig. 2.2.

Stage	Typical age range
Neonatal	6-20 days
Juvenile	11-35 days
Adult	5-10 weeks
Aged	20-24 months

TABLE 2.1: The typical conventions for describing the developmental stages of rats, and the approximate age ranges. Such conventions are, by necessity, slightly arbitrary in nature.

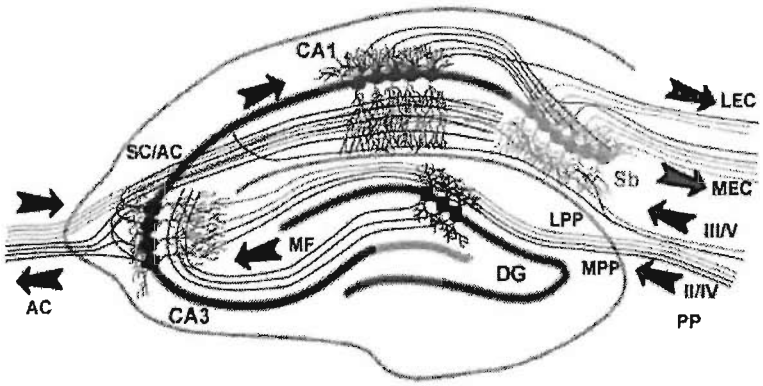

FIGURE 2.2: The hippocampal network. Information flow is principally unidirectional, with input from the entorhinal cortex (EC) via the medial- and lateral perforant paths (MPP and LPP, respectively), forming synapses with the CA3 and dentate gyrus (DG) regions. Mossy fibres (MF) connects the DG to CA3, and the CA3 in turn is connected to the CA1 pyramidal cells via the Schaffer collateral pathway (SC). Additional connectivity between CA3 and the contralateral CA1 region is via the associational commissural pathway (AC). The CA1 neurons also receive input directly from the perforant path, and make connections with the subiculum (Sb). These neurons in turn send axons back into the EC.

Early investigations into long-term plasticity invariably relied upon extracellular stimulation of groups of presynaptic afferents at various frequencies to evoke changes in the level of postsynaptic response. Initiating an extended train of pre- or postsynaptic action potentials is referred to as tetanic stimulation. High frequency tetanic stimulation (HFS) typically refers to stimulation at rates in the region of tens of Hertz, while low frequency tetanic stimulation (LFS) refers to rates of only a few Hertz. A number of experiments also use what is known as a "theta burst" stimulation pattern to induce neuronal plasticity. Such a protocol typically consists of several (10-15) bursts of stimulation, delivered at a low frequency (perhaps 5Hz). A burst will consists of a few (4-5) spikes at a high frequency (say, 100Hz).

In many cases, a variety of pharmacological manipulations are made once the induction of plasticity is reliably demonstrated. Typical treatments include the blocking of ionotropic N-methyl-D-aspartate-receptors (NMDA-receptors) using the NMDA-receptor antagonist AP-5 (or the more active form, D-AP-5). Another, less common manipulation, is the blocking of metabotropic NMDA-receptors (mGluRs) with alpha-methyl-4-carboxyphenylglycine (MCPG). Other treatments include the application of nimodipine to block L-type channels, or the injection of BAPTA or EGTA (fast and slow calcium buffers, respectively). The aim of these pharmacological treatments is to attempt to identify some of the subprocesses required for the induction and expression of the plasticity, perhaps shedding light on its underlying mechanism. For example, both NMDA-receptor-dependent and independent forms of LTD apparently coexist in the rat hippocampus (Oliet et al., 1997), and to observe one but not the other requires the use of an NMDA-receptor antagonist.

Although, in the majority of cases, plasticity is confined to the stimulated inputs, some experimental results suggest that plasticity may spread to other, unstimulated inputs synapsing onto the same target cell (Doyere et al., 1997). This spreading of plasticity to unstimulated inputs is commonly referred to as "heterosynaptic plasticity" in the literature. Here we prefer the terms input-specific (IS) and non-input-specific (NIS). This terminology explicitly acknowledges the fact that the plasticity observed in these experiments is due to changes in collections of input, rather than changes at individual synapses.

We now review a range of experimental results, and summarise their main findings.

2.3.1 Early LTP/LTD Experiments - High and Low Frequency Tetanic Stimulation

Early investigations into long-term plasticity tended to rely upon extracellular stimulation of groups of presynaptic afferents at various frequencies to evoke changes in the level of postsynaptic response. At the time, it was widely believed that information transfer in biological networks took place in a manner based purely on the firing rates on the participating neurons. The presynaptic afferents were therefore always stimulated at some determined frequency, without considering the timing of individual action potentials. Recently, it has been shown that the precise pattern of pre- and postsynaptic spiking, rather than the average rate of spiking, may govern the expression of plasticity, raising the possibility that many early results on LTP and LTD may need careful reinterpretation.

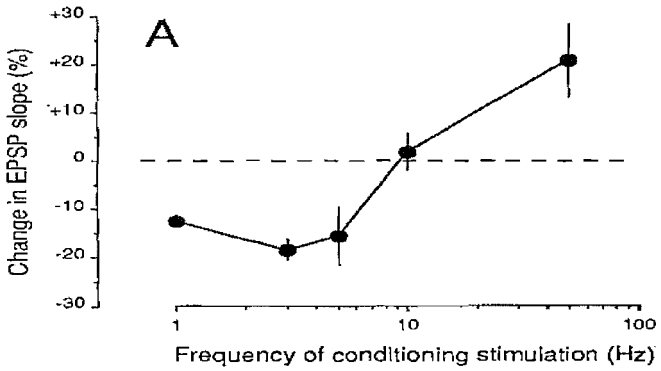

FIGURE 2.3: Input specific LTP and LTD of Schaffer collateral inputs to CA1 rat hippocampus in response to sustained presynaptic stimulation (900 pulses) at various frequencies. As the frequency of stimulation increases, the magnitude of the NMDA-receptor-dependent LTD is reduced, eventually reversing and becoming LTP as a threshold is passed. The data points were recorded 30 minutes after conditioning, with the number of samples, $n \geq 5$, for each point (from Dudek and Bear (1992)).

In the rat hippocampus, groups of Schaffer collateral inputs to CA1 reliably display NMDA-receptor-dependent LTD in an input specific manner, after an initial transient potentiation, by sustained low frequency presynaptic stimulation (Dudek and Bear, 1992). The frequency dependence of this form of plasticity is shown in Fig. 2.3. Increasing the frequency of stimulation saw the magnitude of LTD decrease, eventually reversing and becoming a LTP as a threshold frequency was passed. The crossover point, where depression changed into potentiation, fell at a presynaptic stimulation rate of around 10 Hz. A theta-burst stimulation pattern was shown to trigger LTP in the same connections (Dudek and Bear, 1993). Fig. 2.4 shows the time course of the recorded postsynaptic response to a series of alternating theta-burst and LFS stimulation protocols. The potentiation was reversible, and could be saturated by repeated application of the stimulation protocol.

The developmental dependency of HFS-LTP and LFS-LTD has also been examined. In rats, the magnitude of LTD declines with age (Dudek and Bear, 1992) and, in adult rats, LTD is apparently absent (Wagner and Alger, 1995) but may be uncovered by applying $GABA_A$ anatagonists (a treatment which had no effect on young LTD, possibly indicating a developmental change in the mechanism of plasticity). In aged rats, input specific NMDA-receptor-dependent LTD is once again expressed (Norris et al., 1996). LTP triggered by HFS is apparently uniform across all ages (Norris et al., 1996).

Other, apparently non NMDA-receptor-dependent forms of neuronal plasticity

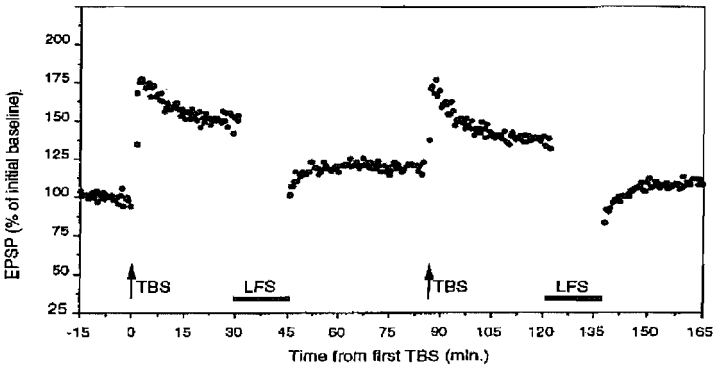

FIGURE 2.4: Time course of the recorded excitatory postsynaptic potential (EPSP) taken from a single experiment using Schaffer collateral inputs to CA1 in the rat hippocampus. Presynaptic stimulation was in the form of low-frequency tetani (LFS) or theta-burst stimulation (TBS). The theta-burst stimulation reliably triggers LTP, which may be reversed by subsequent LFS. Subsequent TBS will repotentiate the same connections, demonstrating that the plasticity was, in this case, reversible. The TBS-LTP could also be saturated by repeated application of the stimulation protocol (from Dudek and Bear (1993)).

have been observed, and appear to coexist alongside NMDA-receptor-dependent types in CA1 hippocampal pyramidal cells (Oliet et al., 1997). In neonatal and juvenile rats, blocking NMDA-receptors (via application of the antagonist D-AP5) gave rise to a long-lasting depression. This depression was, however, dependent on the activation of metabotropic glutamate receptors (mGluRs). This mGluR-LTD was input specific, most easily produced in small inputs, and required at least some postsynaptic depolarisation.

Experiments on hippocampal neurons have also been performed on adult rats in vivo (Doyere et al., 1997). Stimulation of the medial and lateral perforant pathways (MPP and LPP) in an awake animal showed that not all activity dependent changes occurred in an input-specific manner. As Fig. 2.5 shows, weak HFS of the LPP caused those connections to potentiate, a potentiation accompanied by a depression of the unstimulated MPP pathway. The MPP depression could be reversed by subsequent HFS, and there was no correlation between the magnitude of the LPP-LTP and the MPP-LTD. Unlike experiments performed using slices or cultured neurons, the patterns of activity used to induce plasticity were superimposed onto a natural level of background firing. The presence of background firing in experimentally unstimulated inputs may go some way towards explaining the non-input-specific nature of some of these results, as they are not truly "silent". As discussed above, neurons in vivo also have slightly altered conductance levels due to sustained background activity. This may result in small, but significant, differences in synaptic integration.

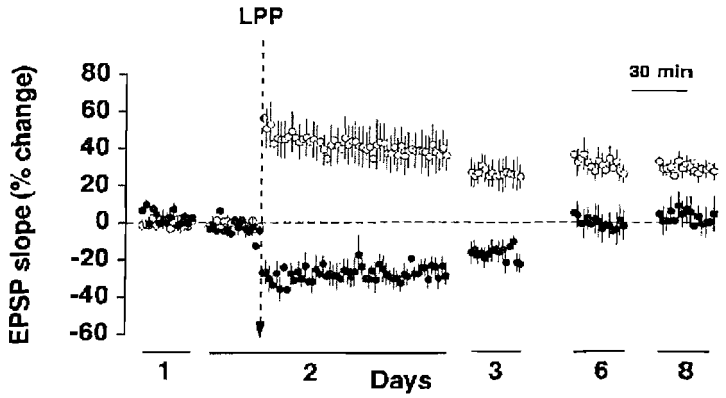

FIGURE 2.5: An in vivo plasticity experiment using adult rat hippocampal neurons. Sustained stimulation of the lateral perforant pathway (LPP, dashed line) results in a robust LTP of the recorded EPSP (open circles) accompanied by LTD of the unstimulated medial perforant pathway (MPP) EPSP (filled circles). Note the decay in amplitude of the remaining plasticity over the course of several days (from Doyere et al. (1997)).

2.3.2 Later LTP/LTD Experiments - Bursts of Pre- and Postsynaptic Action Potentials

Later experiments on LTP/LTD began to show that sustained presynaptic stimulation, in the form of a tetanus, was not always required to trigger afferent plasticity. Using shorter stimulation protocols, the first indications of a time window for interaction of pre- and postsynaptic events began seen in the hippocampus (Debanne et al., 1994), goldfish Mauthner cells (Yang and Faber, 1991), neocortex (Markram et al., 1997), and somatosensory cortex (Egger et al., 1999; Feldman, 2000). Long-term depression has also been shown to operate at inhibitory, GABA-ergic connections in the hippocampus using bursts of presynaptic action potentials (Fitzsimmonds et al., 1997).

In guinea pig hippocampal slices, pairing a presynaptic burst of action potentials with a short period of postsynaptic depolarisation can reliably trigger LTP of CA1 inputs (Gustafsson et al., 1987). The LTP was input specific, NMDA-receptor dependent, had a magnitude dependent on the number of postsynaptic action potentials triggered during the presynaptic burst, and could be occluded by prior tetanisation. Occlusion is often taken as evidence that two protocols evoke the same underlying mechanism, and, by extension, evoke the same type of LTP. The first indication of a time window for interaction of pre- and postsynaptic firing was also shown, with the burst of presynaptic action potentials having to precede the postsynaptic depolarisation by less than 400ms. Decreasing the separation time increased the magnitude of LTP, and the exponential-like dependence on the time difference of pre- and postsynaptic stimulation is shown

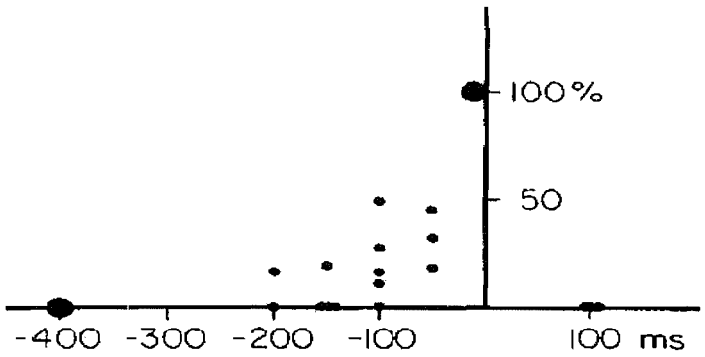

FIGURE 2.6: Timing-dependence of the LTP expressed by inputs to guinea pig CA1 hippocampal neurons. Pairing a presynaptic burst of action potentials with a short period of postsynaptic depolarisation can reliably trigger LTP of CA1 inputs, provided the burst of presynaptic action potentials precedes the postsynaptic depolarisation by less than 400ms. The horizontal axis represents the time delay between pre- and postsynaptic stimulation, with negative values representing the case where presynaptic stimulation precedes postsynaptic stimulation. The vertical axis is the percentage change in recorded EPSP (from Gustafsson et al. (1987)).

in Fig. 2.6. A second protocol, where the order of pairing was reversed, using a 100ms time difference was ineffective in inducing any plasticity.

The timing dependence of LTD initiation was later investigated in hippocampal slice cultures, where Schaffer collateral inputs to CA1 were shown to exhibit associative LTD from repeated LFS when coming 800ms after a 240ms postsynaptic depolarisation (Debanne et al., 1994). This depression was input specific, NMDA-receptor-dependent and could be reversed by subsequent induction of LTP. It could also be induced at previously potentiated connections. At time a difference of 2400ms there was no measurable change. When LTD did occur, the magnitude of depression depended on time between pairing and on duration of depolarising pulse. Reversing the order of pairing, and triggering LFS before the postsynaptic depolarisation, led to LTP. Thus, the same stimulation protocol could lead to potentiation or depression depending on the order of events. Similar protocols have also been shown to trigger LTD in inputs to goldfish Mauthner cells (Yang and Faber, 1991).

In neonatal rats, pyramidal cells in layer five (LV) of the neocortex showed an NMDA-receptor-dependent LTP in response to pairing bursts of four presynaptic action potentials with postsynaptic depolarisation and spiking (Markram et al., 1997). The magnitude of LTP showed a strong frequency dependence, shown in Fig. 2.7, with a sharp onset at 10Hz reaching saturation at 30Hz. LTP could

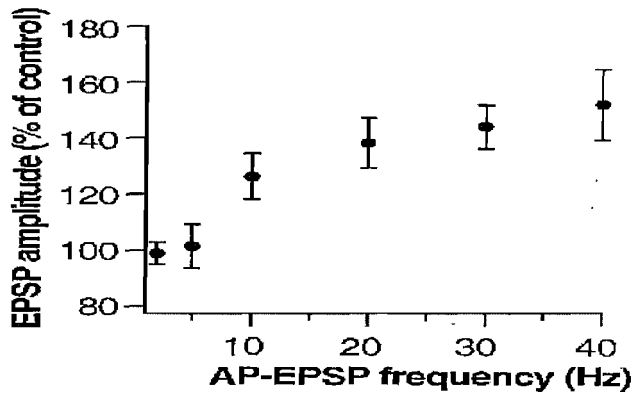

FIGURE 2.7: The effect of pairing bursts of four presynaptic action potentials with postsynaptic depolarisation and spiking in reciprocally coupled pyramidal cells from layer five (LV) of the neocortex of neonatal rats. The LTP induced by this protocol displays a noticeable frequency dependence, with a sharp onset of LTP at 10Hz and reaching saturation at 30Hz (from Markram et al. (1997)).

also be obtained by pairing single action potentials with postsynaptic depolarisation. It was the interplay of pre- and postsynaptic spiking that triggered these changes, as presynaptic bursts or postsynaptic depolarisation alone, or postsynaptic voltage clamping (-30 to -10mV), had no effect. An extension of this experiment using two bidirectionally connected cells showed that it was the timing of the pre- and postsynaptic action potentials that determined the direction of plasticity, and that spiking had to occur in a relatively narrow time window of around 100ms.

LTP and LTD have also been observed in thalamocortical slices from the somatosensory cortex of neonatal rats (Egger et al., 1999). The somatosensory cortex deals with tactile stimuli, with a characteristic barrel structure in L4 processing input from whisker sensory neurons. Simultaneous dual whole-cell recordings were made from spiny L4 neurons, and pre- and postsynaptic spike bursts were paired at different time delays. At delays of less than around 25ms, LTD was elicited in the L4 connections. Reversing the order of stimuli, with the postsynaptic burst occurring before the presynaptic burst, gave exactly the same result. The critical window for interaction of these spike bursts was therefore symmetric, consisting of a single LTD window spanning the origin. No LTP was observed for any timing difference, which is very different to the anti-symmetric curve observed in many other preparations (such as L4-L2/3 connections in S1 (Feldman, 2000), or connections between hippocampal neurons (Bi and Poo, 1998)). As the frequency of spiking within a burst changed, the magnitude of the LTD was altered, at 1Hz a much smaller LTD was observed compared to

10Hz or 20Hz. Pairing using 50Hz-bursts was ineffective in inducing plasticity. The LTD was NMDA-receptor independent, but did depend upon group II mGluRs. As in other experiments, neither pre- nor postsynaptic trains of action potentials were individually effective in triggering plasticity.

The induction of LTD at both glutamatergic and GABA-ergic synapses has been characterised using whole-cell recordings from cultures of embryonic rat hippocampal neurons (Fitzsimmonds et al., 1997). Presynaptic stimulation consisted of 1 second trains of action potentials at 5 Hz. When paired with sustained postsynaptic depolarisation to -50 mV the glutamatergic connections reduced in strength, as evidenced by a reduced EPSC. GABA-ergic connections responded to this protocol in a similar manner, with a persistent reduction in the inhibitory postsynaptic potential (IPSP). The LTD of glutamatergic inputs was abolished by the presence of AP-5, but the GABA-ergic LTD was unaffected. These results suggest that GABA-ergic LTD at hippocampal synapses is NMDA-receptor independent.

2.3.3 Timing-dependence in Neuronal Plasticity

Some of the first indications of timing-dependence of plasticity came from experiments using presynaptic bursts of action potentials (Gustafsson et al., 1987; Debanne et al., 1994), where the presynaptic burst had to arrive within around 400ms of postsynaptic depolarisation to evoke any change. A more detailed examination of this timing-dependence, using single pre- and postsynaptic spikes, soon followed. A critical window for the interaction of pre- and postsynaptic spikes has been observed in a variety of species, including the Mormyrid fish electro-sensory lobe (Bell et al., 1997), at xenopus laevis tadpole retino-tectal synapses (Zhang et al., 1998), between rat hippocampal pyramidal cells (Bi and Poo, 1998), and in the rat somatosensory cortex (Feldman, 2000). These timing experiments generally use pre- and postsynaptic current injections to trigger spiking, removing the need for sustained postsynaptic depolarisation.

One of the first timing experiments used medium ganglion cells from Mormyrid fish electrosensory lobe slices (Bell et al., 1997). These cells are Purkinje-like, receiving many inputs, and display a broad spike in response to depolarisation beyond a threshold. Presynaptic stimulation at fixed times (250ms or 500ms) was paired with a postsynaptic spike evoked at various timings. A window for the interaction of the pre- and postsynaptic spikes was seen, approximately 200ms in width. That is, if the pre- and postsynaptic spikes were separated by more $\pm$100ms no changes occurred. The most striking feature of this time window, shown in Fig. 2.8, was that negative spike timings, where the presynaptic spike followed the postsynaptic spike, gave rise to LTP while positive spike timings let

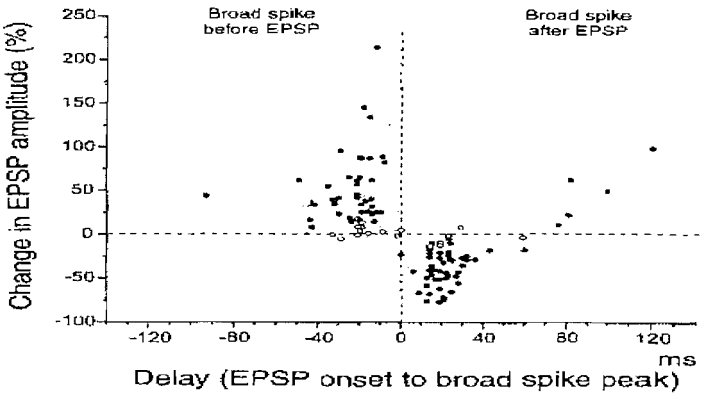

FIGURE 2.8: Percentage change in measured EPSP amplitude plotted against the delay between pre- and postsynaptic stimulation after repeated pairing of pre- and postsynaptic spikes in Purkinje-like medium ganglion cells from Mormyrid fish electrosensory lobe slices. A window for the interaction of the pre- and postsynaptic spikes is clearly seen, of approximately 200ms width. The window is of opposite polarity to similar spike timing windows seen in the hippocampus (Bi and Poo, 1998) and cerebellum (from Bell et al. (1997)).

to LTD. Interestingly this is of opposite polarity to similar spike timing windows seen elsewhere in the hippocampus (Bi and Poo, 1998). There is a sharp transition between the two phases near spike timings of 0ms, approximately 10ms in width. As observed in the majority of plasticity experiments, pre- or postsynaptic synaptic action potentials alone evoked no significant change, showing that it is the interaction of the two that drives changes in connection strength.

Very similar results have been produced using xenopus laevis tadpoles (stage 40-41), where whole cell perforated patch recordings were made from tectal cells in vivo (Zhang et al., 1998). Stimulation of the retina at two distinct sites 50-150 micro metres apart saw additive, convergent input to tectal cells, as is expected at this stage of development. Repeated stimulation of retina followed closely by tectal spiking led to input-specific, NMDA-receptor-dependent LTP. Again, pre- or postsynaptic spiking alone caused no change, and clamping the postsynaptic cell at -70mV blocked all changes. The level of LTP increased with the number of pairings, eventually reaching a saturation level, and was independent of the frequency of pairings or manner of stimulation (such as a theta-burst stimulation). The critical window for interaction of pre- and postsynaptic events is shown in Fig. 2.9. Here, the critical window was 200ms wide, with maximum LTP magnitude (occurring at small spike-time time differences) approximately twice that of the maximum LTD magnitude.

A similar critical window for interaction of pre- and postsynaptic events was

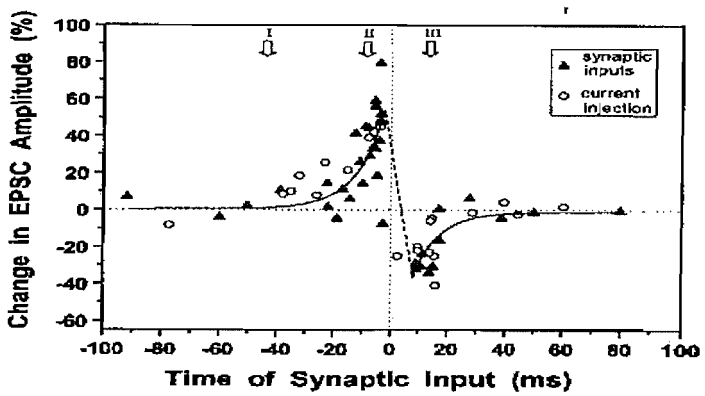

FIGURE 2.9: Plasticity induced by repeated pairing of pre- and postsynaptic spikes across xenopus laevis tadpole retino-tectal connections in vivo. The percentage change in EPSP amplitude, measured 10-30 minutes after conditioning, is shown against the delay between pre- and postsynaptic stimulation. The observed critical window for interaction is very similar to that observed in the hippocampus. Note that the definition of Δt used by the authors is the opposite of the usual convention. This reflects the curve in the y-axis compared to Fig.2.8 and Fig. 2.10 (from Zhang et al. (1998)).

also observed in the experiments of (Bi and Poo, 1998) on embryonic rat hippocampal cell cultures. Whole cell perforated patch recordings were made both the pre- and postsynaptic cells, and both were held in a current clamp to allow spiking to be initiated on demand. A reasonably precise investigation of the timing-dependence of plasticity across a single connection could therefore be made. Positively correlated spiking, where presynaptic spiking precedes postsynaptic spiking, led to LTP. Negatively correlated spiking, where the order is reversed, led to LTD. Both LTP and LTD were NMDA-receptor-dependent. Both positively and negatively correlated spiking were insufficient to elicit plasticity at glutamatergic presynaptic neurons synapsing onto GABA-ergic target cells, suggesting that this form of LTP/LTD is target specific. The potentiating protocol failed for strong inputs, while depression was independent of initial strength. The critical window for interaction is shown in Fig. 2.10. In this experiment, the window was $\pm$40ms in width with a narrow 5ms transition zone. Using nimodipine to block L-type channels, it was shown that this form of LTD required activation of dendritic L-type calcium channels, while LTP is only facilitated by, not dependent on, their activation.

Synapses from L4 to L2/3 neurons from the somatosensory cortex of juvenile rats also display a critical window for the interaction of pre- and postsynaptic spiking (Feldman, 2000). In this study, L4 afferents received extracellular stimulation and whole cell recordings form their target pyramidal cells were made. The connections were thus contained inside a particular column of a barrel. An

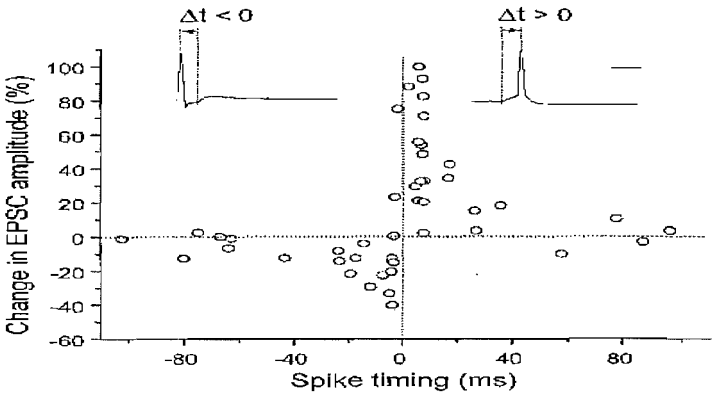

FIGURE 2.10: The critical window for induction of LTP and LTD using repeated pre- and postsynaptic spiking, for embryonic rat hippocampal cell cultures. The percentage change in measured EPSP is plotted against the spike-timing difference (from Bi and Poo (1998)).

asymmetric critical window, of the same polarity as seen in Bi and Poo, was observed but with a dramatically extended LTD phase (more than three times the size of the corresponding LTP phase). This suggests an overall dominance of LTD over LTP, a hypothesis confirmed when uncorrelated pre- and postsynaptic spiking was seen to evoke an overall depression of connections. The NMDA-receptor antagonist AP-5 blocked the original LTP/LTD, but in addition unmasked an NMDA-receptor-independent LTD window at positive spike timings. Throughout this study, inhibition was blocked using $GABA_A$ antagonists. Allowing inhibition did not affect the shape of the LTP/LTD window, but lowered the magnitude of both the potentiating and depressive phases.

2.4 Heterogeneity in Neuronal Plasticity

The complexity of activity-dependent changes is apparent from the variety of experimental results discussed above. Several further issues surround the interpretation of this data. For example, it has been suggested that the history of afferent activity can influence subsequent plasticity expression (Holland and Wagner, 1998), an effect known as meta-plasticity. Several studies have also examined the possibility that individual connections do not all follow the same plasticity rule, but rather form a heterogeneous population (Debanne et al., 1998, 1999). An experiment in which groups of afferents are stimulated will therefore probe only the "average" plasticity rule of that population of inputs. Such an average plasticity is not necessarily respected by individual inputs, but may emerge as a ensemble property of the system. It may also be the case that not every input may actually be capable of undergoing plasticity. More recently, studies of

plasticity across individual inputs have raised the possibility that inputs change their strength in an all–or–none manner (Petersen et al., 1998; O'Connor et al., 2005).

2.4.1 Metaplasticity

A particular stimulation protocol can sometimes have an effect on the plasticity produced by subsequent stimulation. This activity-driven change in the underlying plasticity rule is termed metaplasticity. The effect of prior stimulation on subsequent LTP/LTD induction has been explored in Schaffer collaterals in adult rat hippocapal slices (Holland and Wagner, 1998). The effect of priming was explored by stimulating afferents with a HFS in the presence of D-AP5. The D-AP5 ensured that the NMDA-receptor-dependent LTP that would typically follow such a HFS is suppressed. Priming had a clear effect on subsequent LTD induction, with an the primed input showing a greater degree of LTD compared to an unprimed input. This facilitation could also operate in a non-input-specific manner, where priming of input A caused a greater LFS-LTD in unprimed input, B. This suggests that the presence or absence of priming in other experiments may affect their results, and that meta-plasticity can operate in both input-specific and non-input-specific ways.

2.4.2 Heterogeneity of Individual Inputs

Single connections, rather than groups of afferents, have been studied in neonatal rat hippocampal slice cultures, where recordings were taken from CA3 pyramidal cells (Debanne et al., 1998). Pairing single presynaptic burst with a 240ms post depolarising current pulse, which allowed postsynaptic spiking, led to LTP as previously shown. This LTP is due to a mixture of mossy fibre and CA3 inputs. NMDA-receptor-dependent LTP of CA3-CA3 connections was shown using the same protocol by using sharp microelectrodes to impale two CA3 neurons. Input-specific, NMDA-receptor-dependent LTD of these CA3 pairs could be evoked using LFS. LTP and LTD could be evoked with single pre- and postsynaptic spike pairs, time differences of 0ms led to LTD and time differences of 15ms led to LTP. Pairing short postsynaptic bursts with single presynaptic action potentials 800ms later led to a large LTD, as did pairing a single postsynaptic spike with a single presynaptic spike at 15ms and 70ms, but not 200ms. The magnitude of change was similar in both cases, so the burst of postsynaptic spikes seems to have extended the critical window. In all cases, the changes appeared to be bidirectional with an LTP protocol able to reverse a previously applied LTD protocol.

In the same preparation, pairing single presynaptic action potentials with bursts of postsynaptic action potentials caused LTP, and tetanisation by LFS of presynaptic afferents caused LTD (Debanne et al., 1999). Not every connection displayed plasticity, but this was not due to a lack of postsynaptic NMDA-receptors. There was also some variety in the time course of plasticity, over half of unitary connections showed an initial transient LTP phase which then stabilised at a lower more persistent level. The remainder had no transient phase, and had a sustained LTP level of the same magnitude. A mixture of transient and non-transient connections were found to innervate the same target neuron, so such differences are apparently properties of the inputs themselves. The magnitude of LTP was very variable at unitary connections, ranging from 100% to 650%. The same protocol using groups of afferents triggered LTP ranging from 100% to 250%, reflecting the averaging taking place across plastic and non-plastic afferents. There was a significant inverse correlation between the initial size of the connection and the magnitude of change, with smaller strength connections being strengthened by a greater relative amount. The relationship between initial size and amplitude of plasticity was approximately exponential. LTP induction failed altogether at about a quarter of unitary connections in both CA3-CA1 and CA3-CA3, but in some cases this may have been due to previous saturation of LTP rather than those connections being truly nonplastic. There was no correlation between this failure and initial EPSP size, and some inputs which failed to show LTP still displayed transient potentiation. Both LTP and non-LTP capable connections were made onto same postsynaptic neuron, with the ability to express LTP apparently a property of the presynaptic cell. With double-monosynaptic connections (where single input forms two separate groups of contacts with target) it was possible to get LTP of the long latency response, but no change in the short latency response. This suggests, further, that not all synapses from a single input can display LTP.

2.4.3 All–or–None Potentiation

Physiological neuronal plasticity is generally considered to occur in a graded fashion, with more potentiation being evoked the longer the stimulation protocol continues. Recently, this view has been challenged by the observation that changes in input strength can occur in an all–or–none manner.

All–or–none potentiation has been investigated in hippocampal slices from neonatal rats (Petersen et al., 1998). Recordings were taken from patch clamped CA1 neurons, and Schaffer collaterals from CA3 were stimulated with an extracellular electrode placed in the stratum radiatum. The stimulation protocol consisted

of tetanic stimulation at 1Hz. Pairing this stimulation with postsynaptic depolarisation to 0mV saw a persistent increase in input strength. Repeating this protocol caused a further increase in input strength, as observed in other studies. To investigate the plasticity of single inputs, the magnitude of this stimulation was then reduced until the postsynaptic response disappeared. The magnitude of stimulation was then increased until a postsynaptic EPSC was recorded. This EPSC appeared in an abrupt fashion, and was therefore likely due to the activation of a single presynaptic input rather than a group of inputs. Using this minimal stimulation protocol, the result of pairing with sustained postsynaptic depolarisation to 0mV was explored. The pairing took place in two phases. In the first phase, the putative single input was stimulated 10 times then the change in EPSC assayed. In the second phase, the input received a further 100 pairings. 25 single inputs were explored in turn. The single inputs were either potentiated by the first 10 stimuli, the second 100 stimuli, or not at all. No single input was potentiated by both the 10 and 100 stimuli, and the level of potentiation for inputs potentiated by 10 stimuli was commensurate with the level experienced by inputs potentiated by 100 stimuli. A closer examination of the time course of the potentiation showed it to be step-like, occurring in as little as 10 seconds. Overall, recordings were maintained for 400 seconds. There was considerable variation in the number of stimuli that had to be delivered to a single input before this rapid potentiation occurred, with inputs taking from 3 to 15 stimuli before potentiation occurred. Thus, in contrast to the result of stimulating multiple inputs, the potentiation of this putative single input occurred in an all–or–none fashion. It would appear that such binary changes are not observed when a population of inputs are stimulated, as the individual variations are averaged to produce some total change which is apparently graded.

The ability of depression as well as potentiation to occur in a step-like manner has also been explored in neonatal rat hippocampal slices (O'Connor et al., 2005). Again, recordings were made from CA1 pyramidal cells, and presynaptic stimulation was achieved by placing an extracellular electrode in the stratum radiatum. A similar minimal stimulation protocol to that described above was used to stimulate putative single inputs. In this experiment, the response to changes in stimulation strength and the latency throughout the experiment were also used as indicators that a single input was being stimulated rather than a group of inputs. The LTP protocol consisted of presynaptic stimulation at 10Hz, with every 10th spike was paired with postsynaptic depolarisation to 0mV for 700ms, repeated 40 times. The LTD protocol consisted of the same 1Hz stimulation, but this time every 3rd spike was paired with postsynaptic depolarisation to -55mV, and the pairing was repeated 130 times. Naive inputs either potentiated or depressed depending on the protocol, or did not change. The time

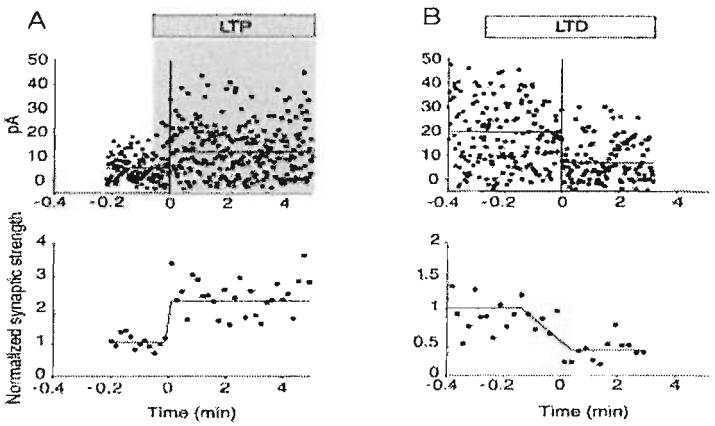

FIGURE 2.11: All–or–none plasticity at putative single inputs in the hippocampus. (Top) EPSC time series during **A** LTP and **B** LTD stimulation protocols. The vertical line shows the estimated onset time. (Bottom) same data as top, but binned into 10 data points per point. The fitted line is a best-fit ramp function, with an onset of approximately 12 seconds (from O'Connor et al. (2005)).

course for the induction of plasticity is shown in Fig. 2.11. Under both potentiation and depression protocols, once plasticity had been initiated a step-like change in input strength was observed, with changes appearing over the course of a few seconds. Initiation of plasticity was quite variable, sometimes taking many pairings to appear. On average, potentiation took 8 pairings to appear (a result consistent with Petersen et al. (1998)), and depression took on average 61 pairings. Recordings were maintained for an average of 10 minutes.

A comparison of inputs which could be potentiated and those which could be depressed showed that inputs existed in apparently two distinct states, of "low" or "high" strength. Inputs of high strength could not potentiate further, but could sometimes depress. Conversely, inputs of low strength could not depress but could sometimes be potentiated. According to this definition, in a naive slice approximately 71% inputs were of low strength, and 21% were of high strength. The value of low and high strengths was different for different inputs, but in all cases only two possible states were available. A further characterisation of the occupancy of these two states in naive slices was performed by comparing the extracellular baseline recordings of naive, saturated potentiated (via theta-burst stimulation) and saturated depressed (via LFS) inputs. The same ratio of approximately 2:1 in favour of low strength inputs was observed. Thus it appears that depression as well as potentiation can occur in a step-like, all–on–none manner. In addition, inputs may operate as binary switches, being in a either a state low strength or a state of high, but not in between. Under such a

paradigm, "potentiation" may be interpreted as the transition of inputs from the low-strength to high-strength state. Likewise, "depression" may be interpreted as the transition from high- to low-strength states.

These results are interesting, but, as always, must be carefully interpreted. In the case of Petersen et al., (1997) the pairing protocol involves clamping the postsynaptic cell to 0mV for a sustained period of time. Clamping at such pathological voltages may strongly influence on the operation of the target neuron and the plasticity it expressed. Although in O'Connor et al., (2005) sustained depolarisation on the scale of 10 or more seconds is avoided, the stimulation protocol still involves a window of depolarisation many times longer than a typical action potential rising phase. In addition, recordings in both experiments were only maintained for around 10 minutes, raising the possibility that the potentiation was transient and not, in fact, LTP. It should be noted, however, that the kinds of protocols used have been shown elsewhere to reliably give rise to long-term changes in input strengths.

2.4.4 Functional differences

Apparently similar forms of neuronal plasticity may, in fact, conceal important differences between plasticity in different brain areas. In neocortical L2/3 neurons, a robust increase in EPSP amplitudes is seen following an LTP protocol (Buonomano, 1999). However, the first 2-3 postsynaptic responses to a presynaptic train were potentiated more than later responses (an effect known as synaptic redistribution). That is, a pair of presynaptic pulses 100 ms apart evoke a similarly enhanced EPSPs, but at shorter intervals the first pulse was preferentially enhanced. Thus, LTP in the neocortex seems to preferentially strengthen early EPSP components. In contrast, this does not appear to be the case in the hippocampus (Buonomano, 1999; Selig et al., 1999). In neonatal rat hippocampal slices, the effect of LTP on the postsynaptic responses to a train of presynaptic action potentials at Schaffer collaterals was examined (Selig et al., 1999). After LTP, all postsynaptic responses were increased by an equal amount, and a similar result was obtained for the induction of LTD with all responses decreasing by an equal amount. Synaptically coupled CA3-CA3 neurons also displayed this uniform increase in postsynaptic response. These differences may reflect important functional differences between different brain areas.

2.5 Summary

Experimental work has explored several types of activity-dependent synaptic plasticity in a range of species, brain areas, and at several stages of development. These experiments are technically very impressive, and have provided a rich library of data upon which theoretical studies can draw. However, many results are often open to at least some degree of interpretation. This is partly due to experimental limitations, and partly due to the complexity of the system under study. For example, it is often the case that the neurons under study are capable of more than one form of plasticity, so that the protocol chosen to elicit measurable changes itself influences the observed result. Experimental results are also often subject to a high level of noise, some of which may be unavoidable (such as that arising from the unreliability of synaptic transmission). Apparently uncontroversial technical details, such as the choice of perfusion medium in slice experiments, can also influence the results. Caution must therefore be taken when attempting to generalise data drawn from different experimental preparations.

A theoretically relevant issue is that the kind of neuronal plasticity observed in experiment may be unrepresentative of that operating in a more natural setting. For example, given the apparent sensitivity of the system under study, and the heterogeneity of the results discussed above, it would be unsurprising if the choice of stimulating protocol, and the environment within which it is administered, had important consequences on the form of plasticity observed. That a plasticity experiment is being performed at all means that the neuronal circuit under study has been disrupted in some manner. Indeed, it is possible that a large proportion of experimental scenarios are rarely experienced under "normal" circumstances, raising the possibility that experimenters are exploring a drastically restricted range of neuronal behaviour. A similar, and equally compelling, argument applies to the various pharmacological manipulations that are routinely undertaken. Experimental results on neuronal plasticity must therefore be taken within their experimental context, being put forth as examples of the kinds of behaviour that a coupled pair of neurons may engage in, rather than an accurate portrayal of their normal operation.

In the context of spike-timing results, although there exists compelling evidence for the importance of individual spikes in driving neuronal plasticity, it is important to acknowledge the inherent limitations of the experimental data and carefully interpret those findings. That initiating an isolated pre- and postsynaptic spike pair leads to the modification of input weights is now an accepted experimental fact. However, from a theoretical point of view, this does not necessarily mean that a spike-based description is the correct level at which to

model the phenomenon. For example, although spike-to-spike interactions may be the basic currency of certain plasticity interactions, some averaging of these interactions will necessarily occur over longer periods of time. The essential phenomenology of the plasticity, in the form of the learning rule that the interactions encode, may therefore be more correctly described at a rate-based level. Thus, although spike-timing plasticity results may "real", in the sense that they are robust and reproducible, a spike-based description is not necessarily required, and in some cases may even be undesirable. We will return to these issues in Chapter 6.

Chapter 3

Rate-based Models of Neuronal Plasticity

In reviewing models of neuronal plasticity, we find it useful to make a distinction between rate-based plasticity rules, where the activity of pre- and postsynaptic neurons is described by the recent time-averaged firing rate, and spike-based plasticity rules, which involve consideration of individual pre- and postsynaptic spiking events. In this Chapter we will discuss rate-based models. In Chapter 4 we extend our discussion to spike-based models.

3.1 Rate-based Learning Rules

Rate-based learning rules take the form of a differential equation describing the rate of change of the input weight vector, $\mathbf{w}(t)$, as a function of the pre- and postsynaptic firing rates. The activity of each neuron is described by a continuous variable. The input vector is often denoted $\mathbf{u}(t)$, and the postsynaptic rate $v(t)$. To simplify the analysis of such models, the firing rates are sometimes defined relative to some background level, and may therefore take negative values. Firing rates are also sometimes expressed as the ratio of this relative firing rate to the average firing rate, in which case they are also dimensionless quantities.

A simple description of the postsynaptic firing rate, v, as a function of the postsynaptic membrane potential, v_m, is

$$v = g(v_m), \tag{3.1}$$

where $g(v_m)$ is some monotonically increasing function. The monotonic character of $g(v_m)$ encodes our expectation that higher levels of postsynaptic depolarisation

leads to higher levels of postsynaptic firing. In the simplest case, $g(v_m) = v_m$. The membrane potential, v_m, may in turn be calculated from the input weights and input activity levels. In most cases, the simple linear sum

$$v_m = \sum_i w_i u_i, \tag{3.2}$$

is an adequate approximation. Thus, we have that

$$v = \mathbf{w}(t).\mathbf{u}(t), \tag{3.3}$$

and the postsynaptic neuron therefore acts as a simple, linear integrator of its input where the postsynaptic response is approximated as the simple sum of the input weights multiplied by their instantaneous firing rate. The vector nature of $\mathbf{u}(t)$ and $\mathbf{w}(t)$ reflects the fact that multiple presynaptic inputs typically converge onto a single postsynaptic target. Although more complicated descriptions may be constructed, this minimalist approach adequately captures the basic features of postsynaptic firing in an analytically tractable manner.

3.2 Hebbian Learning

An early theme in the study of activity-dependent changes in neuronal connectivity was the idea that changes in input weights must be driven by correlations between pre- and postsynaptic firing. The basis of much research into correlation-based plasticity rules has been the Hebb rule (Hebb, 1949). The rule can be paraphrased by stating "If input neuron A regularly contributes to the firing of target neuron B, the connection between the two should be strengthened". Such a learning rule was suggested by Hebb in 1949 to be able to drive the formation of "neuronal assemblies", collections of inter-connected neurons having similar function. At the time, it was a purely theoretical construct, but experimental work has since shown that biological mechanisms exist that are capable of supporting such processes. Hebbian learning is correlation-based in the sense that input A must contribute in some way to the firing of target B.

We make may make a general formulation of a Hebbian learning rule, then extract specific examples from this formulation. We consider local learning rules, so that we restrict the dependence of $dw_i(t)/dt$ to the local variables $w_i(t)$, $u_i(t)$ and $v(t)$. We therefore have that

$$\frac{dw_i}{dt} = G(w_i, u_i, v) \tag{3.4}$$

where G is some unknown function, and we have suppressed the time dependences of w_i, u_i and v for clarity (a convention we follow throughout this discussion). Note that it is possible to incorporate a dependence on the history of the weights in Eqn.3.4 but we do not consider such cases here.

Expanding G around $u_i = v = 0$, we have that

$$\frac{dw_i}{dt} = A(w_i)+B(w_i)v+C(w_i)u_i+D(w_i)vu_i+E(w_i)v^2+F(w_i)u_i^2+O(v^3). \quad (3.5)$$

where $O(v^3)$ represent terms of order 3 or higher. We view this expansion not as some approximation of dw_i/dt but rather as a way of neatly encoding a variety of rate–based learning rules. Thus, our choice to discard higher order terms is is not a statement that these higher order terms are small, but simply that terms of order 3 or higher do not interest us in the class of learning rules we consider. We therefore concentrate on the lower order terms A to F, and simply note that other, more complicated learning rules may be constructed by continuing this expansion.

The coefficients A to F depend on the input weight, w_i, as only u_i and v have been expanded. $A(w_i)$ is a constant term, and does not depend on either pre- or postsynaptic firing. Such a term might be a decay term, proportional to w_i, perhaps included to prevent uncontrolled growth of input weights. $B(w_i)$ and $C(w_i)$ are multiplied by v and u_i, respectively. That is, they represent changes induced by pre- or postsynaptic firing alone. Such changes are usually considered to be non-Hebbian, as they do not require correlations between pre- and postsynaptic firing. $E(w_i)$ and $F(w_i)$ are similarly non-Hebbian, but are multiplied by the square of pre- or postsynaptic firing. $D(w_i)$ is the cross term, which picks up factors of both u_i and v. This term is therefore the most interesting as it describes changes that depend on both pre- and postsynaptic firing.

3.2.1 Basic Hebb-rule

The basic Hebb-rule can be extracted from Eqn. 3.5 by setting $D(w_i) > 0$ and all other coefficients to zero. This produces the rule

$$\frac{dw_i}{dt} = D(w_i)vu_i, \quad (3.6)$$

Setting $D(w_i) = 1/\tau_w \forall i$ for simplicity, the rule may be written in vector form as

$$\tau_w \frac{d\mathbf{w}}{dt} = v\mathbf{u}, \quad (3.7)$$

where τ_w is the learning time constant. The corresponding anti-Hebbian rule may be produced by setting $D(w_i) = -1/\tau_w \forall i$. It can be seen intuitively that, under the basic rate-based Hebb-rule, inputs will strengthen when the product of the pre- and postsynaptic firing rates is high. However, as the rule only describes potentiation, it always leads to uncontrolled growth of synaptic weights. This can readily be seen by examining the derivative of the square of the weight vector, $|\mathbf{w}|^2 = \mathbf{w}.\mathbf{w}$, which gives

$$2\tau_w \mathbf{w}.\frac{d\mathbf{w}}{dt} = v\mathbf{w}.\mathbf{u} \tag{3.8}$$

In the simple case where we set $v = \mathbf{w}.\mathbf{u}$, this becomes $d|\mathbf{w}|^2/dt \propto v^2$, which is always non-negative. Thus, inputs will continue to strengthen under the basic Hebb-rule without bound, a situation commonly referred to as runaway learning.

3.2.2 Bounding

To avoid uncontrolled increase in synaptic strength, upper and lower bounds on input weights may be introduced to the simple Hebb-rule. A "soft" bound may be implemented in the form of a saturating weight dependence in the coefficients of Eqn. 3.5. For example, setting $D(w_i) = D(w_{max} - w_i)$ reduces the level of potentiation and an input weight approaches the upper bound w_{max}. Inputs will therefore not potentiated past w_{max}. A "hard" bound may be implemented by replacing $(w_{max} - w_i)$ with the Heaviside step function $H(w_{max} - w_i)$. Under this formulation, inputs are potentiated normally until they reach the hard upper bound then stop evolving. Numerically, such a bound is usually implemented by forcing any input weight that crosses the bound to become equal to the limiting value.

Lower bounds may be implemented by modifying any negative coefficients in Eqn. 3.5 in a similar manner. Typically, it is common to limit the elements of $\mathbf{w}$ to be greater than zero to reflect the fact that inhibitory and excitatory connections are biologically distinct and one cannot somehow turn into the other.

Although limiting connection strengths by imposing a hard or soft bounds at some upper limit, w_{max}, prevents runaway learning, the dynamics of the underlying rule are unchanged. For example, under the simple Hebb-rule we still have that $d|\mathbf{w}|^2/dt \propto v^2 \geq 0$. The input weights will therefore eventually become saturated, with every input weights being driven to the upper bound, w_{max}.

3.2.3 Pre- and Postsynaptically Gated Rules

The tendency of the simple Hebb-rule to produce uncontrolled potentiation stems from the fact that the rule does not contain any mechanism by which inputs may weaken. It is desirable, therefore, to introduce some form of depression into the Hebb-rule. This depression must not, however, destroy the potentiation when both pre- and postsynaptic firing rates are high. This may be achieved in a number of ways. A simple weight decay term may be introduced by setting $A(w_i) = -Aw_i$, where A is some constant, in Eqn. 3.5. This prevents uncontrolled potentiation by pulling back down those inputs which grow too large. Alternatively, we may impose thresholds on pre- or postsynaptic firing rates below which depression is induced rather than potentiation. These rules are known as pre- or postsynaptically gated rules, depending on which variable acts as the controller. The postsynaptic form may be formulated setting $D(w_i) = 1/\tau_w$ and $C(w_i) = -\theta_v \tau_w$ in Eqn. 3.5 to give, in vector form,

$$\tau_w \frac{d\mathbf{w}}{dt} = (v - \theta_v)\mathbf{u} \tag{3.9}$$

where θ_v is the level of postsynaptic firing above which depression switches to potentiation. A presynaptically gated rule may be produced by setting $D(w_i) = 1/\tau_w$ and $B(w_i) = -\theta_{\mathbf{u}}\tau_w$ in a similar manner, giving

$$\tau_w \frac{d\mathbf{w}}{dt} = v(\mathbf{u} - \theta_{\mathbf{u}}) \tag{3.10}$$

where $\theta_{\mathbf{u}}$ is the presynaptic threshold that determines when an input begins to be strengthened instead of weakened.

Imposing the threshold on postsynaptic firing means that inputs will be modified only if they have non zero activity (an input-specific change), with the direction of the modification is dependent on the level of postsynaptic activity. With a presynaptic threshold, a non-zero postsynaptic activity will weaken even inactive inputs, and the changes are therefore termed non-input-specific. Imposing both thresholds simultaneously leads to the undesirable situation where low pre- and postsynaptic firing rates leads to potentiation, so in most cases a single threshold is used.

3.3 The Covariance Rule

Setting the threshold in either the presynaptically gated or postsynaptically gated rules to the long-time mean firing rate, so that $\theta_v = \langle v \rangle$ or $\theta_{\mathbf{u}} = \langle \mathbf{u} \rangle$,

respectively, and averaging over all training inputs gives a plasticity rule of the form

$$\tau_w \frac{d\mathbf{w}}{dt} = \mathbf{C}\mathbf{w} \tag{3.11}$$

where $C = \langle \mathbf{u}(\mathbf{u}^T - \langle \mathbf{u}^T \rangle) \rangle = \langle \mathbf{u}\mathbf{u}^T \rangle - \langle \mathbf{u} \rangle \langle \mathbf{u}^T \rangle$. Thus, C is the input covariance matrix, $C_{i,j} = \langle u_i u_j \rangle - \langle u_i \rangle \langle u_j \rangle = Cov(u_i, u_j)$. The rule is therefore known as the covariance rule (Sejnowski, 1977). In the case of a postsynaptic threshold, this requires that θ_v change over time, as v is a function of the input strengths and firing rates.

Despite having the same learning rule on average, the choice of threshold has important consequences. Imposing the threshold on postsynaptic firing means that inputs will be modified only if they have non zero activity (an input-specific change), with the direction of the modification is dependent on the level of postsynaptic activity. With a presynaptic threshold, a non-zero postsynaptic activity will weaken even inactive inputs, and the changes are therefore termed non-input-specific.

Covariance rules suffer from the same instabilities as a simple Hebb-rule, despite allowing depression as well as potentiation, due to positive feedback. This can be seen by again examining the rate of change of the square of the input weight vector, $d|\mathbf{w}|^2/dt = 2v(v - \langle v \rangle)$, which is proportional, on average, to the variance of the output $\langle v^2 \rangle - \langle v \rangle^2$. This quantity is always positive, and the weights always grow on average, except in the trivial case when v is constant. In addition, neither of the covariance rules are competitive (see section 3.5).

3.4 The Bienenstock-Cooper-Munro Rule

The postsynaptic covariance rule does not require postsynaptic activity to drive changes in input strengths, only that it is less than the threshold value. Experimentally, pre- or postsynaptic spiking alone is almost always insufficient to induce long-term plasticity (but see Reiter and Stryker (1988)). This observation is incorporated in an alternative learning rule proposed by Bienenstock, Cooper and Munro (Bienenstock et al., 1982). The Bienenstock-Cooper-Munro (BCM) rule, which explicitly requires non-zero pre- and postsynaptic activity levels, takes the form

$$\tau_w \frac{d\mathbf{w}}{dt} = v\mathbf{u}(v - \theta_v) \tag{3.12}$$

where θ_v is the threshold on the postsynaptic firing rate above which potentiation occurs instead of depression. The rule is Hebbian, in the respect that it is the rates of pre- and postsynaptic firing determine if an input is potentiated or depressed. At low presynaptic firing rates inputs tend to be weakened, while at high presynaptic rates inputs tend to be strengthened. A threshold determines the crossover point above which depression turns into potentiation.

In an earlier work (Cooper et al., 1979), a rule was proposed in which the threshold θ_v was fixed. Analysis of this learning rule shows that the dynamics are governed by a fixed-plane where $dw/dt = 0$ (Castellani et al., 1999). This fixed-plane is unstable, so that perturbations about this plane grow with time. In terms of afferent strengths, the phase space is divided into two regions. Above the fixed plane afferents are strengthened continually, and runaway learning is observed. Below the fixed-plane, afferents are continually weakened and fall to zero. Thus, this form of the rule leads to uncontrolled learning in the same way as the simple Hebb-rule.

The BCM-rule introduced the novel idea of allowing the threshold between potentiation and depression to change over time. In the BCM formulation, this threshold is a function of the recent time average of the postsynaptic firing rate. This stabilises the inputs and prevents unbounded growth by allowing the location of the fixed-point to change (Bienenstock et al., 1982; Castellani et al., 1999). It can intuitively be seen that the critical condition is that the threshold grows faster than v as the postsynaptic firing rate becomes large. In one implementation, this is achieved by allowing θ_v to act as a low-pass filtered version of v^2, so that

$$\tau_\theta \frac{d\theta_v}{dt} = v^2 - \theta_v \tag{3.13}$$

where τ_θ sets the time scale for modification of the threshold. The time constant τ_θ must be less τ_w so that θ_v changes faster than v. Fig. 3.1 shows how the learning function of the BCM-rule changes as a function of the recent-time-average of the postsynaptic firing rate. The threshold between potentiation and depression slides to adjust the curve in a dynamic manner.

A sliding threshold also makes the BCM-rule competitive, because a strengthening of some inputs raises the output firing rate which adjusts the potentiation/depression threshold so that other inputs are more likely to weaken. As the postsynaptic firing rate is a function of the input weights, the BCM-rule with its sliding threshold may be considered to be a form of input normalisation.

The BCM-rule has been shown, in simulation, to account for the normal and deprived development of orientation selection and ocular dominance in the cat

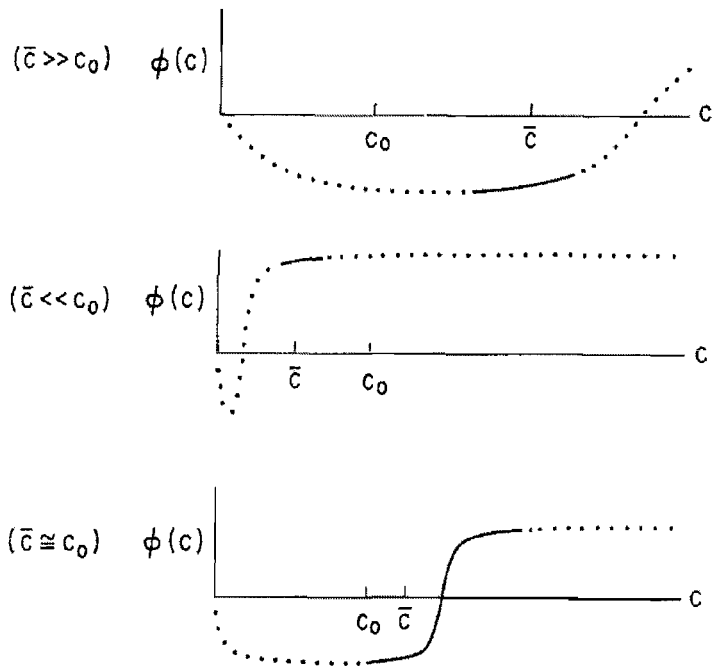

FIGURE 3.1: The learning function, $\phi(c, \bar{c})$, of the Bienenstock-Cooper-Munro rule as a function of postsynaptic firing rate, c, and its recent time average, $\bar{c}$. The threshold between potentiation and depression, θ_m, slides according to difference between the postsynaptic firing rate and some positive constant, c_0. In the top diagram, $\bar{c} >> c_0$, so that postsynaptic firing is much higher than the target level. The learning function is mainly negative, which serves to pull input weights back down and lower the average postsynaptic firing rate. The solid part of the curve is the region around $\bar{c}$, which, from a dynamical point of view, is most relevant. In the central diagram, $\bar{c} << c_0$, so that postsynaptic firing has fallen to a low level. The learning function is mainly positive, particularly the region around $\bar{c}$. Thus, inputs tend to be potentiated and the postsynaptic firing rate will increase. In the lower diagram, $\bar{c} \sim c_0$, and the learning function achieves a rough balance between potentiation and depression. The critical factor in the formulation of the learning function, ϕ, is the nonlinearity that allows the threshold between potentiation and depression to increase or decrease faster than the recent time average of the postsynaptic firing rate, $\bar{c}$ (from Bienenstock et al. (1982)).

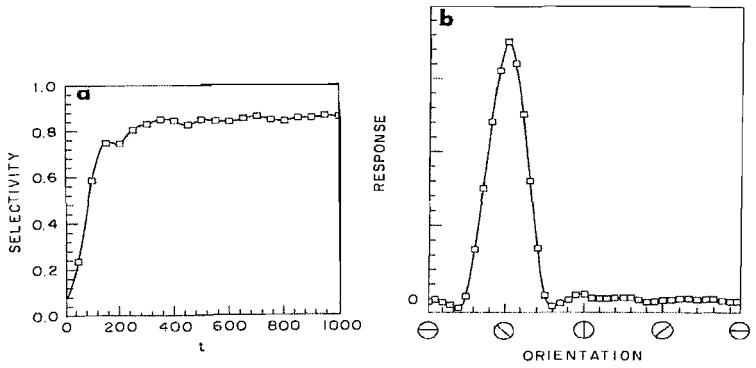

FIGURE 3.2: Evolution of orientation selectivity under the BCM-rule. (**A**) shows the buildup of selectivity in a system presented with circularly symmetrical stimuli. This environment corresponds to orientated bars of light, similar to that used in experimental work. The system starts of in a state of low selectivity, then progressively becomes more selective as time goes on. (**B**) shows the final "tuning curve" for the system. The high level of selectivity indicated in the system is reflected in the large response of the target cell to one particular orientation of stimuli to the exclusion of all others (from Bienenstock et al. (1982)).

primary visual cortex (Bienenstock et al., 1982). Fig. 3.2 shows the development of orientation selectivity under the BCM-rule. Cells become tuned to one particular orientation at the expense of all others (in experimental work, this stimulus is typically an orientated bar of light). The model also reproduces experimental results with restricted monocular input, where only stimuli with a certain orientation are presented. In this case, all visually responsive cells become tuned to the experienced stimulus, or one very close to it.

If input from two eyes is simulated, the BCM-rule can also reproduce various findings on the development and plasticity of ocular dominance. In a simulated normal environment, cells become orientation selective and binocularly driven, with the same orientation selectivity for each eye (Fig. 3.3A). If the input from both eyes is suppressed (binocular deprivation BD, or dark rearing), the cells lose all orientation selectivity, but remain driven by both eyes (Fig. 3.3B). If the input from one eye is suppressed but not the other (monocular deprivation, MD), cells will become orientation selective and monocularly driven, regardless of their initial state (Fig. 3.3C).

3.5 Competition

The situation where all input weights are driven to some upper bound is undesirable for theoretical reasons, as it limits the amount of information that can

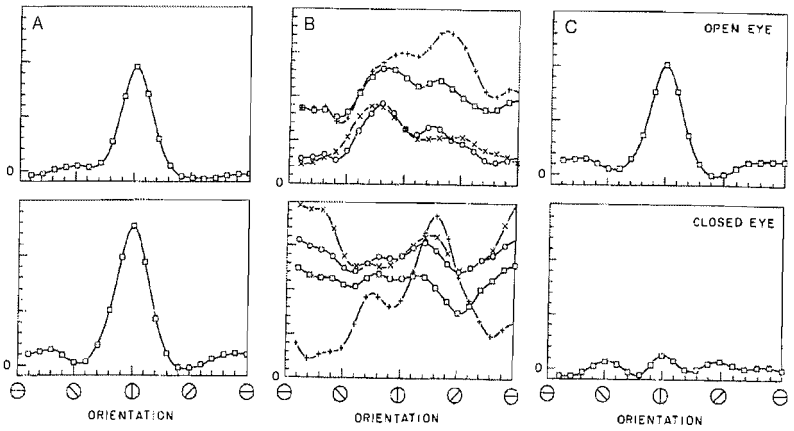

FIGURE 3.3: Simulations of various rearing conditions and their effect on ocular dominance development under the BCM-rule. Upper and lower panels show the response of the simulated target cell to stimuli from the two eyes. (**A**) normal rearing conditions lead to response that is binocular and orientation tuned. (**B**) dark rearing leads to a randomly fluctuating response curve, so that there is no orientation tuning. The cell is, on average, driven binocularly. (**C**) monocular deprivation leads to a monocular, orientation tuned response that completely favours the open eye. (from Bienenstock et al. (1982)).

be stored in the network. To prevent the saturation of input weights, it is useful to introduce an element of competition into the dynamics of a Hebbian learning rule. A learning-rule is said to be competitive if the strengthening of one input leads to the weakening of others.

Competition may be introduced into a Hebbian learning rule in a number of ways. The BCM-rule introduces competition by means of a sliding threshold, which makes it harder for inputs to potentiate if other inputs are already strong. A more direct way of introducing competition is to explicitly add non-Hebbian global terms that depend on the input weights. This typically leads to some form of input-normalisation. Normalisation may be interpreted as the statement that a neuron can support only a fixed total input weight. There are two common choices for the global constraint on total input weights, either a constraint on their sum or a constraint on the sum of their squares. A normalisation constraint may be imposed rigidly at all times, or just once at the end of a period of plasticity.

Input normalisation can drastically alter the outcome of a learning procedure, and different normalisation methods may lead to different outcomes. For example, subtractive normalisation tends to produce afferent distributions where a subset of maximally correlated inputs are saturated to the maximum allowed strength, with all other weights falling to zero (Miller and Mackay, 1994). In contrast to this "sharpening" effect, multiplicative normalisation has been shown to

generate competitive dynamics that lead to "graded" receptive fields, where most mutually correlated inputs are represented (Miller and Mackay, 1994). In both cases, competitive dynamics can be produced capable of segregating afferents, even when those afferents are positively correlated (Elliott, 2003).

In some cases, the mechanism used to introduce competition also serves to stabilise the growth of input weights, but under other plasticity rules, further saturation constraints must be implemented. Stability is generally a desirable feature for a learning rule. With unstable dynamics, any useful information that is learned will only be stored transiently.

3.5.1 Subtractive Normalisation

In this scenario the simple sum of input weights, $\sum_i w_i = \mathbf{n}.\mathbf{w}$ where $\mathbf{n}$ is an N dimensional unit vector, is constrained by adding an extra term to the simple Hebb rule

$$\tau_w \frac{d\mathbf{w}}{dt} = v\mathbf{u} - \frac{v(\mathbf{n}.\mathbf{u})\mathbf{n}}{N} \tag{3.14}$$

where N is the dimensionality of the input. This is called subtractive normalisation as the same amount is subtracted from each input, regardless of its weight. This rigid normalisation keeps the quantity $\mathbf{n}.\mathbf{w}$ fixed. This can be seen by taking the dot product of Eqn. 3.14 with $\mathbf{n}$ to give

$$\tau_w \frac{d\mathbf{n}.\mathbf{w}}{dt} = v\mathbf{n}.\mathbf{u}(1 - \mathbf{n}.\mathbf{n}/N) = 0 \tag{3.15}$$

as $\mathbf{n}.\mathbf{n} = N$. This result can be seen intuitively as the second term in Eqn. 3.14 subtracts the average change in the inputs from each individual input, so that the total change is zero.

Such a rule requires non-local information, in the form of the vector of all input activities, to be available at every input. It is unclear how such non-local information could be acquired in a biological setting, and as such subtractive normalisation is generally considered to be less biologically plausible than other methods of normalisation. A hard lower bound is often imposed in conjunction with subtractive normalisation to prevent weights from becoming negative. A hard upper bound is also often used, to avoid the otherwise common situation where all weights except one are driven to zero. This maximum weight could simply reflect the maximum connection strength sustainable by an individual input. Competition in very strong under a simple Hebb-rule with subtractive normalisation, as weak inputs are depressed by a larger proportion of their size than stronger ones.

3.5.2 Multiplicative Normalisation

Another way of constraining a simple Hebb rule is through multiplicative normalisation. Under multiplicative normalisation, input weights are scaled in a manner proportional to their current strength. Applying this normalisation to the sum of the squares of input weights gives the Oja rule (Oja, 1982), which may be written as

$$\tau_w \frac{d\mathbf{w}}{dt} = v\mathbf{u} - \alpha v^2 \mathbf{w} \tag{3.16}$$

where α is a positive constant. This rule involves pre- and postsynaptic firing rates and the current input weight, all information that is available locally at each input. The stability of the Oja rule can be examined by taking the dot product of this equation with the weight vector w to give

$$\tau_w \frac{d|\mathbf{w}|^2}{dt} = 2v^2(1 - \alpha|\mathbf{w}|^2) \tag{3.17}$$

which indicates that $|\mathbf{w}|^2 \rightarrow 1/\alpha$, preventing runaway growth. The normalisation is thus dynamic, with $|\mathbf{w}|^2$ relaxing to $1/\alpha$ over time, rather than than being rigidly enforced at every step. Again, competition is present as the sum of squares is held fixed. Strengthening of one input will therefore lead to the weakening of another.

3.6 Summary

Rate-based plasticity rules describe the activity of pre- and postsynaptic neurons by the recent time-averaged firing rate. Much of the early work on activity-dependent neuronal plasticity was conducted within such a framework, including that of the Hebb-rule along with its various permutations.

The learning properties of simple Hebbian-rules have been explored in some detail. A common problem amongst early rate-based rules of plasticity is uncontrolled growth of input weights. A variety of means of preventing uncontrolled growth, such as introducing bounds on input weights, have been explored. Normalisation of input weights, often with the aim of introducing competition between inputs, has also been investigated.

The BCM-rule is one of the more widely cited rate-based models of plasticity. The BCM-rule implements competition by means of a sliding threshold between potentiation and depression. The rule has been successfully applied to explain

various developmental phenomena, such as the activity-dependent development of orientation selectivity and binocularity in the developing visual cortex.

Despite these successes, the intrinsic structure of rate-based models means that they are unable to accommodate more recent experimental results which suggest that it is the timing of individual action potentials, rather than their mean rate of arrival, that determines the degree and polarity of change.

Chapter 4

Spike-based Models of Neuronal Plasticity

4.1 Spike-Based Models of Neuronal Plasticity

Recent experimental work has uncovered a form of plasticity, known as spike-timing-dependent-plasticity (STDP), where the relative timing of pre- and post-synaptic action potentials determines the degree and polarity of change in input weights. STDP has been reliably observed in a variety of species, across several brain areas, and at different stages of development. Various spike-based descriptions of plasticity have been formulated based on these results. This is in contrast to the purely rate-based descriptions described in Chapter 3 which do not explicitly represent pre- and postsynaptic spikes.

As before, we denote the presynaptic firing rate vector $\mathbf{u}$, and the postsynaptic firing rate as v. In the context of spike-based descriptions of neuronal activity, the "rate" of firing is interpreted as the recent-time average of the number of spikes arriving per second. In addition, we introduce a new notation to denote particular spike-train patterns. In this notation, we denote presynaptic spikes by π and postsynaptic spikes by p. For example, the 4-spike train pre-post-post-pre would be written simply as $\pi pp\pi$.

A common induction protocol for STDP is pairing of two action potentials, one pre- and one postsynaptic, at various timing differences (Bi and Poo, 1998). To aid our discussion, we define the spike-timing difference, Δt to be the time between pre- and postsynaptic spiking. That is, $\Delta t = t_{post} - t_{pre}$, where t_{pre} and t_{post} are the pre- and postsynaptic spike times, respectively. Thus, $\Delta t > 0$ if the postsynaptic spike follows the presynaptic spike, and $\Delta t < 0$ if this order is reversed. In the literature, these two cases are often referred to as "positively"

and "negatively" correlated firing, respectively. It is usual to constrain the input weight vector, $\mathbf{w}$, by imposing a hard lower bound to prevent its components from becoming negative, $w_i \geq 0 \forall i$. For learning rules which give rise to runaway learning, a hard upper bound, w_{max}, is also usually imposed.

Learning rules based on STDP data have been implemented in several modelling studies. A convenient distinction may be made between phenomenological and biophysical models. Biophysical models attempt to explain experimental observations in terms of the various biological mechanisms thought to underlie plasticity. Phenomenological models place their emphasis on the computational properties of a particular learning rule, standing some way above the actual biological implementation. As a result, phenomenological models are often simpler in nature. We review these studies, and outline the main findings of each implementation.

4.2 Phenomenological STDP Models

We refer to models taking the observed STDP-curve over directly to govern input plasticity as phenomenological STDP models. The learning rule typically takes the form of a biphasic modification window, with the spike-timing difference, Δt, determining the degree and polarity of change. Such a modification is described by the maximum amplitude and characteristic time constant associated with each phase. We denote the potentiating and depressing phase maximum amplitudes as A_+ and A_-, respectively. The plasticity rules we examine are monotonic, in the sense that each phase is non-decreasing in amplitude as the time difference approaches zero. Thus, the maximum amplitude, $A_\pm$, occurs when $\Delta t \to 0$. Similarly, we denote the potentiating and depressing phase characteristic time constants as τ_+ and τ_-, respectively. The precise interpretation of $\tau_\pm$ differs depending on the functional form of the proposed learning rule. For example, under a learning rule composed of two exponential phases (a common approximation of the STDP-curve), $\tau_\pm$ are exponential decay constants.

The phenomenological STDP models reviewed here share several common features. Typically, the STDP-curve is approximated as a biphasic, double exponential. That is, a spike pair of spike-timing difference Δt, adjusts the input weight w according to $w \to w + \Delta w$, with

$$\Delta w = \begin{cases} A_+ f(\Delta t, \tau_+) & \text{for } \Delta t \geq 0 \\ -A_- g(\Delta t, \tau_-) & \text{for } \Delta t < 0 \end{cases}, \tag{4.1}$$

where $A_\pm$ are the plasticity amplitudes and $\tau_\pm$ the characteristic time constants, as defined above. The magnitudes of plasticity, $A_\pm$, represent the change in

input weight due to a single spike pairing. As experimental results invariably rely on multiple spike-pairings to evoke a statistically meaningful change, $A_{\pm}$ are typically estimated by dividing the total change in input weight by the number of spike-pairings. This makes the implicit assumption that multiple spike pairings gave an overall change consistent with a linear sum of pairings under the learning rule.

Learning may be implemented additively (where changes in input weights are independent of their initial value), multiplicatively (where there is some dependence), or by some combination of the two. It is not necessarily the case that potentiation and depression share the same dependence on initial input weight. When learning is multiplicative, the fixed quantities $A_{\pm}$ are replaced by variable quantities $\alpha_{\pm} f(w)$, where $0 \leq \alpha_{\pm} \leq 1$ are the relative changes and $f(w)$ describes the dependence on initial input weight.

4.2.1 The Integrate and Fire Neuron

Phenomenological STDP rules are often explored in the context of a set of plastic excitatory inputs synapsing onto a simple integrate-and-fire target cell. The integrate and fire model generates the explicit postsynaptic spike times that are required to calculate afferent plasticity under spiking rules. This is in contrast to the rate-based approximations discussed in Chapter 3 which described the postsynaptic firing rate as the linear weighed sum of its presynaptic inputs. The integrate and fire neuron captures the basic features of neuronal spiking, such as the non-linear spike generation mechanism, while still having the advantage of being relatively simple in nature compared to more detailed models. The leaky integrate and fire neuron describes the evolution of the membrane potential, V_m, according to

$$\tau_m \frac{dV_m}{dt} = V_{rest} + I_i R - V_m \tag{4.2}$$

where V_{rest} is the resting membrane potential (which may be defined to be zero), τ_m the membrane time constant, I_i an injected current, and R the membrane resistance. Thus, under a constant current injection the membrane potential, V_m, will relax exponentially to $V_{rest} + RI_i$. In the spike-timing plasticity rules examined here the injected current, I_i, is typically zero in which case the membrane potential relaxes to its resting value, V_{rest}.

The influence of presynaptic spiking may be incorporated by adding a synaptic current, which we denote I_s. Labelling the presynaptic inputs with i, we may write the synaptic current I_s as

$$I_s = \sum_i w_i \sum_f I(t - t_i^f) \tag{4.3}$$

where w_i is the efficacy of input i, t_i^f is the time of spike f at input i, and

$$I(t) = g_s(t)\,(E_s - V_m(t)) \tag{4.4}$$

where $g_s(t)$ is the synaptic conductance and E_s is the reversal potential of the relevant ion type. For excitatory synapses this is sodium, and $E_s \sim 0$mV. For inhibitory synapses, this is potassium and $E_s \sim -75$mV. The time course of synaptic conductivity, $g_s(t)$, is typically described as the superposition of exponentials scaled by some maximum conductance value, $\bar{g}$. In this conductance based approach, therefore, input weights and the plasticity amplitudes, $A_\pm$ (or $f(w)$ in the multiplicative case), are described in terms of peak synaptic conductances, measured in pico-Siemens (pS). Postsynaptic spiking is a nonlinear event, defined to occur if $V_m \geq V_t$, where V_t is some defined threshold potential. On initiation of an action potential, V_m is set immediately to some reset potential, V_{reset} which is typically slightly less than the resting value, V_{rest}.

4.2.2 Song et al., (2000)

In the Song model, a biphasic, exponential approximation of the STDP-curve forms the basis of the learning rule (Song et al., 2000). This, it may be argued, is a reasonable approximation of a range of experimental data (Bi and Poo, 1998; Zhang et al., 1998). Parameters are taken mainly from the data of Bi and Poo (1998), but are not in disagreement with the majority of experimental work (but see Bell et al., (1997)) . The Song model is additive, so input weights are adjusted upwards or downwards according to the learning rule by fixed amounts $A_\pm$. That is, when a spike pair occurs with spike-timing difference Δt, the input weight w is adjusted according to $w \rightarrow w + \Delta w$, with

$$\Delta w = \begin{cases} A_+ e^{-\Delta t/\tau_+} & \text{for } \Delta t \geq 0 \\ -A_- e^{\Delta t/\tau_-} & \text{for } \Delta t < 0 \end{cases}, \tag{4.5}$$

where $A_\pm$ are the plasticity amplitudes, and $\tau_\pm$ the exponential decay constants. This learning rule is depicted in Fig. 4.1. A hard lower bound is imposed upon w at zero to prevent input weights from becoming negative. The authors note that, under this particular formulation, there is a requirement for an overall

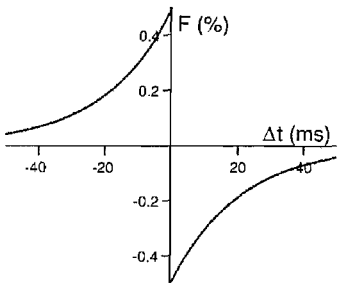

FIGURE 4.1: The biphasic, exponential-like learning rule of Song et. al., (2000). The change in input weight, $F(\Delta t)$, as a function of spike-timing difference, $\Delta t = t_{post} - t_{pre}$, is given as a percentage of the maximum input weight, w_{max} (from Song et al. (2000)).

dominance of depression over potentiation. That is, the condition $\tau_+ A+ < \tau_- A_-$ can empirically demonstrated to be necessary to prevent every input weight from being potentiated uncontrollably. This condition may be interpreted as the integral of the depressive phase being greater than that of the potentiating phase. Thus, the various experimentally observed asymmetries in the parameters (In the hippocampal slice cultures used in Bi and Poo (1998), τ_- is about twice τ_+ and A_+ is about two-and-a-half times A_-) are not strictly essential, only that, overall, depression dominates. The authors choose to set $\tau_+ = \tau_- = 20$ms throughout, in rough agreement with experiment, which leads to the new condition that $A_+ < A_-$. This new condition is duly satisfied by setting the ratio $A_-/A_+ = 1.05$. The learning rule is applied to a system of multiple inputs innervating a single target cell. The target neuron is represented by a simple integrate-and-fire scheme. Input weights are therefore interpreted as conductances, and presynaptic firing triggers excitatory synaptic currents which depolarise the target cell.

The Song model always produces runaway learning. This can intuitively be seen by considering the behaviour of an input which is initially potentiated due to some random fluctuation. After potentiation, the input will influence the postsynaptic integrate-and-fire neuron to a greater degree. As a result, there is a slightly higher probability that further potentiation events will occur. This eventually leads to the situation where one input is sufficiently strong to completely control the postsynaptic response, and as a result it is continually potentiated with all other input weights falling to zero. A hard upper bound is therefore imposed on input weights, w_{max}, chosen to be equal to 0.015.

A value for A_+ was partially derived from experimental work (Bi and Poo, 1998) by making the assumption that the effects of multiple spike-pairings sum linearly. Thus, the observed percentage changes in input weights in the spike-pairing

experiments of Bi and Poo (1998) were simply divided by the number of spike-pairs used in the protocol. A_+ was therefore set to $0.005\,w_{max}$, corresponding to a modification of 0.5% of the maximum input weight (occurring as $\Delta t \rightarrow 0$).

The purely additive nature of the learning rule is a departure from experimental observations, where depression (but not potentiation) apparently occurs in a multiplicative manner (Debanne et al., 1996; Bi and Poo, 1998; Debanne et al., 1999). No temporal restrictions were place on the interaction of pre- and postsynaptic spikes. Thus, in the spike train $\pi ppp\pi$, the outer two presynaptic spikes would individually interact with all three postsynaptic spikes, giving a total of three pre-post and three post-pre spike pairs. As mentioned above, these interactions are assumed to sum linearly. At the time, experimental data had only examined the effect of isolated spike-pairs. The assumption that plasticity sums linearly has since been questioned on both experimental and theoretical grounds (Froemke and Dan, 2002; Izhikevich and Desai, 2003). Although there is no explicit relationship between the plasticity rule and the rate of presynaptic firing, changes in the input firing rate alter the final input weight distribution without destroying the competitive nature of the rule.

Simulations of a single, leaky, integrate-and-fire neuron, innervated by 1000 excitatory and 200 inhibitory inputs, were made. The target neuron had a membrane time constant of $\tau_m = 20$ms, a resting potential of -70mV, and a threshold potential for spike initiation of -54mV. The input resistance was set to 100MΩ, which implies that the maximum input weight, $w_{max} = 0.015$, corresponds to a peak synaptic conductance of 150pS. Excitatory and inhibitory reversal potentials were 0mV and $-70mV$, respectively, and had time courses described by simple exponential functions with time constants $\tau_{ex,in} = 5$ms. Inhibitory inputs were non-plastic, with fixed conductances of 500pS each, and received Poisson spike trains at 10Hz. Excitatory inputs were modified according to the learning rule given above. Although there is no variability in either the level of level of plasticity or the timing of spikes, our own simulations, presented in Chapter 5, show that the introduction of any reasonable level of Gaussian noise to either quantity does not qualitatively affect any of these results.

Under this learning rule, a population of excitatory inputs receiving uncorrelated Poisson spike trains will reach a bimodal equilibrium distribution of input weights (Fig. 4.2a). The weights tend to cluster around the upper and lower bounds, a result independent of the initial distribution provided the postsynaptic neuron is initially firing action potentials. Increasing the presynaptic firing rate shifts this distribution towards the lower bound, a reflection of the highly competitive dynamics (Fig. 4.2b). As a result of this shift, the postsynaptic firing rate is relatively insensitive to changes in the mean presynaptic firing rate (a property referred to as "rate-normalisation" by the authors). The bimodal distribution of

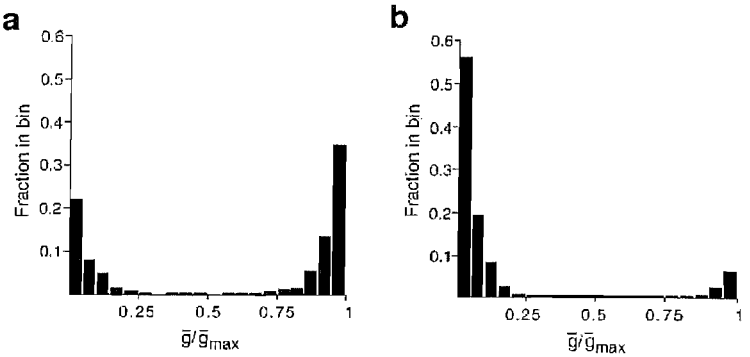

FIGURE 4.2: Equilibrium distribution of input weights reached under the Song learning rule when presynaptic firing is governed by independent Poisson spike-trains. The histogram shows the fraction of input weights falling into different bins, ranging from 0 to w_{max} (referred to as w_{max} in the figure), for input firing rates of **(A)** 10Hz and **(B)** 40Hz. Note the characteristic bimodal distribution, and the shift of this distribution lower input weights as the presynaptic firing rate is increased (from Song et al. (2000)).

input weights, and the highly competitive dynamics, are characteristic properties of the Song model. We return to this particular implementation of an additive STDP-like rule in Chapter 5, where we present a more detailed analysis of its underlying assumptions and explore the learning dynamics more fully in our own simulations.

4.2.3 van Rossum et al., (2000)

van Rossum et al., propose a mixed STDP learning rule, where potentiation is additive but depression is assumed to be multiplicative. A biphasic, exponential learning rule, is postulated, very similar to that of Song et al., (2000), with parameters also taken mainly from Bi and Poo (1998). The experimental asymmetry in the time constants for potentiation and depression were again found to be non-essential, and the authors set $\tau_+ = \tau_-$. Potentiation takes place in an additive manner, whereas depression is multiplicative. That is, the amplitude of a depressive change is proportional to the initial weight of the input. Such a scaling of depression, but not potentiation, with initial input weight has been observed in various experimental preparations (Debanne et al., 1996; Bi and Poo, 1998; Debanne et al., 1999).

Variability in the level of plasticity is implemented in a multiplicative manner. Thus, when a spike pair occurs with spike-timing difference Δt, the input weight

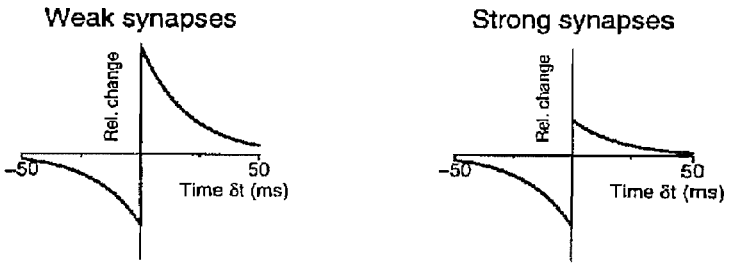

FIGURE 4.3: The biphasic, exponential-like learning rule of van Rossum et. al., (2000). The relative change in input weight as a function of spike-timing difference is shown for **(A)** weak inputs and **(B)** strong inputs. Due to the additive nature of the potentiating phase, stronger inputs are potentiated by a smaller amount relative to their weight. By contrast, the multiplicative depressive phase depresses all inputs equally (from van Rossum et al. (2000)).

w is adjusted according to $w \rightarrow w + \Delta w$, with

$$\Delta w = \begin{cases} (A_+ + xw)e^{-\Delta t/\tau_+} & \text{for } \Delta t \geq 0 \\ (-\alpha_- w + xw)e^{\Delta t/\tau_-} & \text{for } \Delta t < 0 \end{cases}, \tag{4.6}$$

where $\alpha_- = 0.03$ is the relative level of depression, $A_+ = 7$pS the potentiation amplitude, $\tau_\pm = 20$ms the exponential decay constants, and x is a Gaussian distributed random variable of zero mean and standard deviation $\sigma = 0.015$. This multiplicative noise plays an important role in shaping the equilibrium distribution of input weights. With simple additive noise, or in the absence of noise altogether, the final distribution of input weights is significantly different. The interaction of pre- and postsynaptic spikes was subject to a temporal restriction, in that spikes were only allowed to interact with their nearest neighbours. Thus, in the spike train $\pi ppp\pi$, only the first pre-post pair and the last post-pre pair will evoke plasticity. This is in contrast to other simple models of STDP, where spikes are not subject to any temporal constraints (Song et al., 2000). The learning rule is illustrated in Fig. 4.3. The plasticity rule is not dependent on the presynaptic firing rate in any way, and increasing the input firing rates serves only to increase the rate at which potentiation and depression occur. Introducing a presynaptic rate-dependence serves only to adjust the mean input weight and does not affect the stability or competition of the rule.

Activity-dependent scaling (ADS) of input weights can be included to regulate the level of postsynaptic firing, and introduce competitive dynamics. The authors choose to impose a multiplicative ADS that is independent of presynaptic firing rates. Thus, the input weights, w, receive an additional modification term, which

takes the form of an integral controller

$$\frac{dw(t)}{dt} = aw(t)[v_{goal} - v(t)] + bw(t)\int_0^t dt'\,[v_{goal} - v(t')], \tag{4.7}$$

where v_{goal} is the desired level of postsynaptic firing, and a and b are constants. The postsynaptic firing rate $v(t)$ is described as a slowly-varying sensor, receiving injections upon each presynaptic spike, relaxing to zero otherwise

$$\tau_v \frac{dv(t)}{dt} = -v(t) + \sum_i \delta(t - t_i), \tag{4.8}$$

where the t_i are a list of presynaptic spike times, and τ_v is the characteristic time constant associated with the sensor. τ_v is assumed to be large, and set equal to 100 seconds. The integral term in Eqn. (4.7) accumulates error over time, and will become dominant if the postsynaptic firing rate does not approach the desired value, v_{goal}, which was set to 20Hz. The parameters $a = 4\text{x}10^{-5}s^{-1}Hz^{-1}$ and $b = 10^{-7}s^{-2}Hz^{-1}$ were chosen to minimise oscillatory behaviour. The overall effect of this activity-dependent scaling is to adjust the balance of potentiation and depression, moving the mean input weight without affecting the shape of the distribution.

Simulations of a single, leaky, integrate-and-fire neuron, innervated by 100 excitatory and 25 inhibitory inputs, were made. The target neuron had an input resistance of 100MΩ, a membrane time constant of $\tau_m = 20$ms, a resting potential of −60mV, and a threshold potential for spike initiation of −50mV. Excitatory and inhibitory reversal potentials were 0mV and $-70mV$, respectively, and had identical time constants, $\tau_{ex} = \tau_{in} = 5$ms. Inhibitory inputs received Poisson spike trains at 20Hz, and were non-plastic with fixed conductances of 2000pS each. Excitatory inputs received both correlated and uncorrelated spike trains, and were subjected to the learning rule give in Eqn. (4.6).

Initially, excitatory inputs received uncorrelated Poisson spike trains. Prolonged, random stimulation in this manner led to the stable, unimodal distribution of input weights shown in Fig. 4.4. This distribution is relatively insensitive to parameter choices, but was noted to be dependent on the presence of multiplicative noise. With simple additive noise, or in the absence of noise altogether, the final distribution of input weights is much narrower. This distribution may be intuitively understood by considering the behaviour of a single input under this learning rule. With an integrate-and-fire neuron, stronger inputs are more likely to trigger postsynaptic spiking than weaker inputs. Thus, stronger inputs are more likely to precede postsynaptic firing than weaker inputs, which, under the plasticity rule given in Eqn. (4.6), means they are strengthened more often. This tendency for strong inputs to become even stronger is a destabilising force, and

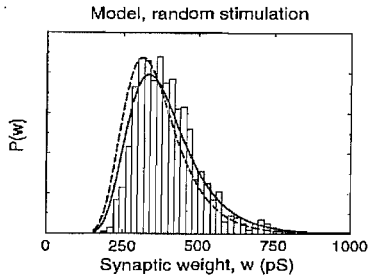

FIGURE 4.4: Equilibrium probability distribution of input weights reached under the van Rossum learning rule when presynaptic firing is governed by independent Poisson spike-trains at 20Hz. The histogram shows the fraction of input weights (written as "synaptic weight" in the figure) falling into different bins when a target integrate-and-fire neuron is simulated. The solid line is the corresponding analytical result. The dashed line is the related analytical result, under the additional assumption stronger inputs do not have an increased probability of triggering postsynaptic spiking (and hence have an equal chance of potentiation as weaker inputs) (from van Rossum et al. (2000)).

would typically lead to runaway learning. However, in this multiplicative model, the reduction in the relative level of potentiation for stronger inputs acts as a stabilising force. The observed unimodal distribution arises due to the balance of these stabilising and destabilising forces. In purely additive models, where the level of potentiation does not decline with increasing input weights, the stabilising force is removed and, unless hard bounds are imposed on input weights, runaway learning occurs (Song et al., 2000).

Correlations amongst presynaptic firing were introduced by randomly generating N Poisson spike trains, then assigning each input a spike train with probability $1/N$ at the start of each time step. With $N < 100$, some inputs will, by definition, be assigned the same spike train. This gives rise to a well defined cross-correlation coefficient of $C(\delta t) = 1/N\delta(\Delta t)$. Inputs firing in groups are more likely to trigger postsynaptic spiking, and thus be potentiated. Accordingly, with more correlated inputs the input weight distribution is shifted towards a higher mean weight (Fig. 4.5). The mean weight of a group of inputs is proportional to their level of input correlation.

Competition is almost absent under this learning rule. This can be seen in simulation by dividing the 100 excitatory inputs into two groups. The first group receives purely uncorrelated Poisson spike trains, the second group initially receives uncorrelated Poisson spike trains but at some point in time begin to be correlated. Both groups initially arrive at a distribution with the same mean weight. After the correlations are introduced into group two, the mean input weight of group two increases, shifting the distribution of weights to the right. The distribution of group one weights is, however, hardly affected by this shift

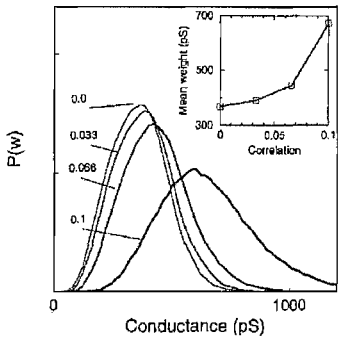

FIGURE 4.5: Introducing correlations between presynaptic spiking leads to an equilibrium probability distribution shifted to higher mean weight, while leaving the overall shape relatively intact. The learning is non-competitive, with the four groups of inputs behaving independently. That is, increases in the mean weight of one group will not cause a subsequent decrease in another. The inset shows the relationship between input correlation parameter and mean equilibrium input weight (from van Rossum et al. (2000)).

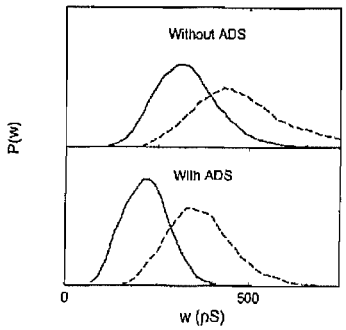

FIGURE 4.6: The effect of activity-dependent scaling (ADS) on the equilibrium distributions of two populations inputs. *Top*: without ADS, introducing correlations into one group (dotted line) causes that group to arrive at a higher mean input weight without affecting the distribution of the uncorrelated group (solid line). The total postsynaptic input is therefore high, and the postsynaptic neuron fires at a high rate. *Bottom*: with ADS, the mean input weight of both groups are scaled down, and the postsynaptic firing rate remains at the desired level, v_{goal} (from van Rossum et al. (2000)).

(Fig. 4.6). This increases the total presynaptic input, and drives the postsynaptic cell to a higher firing rate. The postsynaptic firing rate therefore follows the total input, and is very sensitive to input rate fluctuations.

Competitive dynamics may be introduced by including activity-dependent scaling (ADS), as outlined above. The shape of the input weight distribution, and its stability, are not affected by the introduction of ADS. ADS operates independently of the presynaptic firing rate, whereas the spike-timing learning rate is

accelerated for high-rate inputs. Thus, low rate inputs are effectively governed by ADS and high-rate afferents by spike-timing plasticity.

4.3 Biophysical Models of STDP

An alternative approach to modelling STDP is the attempt to reconstruct the biochemical induction and expression pathways of plasticity across a particular type of connection. Unlike the phenomenological approach, these models commit themselves to a particular mechanistic paradigm. This can limit the generality of any theoretical insights which may be obtained. For example, the theoretical implications of a particular model may only be considered in the context of other connection types if those connections are themselves amenable to a model of similar intrinsic structure.

4.3.1 The Calcium-Control Hypothesis

A large body of evidence suggests that calcium plays a critical role in some forms of neuronal plasticity. In particular, studies on the N-methyl-D-aspartate (NMDA) subtype of ionotropic glutamate receptor have suggested that it might provide a kind of molecular coincidence detection mechanism that could underlie STDP at certain synapses. The idea that NMDA-receptor-dependent calcium dynamics may underlie STDP has become known as the calcium-control hypothesis. This hypothesis may be summarised as follows. Presynaptic glutamate release activates the postsynaptic AMPA- and NMDA-receptors. Activation causes the receptor molecule to undergo a conformational change, allowing the associated ionic channel to open and create a tunnel through the cell membrane. The AMPA-receptors has an ion channel that is largely permeable to sodium and potassium. Presynaptic activation activates the population of postsynaptic AMPA-receptors, and the influx of sodium creates a postsynaptic depolarisation. Although various other processes contribute to this depolarisation, AMPA-receptor-mediated sodium influx is the dominant component. NMDA-receptors also allow the passage of sodium ions, but in addition allow calcium influx. However, the resting NMDA-receptor is blocked by a magnesium ion, which requires postsynaptic depolarisation to clear. Thus, glutamate binding must be coupled with sufficient postsynaptic depolarisation (such as that occurring during an action potential) in order to allow postsynaptic calcium entry. The resulting high level of postsynaptic intracellular calcium is thought to trigger some change that upregulates the strength of the connection between the pre- and postsynaptic neurons.

On the other hand, when the postsynaptic depolarisation caused by the AMPA-receptors is insufficient to trigger an action potential, the magnesium block on the NMDA-receptors is only partially relieved, and the level of postsynaptic calcium is elevated to some sustained intermediate level. Some change then takes place such that the efficacy of the connection between pre- and postsynaptic neurons is decreased.

Thus, the distinct spatio-temporal calcium transient profiles due to the partial or complete activation of NMDA-receptors are thought to trigger distinct signalling cascades leading to either LTP or LTD (Yang et al., 1999). Calcium-activated calmodulin kinase II (CaMKII) may provide the driving mechanism behind this bidirectional plasticity by responding with different kinetics to the two regimes (Soderling, 2000). For example, the carboxyl terminal could promote facilitation of calcium channels and lead to LTP, and the amino terminal could promote inactivation of the same channel and lead to LTD. CaM also participates in multiple local (and hence input specific) signalling complexes. The dynamics and binding of calcium to various proteins, and the consequent induction of LTP or LTD, may therefore reasonably depend in a highly nonlinear way on the relative timing of pre- and postsynaptic action potentials.

The calcium-control hypothesis requires a signal to pass from the soma to the dendrites indicating that an action potential has taken place. It has been suggested that this postsynaptic signal takes the form of a back-propagating action potentials (BPAP). BPAPs have been shown experimentally to modulate LTP induction in a manner consistent with this role (Magee and Johnston, 1997) and a persistent after-hyperpolarisation following a dendritic action potential is a necessary condition for some biophysical models of STDP (Shouval et al., 2002). Other work has shown that at stimulation frequencies below around 10Hz, LTP is not induced (Markram et al., 1997). However, this result may stem, in part, from a lack of BPAP initiation at such low stimulation frequencies. Thus, although the calcium control hypothesis has received considerable experimental support, there remain a number of outstanding questions.

Several models of STDP have been based around a description of calcium dynamics in accordance with the calcium control hypothesis. We review the major findings here.

4.3.2 Senn et al., (2000)

As discussed above, a biophysical model is specific in the sense that it aims to reconstruct the underlying pathway of a particular type of plasticity. Senn et

al., (2000) set out to construct a detailed biophysical model of plasticity expressed by connections between neighbouring, thick-tufted pyramidal cells from layer 5 (L5) of the neocortex. Experimental work based on dual whole-cell voltage recordings from neighbouring L5 pyramidal cells has demonstrated a form of STDP-like input plasticity that depends upon the activation of N-Methyl-D-Aspartate (NMDA) receptors (Markram et al., 1997). This form of plasticity is timing-dependent, input specific, and shows a noticeable frequency-dependence for the onset of afferent modification. In addition, the response of these neocortical neurons to a train of presynaptic action potentials has a characteristic time course, with later spikes in the train evoking a smaller excitatory postsynaptic potential (EPSP) as shown in Fig. 4.7. This effect has been attributed to the depletion of the pool of neurotransmitter vesicles available for discharge, which would be consistent with the hypothesis that a finite number of vesicles are docked and available for release at any one time. Potentiating these inputs causes the early-spike response to increase and the late-spike response to be depressed, an effect known as "synaptic redistribution". Synaptic redistribution is consistent with plasticity being expressed at a presynaptic locus, with, for example, changes in the probability of presynaptic vesicle release underlying the apparent change in input weight. This presynaptic mechanism is in contrast to, for example, a general increase in postsynaptic responsiveness by the insertion or up-regulation of postsynaptic neurotransmitter receptors. Accordingly, Senn et al., (2000) formulate a model of neocortical plasticity, based around the idea that changes in the probability of presynaptic neurotransmitter vesicle release govern changes in input weight. The probability of vesicle release is denoted P_{rel}, and it is assumed to be the product of vesicle discharge probability times the probability that a vesicle is available for release $P_{rel} = P_{dis}P_{ves}$. Changes in the discharge probability, P_{dis}, are assumed to be solely responsible for observed changes in input weight. These changes are assumed to result from two independent second-messenger pathways, one that drives up-regulation after a πp spike pair and one that drives down-regulation following a $p\pi$ pair.

Unlike other models of STDP, changes to input weights in response to pre- and postsynaptic spiking do not occur instantaneously, but rather are assumed to be expressed only after 10-20 minutes have elapsed. In terms of simulation, this means that the modifications arising to P_{dis} due to coincidences of pre- and postsynaptic activity are only realised after the stimulating protocol has finished. That is, changes are recorded and summed linearly $\eta = \sum_i \Delta P^i_{dis}$, where ΔP^i_{dis} is the i'th modification to P_{dis}. The step modification $P_{dis} \rightarrow P_{dis} + \eta$ occurs only after the stimulation patterns have ceased. The consideration of plasticity is limited to that expressed on the time scale of 10-20 minutes. There is no explicit representation of the short-term transient potentiation that typically

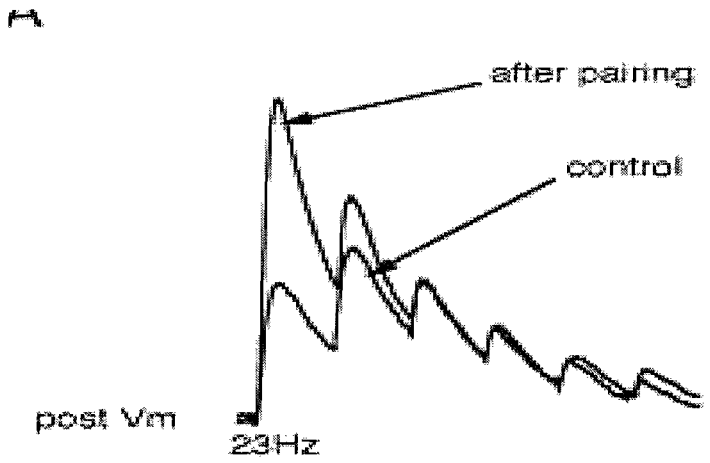

FIGURE 4.7: Effect of repeated pairing of pre- and postsynaptic action potentials on the postsynaptic response to trains of action potentials. The graph shows the postsynaptic response to a 23-Hz train of presynaptic action potentials before and after pairing. After pairing, the early part of the postsynaptic response is potentiated, but the later response remains at the same. The average response of 58 sweeps is shown before and 20 min after pairing (from Markram et al. (1997)).

follows a presynaptic spike-train (post-tetanic potentiation, or PTP), not any consideration of decay of plasticity on time scales greater than one hour.

The central postulate of the model is that postsynaptic NMDA-receptors are responsible for mediating up- and down-regulation of input weight, via these two second messenger pathways. Although this hypothesis is plausible, and it has been experimentally demonstrated that NMDA-receptor activation is necessary for both up- and down-regulation, the exact role of NMDA-receptors in plasticity is still unclear. The NMDA-receptors are assumed to be able to reside in one of three possible states; resting, active, and calcium modulated. There is a fixed population of NMDA-receptors, and the proportion residing in each state is denoted N_r, N_a, and N_c, respectively. The second messenger responsible for triggering the changes in P_{rel} may be in an active state, $B_{u,d}$ or an inactive state $A_{u,d}$, where the subscript u, d denote the up- and downregulating second messengers, respectively.

Under normal circumstances, $P_{ves} = 1$, but, following vesicle release, P_{ves} is set to zero then recovers according to a Poisson process of time constant $\tau_{ves} = 800$ms. This is in accordance with the 'univesicular hypothesis' (Triller and Korn, 1982), which assumes that there is, at most, one vesicle docked and ready for release at any one time. This vesicle depletion model is a stochastic version of the authors' earlier depressing-synapse model (Markram et al., 1997). This stochastic model

replicates the statistics of non-averaged responses seen during experiments involving a single instance of a stimulation pattern. Averaging over repeated trials of the same presynaptic stimulation protocol gives a mean postsynaptic response (in the form of an excitatory postsynaptic potential, or EPSP) which is proportional to P_{rel}.

When presynaptic release of neurotransmitter does occur, for example in response to a presynaptic action potential, it is assumed to immediately diffuse across the synaptic cleft and cause some proportion, r_a, of NMDA-receptors currently in the resting state to make the transition from the resting to the active state, $N_a \rightarrow N_a + r_a N_r$. Neurotransmitter release may also occur spontaneously, triggering an identical set of events. Postsynaptic action potentials are assumed to trigger a significant depolarisation of the dendritic membrane, allowing calcium influx via the population of activated NMDA-receptors. This calcium influx has two effects. Firstly, activation of a proportion, $r_s N_a$, of the potentiating second messenger occurs, $B_u \rightarrow B_u + r_A N_a A_u$, which then diffuse in a retrograde manner to the presynaptic site and triggers up-regulation of η by the amount $r_{dn}\eta[B_d - \theta_d]^+$, where θ_d is the threshold on level of down-regulating messenger activity required for down-regulation, and the quantity $[x]^+ = max[x, 0]$. Secondly, additional calcium influx occurs through voltage-gated calcium channels. This calcium influx causes some proportion, r_c, of NMDA-receptors to make the transition from the resting to the calcium modulated state, $N_c \rightarrow N_c + r_c N_r$. A subsequent presynaptic release of glutamate acts upon these altered-state NMDA-receptors, which activate a proportion, $r_A N_c$, of the down-regulating second messenger, $B_d \rightarrow B_d + r_A N_c A_d$. This down-regulating second messenger also diffuses to the presynaptic site and triggers a down-regulation of η by the amount $r_{up}(1-\eta)[B_u - \theta_u]^+$, where θ_u is the threshold on level of up-regulating messenger activity required for up-regulation.

The NMDA-receptors thus act as coincidence detectors of pre- and postsynaptic spiking. The spike train πp will lead to up-regulation of presynaptic release probability, and the spike train $p\pi$ will lead to a down-regulation. In the absence of further pre- or postsynaptic spiking, the quantities N_a, N_c, and B_u, B_d decay with time constants τ_N and τ_B, respectively. This kinetic scheme may be expressed in differential equation form as follows. First, the evolution of the proportion of NMDA-receptors in the activated, calcium modulated, and resting states (N_a, N_c, and N_r, respectively) as functions of time, t, are given by

$$dN_a/dt = -N_a/\tau_N + r_a N_r \delta(t - t_\pi^i), \tag{4.9}$$

$$dN_c/dt = -N_c/\tau_N + r_c N_r \delta(t - t_p^i), \tag{4.10}$$

$$dN_r/dt = 1 - N_a - N_c, \tag{4.11}$$

where the delta function $\delta(t - t^i_{\pi,p})$ represents the injection received upon the ith pre- or postsynaptic spike, occurring at time $t^i_{\pi,p}$. Similarly, the equations governing the up- and down-regulating second messenger activation states, B_u and B_d, are

$$dB_u/dt = -B_u/\tau_B + r_A N_a(1 - B_u)\delta(t - t^i_\pi), \tag{4.12}$$

$$dB_d/dt = -B_d/\tau_B + r_A N_c(1 - B_d)\delta(t - t^i_\pi). \tag{4.13}$$

The accumulated change in P_{dis}, denoted η, then evolves according to

$$d\eta/dt = r_{up}(1 - \eta)[B^+_u - \theta_u]^+\delta(t - t^i_p) - r_{dn}\eta)[B^+_d - \theta_d]^+\delta(t - t^i_\pi), \tag{4.14}$$

where B^+_u and B^+_d are the values of B_u and B_d immediately after a post- and presynaptic spike

$$B^+_u = r_A N_a(1 - B_u) + B_u, \tag{4.15}$$

$$B^+_d = r_A N_c(1 - B_d) + B_d. \tag{4.16}$$

Finally, the discharge probability, P_{dis} asymptotes to its final value, η, according to

$$dP_{dis}/dt = (\eta - P_{dis})/\tau_\eta, \tag{4.17}$$

with time constant $\tau_\eta = 6.10^5$ms.

Simulations were made of a range of stimulation protocols used in experimental work (Markram et al., 1997). Bursts of pre- and postsynaptic spikes (5 spikes in a 10Hz train, repeated at 2.5Hz) were simulated, with the bursts initiated at timings of $t_{post} - t_{pre} = \pm 100$ms and ± 10ms. In experimental work, spike bursts initiated 100ms apart were unsuccessful at inducing plasticity, and bursts at timings of $+10$ms and -10ms giving potentiation and depression, respectively. Both these findings are reproduced by the model, as shown in Fig. 4.8. Parameters were chosen to fit the experimental data, with NMDA-receptors described by $r_a = 1$, $r_c = 0.5$, $\tau_N = 300$ms, the second messenger dynamics by $r_A = 0.7$, $\tau_B = 600$ms, and the plasticity action by $r_{up} = r_{dn} = 0.1$. This gives good qualitative agreement with the experimental data, but the magnitude of plasticity is only around half that seen experimentally.

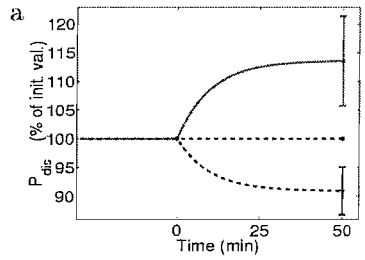

FIGURE 4.8: Asymmetry of modification under the biophysical learning rule of Senn et al., (2000). Simulated results of presenting a pre- and postsynaptic spike train (5 spikes, at 10Hz) at spike timings of $+10ms$ (top trace) and -10ms (bottom trace) are shown. The graph shows the percentage change in P_{dis} as a function of time. When the spike trains were initiated with a time difference of 100ms, no plasticity was induced (middle trace). Parameter values are given in the text, and were set by fitting the simulation results to experimental data (from Senn et al. (2000)).

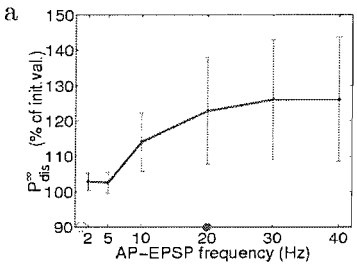

FIGURE 4.9: The frequency-dependence of modification under the learning rule. A sharp onset of plasticity occurs when the spike trains (5 spikes, 2 to 40Hz) reach a frequency of 10Hz. The percentage change in P_{dis} is shown 60 minutes after pairing of pre- and postsynaptic spike bursts at a spike timing of $+2ms$ (from Senn et al. (2000)).

Fig. 4.9 shows the simulated effect of repeating the protocol with a fixed inter-burst delay of 2ms, while varying the spike train frequencies (from 2 to 40Hz). The model reproduces the main characteristics of the frequency-dependent learning curve, with a sharp onset at around 10Hz and a saturation at higher frequencies (Markram et al., 1997). This is expected, as the model is formulated to accommodate this frequency-dependence by explicitly including pair of thresholds, $\theta_{\pm}$, to exclude too much modification at low firing rates. However, the final value for the change in discharge probability, η, was approximately half that seen in experimental work (Markram et al., 1997).

A third experimental protocol was simulated, where spike trains consisted of a variable number of spikes at 20Hz paired at a delay of 2ms. In experimental work, an increased number of spikes in a train reduced the level of LTP expressed

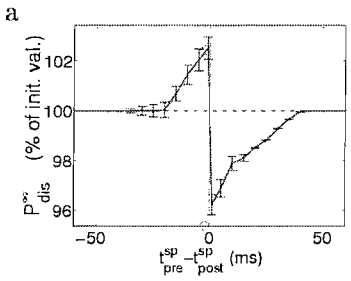

FIGURE 4.10: The STDP-like modification curve. The percentage change in P_{dis} is shown after 50 spike pairings at various spike-timing differences, $\Delta t = t_{post} - t_{pre}$. Although qualitatively similar to the experimentally observed STDP curves, the magnitude of plasticity is only 2% of that typically observed experimentally (from Senn et al. (2000)).

(Markram et al., 1997). In simulation, this behaviour was only approximately captured, with a noticeable departure as the number of spikes became small (although the results still lie within the high standard deviations of the data).

In a number of spike-pairing experimental protocols, a low firing rate is chosen specifically to demonstrate that it is the individual spike pairs that give rise to the observed overall change in input weight, rather than the averaged effect of longer, multi-spike trains (Bi and Poo, 1998). Thus, the experimental protocols typically used to assess the effect of individual spike pairs involve presentation at frequencies lower than that required to elicit significant plasticity under this model. The simulated change in P_{dis} due to a typical STDP protocol are shown in Fig. 4.10. Although the overall shape of the STDP-like curve is qualitatively the same, such protocols give a simulated changes of less than 2% that seen in experiment (Zhang et al., 1998; Bi and Poo, 1998). This is a direct consequence of the thresholds, $\theta_{\pm}$, which are included in the model to accommodate the frequency-dependence of LTP discussed above.

4.3.3 Shouval et al., (2002)

The Shouval model is another biophysical, calcium-based description of neuronal plasticity (Shouval et al., 2002). In the Shouval model, two key assumptions are made. First, that calcium is the primary signal for plasticity, and second, that the NMDA-receptor-dependent calcium influx is the dominant source of this calcium. The authors postulate a learning rule of the form of a differential

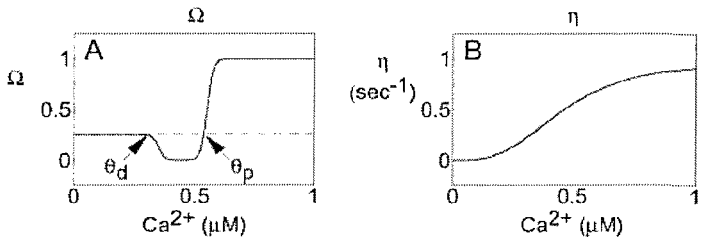

FIGURE 4.11: Plots of (**A**) the learning function, Ω, and (**B**) the learning rate η, as a function of postsynaptic calcium concentration. The learning function is translated directly in changes in input strength. When the calcium concentration is below a threshold, θ_d, the learning function is zero and no plasticity is induced. Between θ_d and θ_p the learning function is negative, representing depression, and above θ_p the function is positive, representing potentiation (from Shouval et al. (2002)).

equation describing the rate of change in input strength of input j, w_j. This differential equation is written as

$$\dot{w}_j = \eta(a_j)\left(\Omega(a_j) - w_j\right), \tag{4.18}$$

where a_j is the level of postsynaptic intracellular calcium, $\Omega(a_j)$ is the learning function and $\eta(a_j)$ the calcium-concentration-dependent learning rate. The second term inside the bracket is a decay term to prevent uncontrolled potentiation. This is not essential and, for example, a hard upper bound on input weights could be used instead.

Plasticity occurs in this model in response to elevation of intracellular calcium levels above normal levels. The learning function Ω describes the exact dependence of this learning on the level of calcium, and is assumed to be the two-phase function shown in Fig 4.11A. The learning function has two explicit thresholds which denote the point at which depression begins, θ_d, and the point at which depression turns into potentiation, θ_p. The dependence of η on a_j prevents oscillations and ensures that input strengths do not decay to zero when a_j returns to basal levels. Thus, we have changes in input weights driven by both changes in the total level of intracellular calcium, through the function Ω, and by the temporal pattern of those changes, through the learning rate η.

Postsynaptic calcium dynamics are assumed to be purely NMDA-receptor-dependent. A standard set of assumptions are made about NMDA-receptor dynamics, and the calcium current, $I(t)$, at some time t following a presynaptic spike is assumed

to be

$$I(t) = PH(v)\left(I_f e^{(-t/\tau_f)} + I_s e^{(-t/\tau_s)}\right), \tag{4.19}$$

where I_f and I_s describe fast and slow components, respectively, with $\tau_f = 50$ms and $\tau_s = 200$ms. P describes the activation profile of the NMDA-receptors, accommodating such factors as the fraction of NMDA-receptors activated by arrival of a presynaptic spike, and $H(v)$ describes the voltage-dependence of the magnesium block. Thus, the calcium influx at time t after some presynaptic spike is the sum of a fast and slow component of NMDA-receptor activation, modulated by the functions P and $H(v)$.

The postsynaptic voltage, v, is assumed to depend on the recent history of presynaptic spiking in two ways. First, there is the usual depolarisation due to AMPA-receptor-dependent sodium influx. Second, there is a much larger depolarisation whenever a postsynaptic spiking occurs, which is assumed to be transmitted throughout the dendritic tree by a backpropagating action potential (BPAP). The time course for this BPAP is critically important for the model. If it is too short, then a pre-post spike pair will only elevate the postsynaptic calcium concentration by a modest amount and trigger a relatively minor increase in input strength. Worse, a post-pre pair will trigger postsynaptic calcium levels that are indistinguishable from presynaptic spiking alone. It is therefore necessary to postulate an extended depolarising tail for the BPAP, that persists well after the PSAP that triggered it. The depolarisation due to a BPAP is therefore modelled as

$$v_{BPAP}(t) = I_f^b e^{(-t/\tau_f^b)} + I_s^b e^{(-t/\tau_s^b)}, \tag{4.20}$$

where I_f^b and I_s^b describe fast and slow components of the BPAP, respectively, with $\tau_f^b = 3$ms and $\tau_s^b = 25$ms. The shape of the BPAP is nearly independent of the slow component provided that the magnitude of the slow component is much smaller than that of the fast component. Thus, although the particular form of BPAP chosen by the authors is motivated by a theoretical requirement of the model rather than a desire to reproduce experimental data in a minimal fashion, the form chosen is not in direct conflict with experimental data (Magee and Johnston, 1997; Larkum et al., 2001). Finally, calcium dynamics are modelled by the simple differential equation

$$\frac{da_j}{dt} = I(t) - \frac{a_j}{\tau_a} \tag{4.21}$$

where $\tau_a = 50$s is the time constant for the decay of the calcium level, representing the various processes by which calcium is returned to basal levels. One of the

key assumptions of the model is that NMDA-dependent influx is the dominant of calcium. Thus, although the consideration of additional, secondary sources of calcium (such as release from intracellular calcium stores) might appear desirable, to first order the learning behaviour exhibited by the model would be unchanged.

Plots of NMDA-receptor activation and postsynaptic depolarisation (panels 1, 3 and 5) and the resulting postsynaptic calcium transient (panels 2, 4 and 6) for various spiking protocols are shown in Fig. 4.12. Presynaptic spiking alone is insufficient to trigger plasticity as the calcium concentration does not exceed the lower threshold, θ_d. A post-pre spike pair causes some intermediate level of calcium elevation to a level between θ_d and θ_p, triggering a decrease in input strength. A pre-post spike pair causes a high level of calcium elevation, exceeding the potentiation threshold θ_p, triggering an increase in input strength. Thus, the basic spike-pairing protocol of a typical STDP experiment will adequately be accounted for (Bi and Poo, 1998).

The need for a calcium-concentration-dependent learning rate can clearly be seen by considering the calcium transient induced by a pre-post spike pair, shown in box 6 of Fig. 4.12. Such a spike pair would be expected to induced potentiation of input strengths (Bi and Poo, 1998; Zhang et al., 1998). Accordingly, the level of potentiation induced by the region above θ_p must be greater than the level of depression induced by the time spent between θ_d and θ_p. This achieved through the variable learning rate, η, which amplifies changes occurring at higher levels of intracellular calcium. Without this variable learning rate, the time spent in the depressing regime would reverse and cancel out any potentiation.

Pairing of presynaptic spikes with sustained postsynaptic depolarisation, such as that performed during voltage-clamp experiments, was simulated. No plasticity was observed when the level of depolarisation was below -65mv. When the depolarisation was between -60mv and -35mv LTD was induced, and when the depolarisation was above -30mv, LTP was induced. As sustained depolarisation paired with presynaptic spiking translates quite straight forwardly to postsynaptic calcium levels, the shape of this dependence is essentially that of the learning function, Ω. Thus, as Ω was formulated to induce zero change at low calcium levels, LTD at intermediate levels, and LTP at high level, this first result is expected.

Simulations of rate-based induction protocols were also performed. Presynaptic activity was set to a rate ranging from 0Hz to 20Hz. Coupled with a simple statistical model of postsynaptic spike generation, this produced a plasticity curve which was, again, similar in shape to the learning function Ω. The cross over point from depression to potentiation was around 8Hz. This curve is similar

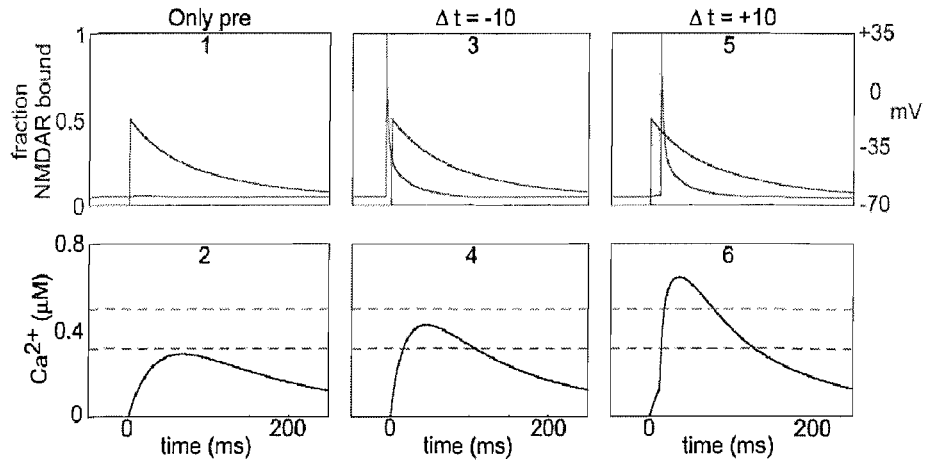

FIGURE 4.12: The postsynaptic response (the NMDA-receptor activation level and the postsynaptic depolarisation), and the resulting calcium transient, for both isolated presynaptic spiking and presynaptic spiking paired with postsynaptic spiking. The leftmost plots are (**1**) the postsynaptic response and (**2**) the corresponding calcium transient due to isolated presynaptic spiking. The fin-shaped trace in (1) is the NMDA-receptor activation level. Postsynaptic depolarisation for this protocol is very small. The calcium transient in (2) is below the lower threshold, θ_d, represented by the lower dotted line and no plasticity is induced. The centre plots are (**3**) the postsynaptic effect and (**4**) the corresponding calcium transient due to postsynaptic spiking followed by presynaptic spiking at a 10ms time difference. Again, the fin-shaped trace in the upper plot is the NMDA-receptor activation levels. This time, a large postsynaptic depolarisation results the postsynaptic action potential. The resulting calcium transient shown in (4) is between the two thresholds, θ_d and θ_p, and depression is therefore induced. The rightmost plots are (**5**) the postsynaptic effect and (**5**) the corresponding calcium transient due to presynaptic spiking followed by postsynaptic spiking, again at a 10ms time difference. The temporal ordering of the fin-shaped, NMDA-receptor activation level and the postsynaptic depolarisation spike is reversed. The resulting calcium transient shown in (6) exceed the upper threshold, θ_p, and potentiation is induced (from Shouval et al. (2002)).

to the ϕ function of the BCM-model discussed in Chapter 3. Thus, the model is able to accommodate simple rate-based induction of plasticity. The authors note that the threshold between potentiation and depression can easily be modified by allowing the parameters describing the NMDA-receptor dynamics to change. This "meta-plasticity" could offer a way to introduce stability to the learning dynamics of this model of STDP by by allowing this threshold to slide in a BCM-like manner.

As is apparent from Fig. 4.12, the Shouval model provides an explanation of the basic spike-pairing results of STDP (Bi and Poo, 1998; Zhang et al., 1998), with spike pairings at 1Hz at time +10ms giving rise to LTP and those at −10ms giving rise to LTD. However, when the full STDP curve is plotted for spiking timings

ranging from −100ms to 100ms, an additional LTD window is observed a high positive spike-timing (that is, when presynaptic spiking precedes postsynaptic spiking). This is not in agreement with the majority of experimental results (Bi and Poo, 1998; Zhang et al., 1998) although one exception does exist where an additional LTD window of this nature was observed (Nishiyama et al., 2000). This second LTD window is due to the intermediate level of calcium elevation that a pre-post spike pair of large spike-timing difference triggers, and, as a result, is a robust prediction of the model. A strong rate-dependence exists in this model, and as the rate of pre- and postsynaptic spiking increases LTP becomes more and more dominant. Pairing of pre- and postsynaptic spikes will therefore induce potentiation, regardless of spike-timing, if the pairing takes place above around 10Hz.

4.3.4 Karmarkar et al., (2002)

In most experimental work, a single LTP and a single LTD window have been observed (Bi and Poo, 1998; Zhang et al., 1998). The prediction of a second LTD window at large pre-post spike timings in simple calcium based models of STDP is inconsistent with this observation. An attempt to remedy this problem by the introduction of a second coincidence detector has been made in the form of the Karmarkar model (Karmarkar and Buonomano, 2002). The model is consistent with the calcium control hypothesis, and incorporates both NMDA-receptor-dependent and voltage-gated-calcium-channel (VGCC) dependent calcium influx. Both NMDA-receptors and VGCCs have been implicated in playing a role in spike-timing plasticity (Bi and Poo, 1998).

NMDA-receptor-dependent calcium influx was assumed to be responsible for driving LTP. Intracellular calcium levels due to the action of NMDA-receptors were assumed to be governed by the differential equation

$$\frac{dCa_{NMDA}}{dt} = a_1 \frac{V_m - 140}{1 + e^{a_2 - V_m}} - \frac{Ca_{NMDA}}{\tau} \tag{4.22}$$

where V_m is the membrane potential, τ is the time constant governing calcium removal, and Ca_{NMDA} is NMDA-receptor-dependent calcium concentration. VGCC-dependent calcium influx was assumed to be responsible for driving LTD. In a similar manner to that of NMDA-receptors, the VGCC-dependent intracellular level was assumed to be governed by the differential equation

$$\frac{dCa_{VGCC}}{dt} = b_1 \frac{V_m - 140}{1 + e^{b_2 - V_m}} - \frac{Ca_{VGCC}}{\tau} \tag{4.23}$$

where Ca_{VGCC} is VGCC-receptor-dependent calcium concentration. The constants a_1, a_2, b_1, b_2 are taken from the literature (see Karmarkar and Buonomano (2002) for details).

The VGCC-dependent calcium influx was assumed to induce LTD via the coincidence of the resulting intracellular calcium elevation with presynaptic glutamate release. Such a process could be governed, biologically, by a metabotropic-glutamate-receptor (mGluR) dependent pathway, which requires postsynaptic calcium entry followed by presynaptic glutamate release. In the model, the integral of the intracellular calcium level across the period where presynaptic glutamate was present is then directly translated into the level of LTD. Similarly, NMDA-receptor calcium influx is integrated and translated directly into the induced level of LTP. As discussed above, NMDA-receptors only permit calcium influx when presynaptic firing is followed by postsynaptic firing due to the requirement for both glutamate activation and the voltage-dependent relief of the magnesium block. Thus, pre-post pairings give rise to LTP via the NMDA-receptor pathway and post-pre pairings give rise to LTD via the VGCC pathway. In both cases, there must be some mechanism that translates the intracellular calcium rise into a graded change in synaptic strength.

The integral values for NMDA-receptor calcium influx and VGCC calcium influx, and the corresponding STDP rule are shown in Fig. 4.13. The STDP rule has two phases, a negative phase representing LTD when presynaptic spiking follows postsynaptic spiking, and a positive phase representing LTP when the order is reversed. There is no additional LTD phase a large pre-post spike timings, as the VGCC-dependent pathway which governs LTD is never activated under a pre-post a spike pair. Thus, the introduction of a second coincidence detector, and the distribution of labour that this allows, produces a calcium-based STDP rule capable of supporting a biphasic learning rule of the form observed in a variety of experimental preparations (Bi and Poo, 1998; Zhang et al., 1998).

4.3.5 Saudargine et al., (2005)

The problem of a second LTD window under pre-post pairings at large spike-timings in a calcium-control model has also been discussed in the Saudargine model (Saudargine and Porr, 2005). In contrast with models of plasticity where changes are dependent only on the absolute level of postsynaptic calcium (Shouval et al., 2002), in the Saudargine model plasticity is directed by the gradient of intracellular calcium changes. The basic principles of the calcium control hypothesis, that it is NMDA-receptor-dependent calcium that drives synaptic plasticity, are then reinterpreted in this context. NMDA-receptor dynamics are

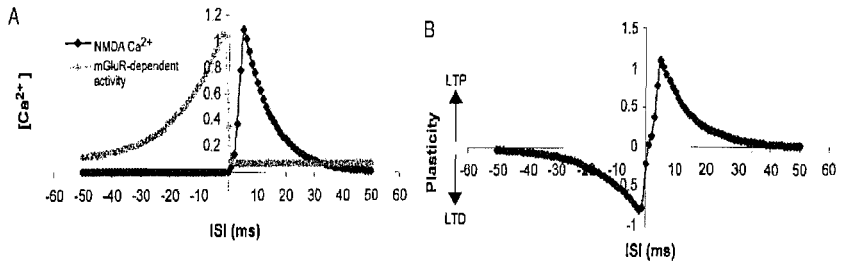

FIGURE 4.13: (**A**)The integral values for the proposed NMDA-receptor LTP pathway (black) and VGCC-mGluR LTD pathway (grey). By introducing two separate coincidence detectors, the two process may be kept distinct. There is no overlap between the LTP process and the LTD as each is only activated by a specific spike-ordering (either pre-post for NMDA-receptor-dependent LTP, or post-pre for VGCC-mGluR-dependent LTD). (**B**) The resulting STDP rule. The rule uses the integral values shown in (A) to determine the level of plasticity directly. The rule is biphasic, with a single LTP and a single LTD phase. The second LTD phase at large pre-post spike timings seen in single-detector models of the calcium control hypothesis is absent (from Karmarkar and Buonomano (2002)).

modelled according to a simple definition standard in the literature (Koch and Segev, 1998). In this scheme, the NMDA conductance, $g(t)$, is written as

$$g(t) = g_{max} \frac{e^{-t/\tau_a} - e^{-t/\tau_b}}{1 + \eta M e^{-\gamma V_m}}, \tag{4.24}$$

where g_{max} is the peak conductance, M the magnesium concentration, and $V_m(t)$ the postsynaptic depolarisation. The time constants $\tau_a \ll \tau_b$ represent slow and fast components of activation, respectively. η and γ are constants of proportionality, and we refer the reader to the original work for further details (Saudargine and Porr, 2005).

The authors assume that both increases (influx) and decreases (elimination) of intracellular calcium depend directly on the postsynaptic membrane voltage. Specifically, the derivative of the postsynaptic voltage, $\dot{V}_m$, is translated directly into influx (positive) and elimination (negative). By using the derivative of V_m, rather than its instantaneous value, the level of intracellular calcium will depend on the temporal profile of postsynaptic depolarisation not simply its magnitude. This approach is different to that of, for example, Shouval et al., (2002) discussed above and, as we will see, prevents the prediction of an additional LTD window at large pre-post spike timings.

Before taking the derivative, the postsynaptic membrane potential is convolved with a low-pass filter. This filter is chosen to reproduce the steep-rise and slow

decay of intracellular calcium transients (Sabatini and Oertner, 2002). The resulting function is termed the "postsynaptic influence", and written

$$F(t) = \frac{d\left(V_m(t)h(t)\right)}{dt} \tag{4.25}$$

where $h(t)$ is the low-pass filter. Defining μ to be some learning rate, the afferent plasticity induced from a spike-pairing of time difference Δt, denoted Δw, is then defined to be

$$\Delta w(\Delta t) = \mu \int_{\Delta t}^{\infty} F(t')G_N(t' + \Delta t) \tag{4.26}$$

for $\Delta t > 0$ (pre-post pairings) or

$$\Delta w(\Delta t) = \mu \int_{\Delta t}^{\infty} F(t' - \Delta t)G_N(t') \tag{4.27}$$

for $\Delta t < 0$ (post-pre pairings), where $G_N(t)$ is the NMDA-receptor conductance. The temporal ordering of pre- and postsynaptic events, which are represented by $G_N(t)$ and $F(t)$, respectively, are implemented in these equations by a temporal shift of Δt in one of the functions.

Thus, the derivative of the filtered postsynaptic membrane potential, $F(t)$, is multiplied by the NMDA-receptor conductance, $G_N(t)$, and integrated over all time $t > \Delta t$. The result is then translated directly into a change in input strength. There is no need to postulate any form of biphasic learning function to account for LTD, as the derivative of the filtered postsynaptic voltage may be positive or negative. Note, however, that this still requires that machinery exist that is capable of translating the postsynaptic voltage derivative, via this integral, into a graded change in input strength.

The resulting plasticity rule produces the biphasic, STDP learning curve shown in Fig. 4.14A. In accordance with the majority of experimental results, the rule predicts a single LTP window and a single LTD window (Bi and Poo, 1998; Zhang et al., 1998). The basic phenomenology of STDP, that isolated spike-pairings can produce bidirectional changes dependent on the precise spike-timing, is therefore reproduced. Increasing the time constant governing the membrane voltage, so that the temporal profile is prolonged, produces a curve where LTD is increasingly dominated by LTP (Fig. 4.14B). Thus, at different synaptic sites the same STDP rule is of potentially different character, depending on local conditions. For example, at synapses further from the soma, where the postsynaptic depolarisation arising from a PSAP is broader and more sustained, the STDP learning

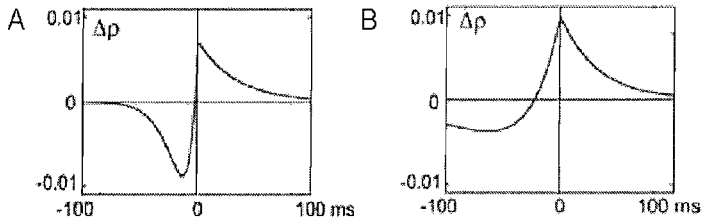

FIGURE 4.14: The STDP-like learning rule of the Saudargine model. The predictions of the model depend critically on postsynaptic membrane potential dynamics. In **(A)** the time constant governing the membrane voltage, τ_m, is set to 10ms. The resulting STDP curve has a roughly equal balance of potentiation and depression, and is in qualitative agreement with the STDP curves observed in experimental work (Bi and Poo, 1998). In **(B)** τ_m = 40ms, resulting in a prolonged, slowly rising depolarisation after a PSAP. This translates into a STDP curve dominated by potentiation. (from Saudargine and Porr (2005)).

rule will be dominated by LTP. A simulated increase in magnesium concentrations suppresses the STDP curve of the Saudargine model but does not change its overall shape.

In summary, the Saudargine model is a calcium-based, biophysical of STDP broadly in line with the calcium-control hypothesis. A central postulate is that postsynaptic calcium concentration depends directly on the postsynaptic membrane potential, V_m. Importantly, this dependence is on the gradient of V_m, rather than its absolute value. When the calcium concentration is translated into an evoked plasticity, a rule is produced that is STDP-like in nature, with a biphasic learning curve. The learning curve does not predict an additional LTD phase at large pre-post spike timings, which is a robust feature of other, purely concentration-based, approaches (Shouval et al., 2002).

4.4 Summary

A great variety of models of activity-dependent neuronal plasticity have been proposed. These range from simple rate-based, rules to more recent, biologically plausible, spike-based ones.

Rate-based plasticity rules describe the activity of pre- and postsynaptic neurons by the recent time-averaged firing rate. Much of the early work on activity-dependent neuronal plasticity was conducted within such a framework, including that of the Hebb-rule along with its various permutations. Rate-based models have been successfully applied to a range of problems, such as the activity-dependent development of orientation selectivity in the developing visual cortex.

However, due to their intrinsic structure, rate-based models are unable to accommodate more recent experimental results which suggest that it is the timing of individual action potentials, rather than their mean rate of arrival, that determines the degree and polarity of change. Spike-based plasticity rules do not suffer from such a limitation as they explicitly consider individual pre- and post-synaptic spiking events. They may be divided into two groups; phenomenological and biophysical. Biophysical models aim to show that particular input plasticity rules can be explained in terms of the complex biological mechanisms thought to underlie plasticity. Such models can often be convoluted, with poorly constrained parameters, and offer little scope for generalisation. Phenomenological models place more emphasis on the computational properties of a particular learning rule, standing some way above the actual biological implementation, and as a result are often simpler in nature. However, by their very nature, these models make the implicit assumption that the STDP-curve is valid across individual afferents, and for individual spike pairs. Although certainly possible, it is not necessarily true that the observed STDP-curve is valid on any other level than this multi-afferent and multi-spike pair level. Indeed, a large proportion of experimental work on STDP has measured the overall plasticity of several afferents in response to multiple spike pairings. We explore the implications of this approach more fully in Chapter 6, where we show that the observed STDP-curve may, indeed, emerge due to the ensemble averaging of individually much simpler changes. We also show that there are several possibilities as to the locus of such an averaging process, so that the averaging may take place at either the input, spike-pair, or synaptic level.

Chapter 5

The Song Model of Neuronal Plasticity

A standard implementation of a phenomenological STDP-like learning rule is the Song model (Song et al., 2000). Although the Song model is based mainly on the data of Bi and Poo (1998), the form of the learning rule is not in direct conflict with the majority of experimental results (but see Bell et al. (1997)). The original work focused on the competitive and stabilising properties of the learning rule when applied to multiple afferents innervating a single target cell. In this Chapter we reproduce these simulations, then explore more fully both the parameter space and the underlying assumptions of the model. This illustrates issues common to several existing formulations of STDP and motivates our own model of STDP presented in Chapter 6.

5.1 Formulation of Model

Experimental work has shown that the activity of an afferent and its target cell can trigger an adjustment of afferent efficacy dependent on the time delay between pre- and postsynaptic firing, a phenomenon known as timing-dependent long-term potentiation STDP. The closer pre- and postsynaptic firing times are, the larger the change in connection strength. In the Song model, the experimentally observed spike-timing dependent plasticity results are taken over directly to form a simple learning rule describing the plasticity of individual inputs. A biphasic, exponential relationship between the timing of pre- and postsynaptic

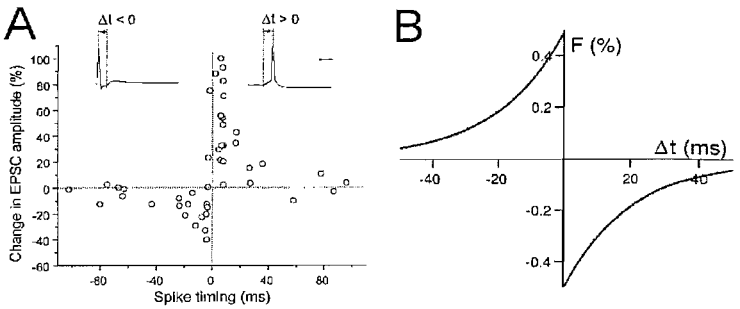

FIGURE 5.1: Timing-dependent plasticity; **A** the STDP curve measured in experimental work on rat hippocampal slice cultures (from Bi and Poo (1998)) and **B** the learning rule of Song et al. (2000) inspired by it (from Song et al. (2000)). The two curves are qualitatively similar, with their apparently opposite polarities simply arising from a difference in the definition of Δt.

spikes and the magnitude of afferent plasticity is postulated. The change in input strength, ΔA, given a spike time interval, t, is given by

$$\Delta A(t) = \begin{cases} g_{max} A_+ e^{-t/\tau_+} & \text{for } t \geq 0 \\ -g_{max} A_- e^{+t/\tau_-} & \text{for } t < 0 \end{cases}, \tag{5.1}$$

where $A_\pm$ are the amplitudes of plasticity, and $\tau_\pm$ the characteristic time scales over which plasticity occurs. The quantity g_{max} represents the maximum input weight. Fig. 5.1A shows the STDP curve from Bi and Poo, (1998) and Fig. 5.1B the resulting learning rule from Song et al., (2000).

The learning rule is constructed by making a simple exponential fit to the experimental data. This guarantees that the characteristic biphasic, exponential form of the STDP curve is reproduced. The parameters of the learning rule may be chosen to give the best possible match to the data. Experimentally, the magnitude of potentiation is greater than that of depression, but the depressive phase has an apparently larger time constant. Integrating the learning curve over all time takes into account the balance of these two factors and gives an indication of whether, overall, potentiation or depression dominates. A negative integral indicate that depression is dominant over potentiation. We show later that in order for the learning rule to exhibit competitive dynamics, this integral must indeed be negative.

In numerical simulations, discrete time steps of 10^{-4}s were used. This value is 1/200th of the STDP-curve decay constants and gives an appropriate temporal resolution. The postsynaptic target cell is modelled as a leaky, integrate-and-fire neuron (Dayan and Abbott, 2001), with 1000 excitatory and 200 inhibitory inputs. The spike-timing plasticity rule is applied to excitatory afferents, which

change their weight according to the correlations between afferent and target cell firing. The inhibitory synapses are non-plastic, with a fixed, uniform efficacy.

Each afferent receives uncorrelated spike trains generated by independent Poisson processes. Poisson spike trains are known to approximate spontaneous activity seen vivo, with a highly variable interspike interval (ISI). However, ISIs in vivo are often correlated and neurons often exhibit bursting behaviour, especially during learning episodes (Bair et al., 1994). Although the excitatory and inhibitory afferents receive spike trains at different rates, the firing rates are uniform within each group.

Although the equations governing the system are developed using a conductance-based approach, the actual implementation treats this conductance as a direct measure of afferent strength. This measure is more conventionally called the "synaptic efficacy", but we make a more careful distinction between synaptic and afferent strengths. It is implicit that afferents make multiple synaptic contacts with target cells, and that the efficacy of these individual contacts is summed in some manner to give the total afferent strength. The Song model is a description of afferent plasticity, and contains no explicit treatment of individual synaptic strengths. We will return to this issue in the discussion. The terms "strength" and "weight" are equivalent and may be used interchangeably, as may the terms "afferent" and "input"'.

Under the integrate-and-fire neuron approximation, the various processes of synaptic transmission which have a finite time course (such as neurotransmitter vesicle binding, diffusion of neurotransmitter across the synaptic cleft, ligand binding and ionic channel dynamics) are typically assumed to be instantaneous. Here, we include an element of Gaussian noise in spike timings to reflect both the noise in synaptic transmission, as well that from other experimental sources.

The plasticity rule is implemented using a set of decaying functions. p_i is associated with the afferents, where $i = 1, 2, 3...$, and p_x associated with the post-synaptic target cells, where $x = 1, 2, 3....$ When a pre- or postsynaptic spike arrives, p_i or p_x is incremented, respectively. In the absence of further firing, the $p_{i,x}$ decay exponentially

$$\tau_{\pm}\frac{dp_{i,x}}{dt} = -p_{i,x} \tag{5.2}$$

where $\tau_{\pm}$ are the decay constants. The use of these functions is a computational device to simplify the implementation of the learning rule. Instead of recording an exponential decaying "interaction function" for each pre- or postsynaptic event, the algorithm simply makes injections to two exponentially decaying functions that combine all the contributions of either presynaptic or postsynaptic

firing into a single value. It can be intuitively seen that this approach only works if evolving a sum of n values according to an exponential decay is equivalent to evolving those n values independently, then summing at the end. It is easy to show this is true by considering a sequence of n presynaptic events at times $(-t_i)$ where $i = 0, ..., n$. Examining the interactions function, p_i, at some later time, $t > 0$, we have that

$$p_i = p_i(-t_i)e^{-(t+t_i)}, \tag{5.3}$$

where $p_i(-t_i)$ is the initial value of function i. The sum of the n functions at time t is then

$$p_1 e^{-(t+t_1)} + p_2 e^{-(t+t_2)} + ... + p_n e^{-(t+t_n)}, \tag{5.4}$$

where we have dropped the time argument of the initial values for clarity. However, instead of evolving n separate functions then summing, we could write

$$p_i = p_i e^{-t} e^{-t_i}, \tag{5.5}$$

so that their sum is

$$(p_1 e^{-t_1} + p_2 e^{-t_2} + ... + p_n e^{-t_n})e^{-t} = (p'_1 + p'_2 + ... + p'_n)e^{-t}. \tag{5.6}$$

Thus, a simple rewriting of the terms allows us calculate sum of primed values once then evolve this sum, rather than evolving and summing all n non-primed values individually. Subsequent presynaptic events are adequately accommodated by making a step-contribution to the sum, which then continues to evolve as a single entity.

This simplification only works if the interaction functions evolve according to an exponential decay, and therefore possess the lack of memory property. If a different interaction function were used this simplification could not be made. We also require that the interactions from multiple spike-pairings are assumed to sum linearly. If this were not the case, the origin of the contributions to the overall interaction function not just their combined value would be important for determining plasticity. This method contrasts to the more intuitive, but computationally intractable, approach where the time step at which each event takes place is recorded and compared to all past events (see below). Under a simple one step Euler method, Eqn. 5.2 becomes

$$p_{i,x} \rightarrow p_{i,x} + \Delta t \frac{\delta p_{i,x}}{\delta t} \tag{5.7}$$

where Δt is the discrete time step. An action potential occurring at input i triggers a postsynaptic depolarisation in proportion to that inputs strength. Two additional plasticity-related computations are also made. Firstly, p_i is incremented by an amount A_+ and secondly, the afferent efficacy of the firing input g_i is reduced by an amount $p_x g_{max}$. When inhibitory inputs fire, they exert an inhibitory influence on the target cell but do not engage in plasticity. Each time a postsynaptic neuron is driven beyond threshold, an action potential is triggered and two plasticity computations occur. Firstly, p_x is decremented by an amount A_-, and secondly, every excitatory input has its efficacy increased by $p_i g_{max}$. At all times, input weights are constrained to be positive, so that $0 \leq g_i \leq g_{max}$.

This computational scheme uses the functions p_i and p_x to encode the two exponential phases of the learning rule. At each time step the five equations describing the evolution of the state variables are numerically integrated using a one-step Euler method (Gerhald and Wheatley, 1994). The order of integration is arbitrary and has no effect on the equilibrium distributions produced. There are no explicit absolute or relative refractory periods. Both the afferents and the target cell are therefore able to fire an action potential at every time step, although this rarely happens in practice.

In summary, the plasticity functions associated with each afferent and target cell receive one time injections whenever an appropriate spike occurs, and then decay according to an STDP-like exponential curve. The functions are recalled at a later time and used to make appropriate changes in afferent efficacy. This slightly esoteric approach is purely a computational device adopted for tractability. Consider, for example, the more intuitive approach of simply recording the firing time of every pre and postsynaptic event and comparing them under the STDP-like learning rule. As interactions between spike pairs are temporally unconstrained, each pre- or postsynaptic spike interacts with all preceding pre- and postsynaptic spikes. Once a simulation begins, the number of pre- and postsynaptic spikes since the start of the simulation increases very rapidly, and the number of interactions increases in a similarly rapid fashion. To run a full simulation would typically require around 10^{12} spike times to be recorded, and an equal number of exponential calculations to be made whenever a new spike occurred. Although possible, in practice this approach is computationally intractable. Introducing temporal constraints, such as a finite temporal window in which spikes interact, would address this issue but only at the expense of introducing additional, arbitrary, constraints into the model.

5.1.1 Parameter Selection

Experimentally, multiple spike pairings (60 pairings at 1Hz) were used to evoke statistically meaningful changes in evoked postsynaptic excitatory potentials (EPSPs) (Bi and Poo, 1998). Due to the variability in initial input strength, this data was presented in the form of a percentage change in input strength. In order to obtain a value for the magnitudes of plasticity, $A_{\pm}$, it is assumed that the effect of multiple spike pairing events sum linearly, so that $A_{\pm}$ can be found by simply dividing the total change by the number of pairings that caused it. The values thus obtained are, however, only approximate $A_{\pm}$ as the additional scaling of all afferent strengths by g_{max} is ignored. The values for $A_{\pm}$ that are used, 0.005 for potentiation and -0.00525 for depression, are therefore only roughly consistent with the experimental data.

The decay constants, $\tau_{\pm}$, determine the size of the time window over which spikes interact and cause a change in afferent efficacies. By making them larger, an afferent can fire further in advance of postsynaptic firing and still enjoy an up regulation of efficacy. The same will be true for depressive pairing when the order is reversed. Making it smaller has the opposite effect, reducing the time window in which pre- and postsynaptic spiking can effectively interact. Experimental data on the width of the STDP window are mixed. Some results have suggested that the two sides are of roughly equal duration (Markram et al., 1997; Bi and Poo, 1998; Zhang et al., 1998), others that the depressive phase is longer (Debanne et al., 1998; Feldman, 2000). As experimental results are drawn from a variety of preparations, this may reflect functional as well as experimental differences. The dynamics of the Song model are insensitive to minor differences in the time constants, $\tau_p m$, as long as the constraint $A_+\tau_+ \leq A_-\tau_-$ is maintained. We set $\tau_{\pm} = 20$ms in accordance with the original paper. The maximum synaptic efficacy, g_{max}, is set to 0.015. This corresponds to a peak synaptic conductance of 150pS, which lies within accepted experimental bounds (Hille, 1992). The inhibitory synapses are non-plastic, with fixed conductances of 0.05.

Noise in the timing of spikes, reflecting both experimental error and variable transmission times, is drawn from a Gaussian distribution with standard deviation of 1ms.

The action potential threshold for the integrate and fire neuron was set to -54mV, resting potential to -70mV, inhibitory and excitatory conductance decay constants to 5ms, and membrane time constant to 20ms. These values are all roughly standard from the literature. The excitatory and inhibitory reversal potentials are set to the 0mV and -70mV, respectively. The reset potential was set to -60mV by matching the neuronal gain of a simple integrate and fire model to match the observed value for pyramidal and sparsely spiny stellate neurons of

the neocortex (Troyer and Miller, 1997). The dynamics of the integrate and fire neuron, when acting as a simple integrator of synaptic input, then give a reasonable approximation to the highly variable ISIs for action potential production as seen in vivo without resorting to balanced inhibition or correlated inputs. Song et al., suggest that this is a desirable property when modelling cortical neurons. However, we show later that the dynamics of the learning rule are relatively insensitive to changes in the reset potential. At the start of each simulation, the target neuron was initially set to its resting potential of -70mV.

5.1.2 Computational Scheme

After initialisation of parameters and clearance of data storage variables and arrays, the following steps are repeated until a fixed amount of simulated time has elapsed.

1. Decay of state variables: State variables g_{ex}, g_{in}, p_x and p_i decay, subject to a lower bound of zero.

2. Compute excitatory input For each excitatory input, Poisson statistics are used to determine if an excitatory afferent fires an action potential or not. If if does, g_{ex} is incremented by the afferents weight and g_i is weakened by m_i, subject to a lower bound of zero. p_i is also incremented by A_+.

3. Compute inhibitory input: For each inhibitory input, Poisson statistics are used to determine if an inhibitory afferent fires an action potential or not. If it does, g_{in} is incremented by the afferents weight.

5. Integrate and fire: Rate of change of membrane potential calculated for the leaky, integrate-and-fire target neuron. The target neuron membrane potential is incremented using a one step Euler method. If the firing threshold is exceeded, a postsynaptic action potential is triggered. v then changes immediately to the reset potential, and m_x is decremented by A_-. The input weights, g_i, are then strengthened by p_i, subject to an upper bound of g_{max}.

5.2 Results

We now reproduce the main findings of the original work on the Song model (Song et al., 2000). In most cases, 1000 seconds of simulated time is sufficient for the distribution of afferent efficacies to reach a dynamic equilibrium. At this point, individual afferents may still move between high and low strength with time, but no change is seen in the overall distribution beyond those associated with the stochastic nature of the simulation. As input firing rates are reduced,

or the ratio A_-/A_+ is decreased, the learning rate naturally falls. This results in a dramatic increase in convergence time, and the simulation must run for longer for the system to reach equilibrium.

5.2.1 Input-Rate Adaptation

The 1000 excitatory synapses were given an initial efficacy of g_{max}. This produces a very high postsynaptic firing rate, with the target neuron firing regularly, as would be expected when averaging the input of a large number of afferents of roughly equal weight. The postsynaptic spiking at this stage is largely independent of individual presynaptic spike timing, so approximately equal numbers of presynaptic spikes will fall either side of the postsynaptic spike. The negative integral of the learning rule STDP-like learning curve means that, on average, the inputs will be weakened. As equilibrium is approached, a more balanced distribution is produced with total afferent input maintaining the postsynaptic neuron very near to its firing threshold. The result is postsynaptic firing that occurs in a manner similar to a Poisson spike train, with postsynaptic spiking occurring in response to random fluctuations in the total afferent input.

As random correlations in input firing are now needed to evoke a postsynaptic spike, more presynaptic events tend to occur before a spike than after. Other afferents continue to weaken in accordance with the overall dominance of depression. Once a group of afferents begin to strengthen they begin to win more and more control of the postsynaptic spiking. The remaining afferents weaken further, eventually losing all control of the postsynaptic cell. The learning rule is therefore competitive, with some afferents increasing in strength at the expense of the others.

Eventually, an equilibrium distribution arises. Afferent strengths fall in the range defined by the hard upper and lower bounds, between zero to g_{max}. Fig. 5.2 shows a histogram of the equilibrium efficacy distribution for excitatory input firing rates of 10Hz, and Fig. 5.3 shows the same for input firing rates of 40Hz. All other parameters were held fixed. We see the characteristic bimodal distribution of this model (Song et al., 2000), with afferent efficacies clustering around the upper and lower bounds. A lower excitatory firing rate causes more efficacies to approach the upper bound, which may be interpreted as a kind of input-rate adaptation where the total synaptic input (the product of afferent efficacy and firing rate, summed over all inputs) is kept roughly constant across different firing rates. As a result, the postsynaptic firing rate is largely insensitive to the presynaptic firing rates. The initial distribution of afferent efficacies has no effect on the equilibrium distribution. This is assuming, of course, that there

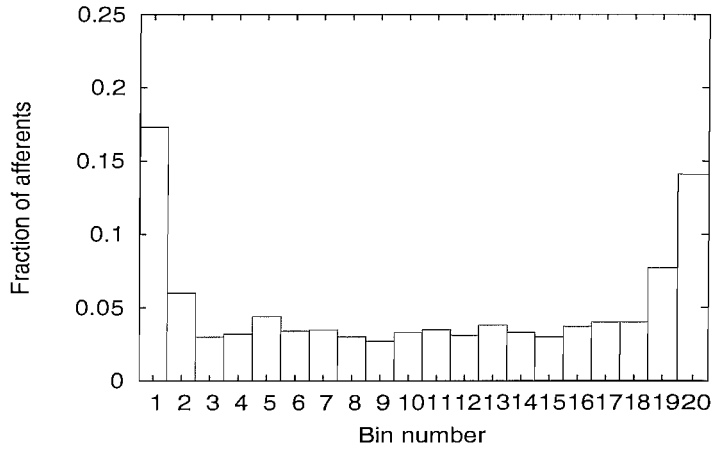

FIGURE 5.2: Equilibrium distribution of afferent efficacies for excitatory input firing rates of 10Hz. The characteristic bimodal distribution of the Song model can clearly be seen. The afferent efficacies lie between zero and $g_{max} = 0.015$, so that each bin represents an interval of 0.00075.

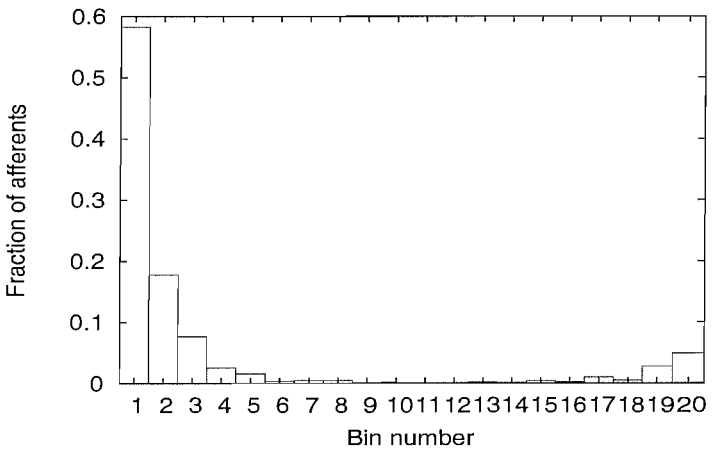

FIGURE 5.3: Equilibrium distribution of afferent efficacies for excitatory input firing rates of 40Hz. The bimodal distribution is shifted towards zero, so that the cluster of inputs with efficacy around g_{max} is much reduced. The afferent efficacies lie between zero and $g_{max} = 0.015$, so that each bin represents an interval of 0.00075. The postsynaptic firing rate is, as a result, almost unchanged (an effect referred to as input-rate-adaptation).

is sufficient initial input to trigger postsynaptic spiking, as in the absence of postsynaptic firing no plasticity is evoked.

The excitatory input firing rate was varied from 10 to 50Hz, with all other parameters held fixed. Fig. 5.4 shows the effect on the postsynaptic firing rate and coefficient of variation. In the figure, the postsynaptic firing rate has been scaled by a factor of 1/10. The postsynaptic firing rate increases by around 1Hz for each

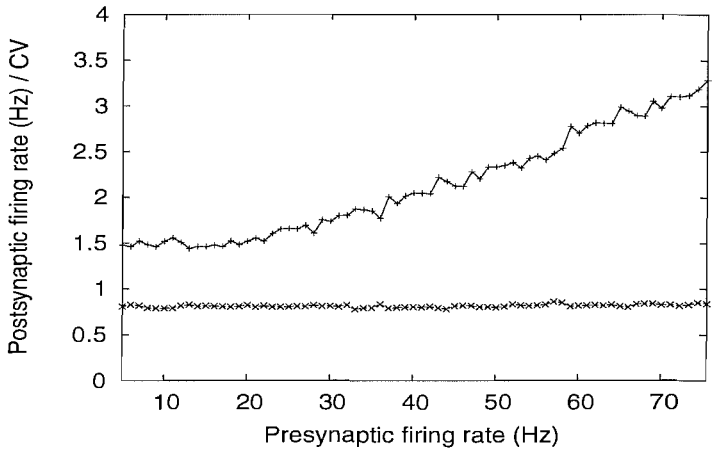

FIGURE 5.4: Effect of changing the presynaptic firing rate on the postsynaptic firing rate (top line) and coefficient of variation (bottom line). The postsynaptic firing rate has been scaled by a factor of 1/10, and the data points are the average of 10 runs.

5Hz increase in presynaptic rates, demonstrating the regulation of postsynaptic firing rates under this learning rule. The coefficient of variation is high, and the neuron remains highly sensitive to random fluctuations in the input firing rates. Fig. 5.5 shows the ratio of total inhibitory to total excitatory synaptic input. There is a slight dominance of inhibition across the whole range of presynaptic firing rates, and this balance of inhibition and excitation contributes to the high variability of postsynaptic spiking. The percentage of strong excitatory afferents (were strong is defined as $A_i \geq 0.8 g_{max}$) decreases with increased firing rate. This provides an additional source of variability in postsynaptic spiking as several excitatory inputs must fire together in order or trigger a postsynaptic action potential.

The ratio of A_-/A_+ determines whether potentiation or depression is dominant (as $\tau_+ = \tau_-$, always). For $A_-/A_+ \geq 1$ the integral of the STDP-like learning curve is negative, and depression is, on average, dominant. We vary the value of the ratio from 1.05 to 1.20. Fig. 5.6 shows the equilibrium postsynaptic firing rate, and coefficient of variation. As the ratio increases, the equilibrium postsynaptic firing rate decreases and the coefficient of variation increases in a non-linear fashion.

As demonstrated in the original work, the bimodal distribution of afferent efficacies can be "filled in" to some degree if g_{max} is made 2.33 times higher, and both A_- and A_+ are made 4 times higher (Song et al., 2000). As Fig. 5.7 shows, this empirical choice of parameters preserves the basic characteristics of the learning rule, such as postsynaptic rate and coefficient of variation regulation, while

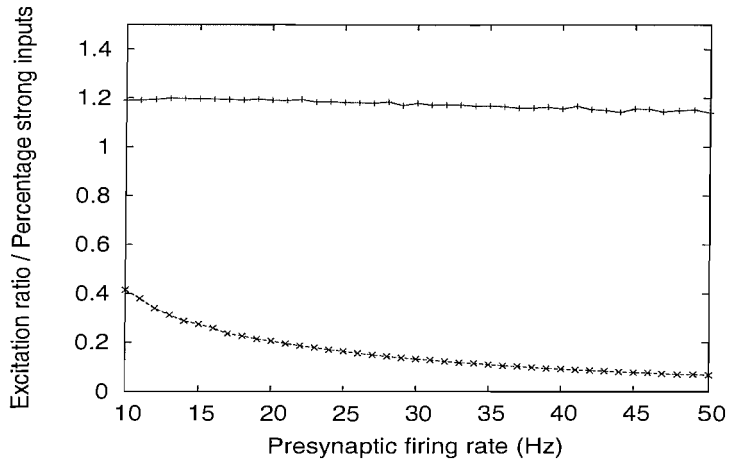

FIGURE 5.5: Effect of changing the presynaptic firing rate on the ratio of total inhibitory to excitatory input (top line) and the percentage of strong inputs (defined as $\geq 0.8 g_{max}$, bottom line). The data points are the average of 10 runs.

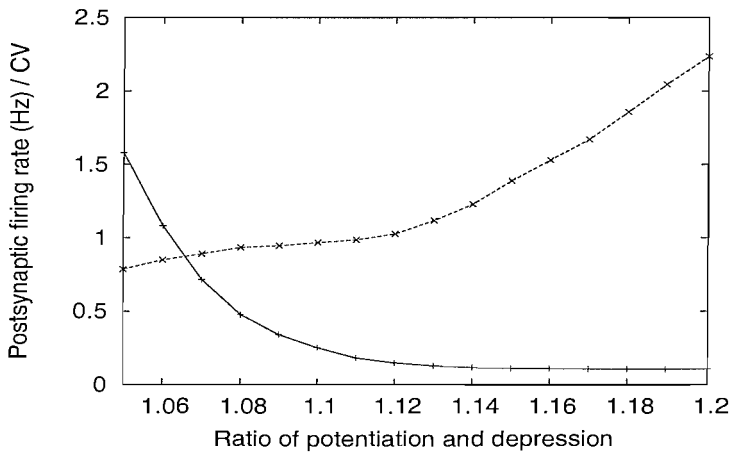

FIGURE 5.6: Effect of changing the dominance of depression over potentiation on the postsynaptic firing rate (decreasing) and coefficient of variation (increasing). The postsynaptic firing rate has been scaled by 1/10, and the data points are the average of 10 runs.

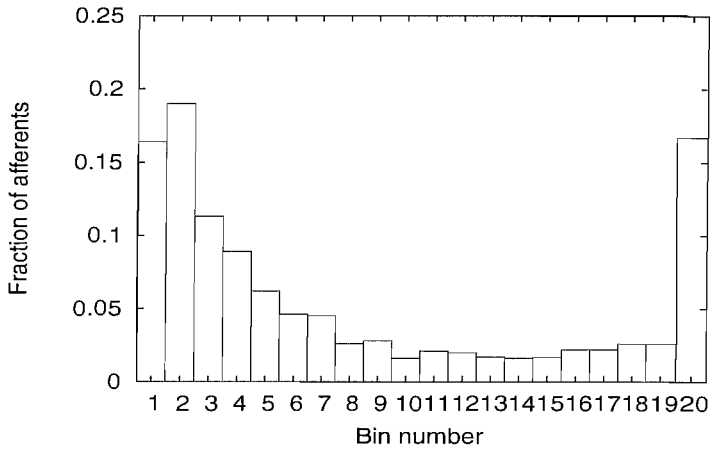

FIGURE 5.7: Equilibrium distribution of afferent efficacies at 10Hz, with g_{max} 2.33 times larger, and A_-/A_+ 4 times larger. The bimodal distribution is still present, but has been "filled-in" to some extent. The afferent efficacies lie between zero and $g_{max} = 0.015$, so that each bin represents an interval of 0.00075.

producing a smoother afferent efficacy distribution. Experimental work measuring the size of miniature EPSPs (which may be interpreted as an indication the afferent efficacy) suggests that such a distribution is more biologically plausible.

5.2.2 Latency Reduction

We now explore the dynamics of the Song rule in response to presynaptic burst firing. We continue to set $A_+ = 0.005$, $A_+/A_- = 1.05$, and $\tau_\pm = 20$ms. Both excitatory and inhibitory afferents fire at a rate of 10Hz, and have initial synaptic weights equal to $0.2g_{max}$. In addition to this background firing, each excitatory afferent was made to fire a single burst of action potentials ten times a second. The bursts consisted of a Poisson spike train at 100Hz lasting for 20ms. These bursts were not synchronous, with each excitatory afferent assigned a fixed random latency, drawn from a Gaussian distribution with zero mean and standard deviation of 15ms. Excitatory afferents with low latencies fired earlier during each bursting phase, and as a result tended to precede the postsynaptic response. Under this learning rule, these low latency afferents are preferentially strengthened. Afferents spiking with higher latency tended to spike after the postsynaptic response, and as a result were weakened. The simulation ran for 1000 seconds of simulated time. This was sufficient to reach an equilibrium in afferent efficacies, as the spread in latencies creates an environment where the final strength of each synapses converges rapidly and predictably.

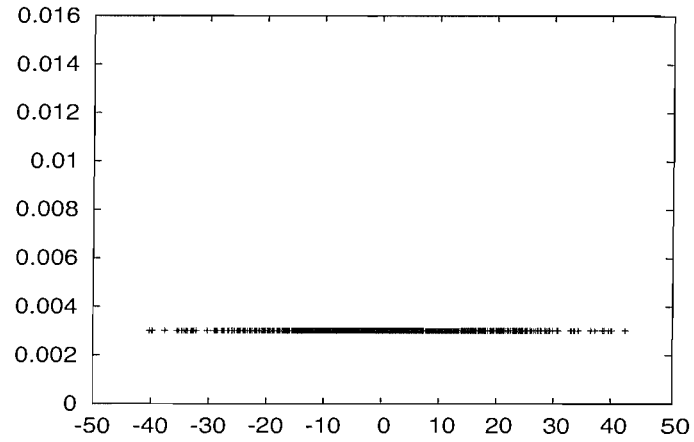

FIGURE 5.8: The initial strengths of the excitatory inputs, as a function of assigned latency.

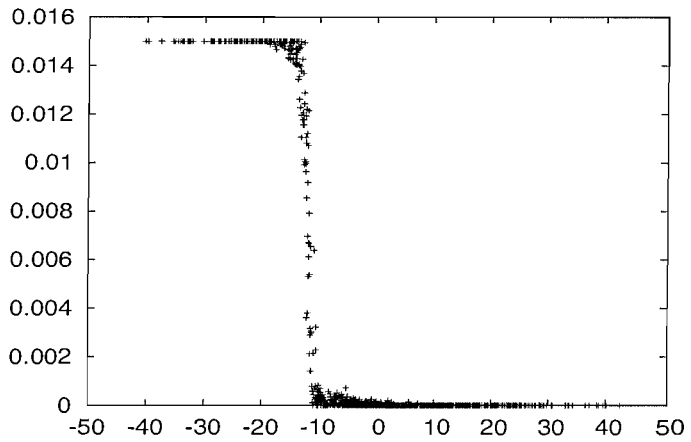

FIGURE 5.9: The equilibrium strengths of the excitatory inputs, as a function of assigned latency. Low latency inputs have invariably been strengthened to g_{max} and high latency inputs weakened to zero. A sharp transition exists between the two groups.

The initial afferent efficacy distribution, as a function of assigned latency, is shown in Fig. 5.8. All afferents have an efficacy of $0.2g_{max}$, and the distribution is Gaussian. Fig. 5.9 shows the equilibrium distribution. Afferents with lower latency have invariably been strengthened to the maximum weight allowed. Higher latency inputs have invariably weakened and remain very small. A sharp transition may be seen between the two populations. Thus, at equilibrium, the target neuron responds much quicker to the onset of afferent bursting, as the low latency afferents effectively control postsynaptic spiking.

5.2.3 The Reset Potential

In the original paper, the reset potential of the integrate and fire neuron was set to -60mV by matching the neuronal gain of a simple integrate and fire model to match the observed value for pyramidal and sparsely spiny stellate neurons of the neocortex (Troyer and Miller, 1997). The authors claim that the dynamics of the integrate and fire neuron, when acting as a simple integrator of synaptic input, then give a reasonable approximation to the highly variable ISIs for action potential production as seen in vivo without resorting to balanced inhibition or correlated inputs. Song et al., suggest that this is a desirable property when modelling cortical neurons. It is important, however, to explore the effect that this somewhat counterintuitive value (a more standard value would be somewhere around -90mv) has on the learning rule. We therefore repeat the input-rate adaptation simulations at 10Hz, using the more standard reset potential of -90mv. We find that the characteristic bimodal distribution of equilibrium afferent efficacies is qualitatively unchanged. The slight differences between the two distributions are due to the stochastic nature of the simulation. This result is expected, as the competitive dynamics that produce the bimodal distribution arise directly from the use of a non–linear integrate–and–fire neuron. As this mechanism is still present, we would expect qualitatively similar dynamics. Certain properties of the postsynaptic neuron, such as the variability of inter–spike–intervals, are, however, influenced by the choice of reset potential Song et al. (2000). We conclude, therefore, that the key properties of the STDP–like learning rule, such as the presence of competitive dynamics, are insensitive to the precise value of the reset potential of the integrate and fire neuron.

5.2.4 Stability

It is important to examine the stability of a learning rule to determine which of its properties are more than just transient behaviours. The characteristic, bimodal, equilibrium distribution of input weights of the Song model, shown in Fig. 5.2 is described in the original paper as being "stable". However, this stability is only in the sense that the overall distribution does not change over time. Individual inputs may still change, moving about within the distribution without affecting its overall shape. It is entirely possible, therefore, for an input of high strength (one that is part of the cluster around g_{max}) to weaken, pass through the intermediate region, and eventually become an input of low strength in the cluster around zero. Fig. 5.11 shows the evolution of five afferents during a simulation lasting 5000 seconds. This duration is approximately five times that required to reach a bimodal distribution of input weights. We see, as expected, that by 1000 seconds all of the afferents have joined one of the two clusters around

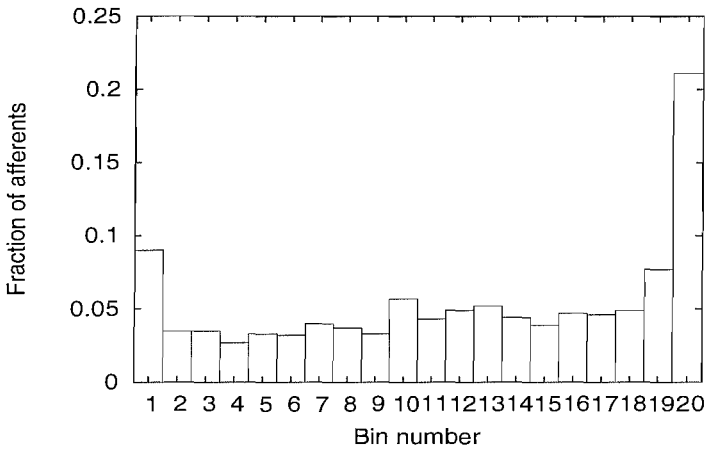

FIGURE 5.10: Equilibrium distribution of afferent efficacies at 10Hz using a reset potential of -90mV. The characteristic bimodal distribution is qualitatively unchanged from the case where the reset potential is -65mV. The afferent efficacies lie between zero and $g_{max} = 0.015$, so that each bin represents an interval of 0.00075.

zero or g_{max}. However, over the next few thousand seconds two of the afferents independently transition from one extremum to the other. A simple numerical investigation confirms that this observation holds for any value of g_{max} in the region of those chosen in the original work. Thus, the bimodal distribution is stable when viewed as a whole but unstable when examined at the level of individual afferents. Scaling the values of $A_{\pm}$ do not affect this result. This is expected, as it can be intuitively seen that it is the non-zero probability of finding afferents between the clusters at zero and g_{max} (the "filled-in" nature of the distribution) that permits gradual transitions between the two extrema. Scaling $A_{\pm}$ leaves this distribution unchanged, altering only the length of time taken by afferents to move around within the distribution. That is, scaling $A_{\pm}$ simply alters the overall learning rate.

Under the parameters chosen in the original paper, then, the Song model is only stable in the sense that the overall distribution of input weights does not change. Individual inputs are not stable as they may flip from a state of high strength to one of low strength over a long enough time period. Even over relatively short periods, inputs are capable of undergoing quite large fluctuations in strength. The segregation of inputs that occurs in the Song model is therefore of a quite different nature to, say, that in the BCM-rule (Bienenstock et al., 1982) in which a single input reliably wins control of the target cell and all other inputs are suppressed. In the BCM-rule, the segregation is stable so that the winning afferent retains control of the target cell and does not subsequently lose control to a different input due to some large fluctuation in strength.

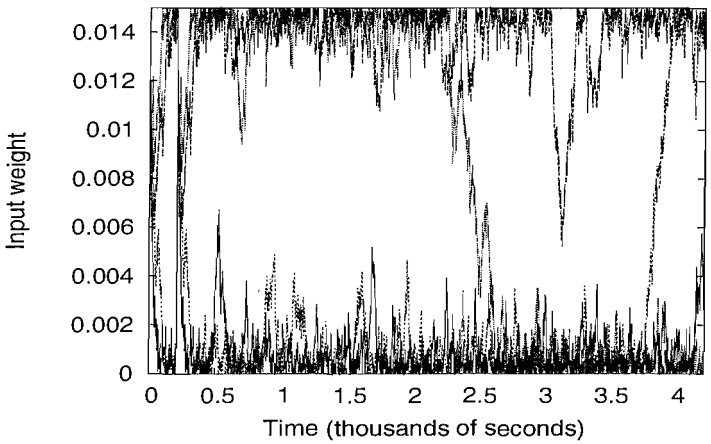

FIGURE 5.11: Evolution of five afferents over 5000 seconds of stimulated time. Afferents initially segregate into two clusters around zero and g_{max}. These two clusters underlie the characteristic bimodal distribution shown in Fig. 5.1 and Fig. 5.2. Subsequently, two afferents transition from one of the extremes to the other. The bimodal distribution is therefore only stable in a global sense, not from the point of view of individual afferents. Parameters are as for Fig. 5.2.

It is possible to stabilise the Song model somewhat by choosing a much higher value for g_{max}. Although this issue was not addressed in the original work, our own simulations show that with a higher g_{max} an equilibrium distribution may be produced where the majority of inputs are clustered around zero with a few, very high strength afferents clustered around g_{max}. Transitions between states of low and high strength then become very unlikely, and individual inputs as well as the overall distribution are stabilised. This occurs because, with a suitably high g_{max}, the bimodal distribution is prevented from being "filled-in", so that very few inputs with weight intermediate between the zero and g_{max} exist.

5.3 Comparison to Original Implementation

Our implementation of the STDP-like learning rule of Song et al., (2000) reproduces all their main findings. Small quantitative differences in the equilibrium distribution of afferent efficacy are to be expected, and are consistent with the stochastic nature of the simulations. The histograms showing the bimodal distribution of input weights characteristic of the Song model are taken from a single simulation, and do not represent data averaged over many repetitions. At best, therefore, a qualitative match between the original paper and our own simulations would be expected. Fig. 5.6 shows the effect of changing the dominance of depression over potentiation (the ratio A_-/A_+). For low values of $A_-/A_+ \leq 1.05$ there were significant quantitative differences when compared

to the original paper, although the qualitative picture is the same (data not shown). It is likely that this numerical difference stems from the fact that we use simulations of fixed duration (around 1000 seconds of simulated time), but the convergence time of the simulations increases dramatically as the integral of the STDP curves approaches zero. Reducing input firing rates also dramatically increases the convergence time, as more simulated time is needed to allow a sufficient number of spike pairings to occur for the system to reach equilibrium. A closer examination of the afferent efficacy distribution when the simulation terminates supports this conclusion. For $A_-/A_+ \leq 1.05$ the efficacy distribution is very similar to the initial conditions (unimodal, and clustered around g_{max}). Running the simulation for longer results in values closer to, but still significantly different to, those presented in the original work. However, in the absence of any discussion of these convergence issues in the original paper, we choose to fix our simulation duration at 1000 seconds and examine the areas of the curve where convergence is more likely, when $A_-/A_+ \geq 1.05$.

5.4 Discussion

The Song model is a phenomenological, additive STDP learning rule formulated by taking a typical experimentally observed STDP curve over directly to govern input plasticity. The proposed learning rule sets out explicitly how pre- and postsynaptic spiking may interact to drive changes in input strength. It has been shown in other modelling studies that the implementation details of a simple STDP-rule can have important consequences for the learning dynamics (van Rossum et al., 2000; Izhikevich and Desai, 2003). It is therefore important to critically evaluate the formulation of a model, and understand fully any assumptions behind it.

5.4.1 Formulation

Experimental results have suggested that the STDP curve has two distinct phases, one for potentiation with positively correlated pre- and postsynaptic spiking, and one for depression with negatively correlated spiking. Approximating the STDP curve with a double-exponential learning rule seems an obvious first choice, and this appears to give a reasonable match to the data of Bi and Poo (1998). However, the infinite tail of the exponential means that, in the context of a learning rule, every presynaptic spike will interact with every earlier postsynaptic spike (albeit to an exponentially diminishing degree). This lack of temporal restriction may appear trivial, but it is just such restrictions that

have been proposed elsewhere to connect simple STDP learning rules with the BCM-theory (Izhikevich and Desai, 2003).

The learning rule is assumed to be of a fixed form, treating each spike pair as an independent event, with the effects of multiple spike-pairs summing linearly. In a situation where three spikes occur, removing any one will not affect the contribution arising from the interaction of the other two in any way. Recent experimental results suggest that real spike pairs do not always sum linearly in such a simple manner (Froemke and Dan, 2002). In the triplet interactions investigated by Froemke and Dan(2002), repeated pairing of two presynaptic action potentials with a single postsynaptic event (60-80 pairings at 0.2Hz) in pyramidal neurons in layer 2/3 of rat visual cortical slices led to a supralinear overall change that cannot be explained under a simple, biphasic, exponential STDP curve. An additional non-linearity, in the form of spike suppression, is required to accommodate these spike-triplet results in simple STDP models such as the Song model.

A sharp transition is assumed to exist between the up and down phases of the learning curve. Although it is indeed possible that the "real" STDP curve has the same underlying form, perhaps masked by the presence of noise, this is certainly not the only possibility. For example, a much smoother transition may occur. The transition region is potentially very important, as this is the region where the largest plasticity events are evoked. This issue of the transition zone becomes even more important when inputs fire bursts of action potentials instead of single spikes as, with bursting inputs, a larger proportion of interactions would be expected to occur in the transition region.

Finally, we note that, in real biological systems, each afferent makes multiple synapses with its target cells. The total afferent efficacy is then the sum of all the individual synapse efficacies, and there is a distinction between individual synaptic and total afferent efficacy. Experimentally, measurements are invariably made of changes in overall input strength, or even of groups of inputs, rather than of individual synapses. The measured STDP rule is therefore also being measured at this input level. Although it is indeed possible that the STDP curve is respected at each individual synapse, it is not the only interpretation of spike-timing results. In the Song model, this distinction between input and synaptic efficacies is not considered. It is implicitly assumed that each afferent supports a single synapse onto the target, and that each is influenced by postsynaptic firing in an identical manner. Each synapse is therefore required to compute the STDP curve, recording pre- and postsynaptic spike-timings and making some graded change in synaptic strength based on their difference. This places considerable computational burden on the synapse, and requires that a range of coincidence detection machinery be postulated and deployed. Thus, although the Song model

is an apparently minimal model of STDP, it nevertheless places considerable computational demands on the synapse.

5.4.2 Simulation

The learning dynamics of any plasticity rule will depend, in part, upon the particular domain in which it is explored. For example, the choice of input firing patterns may strongly influence the observed plasticity. The particular choice of parameters may also have an important influence on the learning behaviours expressed. The majority of parameters in the Song model are poorly constrained by experimental data, and often little justification is given (beyond simple empirical observations) for confining simulations to a particular parameter space.

The parameters $A_{\pm}$, which determine the maximum modification that can occur due to a single spike pair, are poorly constrained by experimental data. The experimental protocols of Bi and Poo (1998), which are typical of this kind of plasticity experiment, used multiple pairings (60 pairings at 1Hz) to evoke statistically reliable changes in evoked EPSPs (a measure of afferent efficacy). Due to the natural variability in initial EPSC strength, these experimental results are often reported as changes in relative afferent strength. In the Song model $A_{\pm}$ describe the maximum *absolute* change induced by a single spike pairing. Values for $A_{\pm}$ are found by making the assumption that multiple spike pairings sum linearly to give an overall change in afferent efficacy, and as a result the change induced by a single spike pair is simply this total change divided by the number of pairings. However, values found by this simple method cannot be directly mapped onto values of $A_{\pm}$, as this ignores the additional scaling factor of g_{max}. Nevertheless, this is precisely what the authors do. In effect, they make the assumption that every afferent involved in the experiment of Bi and Poo had an initial efficacy of $0.62g_{max}$. Given that the results of Bi and Poo are presented as percentage changes precisely due to the variation in initial afferent strength, this is unlikely to be the case. The values for $A_{\pm}$ used, 0.005 for potentiation and -0.00525 for depression, cannot therefore be said to be drawn directly from experimental results, as those results were concerned with relative changes only. At best, it can be said that since A_{+} is positive and A_{-} is negative they are not in direct conflict with the majority of experimental results.

Input spiking is governed by independent Poisson spike trains, which are arguably a reasonable model of spontaneous firing pattern in the biological system. However, in reality such background levels of firing are almost always supplemented by bursting behaviours, especially during learning episodes (Bair et al., 1994). The properties of the learning rule discussed here, such as it competitive dynamics, bimodal distribution of input weights, and input rate-normalisation,

therefore do not require any form of structured input. That is, the inputs will segregate into two groups under purely stochastic input, with each input eventually finding itself randomly placed into one of the two groups. A correlated group of inputs will, of course, be more likely to potentiate as they will tend to precede postsynaptic spiking more than randomly spiking inputs but the correlations are not in themselves a prerequisite for segregation to occur. Thus, though this STDP-like rule may appear to offer an attractive way of generating segregation of two populations of inputs it does so in a very particular manner. This is in contrast with, for example, the BCM-rule, where it is precisely the presence of correlations between two input groups that drives the segregation (Bienenstock et al., 1982).

5.5 Summary

The Song model implements a simple, phenomenological STDP-like learning rule, based around a biphasic, exponential approximation to the experimentally observed STDP curve describing afferent plasticity. The rule leads to a bimodal equilibrium distribution of afferent weights, subject to hard upper and lower bounds on afferent efficacies. This distribution is stable in the sense that its overall shape does not change, but, importantly, individual inputs are not stable and may transition between states of low and high strength. The distributions produced represent a balance of excitatory and inhibitory input which forces the target neuron into a balanced, irregular firing regime where it is sensitive to the timing of input action potentials. This sensitivity on the timing of input spiking leads to competition among afferents. When the inputs fire with different latencies and correlations, the learning rule will selectively strengthen those firing with a short latencies and stronger correlations.

The non-linearity of the spike generation method plays a crucial role in producing the observed learning dynamics. If the chance of spiking increases linearly with membrane voltage then there is no increase in the proportion of inputs firing before a postsynaptic action potential once the system reaches equilibrium. Experimental data only broadly constrains the models parameters. However, the values used are consistent with a range of experimental results, and the equilibrium properties of a Song model are insensitive to the small variations in the models parameters. It is critical, however, that depression dominates potentiation overall (Song et al., 2000).

As the Song model takes the STDP curve over directly to govern input plasticity, it is guaranteed that the model is consistent with those spike-pair results. It

does not, however, generalise well to accommodate other results, such as spike-triplets or quadruplets (Froemke and Dan, 2002). It is, however, possible to make a variety of modifications to the Song model rule in order to account for additional experimental results not considered in the original work. Introducing "spike-suppression", where the plasticity induced by a spike-pair is reduced if preceded by another spike, can accommodate the spike-triplet results of Froemke and Dan (2002).In other work, it has been shown that introducing asymmetric temporal limitation on the interactions of presynaptic spikes can produce a rule that, when averaged over the course of a long spike train, predicts a change in input weight compatible with the BCM model (Izhikevich and Desai, 2003).

5.6 Conclusions

The Song model is a commonly cited example of a phenomenological, additive STDP learning rule. Although the exploration of the Song model has been mostly numerical, several key properties have been noted. In this Chapter, we have reproduced the main findings presented in the original paper (Song et al., 2000), then extended our simulations to further characterise the model. We find that the learning dynamics of the rule are relatively insensitive to the parameters, with the exception that depression is always required to dominate potentiation, on average.

The Song model makes the implicit assumption that STDP is valid at the level of synapses as well as inputs. This view is widely adopted in the literature. However, this is only one possible interpretation of STDP results. We show in Chapter 6 that an alternative view is viable; that the observed change in afferent efficacy arises due to the temporal and spatial average of much simpler changes occurring at the individual synapses comprising that afferent. This ensemble interpretation of STDP leads to a much simpler synaptic rule, which explains a range of spike- and rate-based experimental results while at the same time greatly reducing the computational burden on individual synapses.

Chapter 6

An Ensemble Interpretation of Spike-Timing-Dependent Plasticity

The work in this Chapter is based primarily on the paper "A Synaptic and Temporal Ensemble Interpretation of Neuronal Plasticity" (Appleby and Elliott, 2005).

6.1 The Synaptic Basis of Neuronal Plasticity

The models of activity-dependent plasticity discussed in Chapter 3 and Chapter 4 may be classified as either rate- or timing-based, depending on the underlying description of neuronal firing. A further distinction may be made between phenomenological models and those based on a more biophysical approach.

Rate-based models of plasticity have been successful in explaining certain developmental phenomena, such as the development of orientation selectivity in the visual system (Bienenstock et al., 1982). However, by their very, nature rate-based models are unable to accommodate recent experimental results suggesting that it is the exact timing of pre- and postsynaptic action potentials that determines the degree and polarity of afferent plasticity. More recently, timing-based models have been formulated that consider the interactions of individual pre- and postsynaptic spikes. A variety of phenomenological learning rules have been proposed which take the experimentally observed spike-timing plasticity rule over directly to govern plasticity at individual inputs. In conjunction with certain constraints, this approach can give rise to stable distributions of input weights (Song et al., 2000; van Rossum et al., 2000). Biophysical models have

mainly focused on modelling calcium and NMDA-receptor dynamics according to a "calcium-control hypothesis" (Senn et al., 2000; Shouval et al., 2002; Karmarkar and Buonomano, 2002; Saudargine and Porr, 2005). This approach has yielded models of some complexity which, although reproducing STDP-like plasticity curves, often display behaviours not generally seen experimentally, such as the prediction of an additional second LTD window at large pre-post timing differences (but see Nishiyama et al., 2000). They are also, in general, more difficult to analyse than phenomenological models, and can be sensitive to the choice of parameters.

We note that experimental works invariably measures plasticity across a connection, or group of connections, rather than across individual synapses. As an afferent typically makes multiple synapses onto a target cell, it is possible that the observed STDP plasticity rule exists only the input level, with a different rule governing changes at the synaptic level. That is, the observed plasticity rule may emerge due to the averaging of some other synaptic rule rather than actually being implemented at any individual synapse. In this Chapter, we propose such a rule, in the form of a three-state synaptic switch that governs plasticity at individual synapses, and find that this simple rule can explain a variety of rate- and timing-based plasticity results.

6.2 Basic Formulation

We now construct an activity-dependent synaptic plasticity rule that governs changes at individual synapses in response to pre- and postsynaptic spiking. We propose that a positively correlated spike pair will potentiate a given synapse by a fixed amount A_+, subject only to the requirement that the postsynaptic spike occurs within a finite time window relative to the presynaptic spike. Outside this time window, the postsynaptic spike does not evoke any change in synaptic strength. The duration of this time window is not fixed, but is taken to be a stochastic quantity governed by some probability distribution. This simple modification rule could be embodied by some biological, synaptic switch mechanism. The arrival of a presynaptic spike activates some process that elevates the synapse into a different functional state. The arrival of a postsynaptic spike while this process is still active, and the synapse is still in the elevated state, induces potentiation of the synapse by a fixed amount A_+. The postsynaptic spike is also taken to deactivate the process. In the absence of postsynaptic firing, the process will naturally deactivate in a stochastic, random manner, and subsequent postsynaptic spiking will not evoke a change in synaptic strength unless preceded by further presynaptic spiking.

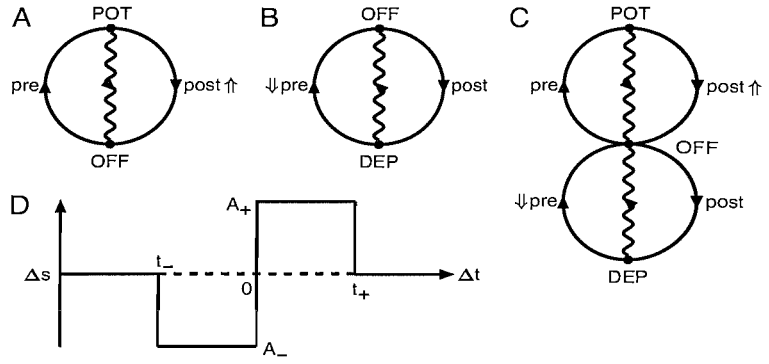

FIGURE 6.1: The simplest, self-consistent forms that the proposed synaptic switch can take, and its resulting plasticity rule. **(A)** shows the states and transitions that must exist in a switch accommodating the timing-dependent induction of LTP due to pre- and postsynaptic spiking; **(B)** same as (A) but for the induction of LTD; **(C)** a unified, 3-state synaptic switch that can exhibit both LTP and LTD; **(D)** change in synaptic strength evoked under the unified 3-state switch for a representative spike pair at various spike timings. In this example, the switch was activated to the POT or DEP state and remained there for some arbitrary time $t_{\pm}$. Although we show the case where $t_+ = t_-$, these times would, in general, be different and would also vary from trial to trial. The arrows $\Uparrow$ and $\Downarrow$ indicate the induction of potentiation and depression, respectively.

We label the resting state of the synapse as the OFF state. The elevation of the synapse into a different functional state, due to the arrival of a presynaptic spike, is represented by a transition to a different state that we label the POT state. While in the POT state, additional presynaptic spiking has no further effect. If, on the other hand, a postsynaptic spike occurs while the switch is in the POT state, then the switch is immediately returned to the OFF state via the transition $POT \rightarrow OFF$. This transition is defined to induce an associated potentiation of synaptic strength of A_+. In the absence of further spiking, the switch will move from the POT state back to the OFF state in a stochastic manner, governed by some probability distribution. We refer to transitions triggered by pre- or postsynaptic spiking as active transitions, while those that occur stochastically are referred to as passive transitions. This abstracted rule is represented in Fig. 6.1A, with semicircles representing active transitions and wavy lines representing passive transitions. The active transition $POT \rightarrow OFF$ is the only transition capable of inducing potentiation of synaptic strength.

To account for depression of synaptic strength under negatively correlated spiking, a second synaptic switch is postulated, see Fig. 6.1B. This switch behaves in a very similar manner to the one just described, except here postsynaptic spiking triggers the initial active transition to a different functional state $OFF \rightarrow DEP$. When in the DEP state, further postsynaptic spiking has no effect, but a presynaptic spike will trigger the active transition $DEP \rightarrow OFF$ with an associated

decrease in synaptic strength, by an amount A_-. The stochastic, passive transition $DEP \rightarrow OFF$ returns the switch to the OFF state in the absence of further spiking.

These two switch mechanisms adequately describe a step change in synaptic strength in response to positively or negatively correlated spiking. The two switches could exist independently, but we can unify them into a single, three-state synaptic switch, see Fig. 6.1C. The parameters of this three-state switch are not necessarily symmetric, which in biological terms reflects the possible independence of the processes activated by pre- or postsynaptic firing when the synapse is in the OFF state. The modification induced by a representative spike pairing at various spike time differences is plotted in Fig. 6.1D. The switch rule gives rise to a modification in accordance with a two-step function of fixed step heights, $A_\pm$. The two random step widths, $t_\pm$, are governed by the probability distributions that describe the passive transitions $POT \rightarrow OFF$ and $DEP \rightarrow OFF$. If this hypothetical spike pairing were repeated, then the widths $t_\pm$ would likely take different values, giving rise to a different "critical window". It is important for our model that the magnitude of synaptic plasticity, represented by the heights of the two step functions, is not dependent on the difference in spike times. The level of coincidence detection required is therefore minimal, as the synapse is required only to record the occurrence of a pre- or postsynaptic spike, not the precise time of occurrence. Although additional states and transitions may freely be added, we find that this simple, three-state switch is all that is required to reproduce a variety of STDP results.

We do not propose that synapses themselves are discretised, only that their potential to undergo plasticity is determined by some controller mechanism that has 3 distinct states. Synaptic strengths are therefore continuous quantities which are adjusted in finite amounts in response to transitions in the controller switch. There is thus no intrinsic upper bound on synaptic strengths.

We note that the active transitions $OFF \rightarrow POT$ and $OFF \rightarrow DEP$ could be made stochastic, perhaps occurring with some fixed probability when a pre- or postsynaptic event occurs. Such a modification would adequately accommodate various sources of noise such as unreliability in synaptic transmission or failure of action potentials to propagate. However, such a modification simply alters the effective learning rates of the two halves of the switch, a change which can be absorbed into a scaling of $A_\pm$.

We assume that an afferent makes multiple synapses onto a target cell. The overall strength of the connection between the two cells is defined, for simplicity and according to the usual convention, to be the linear sum of each individual synaptic strength. The synapses are treated independently, which, due to the

stochastic nature of the synaptic modification rule, means that the synapses comprising a connection will often be in different states. It is therefore the spatial average over synapses, and the temporal average over spike pairs, that determines the overall change in connection strength.

For notational convenience throughout the analysis, we denote a presynaptic spike by the symbol "π" and a postsynaptic spike by the symbol "p". Pre- and postsynaptic firing are assumed to be independent Poisson processes with rates λ_π and λ_p, respectively, and we set $\beta = \lambda_\pi + \lambda_p$. Because they are independent, the combined pre- and postsynaptic spike sequences form a single Poisson process of overall rate β. For a Poisson process of rate λ, the inter-event time (the "waiting time") is an exponentially-distributed random variable with parameter λ, and thus, in particular, the waiting time between any two spikes is exponentially distributed with parameter β, and has the probability density function $f_T(t) = \beta e^{-\beta t}$. For any given spike in the combined train, the probability that it is presynaptic is λ_π/β and the probability that it is postsynaptic is λ_p/β.

6.3 The 2-spike Response Function

Here we restrict our analysis to spike trains consisting of two spikes only, so that a two-spike train can manifest itself as one of four possible sequences: $\pi\pi$, πp, $p\pi$ or pp. The probability of observing a particular spike pattern ij, where $i,j \in \{\pi,p\}$, is then just $p_{ij} = \lambda_i\lambda_j/\beta^2$. Longer spike trains are investigated numerically in Results. Despite the more complicated nature of the higher order interactions between multiple spikes, our results for longer spike trains share characteristics similar to those for the two-spike case.

Under a specific spike pattern, modification of synaptic strength may or may not occur, depending on the state of the switch when the second spike arrives. We therefore seek an expression for the expected change in synaptic efficacy induced by a single spike pair under our switch rule. The spike patterns $\pi\pi$ and pp cannot cause a change in synaptic strength under our switch rule, so we need only to consider the πp and $p\pi$ patterns. Consider the πp pattern. The initial presynaptic spike triggers the active transition $OFF \rightarrow POT$. The switch remains in the POT state until either the arrival of the postsynaptic spike or the occurrence of a stochastic, passive transition. In either case, the switch will be returned to the OFF state, but the active transition triggered by the postsynaptic spike will also induce a change in synaptic strength. We therefore have a sum of four integrals. However, only two of these integrals make non-zero contributions to ΔS (those corresponding to πp and $p\pi$) and we may discard the remaining two. Performing this integrals is, for the two-spike trains considered

here, relatively straightforward. Consider the term arising from πp. We have that

$$\mathbf{E}[\Delta S^{\pi p}] = p_{\pi p} \int_0^\infty dt_1 \, f_T(t_1) \int_0^\infty dt_2 \, f_T(t_2) \Delta S_{\pi p}(t_2). \tag{6.1}$$

Writing this out explicitly

$$\mathbf{E}[\Delta S^{\pi p}] = \frac{\lambda_\pi \lambda_p}{\beta^2} \int_0^\infty dt_1 \beta e^{-\beta t_1} \int_0^\infty dt_2 \beta e^{-\beta t_2} \left[A_+ e^{-\lambda_+ t_2} \sum_{i=0}^{n-1} \frac{(t_2 \lambda_+)^i}{i!} \right], \tag{6.2}$$

where the integral over t_1 evaluates to unity, representing the arbitrary timing of the initial π spike. A change of variables to $y = t_2(\beta + \lambda_+)$, yields

$$\mathbf{E}[\Delta S^{\pi p}] = A_+ \frac{\lambda_\pi \lambda_p}{\beta^2} \frac{\beta}{\beta + \lambda_+} \sum_{i=0}^{n-1} \frac{\lambda_+^i}{i!} \frac{1}{(\beta + \lambda_+)^i} \int_0^\infty dy e^{-y} y^i. \tag{6.3}$$

We recognise the integral over dy as the integral definition of the gamma function, which evaluates to $i!$. Thus,

$$\mathbf{E}[\Delta S^{\pi p}] = A_+ \frac{\lambda_\pi \lambda_p}{\beta^2} \frac{\beta}{\beta + \lambda_+} \sum_{i=0}^{n-1} \frac{1}{(\beta + \lambda_+)^i}. \tag{6.4}$$

Explicitly summing this geometric sum, we arrive at

$$\mathbf{E}[\Delta S^{\pi p}] = \frac{\lambda_\pi \lambda_p}{\beta^2} A_+ \left[1 - \frac{1}{(1 + \lambda \tau_+)^{n_+}} \right]. \tag{6.5}$$

An identical process leads to the corresponding expression for $\mathbf{E}[\Delta S^{p\pi}]$. Summing these two terms gives

$$\mathbf{E}[\Delta S] = \frac{\lambda_\pi \lambda_p}{\beta^2} \left\{ A_+ \left[1 - \frac{1}{(1 + \lambda \tau_+)^{n_+}} \right] - A_- \left[1 - \frac{1}{(1 + \lambda \tau_-)^{n_-}} \right] \right\}. \tag{6.6}$$

Defining

$$K_\pm^1(\lambda) = 1 - \frac{1}{(1 + \lambda \tau_\pm)^{n_\pm}}, \tag{6.7}$$

we then obtain

$$\mathbf{E}[\Delta S] = \frac{\lambda_\pi \lambda_p}{\beta^2} \left[A_+ K_+^1(\beta) - A_- K_-^1(\beta) \right] \tag{6.8}$$

as the expected synaptic change arising from any two-spike sequence.

This equation is an analytical expression for the expected change in synaptic efficacy induced by a two-spike train at given pre- and postsynaptic firing rates, λ_π and λ_p. In the limit of large λ_π and λ_p in Eq. (6.8), we have that $\mathbf{E}[\Delta S] \propto (A_+ - A_-)$. The sign of this expression indicates whether potentiation or depression of synaptic strengths is expected for high pre- and postsynaptic firing rates. Experimental work on rate-based LTP shows that high pre- and postsynaptic firing rates generally lead to LTP (Sjöström et al., 2001). This requires that $\mathbf{E}[\Delta S] > 0$ for large λ_π and λ_p, i.e. $A_+ > A_-$. However, to be able to generate competitive dynamics, we also require a depressive phase where $\mathbf{E}[\Delta S] < 0$, otherwise synapses can never weaken on average. Putting $\lambda_p = \lambda_\pi$, and maintaining the requirement that $A_+ > A_-$, a sufficient condition is that $\partial \mathbf{E}[\Delta S]/\partial \lambda_\pi|_{\lambda_\pi,\lambda_p=0} < 0$. As $\mathbf{E}[\Delta S]|_{\lambda_\pi,\lambda_p=0} = 0$, this guarantees the presence of a depressive region. This produces a second constraint,

$$\gamma = \frac{A_+ n_+ \tau_+}{A_- n_- \tau_-} < 1, \tag{6.9}$$

which we interpret as depression dominating over potentiation. An identical condition has been observed, but not mathematically derived, for simulations of exponential-like STDP plasticity rules in the context of generating dynamics that give rise to bimodal synaptic distributions (Song et al., 2000). Empirical work has also shown it to be a requirement for a BCM-like learning rule to emerge on average from such rules (Izhikevich and Desai, 2003). Here, we have shown that requiring our switch-rule to maintain, on average, a BCM-like learning rule leads to mathematically derivable constraints. Whether or not the presence of a depressive regime, for which $\gamma < 1$, guarantees the presence of competitive dynamics in our switch rule is an issue that we shall explore in Chapter 8.

We set $A_+ = 1$ and $A_- = 0.95$, in accordance with the condition that $A_+ > A_-$, and choose $n_+ = n_- = 3$, as before. Setting $\tau_- = 20$ms and choosing γ determines the remaining parameter τ_+. We also assume that the postsynaptic firing rate λ_p is linearly related to the presynaptic rate λ_π once it exceeds a value Θ, so that

$$\lambda_p = \begin{cases} \lambda_\pi - \eta & \text{for } \lambda_\pi \geq \Theta \\ 0 & \text{for } \lambda_\pi < \Theta \end{cases}, \tag{6.10}$$

and we set $\Theta = 5$Hz. Varying the presynaptic firing rate λ_π produces the family of curves shown in Fig. 6.3 for different values of γ. When $\gamma < 1$ we observe that the behaviour is qualitatively BCM-like, with a depressive phase at low presynaptic firing rates followed by a transition to potentiation as a threshold is passed. We find exact agreement, presented in Results, between this analytical result

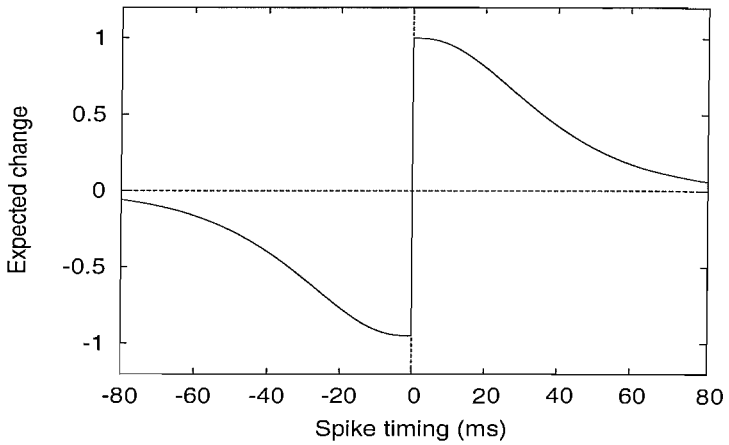

FIGURE 6.2: The expected change in synaptic strength as a function of the spike time difference t, from Eq. (6.8).

and numerical simulation of two-spike trains. For longer spike trains, Eq. (6.8) represents the expected change induced by any pair of spikes, not the expected change induced by all spike interactions in these processes. Nevertheless, for longer spike trains the total change induced by multiple spike interactions is of a qualitatively similar character, as we will show in Results. An important feature of the BCM model (Bienenstock et al., 1982) is the sliding of the potentiation threshold in response to changes in the postsynaptic firing rate (Kirkwood et al., 1996; Philpot et al., 2003). As the analytical expression shows, a threshold emerges from our model which is a function of various, easily modifiable parameters. Allowing some of these parameters to depend on the recent time average of postsynaptic firing, in a manner similar to other modelling approaches, would capture the sliding threshold of the BCM rule in a satisfactory way.

6.4 Numerical Results

We now turn to numerical simulation to study the behaviour of a single afferent innervating a single target cell. The connection between the afferent and target cells is assumed to comprise of multiple synapses, which individually obey the stochastic switch rule set out above. The STDP curve governing changes in overall connection strength emerges from the averaged effect of our synaptic switching rule. This averaging process can take place over multiple synapses or, equivalently, multiple spike pairings. We choose to simulate ten synapses per afferent. This is partly so that an averaging process can be observed even with single spike pairs, but also to show that the synapses comprising a connection can often be in different states and undergo different modifications while still giving

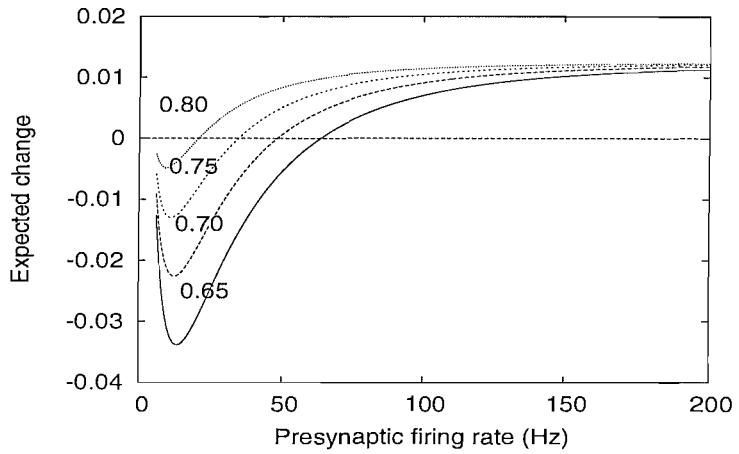

FIGURE 6.3: The expected change in synaptic strength due to a single pair of spikes, for values of γ shown attached to each curve. The pre- and postsynaptic cells fire according to a Poisson process.

rise to the STDP curve when viewed as an ensemble. As the stimulation protocols used involve many spike pairings, the simulations can, in fact, be repeated with just one synapse. The averaging process then occurs at this one synapse over many events, and the results are qualitatively similar.

A typical experimental protocol relies on evoking pre- and postsynaptic action potentials in synaptically coupled cells. In both cases, the normal function of a synapse as a propagator of neuronal activity is suppressed, with external current injections typically used to achieve spiking on demand. In our simulations, as in the experimental procedures, afferent and target cell spiking is assumed to be driven by an external force. Presynaptic spiking does not contribute to postsynaptic spiking in any way, and simulation of any kind of integrate-and-fire target cell is not required.

Due to a high level of variability, the majority of experimental data on STDP describes relative changes in connection strength. Multiple spike pairings are needed to evoke a statistically significant change in overall connection strength. We adopt a similar approach by defining the combined initial synaptic strength of the input afferent to be equal to 1, and then scaling the magnitudes of synaptic plasticity, A_+ and A_-, to reproduce the measured relative change in overall connection strength under a particular experimental protocol.

6.4.1 Spike based results

In order to examine the timing dependence of our rule, we implement a particular experimental protocol that has been shown to evoke STDP-like changes in

embryonic rat hippocampal cultures (Bi and Poo, 1998). This protocol is typical of timing-based LTP experiments and our parameters are chosen to reflect the main features of these data. As described above, we set $n_+ = n_- = 3$, $A_+ = 1.00$, $A_- = 0.95$, and $\tau_- = 20$ms. Choosing $\gamma = 0.70$ generates a value for $\tau_+ = \gamma A_- \tau_- / A_+ \simeq 13$ms. The magnitudes of synaptic plasticity, A_+ and A_-, set the overall scale for synaptic modifications. To match the experimental data, we require that the maximum possible relative change in overall connection strength evoked by 60 spike pairings is approximately ± 1. We therefore require a scaling factor of 60 to be applied to the magnitudes of synaptic modification, and we set $A_+ = 1.00/60$ and $A_- = 0.95/60$ accordingly. This scaling has no other effect beyond producing a simulated change in overall connection strength equal to the experimentally observed value, and the dynamics of the synaptic switch are unchanged. Noise in the timing of spikes, reflecting both experimental error and variable transmission times, is drawn from a Gaussian distribution with standard deviation of 1ms. The spike pairing protocol consists of 60 pairings at 1Hz applied at time differences ranging from -80ms to $+80$ms (Bi and Poo, 1998). The averaging of the synaptic modification rule over multiple synapses and pairings gives an overall change in connection strength that has two exponential-like phases, plotted in Fig. 6.4. This change is consistent with experimental data (Bi and Poo, 1998), with polarity depending only on the signs of $A_\pm$. These simulation results agree with the analytical expressions, and are qualitatively unchanged for any realistic level of temporal Gaussian noise with a standard deviation $\sigma \leq 10$ms. The parameters $\tau_\pm$ determine the width of the temporal window — as $\tau_\pm$ increase, spike pairings at greater time differences begin to evoke a significant change in synaptic efficacy. As discussed above, $A_\pm$ determine the maximum amplitude of plasticity, evoked when the time difference is very small. Adjustment of these parameters produces a family of STDP curves that can reproduce a variety of experimental results without altering the basic characteristics, such as the exponential-like slopes, of the curve.

A number of more complicated spike patterns, including triplets and quadruplets, have been explored in experimental preparations (Froemke and Dan, 2002). In the case of spike triplets, the experimental protocols are very similar to that of spike-pairings, but instead of one pre- and one postsynaptic spike, an additional third spike (either pre- or postsynaptic) is introduced. We reproduce the experimental protocol exactly, repeating a particular stimulation pattern 60 times at 0.2Hz. The parameters are the same as for the spike pairing simulations, and the results are set out in Table 6.1. The simulated results for the spike-triplet protocols are in close agreement with experiment (Froemke and Dan, 2002), a result that cannot be reproduced under other modification rules without additional constraints on spike interaction such as spike suppression (Froemke and Dan, 2002). Under earlier models of STDP, spike triplets were treated as two separate

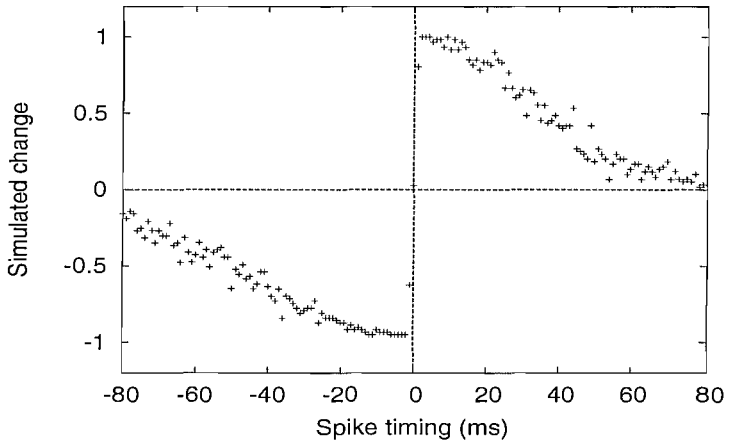

FIGURE 6.4: Simulated total change in overall connection strength as spike-timing varies.

spike pairings that individually obeyed a STDP-like modification curve. Their linear addition gives a predicted change that is not in agreement with experimental data. Here, the switch provides a mechanism, in the form of the passive transitions to the OFF state, under which a triplet can evoke a change with a sign opposite to that predicted by such linear addition of pairings respecting the STDP curve. With two presynaptic spikes and one postsynaptic spike, the first presynaptic spike moves the switch into the POT state. If the postsynaptic event occurs in a timely fashion, it will move the switch back to the OFF state and trigger an increase in synaptic strength. In this case, the second presynaptic event will only move the switch to the POT state, and this does not trigger any change in synaptic strength. If, however, the postsynaptic events occurs too late and the switch has already returned to the OFF state via a passive transition, then the switch will instead be moved to the DEP state. In this case, the second presynaptic event moves the switch back to OFF and triggers a depression of synaptic efficacy. It is the choice of parameters describing the switch, $A_{\pm}$, $n_{\pm}$, and $\tau_{\pm}$, that determines the average outcome for a given protocol. In fact, the parameters chosen roughly to reflect simple spike pairing results are sufficient to accommodate spike triplets. This explanation of triplet interactions emerges as a natural consequence of the switch rule, with no need for modifications or additional constraints.

In the case of spike quadruplets, comprised of two pre- and two postsynaptic spikes, the switch rule leads to potentiation under both of the protocols set out in Table 6.1. This is not in agreement with the experimental results, where the second quadruplet protocol leads to depression.

Pattern	Timing (ms)	Experiment	Simulation
$\pi p \pi$	2.6/6.0	$\uparrow$	+1.00
$p \pi p$	6.5/0.5	$\downarrow$	−0.94
$\pi p p \pi$	8.8/10.6/9.6	$\uparrow$	+0.03
$p \pi \pi p$	7.9/9.6/9.0	$\downarrow$	+0.03

TABLE 6.1: The experimental (Froemke and Dan, 2002) and simulated effect of spike triplets and quadruplet. The first column gives the spiking patterns (a presynaptic spike is denoted by π, and a postsynaptic spike by p). The second column gives the spike time differences for the patterns. The third column gives an indication of the experimental measurement (Froemke and Dan, 2002), upward arrows indicate potentiation, downward arrows indicate depression. The fourth column gives our simulated results.

It is possible to accommodate the quadruplet results by replacing the active transitions $POT \rightarrow OFF$ and $DEP \rightarrow OFF$ with active $POT \rightarrow POT$ and $DEP \rightarrow DEP$ transitions, respectively. Such a modification does not disrupt the spike-pair results by the following argument. The initial spike triggers an active transition $OFF \rightarrow POT$ or $OFF \rightarrow DEP$. If the second spike arrives before the switch decays back to OFF then plasticity will be induced. Whether this spike also induces a transition $POT \rightarrow OFF$ or $POT \rightarrow POT$ is, from a plasticity point of view, irrelevant as there are no further spikes and, by definition, there can be no further plasticity events. Under a spike triplet this is not the case. After a plasticity event, the synapse may undergo further plasticity dependent on the termini of the $POT \rightarrow OFF$ and $DEP \rightarrow OFF$ transitions. However, we find that the contribution from such processes is small, and that the spike-triplet results are also largely unaffected by this proposed modification. We are thus free to modify such transitions without greatly affecting pair or triplet results. However, such a modification of our switch rule destroys the stability of the rate-based limit (unpublished results), and this seems to be a high price to pay to account for results whose significance is currently unclear.

6.4.2 Rate based results

Induction of LTP using a rate-based protocol was simulated by driving the presynaptic cell at a fixed frequency (ranging from 0-200Hz) governed by a Poisson process, as set out above. The postsynaptic cell fires in a similar, Poisson manner, with frequency given by Eq. (6.10). This suppresses postsynaptic firing at very low presynaptic rates, as would be expected in a real system with many inputs. The cells are decoupled in the sense that presynaptic firing does not influence the postsynaptic cell membrane potential in any way, thus reproducing

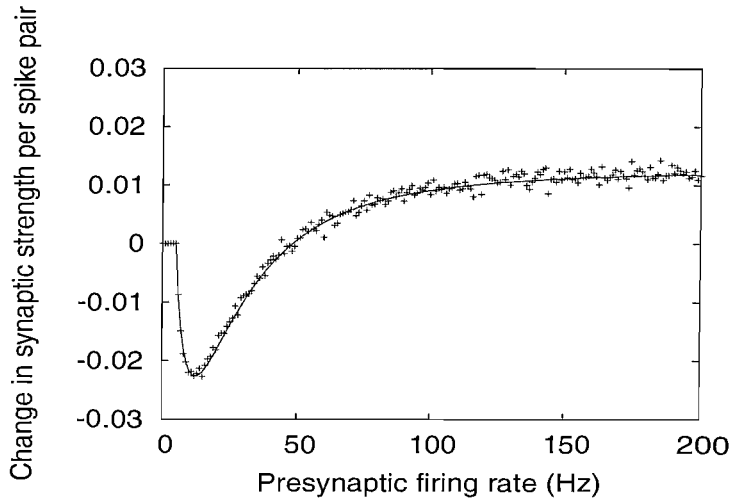

FIGURE 6.5: Change in overall connection strength, per pair of spikes, for simulated two-spike trains, as a function of presynaptic firing rate. The solid line shows the corresponding analytical result.

a typical experimental protocol in which two cells are held in current clamps and current injections are used to induce spiking (Bi and Poo, 1998). The parameters are identical to the spike-pairing simulations. In order to verify our analytical results, we first simulate the effect of pairs of spikes by truncating the simulation after the first pair of events. Each pair of events can therefore consist of either two presynaptic spikes, two postsynaptic spikes, or one of each. The average, overall connection change after a total of 10^6 total spikes is shown in Fig. 6.5, as a function of the presynaptic firing rate. We see exact agreement between the simulated two-spike interaction and the analytically derived result, Eq. (6.8).

The two-spike results consider spike trains containing exactly two events, the interactions of which gives rise to a BCM-like change in the overall connection strength. An identical change will arise from the interaction of any pair of spikes in a train provided that the synapse is in the OFF state. However, when longer pre- and postsynaptic spike trains are considered, further spike interactions may occur and it is important to show that qualitative form of the two-spike learning rule is unchanged by these higher order corrections. We therefore simulate longer spike trains, of 50 and 100 spikes. The total change in overall connection strength then arises from a summation of the many individual transitions that occur. The simulated overall connection change per spike pair, averaged over many such trains, is shown in Fig. 6.6, allowing a direct comparison to the two-spike train result, which is also shown. An averaged, BCM-like plasticity rule emerges in all cases. The 50- and 100-spike trains give an average change per spike pair that is different from that calculated for simple two-spike trains, reflecting the influence of higher order interactions between the spikes, but the results are nevertheless qualitatively similar. The higher order interactions are of smaller

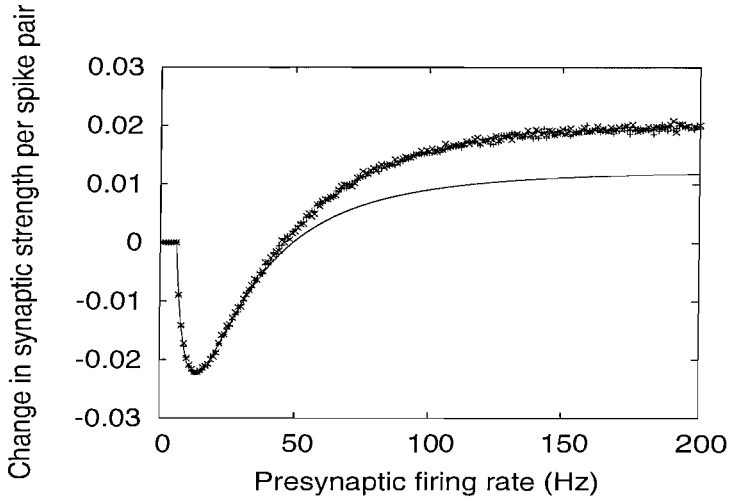

FIGURE 6.6: Change in overall connection strength, per pair of spikes, for simulated trains of 50 spikes (vertical crosses) and 100 spikes (diagonal crosses), as a function of presynaptic firing rate. The change per pair of spikes for a 50-spike train is almost identical to that for the 100-spike train. Also shown for comparison is the analytical two-spike result, represented by the solid line.

and smaller significance, with the average change in connection strength per pair of spikes converging to a limiting value as the number of spikes in a train grows large. Thus, as intuitively expected, if we were to plot the total change after the 100-spike train, it would simply be twice the total change after the 50-spike train.

6.5 Discussion

In this Chapter, we have shown that our proposed synaptic switch rule is a viable model of STDP. The model is susceptible to some degree of understanding and analysis, and very much reduces the computational demands placed on synapses. We have postulated that an individual synapse, when presented with a pre- and postsynaptic spike pair, adjusts its synaptic strength by a constant positive or negative jump, or does not change its strength at all. Our hypothetical synapse is therefore only required to record the occurrence of a pre- or postsynaptic event and adjust its strength by a fixed amount if an appropriate spike is generated in a timely fashion. Provided that the synapse destroys this record in some stochastic manner, so that the trace is short-lived, we have shown that we can derive the STDP rule directly. Thus, a simple synaptic modification rule can indeed give rise directly to a much more complex STDP rule. In this interpretation, the STDP rule can be viewed as an average, ensemble, emergent property of neurons, where the average is over either multiple synapses or multiple spike pairings (or both).

Our switch rule also provides a mechanism by which spike triplets may naturally give rise to an overall change in connection strength similar to that observed in experiment. Under typical STDP models (Song et al., 2000) this result can be achieved by introducing additional non-linearities, such as spike-suppression (Froemke and Dan, 2002). In our model, once a synapse is in, say, the POT state, this state cannot be changed by further presynaptic spikes until the synapse returns to the OFF state. In a strict sense, the synapse suppresses the effect of these subsequent presynaptic spikes, but this non-linearity is of a rather different form from that proposed by Froemke and Dan (2002), in which the values of A_+ and A_- are scaled depending on the spike history.

We have also shown that a BCM-like, rate-based rule (rather than spike-based rule) can be formally derived from our model. This derivation requires no further assumptions, and produces two explicit constraints on the choice of parameters. First, in order to generate LTP when both pre- and postsynaptic firing rates are high, we require that $A_+ > A_-$. That is, the level of potentiation induced under our switch rule by a presynaptic spike followed by a postsynaptic spike must be greater than the level of depression induced when the spike order is reversed. Second, in order to generate LTD at low firing rates, we require that $\gamma = (A_+ n_+ \tau_+)/(A_- n_- \tau_-) < 1$. This may be interpreted as depression dominating over potentiation. An identical condition has been observed to be a requirement for generating bimodal synaptic distributions under a simple, additive STDP rule (Song et al., 2000).

In the analysis presented here, we deconstruct a 2-spike train into all possible combinations of spikes, then calculate the expected synaptic response. Summing the expected response over all possible combinations gives Eqn. 6.8: the expected response to an average 2-spike train. In a 2-spike train, the number of possible combinations of spikes, and the way in which those combinations can interact, is relatively restricted. For 3-, 4-, or 5-spike trains the number of spike combinations, and the number of possible interactions, is considerably increased. Thus, the direct approach to calculating the spike-response function used in this Chapter, although sufficient for small numbers of spikes, rapidly becomes cumbersome as the number of spikes increases. We therefore desire a more general approach that allows us to calculate explicitly any n-spike response function. We will investigate such a general approach in Chapter 7.

That a BCM-like form may be obtained from the 2-spike form of our switch rule does not necessarily mean that the learning dynamics characteristic of the BCM-rule are also reproduced. The BCM-rule has several desirable properties as a model of plasticity, such as the presence of stable, competitive dynamics (Bienenstock et al., 1982). Such learning dynamics are necessary to explain phenomena such as the formation of orientation selectivity or ocular dominance

columns. A more detailed examination of the 2-spike rule, and a characterisation of its learning dynamics, is therefore necessary. We will explore the learning dynamics of the 2-spike rule in more detail in Chapter 8, and examine the differences between the 2-spike response function derived here and the more general n-spike response functions derived in Chapter 7.

6.6 Conclusions

Under our proposed switch rule, a variety of experimental results on spike-timing-dependent plasticity can emerge from temporal and spatial averaging over multiple synapses and multiple spike pairings. In particular, a critical window for the interaction of pre- and postsynaptic spikes emerges as an ensemble property of the collective system, with individual synapses exhibiting only a minimal form of spike coincidence detection. In addition, a Bienenstock-Cooper-Munro-like, rate-based plasticity rule emerges directly from such a model. This demonstrates that two, apparently separate, forms of neuronal plasticity can emerge from a much simpler rule governing the plasticity of individual synapses.

Chapter 7

Multi-Spike Interactions in the Switch Rule

The work in this Chapter is based primarily on the paper "Multi-Spike Interactions in a Stochastic Model of Spike Timing Dependent Plasticity: Derivation of Learning Rules" Appleby and Elliott (2005) (submitted).

7.1 Introduction

In the previous Chapter, we presented a stochastic, ensemble-based view of STDP in which STDP is an emergent property of neurons at the temporal or synaptic ensemble level (Appleby and Elliott, 2005). We showed that it is possible to build a model in which single synapses respond to spike pairs by adjusting their synaptic efficacies in a fixed, all–or–none fashion (cf. Petersen et al., 1998; O'Connor, Wittenberg, and Wang, 2005). If the second spike arrives within a time window that is stochastically determined, then a fixed level of potentiation or depression will occur, but if the second spike arrives too late, then no change in efficacy will take place. Critically, any changes that do take place do not depend on the time of the second spike relative to the first. Because the window size is stochastic, different synapses can respond to the same spike pairing differently, and the same synapse can respond to multiple presentations of the same pairing differently. Only at this temporal or synaptic ensemble level does the overall synaptic connection between a pair of neurons exhibit STDP. Due to the overall structure of the proposed model, we found that the experimental data on spike triplet interactions (Froemke and Dan, 2002) were automatically reproduced, and we also showed that a BCM-like learning rule can emerge from the 2-spike interactions.

Here, we derive exact, analytic results for the general multi-spike or n-spike interaction functions, for any $n \geq 2$, averaged over all possible sequences of n spikes in our model. This is achieved by extending our previous model slightly, and then finally taking a limit that reduces it to its original form. Although we derive exact results for the n-spike interaction function, we show that the asymptotic, large n form of this function is of particular interest, not least because the finite n form rapidly converges to the asymptotic limit, so that the limit is achieved within just a few tens of spikes. We derive these results here in preparation for a later analysis of the differences between the 2-spike and the general n-spike, $n > 2$, interaction functions, this analysis showing that, at least in our model, the 2-spike interaction function's form is entirely different from the general form observed for all other n-spike interaction functions.

7.2 Preliminaries

It is convenient before deriving general results for ΔS_n to derive some preliminary results that will be of use throughout the remainder of this Chapter. For the moment, we do not consider resetting.

Consider first a sequence of n identical spikes, say π spikes, followed by a different spike, say a p spike, giving a total of $n+1$ spikes, the first spike occurring at time t_0 and moving the synapse into the POT state, and let the subsequent interspike intervals be $t_1, \ldots, t_n$, so that t_i, $i > 0$, is the time between the ith spike and $(i+1)$th spike. The time of the ith spike is then $\sum_{j=0}^{i-1} t_j$, any i. Let the multi-argument function $P_{ON}^{+}(t_1, \ldots, t_n)$ be the probability that the synapse is in the POT state (P_{ON}^{-} for the DEP state and p spikes) at time $\sum_{i=0}^{n} t_i$ when the final spike arrives; clearly $P_{ON}^{\pm}$ do not depend on the time of the first spike, t_0, as this merely moves the synapse from the OFF state to the POT or DEP states. If, when the second spike arrives, the synapse is OFF due to a stochastic decay, then this second spike moves the synapse back to POT, and the probability, given OFF at time $t_0 + t_1$, that the synapse is in the POT state at time $\sum_{i=0}^{n} t_i$ is just $P_{ON}^{+}(t_2, \ldots, t_n)$. The probability of being in the OFF state at time $t_0 + t_1$ is just $1 - P^{+}(t_1)$. The probability of being in the POT state at time $t_0 + t_1$ is obviously $P^{+}(t_1)$, so we can write

$$\begin{aligned} P_{ON}^{+}(t_1, \ldots, t_n) &= [1 - P^{+}(t_1)] P_{ON}^{+}(t_2, \ldots, t_n) \\ &+ P^{+}(t_1) P_{ON}^{+}(t_2, \ldots, t_n | POT \text{ at } t_0 + t_1), \end{aligned} \tag{7.1}$$

where $P_{ON}^{+}(t_2, \ldots, t_n | POT \text{ at } t_0 + t_1)$ is the conditional probability that the synapse is in the POT state at time $\sum_{i=0}^{n} t_i$ given that it was still in the POT

state when the second spike arrived at time $t_0 + t_1$. Thus, we condition on whether or not the synapse is still in the POT state at time $(t_0 + t_1)$. Note that $P^+_{ON}(t)$ represents the total probability of being in the POT state at some later time, t. This accounts for all possible processes which terminate in POT at time t, including processes which involve transitions to OFF then back to POT during the intermediate time. It contrasts with the function $P^+(t_0 + t_1)$ which is the probability that the synapse is still in the POT state at time $(t_0 + t_1)$ without ever having returned to the OFF state in the intervening time. We now repeat the above argument, conditioning on whether the synapse is still in the POT state when the third spike arrives at time $t_0 + t_1 + t_2$ given that it was still in the POT state when the second spike arrived at time $t_0 + t_1$. Suppose that the synapse is in the OFF state at time $t_0 + t_1 + t_2$ given this sequence. Then the third spike sends the synapse back to POT. The probability that the synapse is in the OFF state at time $t_0 + t_1 + t_2$ given that it was in the POT state at time $t_0 + t_1$ is $1 - P^+(t_1 + t_2|POT \text{ at } t_0 + t_1)$, and the probability that it is still in the POT state is $P^+(t_1 + t_2|POT \text{ at } t_0 + t_1)$. Hence, we can now write

$$\begin{aligned} P^+_{ON}(t_1, \ldots, t_n) &= [1 - P^+(t_1)]P^+_{ON}(t_2, \ldots, t_n) \\ &+ [P^+(t_1) - P^+(t_1 + t_2)]P^+_{ON}(t_3, \ldots, t_n) \\ &+ P^+(t_1 + t_2)P^+_{ON}(t_3, \ldots, t_n|POT \text{ at } t_0 + t_1 + t_2), \end{aligned} \tag{7.2}$$

where we have used the definition of the conditional probability,

$$\begin{aligned} P^+(t_1 + t_2|POT \text{ at } t_0 + t_1) &= \frac{P^+(POT \text{ at } t_0 + t_1 + t_2 \;\&\; POT \text{ at } t_0 + t_1)}{P^+(POT \text{ at } t_0 + t_1)} \\ &= \frac{P^+(t_1 + t_2)}{P^+(t_1)}. \end{aligned} \tag{7.3}$$

Thus, repeating the above argument for all subsequent spikes, we see that

$$\begin{aligned} P^+_{ON}(t_1, \ldots, t_n) &= [1 - P^+(t_1)]P^+_{ON}(t_2, \ldots, t_n) \\ &+ [P^+(t_1) - P^+(t_1 + t_2)]P^+_{ON}(t_3, \ldots, t_n) \\ &+ \cdots \\ &+ [P^+(t_1 + \cdots + t_{i-1}) - P^+(t_1 + \cdots + t_i)]P^+_{ON}(t_{i+1}, \ldots, t_n) \\ &+ \cdots \\ &+ [P^+(t_1 + \cdots + t_{n-2}) - P^+(t_1 + \cdots + t_{n-1})]P^+_{ON}(t_n) \\ &+ P^+(t_1 + \cdots + t_n), \end{aligned} \tag{7.4}$$

where clearly $P^+_{ON}(t) \equiv P^+(t)$.

The conditional expectation value for the extent of potentiation or depression induced by this sequence of n identical spikes followed by a different spike is just $+A_+P^+_{ON}(t_1,\ldots,t_n)$ or $-A_-P^-_{ON}(t_1,\ldots,t_n)$, respectively. As with the 2-spike calculation, we convert these into unconditional expectations by integrating out the interspike intervals according to their pdfs and weighting by the spike sequences' probabilities. We do this for the general distribution P_{ON}, not specifying whether P^+_{ON} or P^-_{ON}, and hence we do not at this point weight by the spike sequence probabilities as these differ for the two possible cases. Because the combined pre- and postsynaptic spike sequence is Poisson with rate $\beta = \lambda_\pi + \lambda_p$, fortunately the interspike pdf is always the same, although in general, for non-Poisson processes, this will not be true.

We take Eq. (7.4) and integrate $t_1,\ldots,t_n$ over their entire ranges, $t_i \in [0,\infty)$, having weighted by $\beta^n \exp(-\beta\sum_{j=1}^n t_j)$. Notice that the integration over the first spike time, t_0, always produces a result of unity, by the definition of the pdf. Defining

$$T_l(\beta) = \beta^l \int_0^\infty dt_1 \cdots \int_0^\infty dt_l \exp(-\beta \sum_{j=1}^{l} t_j) P_{ON}(t_1,\ldots,t_l), \tag{7.5}$$

and

$$K_l(\beta) = \beta^l \int_0^\infty dt_1 \cdots \int_0^\infty dt_l \exp(-\beta \sum_{j=1}^{l} t_j) P(t_1 + \cdots + t_l), \tag{7.6}$$

where $P(t_1 + \cdots + t_l)$ is the conditional probability of the switch still being in the POT or DEP state at time $t_1 + \cdots + t_l$ without having returned to OFF in the intervening time. Both equations apply for $l \geq 1$. We then have

$$\begin{aligned} T_n(\beta) &= T_{n-1}(\beta)[1 - K_1(\beta)] \\ &+ T_{n-2}(\beta)[K_1(\beta) - K_2(\beta)] \\ &+ \cdots \\ &+ T_{n-i}(\beta)[K_{i-1}(\beta) - K_i(\beta)] \\ &+ \cdots \\ &+ T_1(\beta)[K_{n-2}(\beta) - K_{n-1}(\beta)] \\ &+ K_n(\beta), \end{aligned} \tag{7.7}$$

or, defining $K_0(\beta) \equiv 1$, in a more compact notation, dropping the arguments on the functions T_i and K_i for notational convenience, we have

$$T_n = K_n + \sum_{i=1}^{n-1} T_{n-i}(K_{i-1} - K_i). \tag{7.8}$$

Thus, $T_l(\beta)$ denotes the probability that the switch is in an activated state (either POT or DEP) after l subsequent spikes, in any possible way. This includes transitions back to OFF and subsequent reactivation. Thus, the probability T_l focuses only on the *final* state of the switch, and is made up of a sum of many individual processes that arrive in POT or DEP after l spikes. K_l denotes the probability that the switch is in the POT or DEP state after l subsequent spikes, without having returned to OFF. That is, K_l excludes those processes which involved a transition to OFF followed by reactivation.

Unless explicitly indicated otherwise in what follows, the argument of the T_i's or K_i's is always β. Using the gamma pdfs for $P(t)$ above, it is a straightforward although tedious matter to show that

$$K_l^{\pm}(\beta) = \left(\frac{\tau_{\pm}\beta}{1+\tau_{\pm}\beta}\right)^l \sum_{i=0}^{n_{\pm}-1} \binom{i+l-1}{l-1} \frac{1}{(1+\tau_{\pm}\beta)^i}. \tag{7.9}$$

The K_i's will play an important role throughout. It follows from the standard properties of the binomial coefficients that the K_i's are ordered, so that $K_{l+m} \leq K_l K_m$.

Returning to Eq. (7.8), defining $T_0(\beta) \equiv 0$ and $\Delta T_i = T_i - T_{i-1}$, $i > 0$, we may rewrite this equation in the more elegant form

$$\sum_{j=0}^{n-1} K_j \Delta T_{n-j} = K_n, \tag{7.10}$$

for $n \geq 1$. Direct calculation of the first few ΔT_i's reveals that

$$\begin{aligned}
\Delta T_1 &= K_1, \\
\Delta T_2 &= K_2 - K_1^2, \\
\Delta T_3 &= K_3 - 2K_2K_1 + K_1^3, \\
\Delta T_4 &= K_4 - 2K_3K_1 + 3K_2K_1^2 - K_2^2 - K_1^4, \\
\Delta T_5 &= K_5 - 2K_4K_1 + 3K_3K_1^2 - 2K_3K_2 - 4K_2K_1^3 + 3K_2^2K_1 + K_1^5.
\end{aligned}$$

We see that, perhaps as expected, the expression for ΔT_l involves multiples of powers of the K_i's that form the integer partitions of l. It is hardly surprising that integer partitions should appear in this calculation, because T_l is essentially combinatoric in nature, determined by traversing a binary decision tree governed by conditioning on the synapse's state at each spike arrival time (or node in the tree). Let $p_l = \{a_1, \ldots, a_l\}$ be a partition of the integer l, so that $\sum_{i=1}^{l} i a_i = l$, into k parts, so that $\sum_{i=1}^{l} a_i = k$. Then we observe that

$$\Delta T_l = -\sum_{k=1}^{l} \sum_{\{p_l \vdash k\}} \frac{k!}{a_1! \cdots a_l!} (-1)^k K_1^{a_1} \cdots K_l^{a_l}, \tag{7.11}$$

where $\sum_{\{p_l \vdash k\}}$ means a sum over all partitions of l into k parts. We also have the standard result that

$$\sum_{\{p_l \vdash k\}} \frac{k!}{a_1! \cdots a_l!} = \begin{pmatrix} l-1 \\ k-1 \end{pmatrix}. \tag{7.12}$$

Such identities as these are standard in the theory of combinatorics and we shall require a similar identity to that for ΔT_l later when finding a generating function for the terms that appear in ΔS_n.

7.3 Extensions

We now extend the model to a slightly more general form that will facilitate the derivation of general multi-spike interaction functions. In the basic formulation, we assumed that if a synapse is in the POT (DEP) state, then further presynaptic (postsynaptic) spiking has no effect. We now propose, instead, that after r additional presynaptic (postsynaptic) spikes, the stochastic processes governing the return of the synapse from the POT (DEP) state to the OFF state are reset, provided that the synapse has not returned to the OFF state at any time after the first spike moved it into the POT (DEP) state. The overall result of this is to increase the likelihood that potentiation or depression will occur by essentially discarding the spike history of the synapse after $r + 1$ identical spikes and starting afresh. For $n_\pm = 1$, however, the exponential distribution is memoryless, so such a system should exhibit dynamics independent of the value of r. Strictly, we should allow differing numbers of spikes to reset the POT and DEP states, but it is easy to generalise our results to this case. Furthermore, we are particularly interested in the properties of the $r = 1$ and $r \to \infty$ models, the latter being equivalent to the original form of our model. Arguably, these $r = 1$ and $r \to \infty$ models are more biologically plausible than models with other values of r: either a spike of the appropriate sort resets the stochastic process or it does

not. Other values of r would require us to postulate some form of spike counting machinery at the synapse which, in the spirit of reducing the demands placed on the synapse as far as possible, we consider to be less likely. Issues such as the finite transmission probability of synapses, or action potential propagation failure, do not affect this argument as inclusion of such processes would only change the effective learning rate, a change that would be absorbed into a redefinition of $A_{\pm}$. We therefore assume that r additional spikes of the appropriate kind reset both stochastic processes for $DEP \rightarrow OFF$ and $POT \rightarrow OFF$, and we refer to this generalised model as the r-reset model. We call the 1-reset model the resetting model and the ∞-reset model the non-resetting model.

We may consider how the possibility of spikes resetting the stochastic processes affects the above analysis. In the r-reset model, if $r+1 < n$, then after the $(r+1)$th identical spike, the process, if it has never de-activated during these identical spike events, is reset. This has the effect of making subsequent probabilities independent of any of the events that occurred before the $(r+1)$th identical spike arrived, and causes subsequent conditional probabilities to cease to depend on times earlier than $t_0 + \cdots + t_r$.

For concreteness, consider $r = 1$. In moving between Eqs. (7.1) and (7.2), we had to expand $P^+_{ON}(t_2, \ldots, t_n | POT \text{ at } t_0 + t_1)$ by conditioning on the state of the synapse when the third spike arrived. However, for $r = 1$, the second π spike at time $t_0 + t_1$ resets the stochastic process, because it is given that the synapse has stayed in the POT state since the first π spike at time t_0 moved the synapse into the POT state. Thus, subsequent changes in synaptic strength are independent of the fact that the synapse did not move to the OFF state between the first and second π spikes, because the second π spike restarted the stochastic process afresh. Thus,

$$P^+_{ON}(t_2, \ldots, t_n | POT \text{ at } t_0 + t_1) = P^+_{ON}(t_2, \ldots, t_n) \quad \text{only for } r = 1 \tag{7.13}$$

Hence, Eq. (7.1) becomes $P^+_{ON}(t_1, \ldots, t_n) = P^+_{ON}(t_2, \ldots, t_n)$ and of course we finally arrive at $P^+_{ON}(t_1, \ldots, t_n) = \cdots = P^+_{ON}(t_n) \equiv P^+(t_n)$. This equation means that, for $r = 1$, the final spike in a long sequence of identical spikes is the only spike that counts. Hence, for $r = 1$, we have that $T_l = T_1 = K_1$, any l, by the definition of the T_i's in Eq. (7.5). Putting $T_l = K_1$ into Eq. (7.8), which we now regarded as an equation defining the K_i's in terms of the T_i's, it is easy to see that, for this $r = 1$ case, $K_l = K_1^l$. Thus, we extract the 1-reset model from the non-resetting model by replacing K_l by K_1^l in any equations in which it occurs.

We can see this $K_l \rightarrow K_1^l$ rule in another way. If we are myopic and do not notice the obvious cancellation between terms in Eq. (7.1) for $r = 1$, then we

would condition on the third spike as usual, but using $P^+_{ON}(t_2,\ldots,t_n)$ instead of $P^+_{ON}(t_2,\ldots,t_n|POT \text{ at } t_0+t_1)$. Thus, Eq. (7.2) would become, for $r=1$,

$$\begin{aligned} P^+_{ON}(t_1,\ldots,t_n) &= [1-P^+(t_1)]P^+_{ON}(t_2,\ldots,t_n) \\ &+ [P^+(t_1)-P^+(t_1)P^+(t_2)]P^+_{ON}(t_3,\ldots,t_n) \\ &+ P^+(t_1)P^+(t_2)P^+_{ON}(t_3,\ldots,t_n). \end{aligned} \tag{7.14}$$

Comparing this to Eq. (7.2), we see that $P(t_1+t_2)$ is replaced by the product $P(t_1)P(t_2)$, and, in general, after further conditioning, it is easy to see that $P(t_1+\cdots+t_l)$ is replaced by $\prod_{i=1}^{l}P(t_i)$ in the $r=1$ model. Replacing $P(t_1+\cdots+t_l)$ by this product in the definition of the K_i's in Eq. (7.6), we again see that $K_l \to K_1^l$. Making this substitution in Eq. (7.11), we obtain

$$\Delta T_l \to -K_1^l \sum_{k=1}^{l}(-1)^k \begin{pmatrix} l-1 \\ k-1 \end{pmatrix} \equiv 0, \tag{7.15}$$

for $l>1$, by the binomial theorem. Thus, we see again that, for $r=1$, $T_l=K_1$, any l. These results, in fact, make the 1-reset model particularly simple to analyse, with simple expressions available for ΔS_n, any n.

Using similar arguments to these for the 1-reset model, it is easy to see that the r-reset model replaces $P(t_1+\cdots+t_{lr+s})$, $s<r$ and $l>0$, with

$$\begin{aligned} & P(t_1+\cdots+t_r) \\ \times\ & P(t_{r+1}+\cdots+t_{2r}) \\ \times\ & \cdots \\ \times\ & P(t_{(i-1)r+1}+\cdots+t_{ir}) \\ \times\ & \cdots \\ \times\ & P(t_{(l-1)r+1}+\cdots+t_{lr}) \\ \times\ & P(t_{lr+1}+\cdots+t_{lr+s}), \end{aligned}$$

where of course the last term is absent if $s=0$. After integrating out the spike times, the r-reset model therefore replaces K_{rl+s} with $K_r^l K_s$, so any results that we obtain for the non-resetting model can be converted into those for the r-reset model by the substitution $K_{rl+s} \to K_r^l K_s$. In terms of the T_i's, the conditioning

argument that led to Eq. (7.8) terminates with a cancellation at the rth step, giving a truncated form of the recurrence relation in Eq. (7.10), so that

$$\sum_{j=0}^{r-1} K_j \Delta T_{n-j} = 0, \tag{7.16}$$

for $n \geq r$. For $r = 2$, we have the particularly simple result $\Delta T_{n+2} = (-K_1)^n \Delta T_2$, $n \geq 0$, with $\Delta T_1 = K_1$.

In general, the finite r models are useful precisely because they limit multi-spike interactions to at most $r + 1$ spikes, and this enables us to set up a recurrence relation in ΔS_n that is always of order $r + 1$.

7.4 Multi-Spike Recurrence Relations

We previously derived the expected response to an average 2-spike pairing

$$\Delta S_2 = \sum_{i,j \in \{\pi,p\}} \lambda_i' \lambda_j' \int_0^\infty dt_0 f_T(t_0) \int_0^\infty dt_1 f_T(t_1) \Delta S_{ij}(t_1), \tag{7.17}$$

In this Chapter, we shall principally be concerned with the derivation of an expression for ΔS_n, the expected response to an average n-spike sequence. Our strategy in much of what follows is to extract n-spike results for the small r models, because r limits the possible extent of spike interactions in a long sequence of spikes. The behaviour for large r will then be clear and the results will fall out directly.

By considering the effect of adding m additional spikes in front of this n spike train, we may derive a recurrence relation describing how the expected change in synaptic strength for the $n + m$ spike train, ΔS_{n+m}, is related to that of the n spike train, ΔS_n. We first develop recurrence relations for the 1-,2- and 3-reset models and then, having developed intuitions based on these small r cases, proceed to derive the results for the arbitrary r-reset model.

7.4.1 1-Reset Model

Let $\Sigma_n(t_0, t_1, \ldots, t_{n-1})$ denote an arbitrary sequence of n spikes with interspike intervals t_i, $i > 0$, and t_0 the time of the first spike, and let the function $W_{\Sigma_n}(t_0, \ldots, t_{n-1})$ denote the pdf and spike probability weighting for this sequence that is required to transform conditional expectation values into unconditional expectation values. Let $\Delta S_{\Sigma_n}(t_1, \ldots, t_{n-1})$ be the conditional expectation value for the change in synaptic strength induced by the sequence Σ_n. This

conditional expectation value depends only of the interspike intervals. It is not a function of the first spike time, t_0, as the first spike always moves the switch into the POT or DEP state and starts the stochastic process governing the decay back to OFF.

Consider a sequence of $n+2$ spikes, the first two of which are either both π spikes or both p spikes. For the $r = 1$ model, the first spike is essentially irrelevant because the second, identical spike either re-activates the same process if the synapse is in the OFF state when it arrives at time $t_0 + t_1$, or resets the process if the synapse is still in the POT state. We therefore see that

$$\Delta S_{\pi\pi\Sigma_n}(t_1, \ldots, t_{n+1}) = \Delta S_{\pi\Sigma_n}(t_2, \ldots, t_{n+1}), \tag{7.18}$$

$$\Delta S_{pp\Sigma_n}(t_1, \ldots, t_{n+1}) = \Delta S_{p\Sigma_n}(t_2, \ldots, t_{n+1}). \tag{7.19}$$

Now consider the spike sequence $\pi p\Sigma_n(t_0, \ldots, t_{n+1})$. The first π spike activates the synapse into the POT state. If the synapse has decayed back to the OFF state before the second, p spike arrives, then that first spike is irrelevant and we may disregard it entirely. If, on the other hand, the synapse is still in the POT state when the p spike arrives, then the p spike will return the synapse to the OFF state and induce a potentiation by an amount A_+. The Markovian property associated with a return to OFF means that the remainder of the spike sequence $\Sigma_n(t_2, \ldots, t_{n+1})$ contributes to $\Delta S_{\pi p\Sigma_n}(t_1, \ldots, t_{n+1})$ independently of these first two spikes. Hence, we have that

$$\begin{aligned}\Delta S_{\pi p\Sigma_n}(t_1, \ldots, t_{n+1}) &= [1 - P^+(t_1)]\Delta S_{p\Sigma_n}(t_2, \ldots, t_{n+1}) \\ &+ P^+(t_1)[\Delta S_{\Sigma_n}(t_3, \ldots, t_{n+1}) + A_+].\end{aligned} \tag{7.20}$$

The first term is the case where the synapse returns to OFF before the p spike arrives. The second term is the case where the synapse is still in the POT state when the p spike arrives. An identical argument leads to

$$\begin{aligned}\Delta S_{p\pi\Sigma_n}(t_1, \ldots, t_{n+1}) &= [1 - P^-(t_1)]\Delta S_{\pi\Sigma_n}(t_2, \ldots, t_{n+1}) \\ &+ P^-(t_1)[\Delta S_{\Sigma_n}(t_3, \ldots, t_{n+1}) - A_-].\end{aligned} \tag{7.21}$$

Defining

$$\Delta S_{\Sigma_n} = \int_0^\infty dt_0 \cdots \int_0^\infty dt_{n-1} W_{\Sigma_n}(t_0, \ldots, t_{n-1}) \Delta S_{\Sigma_n}(t_1, \ldots, t_{n-1}), \quad (7.22)$$

where the absence of temporal arguments on ΔS_{Σ_n} indicates that interspike intervals have been integrated out according to their pdfs and spike probabilities included. Eqs. (7.18) and (7.19) therefore become

$$\Delta S_{\pi\pi\Sigma_n} = \lambda'_\pi \Delta S_{\pi\Sigma_n}, \quad (7.23)$$
$$\Delta S_{pp\Sigma_n} = \lambda'_p \Delta S_{p\Sigma_n}, \quad (7.24)$$

where the factors of λ'_π and λ'_p arise because the dependence of $W_{\pi\pi\Sigma_n}$ or $W_{pp\Sigma_n}$ on the first spike factors out and the remainder of $W_{\pi\Sigma_n}$ or $W_{p\Sigma_n}$ is absorbed into the definition of $\Delta S_{\pi\Sigma_n}$ or $\Delta S_{p\Sigma_n}$. Similarly,

$$\Delta S_{\pi p\Sigma_n} = \lambda'_\pi (1 - K_1^+) \Delta S_{p\Sigma_n} + \lambda'_\pi \lambda'_p K_1^+ (\Delta S_{\Sigma_n} + A_+ F_{\Sigma_n}), \quad (7.25)$$
$$\Delta S_{p\pi\Sigma_n} = \lambda'_p (1 - K_1^-) \Delta S_{\pi\Sigma_n} + \lambda'_\pi \lambda'_p K_1^- (\Delta S_{\Sigma_n} - A_- F_{\Sigma_n}), \quad (7.26)$$

where

$$F_{\Sigma_n} = \int_0^\infty dt_2 \cdots \int_0^\infty dt_{n+1} W_{\Sigma_n}(t_2, \ldots, t_{n+1}) = {\lambda'_\pi}^m {\lambda'_p}^{n-m} \quad (7.27)$$

if the spike sequence Σ_n contains m π spikes and $n - m$ p spikes.

We now sum over all possible spike sequences Σ_n, of which there are 2^n. Let $\sum_{\{\Sigma_n\}}$ denote this sum and define $\Delta S_n = \sum_{\{\Sigma_n\}} \Delta S_{\Sigma_n}$, $\Delta S^\pi_{n+1} = \sum_{\{\Sigma_n\}} \Delta S_{\pi\Sigma_n}$, $\Delta S^{\pi\pi}_{n+2} = \sum_{\{\Sigma_n\}} \Delta S_{\pi\pi\Sigma_n}$, etc., in an obvious notation. Then we have that

$$\Delta S^{\pi\pi}_{n+2} = \lambda'_\pi \Delta S^\pi_{n+1}, \quad (7.28)$$
$$\Delta S^{\pi p}_{n+2} = \lambda'_\pi (1 - K_1^+) \Delta S^p_{n+1} + \lambda'_\pi \lambda'_p K_1^+ (\Delta S_n + A_+), \quad (7.29)$$
$$\Delta S^{p\pi}_{n+2} = \lambda'_p (1 - K_1^-) \Delta S^\pi_{n+1} + \lambda'_\pi \lambda'_p K_1^- (\Delta S_n - A_-), \quad (7.30)$$
$$\Delta S^{pp}_{n+2} = \lambda'_p \Delta S^p_{n+1}, \quad (7.31)$$

where we have used the binomial theorem, $\sum_{\{\Sigma_n\}} F_{\Sigma_n} \equiv 1$. Since $\Delta S^\pi_{n+2} \equiv \Delta S^{\pi\pi}_{n+2} + \Delta S^{\pi p}_{n+2}$ and $\Delta S^p_{n+2} \equiv \Delta S^{p\pi}_{n+2} + \Delta S^{pp}_{n+2}$ (by summing over all possible

second spikes), and $\Delta S_n \equiv \Delta S_n^{\pi} + \Delta S_n^{p}$, we finally obtain the coupled recurrence relation

$$\begin{aligned}\mathbf{\Delta S}_{n+2} &= \lambda'_\pi \lambda'_p \begin{pmatrix} +A_+ K_1^+ \\ -A_- K_1^- \end{pmatrix} + \lambda'_\pi \lambda'_p \begin{pmatrix} K_1^+ & K_1^+ \\ K_1^- & K_1^- \end{pmatrix} \mathbf{\Delta S}_n \\ &+ \begin{pmatrix} \lambda'_\pi & \lambda'_\pi (1 - K_1^+) \\ \lambda'_p (1 - K_1^-) & \lambda'_p \end{pmatrix} \mathbf{\Delta S}_{n+1}, \end{aligned} \tag{7.32}$$

where $\mathbf{\Delta S}_n = (\Delta S_n^{\pi}, \Delta S_n^{p})^T$, the superscript T denoting the transpose. This coupled recurrence relation represents our final form for the 1-reset model.

7.4.2 2-Reset Model

For the 1-reset model, we had to look at specific spike patterns extending out to 2 spikes. For the 2-reset model, we must go out to 3 spikes, because only the third spike in a sequence of three identical spikes possesses the capacity to reset the stochastic processes. Consider $\Delta S_{\pi\pi\pi\Sigma_n}(t_1, \ldots, t_{n+2})$. As usual, we condition on the state of the synapse as the second and third π spikes arrive. If the synapse is in the OFF state when the second π spike arrives, we obtain a subsequent change governed by $\Delta S_{\pi\pi\Sigma_n}(t_2, \ldots, t_{n+2})$. If the synapse is still in the POT state when the second spike arrives, then we look at the state when the third spike arrives. If the synapse is in the OFF state at this time, we obtain a subsequent change governed by $\Delta S_{\pi\Sigma_n}(t_3, \ldots, t_{n+2})$. However, if the synapse is still in the POT state, then the third spike resets the process, and subsequent changes are still governed by $\Delta S_{\pi\Sigma_n}(t_3, \ldots, t_{n+2})$. Hence, two conditioning steps lead to

$$\begin{aligned}\Delta S_{\pi\pi\pi\Sigma_n}(t_1, \ldots, t_{n+2}) &= P^+(t_1) \Delta S_{\pi\pi\Sigma_n}(t_2, \ldots, t_{n+2} | POT \text{ at } t_0 + t_1) \\ &+ [1 - P^+(t_1)] \Delta S_{\pi\pi\Sigma_n}(t_2, \ldots, t_{n+2}), \end{aligned} \tag{7.33}$$

where the conditional change

$$\begin{aligned}&\Delta S_{\pi\pi\Sigma_n}(t_2, \ldots, t_{n+2} | POT \text{ at } t_0 + t_1) \\ &= P^+(t_1 + t_2 | POT \text{ at } t_0 + t_1) \Delta S_{\pi\Sigma_n}(t_3, \ldots, t_{n+2}) \\ &\quad + [1 - P^+(t_1 + t_2 | POT \text{ at } t_0 + t_1)] \Delta S_{\pi\Sigma_n}(t_3, \ldots, t_{n+2}) \\ &= \Delta S_{\pi\Sigma_n}(t_3, \ldots, t_{n+2}), \end{aligned} \tag{7.34}$$

so that

$$\begin{aligned}\Delta S_{\pi\pi\pi\Sigma_n}(t_1,\ldots,t_{n+2}) &= P^+(t_1)\Delta S_{\pi\Sigma_n}(t_3,\ldots,t_{n+2}) \\ &+ [1-P^+(t_1)]\Delta S_{\pi\pi\Sigma_n}(t_2,\ldots,t_{n+2}). \quad (7.35)\end{aligned}$$

This cancellation between the conditional probabilities in the final conditioning step is characteristic of all the r-reset models and is due precisely to the stochastic process resetting, as we saw in the derivation of the recurrence relation for the T_i's above. Similarly, we obtain

$$\begin{aligned}\Delta S_{\pi\pi p\Sigma_n}(t_1,\ldots,t_{n+2}) &= P^+_{ON}(t_1,t_2)[\Delta S_{\Sigma_n}(t_4,\ldots,t_{n+2})+A_+] \\ &+ [1-P^+_{ON}(t_1,t_2)]\Delta S_{p\Sigma_n}(t_3,\ldots,t_{n+2}), \quad (7.36)\\ \Delta S_{\pi p\pi\Sigma_n}(t_1,\ldots,t_{n+2}) &= P^+(t_1)[\Delta S_{\pi\Sigma_n}(t_3,\ldots,t_{n+2})+A_+] \\ &+ [1-P^+(t_1)]\Delta S_{p\pi\Sigma_n}(t_2,\ldots,t_{n+2}), \quad (7.37)\\ \Delta S_{\pi pp\Sigma_n}(t_1,\ldots,t_{n+2}) &= P^+(t_1)[\Delta S_{p\Sigma_n}(t_3,\ldots,t_{n+2})+A_+] \\ &+ [1-P^+(t_1)]\Delta S_{pp\Sigma_n}(t_2,\ldots,t_{n+2}), \quad (7.38)\end{aligned}$$

and the remaining four expressions with a leading p spike follow by symmetry.

Integrating out the spike times and summing over all spike sequences Σ_n as for the 1-reset model, we have

$$\begin{aligned}\Delta S^{\pi\pi\pi}_{n+3} &= \lambda'_\pi(1-K_1^+)\Delta S^{\pi\pi}_{n+2}+\lambda'^{\,2}_\pi K_1^+\Delta S^{\pi}_{n+1}, \quad (7.39)\\ \Delta S^{\pi\pi p}_{n+3} &= \lambda'^{\,2}_\pi(1-T_2^+)\Delta S^{p}_{n+1}+\lambda'^{\,2}_\pi\lambda'_p T_2^+(\Delta S_n+A_+), \quad (7.40)\\ \Delta S^{\pi p\pi}_{n+3} &= \lambda'_\pi(1-T_1^+)\Delta S^{p\pi}_{n+2}+\lambda'_\pi\lambda'_p T_1^+(\Delta S^{\pi}_{n+1}+\lambda'_\pi A_+), \quad (7.41)\\ \Delta S^{\pi pp}_{n+3} &= \lambda'_\pi(1-T_1^+)\Delta S^{pp}_{n+2}+\lambda'_\pi\lambda'_p T_1^+(\Delta S^{p}_{n+1}+\lambda'_p A_+). \quad (7.42)\end{aligned}$$

Since $\Delta S^{\pi\pi}_{n+3}=\Delta S^{\pi\pi\pi}_{n+3}+\Delta S^{\pi\pi p}_{n+3}$ and $\Delta S^{\pi p}_{n+3}=\Delta S^{\pi p\pi}_{n+3}+\Delta S^{\pi pp}_{n+3}$, etc., we obtain

$$\begin{aligned}\Delta S^{\pi\pi}_{n+3} &= \lambda'^{\,2}_\pi(1-T_2^+)\Delta S^{p}_{n+1}+\lambda'^{\,2}_\pi\lambda'_p T_2^+(\Delta S_n+A_+) \\ &+ \lambda'_\pi(1-K_1^+)\Delta S^{\pi\pi}_{n+2}+\lambda'^{\,2}_\pi K_1^+\Delta S^{\pi}_{n+1}, \quad (7.43)\\ \Delta S^{\pi p}_{n+3} &= \lambda'_\pi(1-T_1^+)\Delta S^{p}_{n+2}+\lambda'_\pi\lambda'_p T_1^+(\Delta S_{n+1}+A_+). \quad (7.44)\end{aligned}$$

The trick to simplify these equations further is to observe that $\Delta S^{\pi\pi}_{n+2}=\Delta S^{\pi}_{n+2}-\Delta S^{\pi p}_{n+2}$ and use Eq. (7.44) to rewrite $\Delta S^{\pi p}_{n+2}$. Since this equation has no terms in ΔS_n, this trick does not increase the order of the resulting recurrence relation, but does reduce it from a coupled 4×4 matrix to a more manageable, coupled

2×2 matrix system. Hence,

$$\begin{aligned}\Delta S^{\pi}_{n+3} &= {\lambda'_\pi}^2\lambda'_p[T_2^+ - T_1^+(1-K_1^+)](\Delta S_n + A_+) \\ &+ {\lambda'_\pi}^2[1 - T_2^+ - (1-T_1^+)(1-K_1^+)]\Delta S^p_{n+1} + {\lambda'_\pi}^2 K_1^+ \Delta S^{\pi}_{n+1} \\ &+ \lambda'_\pi\lambda'_p T_1^+(\Delta S_{n+1} + A_+) + \lambda'_\pi(1-T_1^+)\Delta S_{n+2}. \qquad (7.45)\end{aligned}$$

A symmetrical version of this equation holds for ΔS^p_{n+3}. We observe that $T_2^+ - T_1^+(1-K_1^+) = K_2^+$ and $1 - T_2^+ - (1-T_1^+)(1-K_1^+) = K_1^+ - K_2^+$, and the overall multiplier of A_+ is $\lambda'_p(\lambda'_\pi K_1^+ + {\lambda'_\pi}^2 K_2^+)$. So, defining $Q_2^+ = \lambda'_\pi K_1^+ + {\lambda'_\pi}^2 K_2^+$ and $Q_2^- = \lambda'_p K_1^- + {\lambda'_p}^2 K_2^-$, we obtain the coupled system

$$\begin{aligned}\mathbf{\Delta S}_{n+3} &= \begin{pmatrix} +\lambda'_p A_+ Q_2^+ \\ -\lambda'_\pi A_- Q_2^- \end{pmatrix} + \lambda'_\pi\lambda'_p \begin{pmatrix} \lambda'_\pi K_2^+ & \lambda'_\pi K_2^+ \\ \lambda'_p K_2^- & \lambda'_p K_2^- \end{pmatrix} \mathbf{\Delta S}_n \\ &+ \begin{pmatrix} \lambda'_\pi K_1^+ & \lambda'_\pi K_1^+ - {\lambda'_\pi}^2 K_2^+ \\ \lambda'_p K_1^- - {\lambda'_p}^2 K_2^- & \lambda'_p K_1^- \end{pmatrix} \mathbf{\Delta S}_{n+1} \\ &+ \begin{pmatrix} \lambda'_\pi(1-K_1^+) & \lambda'_\pi(1-K_1^+) \\ \lambda'_p(1-K_1^-) & \lambda'_p(1-K_1^-) \end{pmatrix} \mathbf{\Delta S}_{n+2}. \qquad (7.46)\end{aligned}$$

This coupled recurrence relation represents our final form for the 2-reset model.

7.4.3 3-Reset Model

Now we go out to the first four spikes in the sequence because the fourth identical spike possesses the capacity to reset the stochastic processes. Three conditioning steps and the usual weighting by W give

$$\Delta S^{\pi\pi\pi\pi}_{n+4} = \lambda'_\pi(1-K_1^+)\Delta S^{\pi\pi\pi}_{n+3} + {\lambda'_\pi}^2(K_1^+ - K_2^+)\Delta S^{\pi\pi}_{n+2} + {\lambda'_\pi}^3 K_2^+ \Delta S^{\pi}_{n+1}. \quad (7.47)$$

The four equations for $\Delta S^{\pi p\pi\pi}_{n+4}$, $\Delta S^{\pi p\pi p}_{n+4}$, $\Delta S^{\pi pp\pi}_{n+4}$ and $\Delta S^{\pi ppp}_{n+4}$ all sum to give

$$\Delta S^{\pi p}_{n+4} = \lambda'_\pi\lambda'_p T_1^+(\Delta S_{n+2} + A_+) + \lambda'_\pi(1-T_1^+)\Delta S^p_{n+3}, \qquad (7.48)$$

while those for $\Delta S^{\pi\pi p\pi}_{n+4}$ and $\Delta S^{\pi\pi pp}_{n+4}$ sum to give

$$\Delta S^{\pi\pi p}_{n+4} = {\lambda'_\pi}^2\lambda'_p T_2^+(\Delta S_{n+1} + A_+) + {\lambda'_\pi}^2(1-T_2^+)\Delta S^p_{n+2}, \qquad (7.49)$$

and we also have

$$\Delta S^{\pi\pi\pi p}_{n+4} = {\lambda'_\pi}^3\lambda'_p T_3^+(\Delta S_n + A_+) + {\lambda'_\pi}^3(1-T_3^+)\Delta S^p_{n+1}. \qquad (7.50)$$

In these four equations, for reasons that will become clear in the r-reset model, we maintain a rigorous distinction between T_1^+ and K_1^+ even though they are identical. We see that these last three equations are, in fact, instances of the

generic equation

$$\Delta S_m^{\pi^i p} = {\lambda'_\pi}^i \lambda'_p T_i^+ (\Delta S_{m-i-1} + A_+) + {\lambda'_\pi}^i (1 - T_i^+) \Delta S_{m-i}^p, \tag{7.51}$$

for $m > i$, where $\pi^i p$ means i π spikes followed by a p spike. This equation follows easily, in fact, from conditioning on whether the synapse is in the POT state when the first p spike arrives, leading to the T_i^+ term, or in the OFF state when the first p spike arrives, leading to the $1 - T_i^+$ term, and then, as usual, integrating out interspike intervals and summing over all subsequent spikes.

By observing that

$$\Delta S_m^\pi = \Delta S_m^{\pi^i} + \sum_{j=1}^{i-1} \Delta S_m^{\pi^j p}, \tag{7.52}$$

we can use this equation to eliminate all occurrences of $\Delta S_m^{\pi^i}$ terms, $i > 1$, from Eq. (7.47) and then use Eq. (7.51) to eliminate all occurrences of $\Delta S_m^{\pi^i p}$ terms, $i > 0$, leaving terms in ΔS_m^π and ΔS_m^p only. Defining $Q_3^+ = \lambda'_\pi K_1^+ + {\lambda'_\pi}^2 K_2^+ + {\lambda'_\pi}^3 K_3^+$, after some algebra, we finally obtain

$$\begin{aligned} \Delta S_{n+4}^\pi &= \lambda'_\pi (1 - K_1^+) \Delta S_{n+3} + (\lambda'_\pi K_1^+ - {\lambda'_\pi}^2 K_2^+) \Delta S_{n+2} + {\lambda'_\pi}^2 K_2^+ \Delta S_{n+1} \\ &- {\lambda'_\pi}^3 K_3^+ \Delta S_{n+1}^p + \lambda'_p {\lambda'_\pi}^3 K_3^+ \Delta S_n + \lambda'_p A_+ Q_3^+, \end{aligned} \tag{7.53}$$

with a similar equation for ΔS_{n+4}^p. It is convenient to define $\tilde{K}_i^+ = {\lambda'_\pi}^i K_i^+$ and $\tilde{K}_i^- = {\lambda'_p}^i K_i^-$. Then $Q_i^\pm = \sum_{j=1}^i \tilde{K}_j^\pm$ and we may finally write

$$\begin{aligned} \mathbf{\Delta S}_{n+4} &= \begin{pmatrix} +\lambda'_p A_+ Q_3^+ \\ -\lambda'_\pi A_- Q_3^- \end{pmatrix} + \begin{pmatrix} \lambda'_p \tilde{K}_3^+ & \lambda'_p \tilde{K}_3^+ \\ \lambda'_\pi \tilde{K}_3^- & \lambda'_\pi \tilde{K}_3^- \end{pmatrix} \mathbf{\Delta S}_n \\ &+ \begin{pmatrix} \tilde{K}_2^+ & \tilde{K}_2^+ - \tilde{K}_3^+ \\ \tilde{K}_2^- - \tilde{K}_3^- & \tilde{K}_2^- \end{pmatrix} \mathbf{\Delta S}_{n+1} \\ &+ \begin{pmatrix} \tilde{K}_1^+ - \tilde{K}_2^+ & \tilde{K}_1^+ - \tilde{K}_2^+ \\ \tilde{K}_1^- - \tilde{K}_2^- & \tilde{K}_1^- - \tilde{K}_2^- \end{pmatrix} \mathbf{\Delta S}_{n+2} \\ &+ \begin{pmatrix} \lambda'_\pi - \tilde{K}_1^+ & \lambda'_\pi - \tilde{K}_1^+ \\ \lambda'_p - \tilde{K}_1^- & \lambda'_p - \tilde{K}_1^- \end{pmatrix} \mathbf{\Delta S}_{n+3}. \end{aligned} \tag{7.54}$$

This coupled recurrence relation represents our final form for the 3-reset model.

7.4.4 r-Reset Model

Having built some intuitions for small r, we can now proceed to the general r case. The standard conditioning argument gives

$$\Delta S_{n+r+1}^{\pi^{r+1}} = \sum_{i=0}^{r-1} (\lambda'_\pi \tilde{K}_i^+ - \tilde{K}_{i+1}^+) \Delta S_{n+r-i}^{\pi^{r-i}} + \tilde{K}_r^+ \Delta S_{n+1}^\pi, \tag{7.55}$$

where $\tilde{K}_0^{\pm} \equiv 1$, and we still have the generic results in Eqs. (7.51) and (7.52). It is, of course, Eq. (7.55) that always ensures that the resulting recurrence relation in ΔS_m^{π} and ΔS_m^{p} is of order $r+1$ and hence a finite system of coupled equations. Pulling out the $i = r-1$ term from the sum and then writing $j = i+1$, we have the slightly more convenient form

$$\Delta S_{n+r+1}^{\pi^{r+1}} = \sum_{j=1}^{r-1} (\lambda_\pi' \tilde{K}_{j-1}^{+} - \tilde{K}_j^{+}) \Delta S_{n+r+1-j}^{\pi^{r+1-j}} + \lambda_\pi' \tilde{K}_{r-1}^{+} \Delta S_{n+1}^{\pi}. \tag{7.56}$$

Using Eq. (7.52) to eliminate terms in $\Delta S_m^{\pi^i}$, $i > 1$, and then Eq. (7.51) to eliminate terms in $\Delta S_m^{\pi^i p}$, $i > 0$, we obtain the apparently rather unpromising equation

$$\Delta S_{n+r+1}^{\pi} = X_1 + \lambda_p' A_+ X_2 + \lambda_p' X_3 + X_4, \tag{7.57}$$

where

$$X_1 = \lambda_\pi' \tilde{K}_{r-1}^{+} \Delta S_{n+1}^{\pi} + \sum_{j=1}^{r-1} (\lambda_\pi' \tilde{K}_{j-1}^{+} - \tilde{K}_j^{+}) \Delta S_{n+r+1-j}^{\pi}, \tag{7.58}$$

$$X_2 = \sum_{j=1}^{r} {\lambda_\pi'}^{j} T_j^{+} - \sum_{j=1}^{r-1} \sum_{i=1}^{r-j} {\lambda_\pi'}^{i+j} (K_{j-1}^{+} - K_j^{+}) T_i^{+}, \tag{7.59}$$

$$X_3 = \sum_{j=1}^{r} {\lambda_\pi'}^{j} T_j^{+} \Delta S_{n+r-j} - \sum_{j=1}^{r-1} \sum_{i=1}^{r-j} {\lambda_\pi'}^{i+j} (K_{j-1}^{+} - K_j^{+}) T_i^{+} \Delta S_{n+r-i-j}, \tag{7.60}$$

$$X_4 = \sum_{j=1}^{r} {\lambda_\pi'}^{j} (1 - T_j^{+}) \Delta S_{n+r+1-j}^{p} - \sum_{j=1}^{r-1} \sum_{i=1}^{r-j} {\lambda_\pi'}^{i+j} (K_{j-1}^{+} - K_j^{+})(1 - T_i^{+}) \Delta S_{n+r+1-i-j}^{p}. \tag{7.61}$$

However, the quantities X_2, X_3 and X_4 simplify dramatically. Consider, for example, X_3:

$$\begin{aligned} X_3 &= \sum_{j=1}^{r} {\lambda_\pi'}^{j} T_j^{+} \Delta S_{n+r-j} - \sum_{j=1}^{r-1} \sum_{i=1}^{r-j} {\lambda_\pi'}^{i+j} (K_{j-1}^{+} - K_j^{+}) T_i^{+} \Delta S_{n+r-i-j} \\ &= \sum_{j=1}^{r} {\lambda_\pi'}^{j} T_j^{+} \Delta S_{n+r-j} - \sum_{k=2}^{r} {\lambda_\pi'}^{k} \Delta S_{n+r-k} \sum_{l=1}^{k-1} (K_{l-1}^{+} - K_l^{+}) T_{k-l}^{+} \\ &= \lambda_\pi' T_1^{+} \Delta S_{n+r-1} + \sum_{j=2}^{r} {\lambda_\pi'}^{j} \Delta S_{n+r-j} \left[T_j^{+} - \sum_{i=1}^{j-1} (K_{i-1}^{+} - K_i^{+}) T_{j-i}^{+} \right] \\ &\equiv \lambda_\pi' T_1^{+} \Delta S_{n+r-1} + \sum_{j=2}^{r} {\lambda_\pi'}^{j} \Delta S_{n+r-j} K_j^{+} \\ &= \sum_{j=1}^{r} \tilde{K}_j^{+} \Delta S_{n+r-j}, \end{aligned} \tag{7.62}$$

where the penultimate line follows from the recurrence relation defining T_j^+ in Eq. (7.8). This was the reason for rigorously distinguishing between T_1 and K_1, even though they are identical, so that we could immediately use Eq. (7.8) without any further manipulation. The key step in this simplification is to rewrite the double sum by defining $i+j$ to be k and summing over its range. Identical manipulations reduce X_2 and X_4, giving

$$X_2 = \sum_{j=1}^{r} \tilde{K}_j^+ \equiv Q_r^+, \tag{7.63}$$

$$X_4 = \sum_{j=1}^{r} (\lambda'_\pi \tilde{K}_{j-1}^+ - \tilde{K}_j^+) \Delta S^p_{n+r+1-j}. \tag{7.64}$$

Thus, finally we have

$$\begin{aligned} \Delta S^\pi_{n+r+1} &= \lambda'_p A_+ Q_r^+ + \lambda'_\pi \tilde{K}_{r-1}^+ \Delta S^\pi_{n+1} + \sum_{j=1}^{r-1} (\lambda'_\pi \tilde{K}_{j-1}^+ - \tilde{K}_j^+) \Delta S^\pi_{n+r+1-j} \\ &+ \lambda'_p \sum_{j=1}^{r} \tilde{K}_j^+ \Delta S_{n+r-j} + \sum_{j=1}^{r} (\lambda'_\pi \tilde{K}_{j-1}^+ - \tilde{K}_j^+) \Delta S^p_{n+r+1-j}. \end{aligned} \tag{7.65}$$

This expression is valid for all $r \geq 1$. After a little further manipulation, we obtain, for $r \geq 2$,

$$\begin{aligned} \Delta S^\pi_{n+r+1} &= \lambda'_p A_+ Q_r^+ + \lambda'_p \tilde{K}_r^+ \Delta S_n + \tilde{K}_{r-1}^+ \Delta S_{n+1} - \tilde{K}_r^+ \Delta S^p_{n+1} \\ &+ \sum_{j=1}^{r-2} (\tilde{K}_j^+ - \tilde{K}_{j+1}^+) \Delta S_{n+r-j} + (\lambda'_\pi - \tilde{K}_1^+) \Delta S_{n+r}. \end{aligned} \tag{7.66}$$

To obtain the correct expression for $r = 1$, we must replace the terms in ΔS_{n+1} and ΔS^p_{n+1} by $\lambda'_\pi \Delta S_{n+1} - \tilde{K}_1^+ \Delta S^p_{n+1}$. A similar expression follows by symmetry for ΔS^p_{n+r+1}, so that in matrix form, for $r \geq 2$, we have

$$\begin{aligned} \mathbf{\Delta S}_{n+r+1} &= \begin{pmatrix} +\lambda'_p A_+ Q_r^+ \\ -\lambda'_\pi A_- Q_r^- \end{pmatrix} + \begin{pmatrix} \lambda'_p \tilde{K}_r^+ & \lambda'_p \tilde{K}_r^+ \\ \lambda'_\pi \tilde{K}_r^- & \lambda'_\pi \tilde{K}_r^- \end{pmatrix} \mathbf{\Delta S}_n \\ &+ \begin{pmatrix} \tilde{K}_{r-1}^+ & \tilde{K}_{r-1}^+ - \tilde{K}_r^+ \\ \tilde{K}_{r-1}^- - \tilde{K}_r^- & \tilde{K}_{r-1}^- \end{pmatrix} \mathbf{\Delta S}_{n+1} \\ &+ \sum_{j=1}^{r-2} \begin{pmatrix} \tilde{K}_j^+ - \tilde{K}_{j+1}^+ & \tilde{K}_j^+ - \tilde{K}_{j+1}^+ \\ \tilde{K}_j^- - \tilde{K}_{j+1}^- & \tilde{K}_j^- - \tilde{K}_{j+1}^- \end{pmatrix} \mathbf{\Delta S}_{n+r-j} \\ &+ \begin{pmatrix} \lambda'_\pi - \tilde{K}_1^+ & \lambda'_\pi - \tilde{K}_1^+ \\ \lambda'_p - \tilde{K}_1^- & \lambda'_p - \tilde{K}_1^- \end{pmatrix} \mathbf{\Delta S}_{n+r}, \end{aligned} \tag{7.67}$$

where, of course, $\sum_{j=1}^{r-2} \equiv 0$ for $r = 2$.

Notice that the equations for ΔS^π_{n+r+1} and ΔS^p_{n+r+1} involve terms in ΔS_m only,

except for the single terms ΔS^{π}_{n+1} and ΔS^{p}_{n+1}. These two terms thus prevent us from simply adding the equations for ΔS^{π}_{n+r+1} and ΔS^{p}_{n+r+1} and hence obtaining a single, uncoupled $(r+1)$th order recurrence relation for ΔS_m. The above matrix recurrence relation represents our final form and we now turn to obtaining the exact solution for $r = 1$ and the asymptotic, large n solution for $r > 1$.

7.5 Solution of Recurrence Relations

We now solve the 1-reset model exactly for any n, then determine the leading order, asymptotic solutions for large n for the r-reset, $r > 1$, model. Finally we perform some consistency checks on these various results.

7.5.1 Exact Solution of the 1-Reset Model

The 1-reset model has the second order recurrence relation given by

$$\begin{aligned} \mathbf{\Delta S}_{n+2} &= \begin{pmatrix} +\lambda'_p A_+ Q_1^+ \\ -\lambda'_\pi A_- Q_1^- \end{pmatrix} + \begin{pmatrix} \lambda'_p \tilde{K}_1^+ & \lambda'_p \tilde{K}_1^+ \\ \lambda'_\pi \tilde{K}_1^- & \lambda'_\pi \tilde{K}_1^- \end{pmatrix} \mathbf{\Delta S}_n \\ &+ \begin{pmatrix} \lambda'_\pi & \lambda'_\pi - \tilde{K}_1^+ \\ \lambda'_p - \tilde{K}_1^- & \lambda'_p \end{pmatrix} \mathbf{\Delta S}_{n+1}. \end{aligned} \tag{7.68}$$

We are not interested in how the separate components ΔS^{π}_n and ΔS^{p}_n evolve but only in their sum, $\Delta S_n = (1,1)^T \cdot \mathbf{\Delta S}_n = \Delta S^{\pi}_n + \Delta S^{p}_n$. We have that

$$\begin{aligned} \Delta S_{n+2} &= (\lambda'_p A_+ Q_1^+ - \lambda'_\pi A_- Q_1^-) + (\lambda'_p \tilde{K}_1^+ + \lambda'_\pi \tilde{K}_1^-) \Delta S_n \\ &+ \Delta S_{n+1} - (\tilde{K}_1^+ \Delta S^{p}_{n+1} + \tilde{K}_1^- \Delta S^{\pi}_{n+1}) \end{aligned} \tag{7.69}$$

and also

$$\begin{aligned} \tilde{K}_1^+ \Delta S^{p}_{n+2} + \tilde{K}_1^- \Delta S^{\pi}_{n+2} &= (\lambda'_p A_+ \tilde{K}_1^- Q_1^+ - \lambda'_\pi A_- \tilde{K}_1^+ Q_1^-) + \tilde{K}_1^+ \tilde{K}_1^- \Delta S_n \\ &+ (\lambda'_p \tilde{K}_1^+ + \lambda'_\pi \tilde{K}_1^-) \Delta S_{n+1} - \tilde{K}_1^+ \tilde{K}_1^- \Delta S_{n+1}. \end{aligned} \tag{7.70}$$

At the cost of increasing the order of the recurrence relation, we can thus obtain a relation purely in ΔS_m, to give

$$(\Delta S_{n+3} - \Delta S_{n+2}) - \tilde{K}_1^+ \tilde{K}_1^- (\Delta S_{n+1} - \Delta S_n) = \tilde{A}_1, \tag{7.71}$$

where we have defined the inhomogeneous part of this equation, $\tilde{A}_1$, by

$$\tilde{A}_1 = \lambda'_p A_+ (1 - \tilde{K}_1^-) Q_1^+ - \lambda'_\pi A_- (1 - \tilde{K}_1^+) Q_1^-. \tag{7.72}$$

The characteristic equation for the homogeneous part of the equation has the three roots $+1$ and $\pm\sqrt{\tilde{K}_1^+\tilde{K}_1^-}$. Since $0 \leq K_1^{\pm} < 1$ for finite β and $0 \leq \lambda'_\pi$, $\lambda'_p \leq 1$, these last two roots are real and have moduli less than unity. The general solution of this equation is then

$$\Delta S_n = n\frac{\tilde{A}_1}{1-\tilde{K}_1^+\tilde{K}_1^-} + \left(\tilde{K}_1^+\tilde{K}_1^-\right)^{n/2}[D_1 + (-1)^n D_2] + D_3, \tag{7.73}$$

where the constants D_1, D_2 and D_3 can be determined by requiring that

$$\begin{aligned}
\Delta S_0 &= 0, \\
\Delta S_1 &= 0, \\
\Delta S_2 &= \lambda'_p A_+ \tilde{K}_1^+ - \lambda'_\pi A_- \tilde{K}_1^-.
\end{aligned}$$

Because $\tilde{K}_1^+\tilde{K}_1^- < 1$, the asymptotic, large n form of this solution is then just

$$\frac{1}{n}\Delta S_n \sim \frac{\tilde{A}_1}{1-\tilde{K}_1^+\tilde{K}_1^-}. \tag{7.74}$$

The solution for large enough n thus essentially scales linearly with n, and this is exactly what we should expect: a typical train of $2n$ spikes should induce, on average, twice as much change in synaptic strength as a typical train of n spikes, for n sufficiently large. Of course, real neurons would not be expected to scale in this manner, because synaptic strengths are presumably constrained and thus saturate at upper or lower limits. In the general r-reset model, we shall deploy this scaling argument to discard the non-unity solutions of the associated characteristic equation, leaving only the asymptotic, scaling solution.

It is worth using the recurrence relation directly to compute the first few ΔS_n, in order to develop an intuition for how the terms develop. Although Eq. (7.73) of course produces these solutions, they are rather opaque. Defining ΔS_n^+ to be only the A_+-dependent part of the solution, with ΔS_n^-, the A_--dependent part, following by symmetry, we can directly compute the following:

$$\begin{aligned}
\Delta S_2^+ &= \lambda'_p A_+ \tilde{K}_1^+ [1], \\
\Delta S_3^+ &= \lambda'_p A_+ \tilde{K}_1^+ \left[2 - 1\tilde{K}_1^-\right], \\
\Delta S_4^+ &= \lambda'_p A_+ \tilde{K}_1^+ \left[3 - 2\tilde{K}_1^- + 1(\tilde{K}_1^-\tilde{K}_1^+)\right], \\
\Delta S_5^+ &= \lambda'_p A_+ \tilde{K}_1^+ \left[4 - 3\tilde{K}_1^- + 2(\tilde{K}_1^-\tilde{K}_1^+) - 1\tilde{K}_1^-(\tilde{K}_1^-\tilde{K}_1^+)\right], \\
\Delta S_6^+ &= \lambda'_p A_+ \tilde{K}_1^+ \left[5 - 4\tilde{K}_1^- + 3(\tilde{K}_1^-\tilde{K}_1^+) - 2\tilde{K}_1^-(\tilde{K}_1^-\tilde{K}_1^+) + 1(\tilde{K}_1^-\tilde{K}_1^+)^2\right].
\end{aligned}$$

The trend is absolutely clear, and in fact, for the general r-reset model, we shall prove such a trend, although the terms are rather more complicated in that case.

Indeed, it is one of the virtues of the 1-reset model, mathematically speaking, that its solutions are so transparent, allowing intuition to develop for the more difficult, general r model.

We can relate the solution obtained by iterating the recurrence relation above to that obtained by a direct, general solution of the recurrence relation, Eq. (7.73). For $n = 2m + 1$, $m \geq 1$, the general expression for ΔS_n^+ is seen to be

$$\Delta S_{2m+1}^+ = \lambda_p' A_+ \tilde{K}_1^+ \sum_{i=0}^{m-1} \left[2(m-i)(1-\tilde{K}_1^-) + \tilde{K}_1^- \right] \left(\tilde{K}_1^+ \tilde{K}_1^- \right)^i , \qquad (7.75)$$

and for $n = 2(m + 1)$, $m \geq 0$, we have

$$\Delta S_{2(m+1)}^+ = \lambda_p' A_+ \tilde{K}_1^+ \sum_{i=0}^{m} \left\{ [2(m-i)+1](1-\tilde{K}_1^-) + \tilde{K}_1^- \right\} \left(\tilde{K}_1^+ \tilde{K}_1^- \right)^i . \qquad (7.76)$$

These two expressions should really be regarded as a conjecture for the general form of the solution, since they are based on a generalisation from the small n case only. We can directly evaluate the sums in Eqs. (7.75) and (7.76) and obtain expressions that agree exactly with the A_+-dependent part of Eq. (7.73) after an evaluation of $D_1 - D_2$, $D_1 + D_2$ and D_3, thus confirming the conjecture. We therefore see that although the exact solution in Eq. (7.73) was obtained very easily and is mathematically very simple, its structure actually obscures the regularity in the expansion producing Eqs. (7.75) and (7.76).

7.5.2 Asymptotic Solution of the r-Reset Model

As for the 1-reset model, we can write

$$\begin{aligned} \Delta S_{n+r+1} &= (\lambda_p' A_+ Q_r^+ - \lambda_\pi' A_- Q_r^-) + (\lambda_p' \tilde{K}_r^+ + \lambda_\pi' \tilde{K}_r^-) \Delta S_n \\ &+ (\tilde{K}_{r-1}^+ + \tilde{K}_{r-1}^-) \Delta S_{n+1} + (1 - \tilde{K}_1^+ - \tilde{K}_1^-) \Delta S_{n+r} \\ &+ \sum_{j=1}^{r-2} (\tilde{K}_j^+ + \tilde{K}_j^- - \tilde{K}_{j+1}^+ - \tilde{K}_{j+1}^-) \Delta S_{n+r-j} \\ &- (\tilde{K}_r^+ \Delta S_{n+1}^p + \tilde{K}_r^- \Delta S_{n+1}^\pi) \end{aligned} \qquad (7.77)$$

and

$$\begin{aligned}
&\tilde{K}_r^+ \Delta S^p_{n+r+1} + \tilde{K}_r^- \Delta S^\pi_{n+r+1} \\
&= (\lambda_p' A_+ \tilde{K}_r^- Q_r^+ - \lambda_\pi' A_- \tilde{K}_r^+ Q_r^-) + \tilde{K}_r^+ \tilde{K}_r^- \Delta S_n \\
&+ (\tilde{K}^+_{r-1} \tilde{K}_r^- + \tilde{K}^-_{r-1} \tilde{K}_r^+ - \tilde{K}_r^+ \tilde{K}_r^-) \Delta S_{n+1} \\
&+ \sum_{j=1}^{r-2} \left[\tilde{K}_r^- (\tilde{K}_j^+ - \tilde{K}_{j+1}^+) + \tilde{K}_r^+ (\tilde{K}_j^- - \tilde{K}_{j+1}^-) \right] \Delta S_{n+r-j} \\
&+ (\lambda_p' \tilde{K}_r^+ + \lambda_\pi' \tilde{K}_r^-) \Delta S_{n+r} - (\tilde{K}_1^+ \tilde{K}_r^- + \tilde{K}_1^- \tilde{K}_r^+) \Delta S_{n+r}. \qquad (7.78)
\end{aligned}$$

Hence, defining

$$\tilde{A}_r = \lambda_p' A_+ (1 - \tilde{K}_r^-) Q_r^+ - \lambda_\pi' A_- (1 - \tilde{K}_r^+) Q_r^-, \qquad (7.79)$$

after some algebra we have

$$\begin{aligned}
&\Delta S_{n+2r+1} - \Delta S_{n+2r} \\
&= \tilde{A}_r + \tilde{K}_r^+ \tilde{K}_r^- (\Delta S_{n+1} - \Delta S_n) \\
&- \sum_{j=1}^{r-1} (\tilde{K}_j^+ + \tilde{K}_j^-)(\Delta S_{n+2r+1-j} - \Delta S_{n+2r-j}) \\
&+ \sum_{j=1}^{r-1} (\tilde{K}_j^+ \tilde{K}_r^- + \tilde{K}_j^- \tilde{K}_r^+)(\Delta S_{n+r+1-j} - \Delta S_{n+r-j}), \qquad (7.80)
\end{aligned}$$

where we have written this equation so that it is clear that there is a $+1$ solution of the characteristic equation for the homogeneous form of the recurrence relation. With $\sum_{j=1}^{r-1} \equiv 0$ for $r = 1$, this equation is valid for all $r \geq 1$.

We define $U_n = \Delta S_{n+1} - \Delta S_n$, with $\Delta S_1 \equiv 0$, so that $\Delta S_{n+1} = \sum_{i=1}^{n} U_i$. Eq. (7.80) then implies that the U_i's satisfy the recurrence relation

$$\sum_{j=0}^{r} \Lambda_j U_{n+2r-j} - \sum_{j=0}^{r} (\tilde{K}_j^+ \tilde{K}_r^- + \tilde{K}_j^- \tilde{K}_r^+) U_{n+r-j} = \tilde{A}_r - \tilde{K}_r^+ \tilde{K}_r^- U_n, \qquad (7.81)$$

where $\Lambda_j = \tilde{K}_j^+ + \tilde{K}_j^-$, $j \geq 1$, and we define $\Lambda_0 \equiv 1$. These U_i's will play an important role later in the construction of a generating function for ΔS_n. The characteristic equation for the homogeneous part of Eq. (7.81) is, after a little re-arrangement

$$\theta^{2r} + \sum_{j=1}^{r-1} \Lambda_j \theta^{2r-j} - \sum_{j=1}^{r-1} (\tilde{K}_j^+ \tilde{K}_r^- + \tilde{K}_j^- \tilde{K}_r^+) \theta^{r-j} - \tilde{K}_r^+ \tilde{K}_r^- = 0. \qquad (7.82)$$

There is no θ^r term. This equation has $2r$ roots. Call them θ_i, $i = 1, \ldots, 2r$. Of course, the roots of the characteristic equation for the homogeneous part of

Eq. (7.80) are just the θ_i's supplemented with the additional, known unity root. For $r = 1$, we saw that $|\theta_i| < 1$, all i, leading to the observed scaling behaviour. In principle, we could find the θ_i's for $r = 2$, but the expressions are messy. Furthermore, we would still be left with all the $r > 2$ cases, for which closed-form expressions for the roots in general likely do not exist. In order to be able to understand the asymptotic behaviour of ΔS_n for general r, we need to bound $|\theta_i|$, all i. The standard bounding theorems (e.g., Gerschgorin's circle theorem applied to the companion matrix of Eq. (7.82)) do not appear to bound the $|\theta_i|$'s sufficiently strongly for our purposes. We therefore adopt a different approach to show that $|\theta_i| \leq 1$, all i.

We can, in fact, prove that ΔS_n grows no faster than linearly in n in the general r-reset model without having to check by explicit calculation that the roots θ_i do not have moduli greater than unity. Consider a sequence of spikes $\Sigma_n(t_0, \ldots, t_{n-1})$ containing m π spikes and $n - m$ p spikes. Each π spike could induce a change of $-A_-$ in synaptic efficacy if it triggers a $DEP \rightarrow OFF$ transition; otherwise it induces no change. Similarly, each p spike could induce a change of $+A_+$ if it triggers a $POT \rightarrow OFF$ transition; otherwise it induces no change. Hence, $\Delta S_{\Sigma_n}(t_0, \ldots, t_{n-1})$ is bounded,

$$-mA_- \leq \Delta S_{\Sigma_n}(t_0, \ldots, t_{n-1}) \leq +(n-m)A_+. \tag{7.83}$$

We integrate out the interspike intervals, weight by the probability of the particular spike sequence Σ_n, and sum over all such sequences to obtain

$$-A_- \sum_{m=0}^{n} \binom{n}{m} m {\lambda'_\pi}^m {\lambda'_p}^{n-m} \leq \Delta S_n \leq +A_+ \sum_{m=0}^{n} \binom{n}{m} (n-m) {\lambda'_\pi}^m {\lambda'_p}^{n-m}, \tag{7.84}$$

or

$$-\lambda'_\pi A_- \leq \frac{1}{n} \Delta S_n \leq +\lambda'_p A_+. \tag{7.85}$$

Thus, ΔS_n is bounded from both above and below by a linear function of n and so cannot grow faster than linearly in n. All the roots θ_i must therefore have moduli not exceeding unity.

We can dispense with the roots θ_i satisfying the strict inequality $|\theta_i| < 1$ because we know that the particular solution of Eq. (7.80) induces at least a linear term in n (because of the known, unity root), and this will dominate these exponentially decaying θ_i^n terms. We are left with any possible roots satisfying $|\theta_i| = 1$. If these are unrepeated, then, again, the particular solution dominates for large enough n. We must therefore examine the possibility of repeated, unit modulus roots, including the supplemental, known unity root.

By applying Descartes' rule of signs to Eq. (7.82), we see that there is precisely one positive (real) root. Applying the rules of synthetic division to obtain an upper bound on the location of this root, we find that it is less than unity. The known, unity root of the characteristic equation for ΔS_n is therefore not repeated, and so does not contribute any linear (or higher) terms to ΔS_n, which would affect the asymptotic solution. Moreover, we know that were this unity root repeated m times, $m > 1$, the particular solution of the recurrence relation would have to grow like n^m, and this can be excluded by the bounding argument above. We are then left with the possibility of repeated $|\theta_i| = 1$, $\theta_i \neq 1$ roots. We know that such roots cannot be repeated three or more times, since otherwise they would contribute quadratic or higher terms to ΔS_n, violating the bound above. Thus, if there are any repeated $|\theta_i| = 1$, $\theta_i \neq 1$ roots, then they can only be repeated twice. However, such twice-repeated roots would contribute linearly-growing, oscillatory terms to ΔS_n, and, given the construction of our model, we can exclude this possibility on the grounds of incompatibility with the model's dynamics.

In summary, we have that $|\theta_i| \leq 1$, all i, and if there are any roots $\theta_i \neq 1$ satisfying $|\theta_i| = 1$, then they are not repeated. In fact, numerically, it appears that all the θ_i satisfy the strict inequality $|\theta_i| < 1$, so that unit modulus roots (except the singleton, known unity root) appear to be entirely excluded, except in certain, limiting cases. So as not to disrupt the flow of the Chapter, we explore in the appendix particular cases in which the roots can be calculated exactly.

The linear boundedness of ΔS_n derived in Eq. (7.85) proves the linear scaling behaviour of ΔS_n in general, and we have argued that the dominant, linear term in this scaling behaviour must derive only from the particular solution of the recurrence relation. Thus, we finally obtain the asymptotic, large n solution

$$\frac{1}{n}\Delta S_n \sim \frac{\lambda_p' A_+(1-\tilde{K}_r^-)Q_r^+ - \lambda_\pi' A_-(1-\tilde{K}_r^+)Q_r^-}{(1-\tilde{K}_r^+)(1-\tilde{K}_r^-) + (1-\tilde{K}_r^-)Q_r^+ + (1-\tilde{K}_r^+)Q_r^-}, \tag{7.86}$$

valid for any r. We will provide two different derivations of this asymptotic solution, one from a direct construction of a generating function for ΔS_n. These approaches do not involve any discussion of the roots θ_i, and therefore these alternative derivations can be viewed as independent, although somewhat indirect proofs that $|\theta_i| \leq 1$, all i. We shall see later that the asymptotic solution is pretty good even for $n \approx 10$.

The formal limit $r \to \infty$ is, of course, beset by various technical difficulties associated with the limit and the resulting sum over a potentially infinite number of polynomial roots. However, in the usual spirit of applied mathematics, we shall ignore these possible difficulties and just assume that the limit of the finite r asymptotic solution is the asymptotic solution of the non-resetting model (and

confirm this numerically). Now, $\tilde{K}_r \to 0$ as $r \to \infty$, and we require $Q^{\pm} = \lim_{r\to\infty} Q_r^{\pm}$. We have, for example,

$$\begin{aligned} Q^+ &= \sum_{l=1}^{\infty} {\lambda'_\pi}^l K_l^+(\beta) \\ &= \sum_{l=1}^{\infty} \left(\frac{\tau_+\lambda_\pi}{1+\tau_+\beta}\right)^l \sum_{i=0}^{n_+-1} \left(\begin{array}{c} i+l-1 \\ l-1 \end{array}\right) \frac{1}{(1+\tau_+\beta)^i} \\ &= \sum_{i=0}^{n_+-1} \frac{1}{(1+\tau_+\beta)^i} \sum_{l=1}^{\infty} \left(\begin{array}{c} i+l-1 \\ l-1 \end{array}\right) \left(\frac{\tau_+\lambda_\pi}{1+\tau_+\beta}\right)^l \\ &= \frac{\tau_+\lambda_\pi}{1+\tau_+\beta} \sum_{i=0}^{n_+-1} \frac{1}{(1+\tau_+\beta)^i} \left(1-\frac{\tau_+\lambda_\pi}{1+\tau_+\beta}\right)^{-(i+1)} \\ Q^+ &= \frac{\lambda_\pi}{\lambda_p} K_1^+(\lambda_p), \end{aligned} \tag{7.87}$$

where we summed the infinite series by recognising it as a binomial expansion for negative powers. Similarly $Q^- = \frac{\lambda_p}{\lambda_\pi} K_1^-(\lambda_\pi)$, so, for the non-resetting or $r \to \infty$ model, we obtain the particularly simple expression

$$\frac{1}{n}\Delta S_n \sim \frac{\lambda'_\pi A_+ K_1^+(\lambda_p) - \lambda'_p A_- K_1^-(\lambda_\pi)}{1 + \frac{\lambda_p}{\lambda_\pi} K_1^-(\lambda_\pi) + \frac{\lambda_\pi}{\lambda_p} K_1^+(\lambda_p)}. \tag{7.88}$$

Notice that the limits $\lambda_\pi \to 0$ or $\lambda_p \to 0$ are well defined since $K_1(\lambda) \to 0$ as $\lambda \to 0$ at least as fast as λ.

7.5.3 Consistency Checks on Solutions

According to the arguments of section 7.2, we can obtain the r-reset model by the general substitution $K_{lr+s} \to K_r^l K_s$, $s < r$ and l some non-negative integer. More particularly, the $(r_1 r_2)$-reset model should reduce to the r_1-reset model under the global replacement $K_{lr_1+s} \to K_{r_1}^l K_s$, $s < r_1$. We check that this is actually the case for the asymptotic $(r_1 r_2)$-reset result,

$$\frac{1}{n}\Delta S_n \sim \frac{\lambda'_p A_+ (1-\tilde{K}^-_{r_1r_2}) Q^+_{r_1r_2} - \lambda'_\pi A_- (1-\tilde{K}^+_{r_1r_2}) Q^-_{r_1r_2}}{(1-\tilde{K}^+_{r_1r_2})(1-\tilde{K}^-_{r_1r_2}) + (1-\tilde{K}^-_{r_1r_2}) Q^+_{r_1r_2} + (1-\tilde{K}^+_{r_1r_2}) Q^-_{r_1r_2}}. \tag{7.89}$$

We have, dropping $\pm$ superscripts for notational convenience,

$$\begin{aligned} Q_{r_1r_2} &= \sum_{i=1}^{r_1r_2} \tilde{K}_i \\ &= (\tilde{K}_1 + \cdots + \tilde{K}_{r_1}) + (\tilde{K}_{r_1+1} + \cdots + \tilde{K}_{2r_1}) + (\tilde{K}_{2r_1+1} + \cdots + \tilde{K}_{3r_1}) \\ &\quad + \cdots + (\tilde{K}_{(r_2-1)r_1+1} + \cdots + \tilde{K}_{r_1r_2}) \\ &\to (\tilde{K}_1 + \cdots + \tilde{K}_{r_1})(1 + \tilde{K}_{r_1} + \tilde{K}_{r_1}^2 + \cdots + \tilde{K}_{r_1}^{r_2-1}) \\ &= Q_{r_1}\frac{1-\tilde{K}_{r_1}^{r_2}}{1-\tilde{K}_{r_1}}. \end{aligned} \tag{7.90}$$

Of course, $\tilde{K}_{r_1r_2} \to \tilde{K}_{r_1}^{r_2}$ in Eq. (7.89), so

$$\begin{aligned} \frac{1}{n}\Delta S_n &\to \frac{\lambda'_p A_+ \frac{Q^+_{r_1}}{1-\tilde{K}^+_{r_1}} - \lambda'_\pi A_- \frac{Q^-_{r_1}}{1-\tilde{K}^-_{r_1}}}{1 + \frac{Q^+_{r_1}}{1-\tilde{K}^+_{r_1}} + \frac{Q^-_{r_1}}{1-\tilde{K}^-_{r_1}}} \\ &= \frac{\lambda'_p A_+ (1-\tilde{K}^-_{r_1})Q^+_{r_1} - \lambda'_\pi A_- (1-\tilde{K}^+_{r_1})Q^-_{r_1}}{(1-\tilde{K}^+_{r_1})(1-\tilde{K}^-_{r_1}) + (1-\tilde{K}^-_{r_1})Q^+_{r_1} + (1-\tilde{K}^+_{r_1})Q^-_{r_1}}. \end{aligned} \tag{7.91}$$

So, indeed, the (r_1r_2)-reset model's asymptotic solution reduces to the r_1-reset model's asymptotic solution under the appropriate substitution.

For $n_\pm = 1$, the memorylessness of the exponential distribution requires that all the r-reset models reduce to the 1-reset model. We have that

$$\begin{aligned} K_l^\pm(\beta) &= \left(\frac{\tau_\pm\beta}{1+\tau_\pm\beta}\right)^l \sum_{i=0}^{n_\pm-1} \binom{i+l-1}{l-1} \frac{1}{(1+\tau_\pm\beta)^i} \\ &= \left(\frac{\tau_\pm\beta}{1+\tau_\pm\beta}\right)^l, \end{aligned} \tag{7.92}$$

for $n_\pm = 1$. But $K_1^\pm(\beta) = \tau_\pm\beta/(1+\tau_\pm\beta)$ for $n_\pm = 1$, so $K_l^\pm(\beta) = \left[K_1^\pm(\beta)\right]^l$, and to reduce the l-reset model to the 1-reset model, we would have to make precisely this substitution, $K_l \to K_1^l$, throughout. For $n_\pm = 1$, this replacement holds in virtue of the strict identity $K_l = K_1^l$ in this case, and hence, for $n_\pm = 1$, all r-reset asymptotic solutions reduce to the 1-reset model's asymptotic solution, as required.

The consistency checks on the asymptotic solutions do not amount to consistency checks on the full solutions. Although we have not yet derived the full solution for the general r-reset model, we can in principle check that the recurrence relation for the (r_1r_2)-reset model reduces to that for the r_1-reset model under the appropriate substitution. As an example, we show that the 1-reset model's recurrence relation is identical to that for the 2-reset model's recurrence relation after the replacement of K_2 by K_1^2 in the latter. Of course, the 2-reset model's

equation is third order, while that for the 1-reset model is second order. Thus, we must use the 1-reset model's recurrence relation to take a further step, generating a third order equation. So, in Eq. (7.68), we replace n by $n+1$ and then write

$$\begin{pmatrix} \lambda'_\pi & \lambda'_\pi - \tilde{K}_1^+ \\ \lambda'_p - \tilde{K}_1^- & \lambda'_p \end{pmatrix} = \begin{pmatrix} \lambda'_\pi - \tilde{K}_1^+ & \lambda'_\pi - \tilde{K}_1^+ \\ \lambda'_p - \tilde{K}_1^- & \lambda'_p - \tilde{K}_1^- \end{pmatrix} + \begin{pmatrix} \tilde{K}_1^+ & 0 \\ 0 & \tilde{K}_1^- \end{pmatrix}, \tag{7.93}$$

this being the matrix multiplying the $\mathbf{\Delta S}_{n+2}$ term. We then substitute Eq. (7.68) (at order $n+2$, not $n+3$) back into the second term on the right-hand-side of Eq. (7.93) and, after a little re-arrangement, obtain

$$\begin{aligned} \mathbf{\Delta S}_{n+3} &= \begin{pmatrix} +\lambda'_p A_+ \left[\tilde{K}_1^+ + (\tilde{K}_1^+)^2\right] \\ -\lambda'_\pi A_- \left[\tilde{K}_1^- + (\tilde{K}_1^-)^2\right] \end{pmatrix} \\ &+ \begin{pmatrix} \lambda'_p(\tilde{K}_1^+)^2 & \lambda'_p(\tilde{K}_1^+)^2 \\ \lambda'_\pi(\tilde{K}_1^-)^2 & \lambda'_\pi(\tilde{K}_1^-)^2 \end{pmatrix} \mathbf{\Delta S}_n \\ &+ \begin{pmatrix} \tilde{K}_1^+ & \tilde{K}_1^+ - (\tilde{K}_1^+)^2 \\ \tilde{K}_1^- - (\tilde{K}_1^-)^2 & \tilde{K}_1^- \end{pmatrix} \mathbf{\Delta S}_{n+1} \\ &+ \begin{pmatrix} \lambda'_\pi - \tilde{K}_1^+ & \lambda'_\pi - \tilde{K}_1^+ \\ \lambda'_p - \tilde{K}_1^- & \lambda'_p - \tilde{K}_1^- \end{pmatrix} \mathbf{\Delta S}_{n+2}, \end{aligned} \tag{7.94}$$

which is precisely the 2-reset model's matrix recurrence relation under the substitution $K_2 \to K_1^2$. Hence, the solutions (asymptotic or otherwise) of the 2-reset model must reduce to the solutions of the 1-reset model under the appropriate substitution.

7.6 1-Transition Processes

We now seek to provide a deeper understanding of the asymptotic solutions derived above. This will then permit us to develop a generating function for ΔS_n in the next section.

We rewrite the general, asymptotic solution in the slightly more transparent form

$$\frac{1}{n}\Delta S_n \sim \frac{\mathcal{T}_1}{\mathcal{N}_1}, \tag{7.95}$$

where

$$\mathcal{T}_1 = \lambda'_p A_+ \frac{Q_r^+}{1-\tilde{K}_r^+} - \lambda'_\pi A_- \frac{Q_r^-}{1-\tilde{K}_r^-}, \tag{7.96}$$

$$\mathcal{N}_1 = 1 + \frac{Q_r^+}{1-\tilde{K}_r^+} + \frac{Q_r^-}{1-\tilde{K}_r^-}, \tag{7.97}$$

and we work, henceforth, for convenience, only with the non-resetting, $r \to \infty$ model, so that

$$\mathcal{T}_1 = \lambda'_p A_+ Q^+ - \lambda'_\pi A_- Q^-, \tag{7.98}$$

$$\mathcal{N}_1 = 1 + Q^+ + Q^-. \tag{7.99}$$

We specifically do not sum the terms in $Q^\pm$, obtaining, for example, $Q^+ = \frac{\lambda_\pi}{\lambda_p} K_1(\lambda_p)$, as this obscures the reduction of this model to the r-reset model and also renders opaque a natural interpretation of the terms in $\mathcal{T}_1$ and $\mathcal{N}_1$. Notice that, for the reduction to the r-reset model, we have, for example,

$$Q^+ = \sum_{i=1}^{\infty} \tilde{K}_i^+ \to Q_r^+ \left[1 + \tilde{K}_r^+ + (\tilde{K}_r^+)^2 + \cdots\right] = \frac{Q_r^+}{1 - \tilde{K}_r^+}, \tag{7.100}$$

so we see immediately how terms like $Q_r^+/(1 - \tilde{K}_r^+)$ arise from Q^+ and therefore no longer need to consider the general r-reset model.

We define $\mathcal{T}_1^+ = +\lambda'_p A_+ Q^+$ and $\mathcal{T}_1^- = -\lambda'_\pi A_- Q^-$, so that $\mathcal{T}_1 = \mathcal{T}_1^+ + \mathcal{T}_1^-$. What processes give rise to the terms that occur, for example, in $\mathcal{T}_1^+$? We have that

$$\mathcal{T}_1^+ = \lambda'_p A_+ \sum_{i=1}^{\infty} {\lambda'_\pi}^i K_i^+. \tag{7.101}$$

We know that a K_i^+ term arises from a sequence of i π spikes in which the synapse never returns to the OFF state once placed in the POT state by the first π spike. That is, K_i^+ arises from the probability that the synapse never returns to the OFF state during the whole time interval $[t_0, \sum_{j=0}^{i} t_j)$, which is $P^+(t_1 + \cdots + t_i)$. If the synapse is in the POT state when a final p spike arrives at time $\sum_{j=0}^{i} t_j$, then a change in synaptic strength by an amount A_+ will occur. Thus, the term $\lambda'_p A_+ \tilde{K}_i^+$ is obtained from the spike sequence $\pi^i p$ in which no stochastic transitions back to OFF ever occur and in which the only transition back to OFF is induced by the final p spike. An identical argument holds for the term $-\lambda'_\pi A_i \tilde{K}_i^-$ for the $p^i\pi$ sequence. $\mathcal{T}_1^+$ is thus obtained by summing over all spike sequences $\pi p, \pi^2 p, \pi^3 p, \ldots$, in which there is only one transition back to the OFF state, inducing potentiation by an amount A_+. Similarly for $\mathcal{T}_1^-$. We refer to such processes as 1-transition processes. $\mathcal{T}_1$ is therefore the expected change in synaptic strength induced by all 1-transition processes. Note that if we consider a train of spikes, which may be decomposed into a train of potentiating or depressing 1-transition processes, the final potentiating or depressing transition need not necessarily occur at the very end of the train. That is, further spikes and transitions may freely occur after the final potentiating or depressing transition to OFF, provided that any resulting transitions to OFF do not induce any change in synaptic strength.

Another 1-transition process is of the form $\pi^i\gamma$, where the γ represents the first stochastic transition back to the OFF state. Of course, such a process induces no change in synaptic strength and so does not contribute to $\mathcal{T}_1$. What is the expected number of spikes in all possible 1-transition processes, $\pi^i p$, $p^i\pi$, $\pi^i\gamma$ and $p^i\gamma$, any $i > 0$? The expected number of spikes in all the sequences $\pi^i p$ is just

$$\mathcal{N}_p^+ = \lambda'_p \sum_{i=1}^{\infty}(i+1)\tilde{K}_i^+, \tag{7.102}$$

and similarly for $p^i\pi$ sequences we have

$$\mathcal{N}_\pi^- = \lambda'_\pi \sum_{i=1}^{\infty}(i+1)\tilde{K}_i^-. \tag{7.103}$$

For the sequence $\pi^i\gamma$, the probability that the first stochastic transition back to the OFF state occurs after the ith π spike is $P^+(t_1+\cdots+t_{i-1})-P^+(t_1+\cdots+t_i)$, $i > 1$, leading to $K_{i-1}^+ - K_i^+$, with $K_0^+ \equiv 1$, $i > 0$. Hence, the expected number of spikes in all the sequences $\pi^i\gamma$ is

$$\mathcal{N}_\gamma^+ = \sum_{i=1}^{\infty} i(\lambda'_\pi \tilde{K}_{i-1}^+ - \tilde{K}_i^+) \tag{7.104}$$

and in all $p^i\gamma$ sequences, we have

$$\mathcal{N}_\gamma^- = \sum_{i=1}^{\infty} i(\lambda'_p \tilde{K}_{i-1}^- - \tilde{K}_i^-). \tag{7.105}$$

Now,

$$\begin{aligned}
\mathcal{N}_p^+ + \mathcal{N}_\gamma^+ &= \sum_{i=1}^{\infty}\left[i\lambda'_\pi \tilde{K}_{i-1}^+ - i\tilde{K}_i^+ + \lambda'_p(i+1)\tilde{K}_i^+\right] \\
&= \lambda'_\pi + \sum_{i=1}^{\infty}\left[(i+1)\tilde{K}_i^+ - i\tilde{K}_i^+\right] \\
&= \lambda'_\pi + Q^+.
\end{aligned} \tag{7.106}$$

Similarly $\mathcal{N}_\pi^- + \mathcal{N}_\gamma^- = \lambda'_p + Q^-$, and so the expected number of spikes in all 1-transitions is just $1 + Q^+ + Q^-$, which is precisely $\mathcal{N}_1$, the denominator is the asymptotic expression for ΔS_n. $\mathcal{T}_1$ is therefore the expected change in synaptic strength in all 1-transition processes, and $\mathcal{N}_1$ is the expected number of spikes in all 1-transition processes.

Consider some general sequence of n spikes. This sequence will decompose, by definition, into a chain of successive 1-transition processes defined by the synapse returning to the OFF state at the end of each 1-transition process. How many 1-transition processes do we expect, on average, to occur in a sequence of n spikes?

Because the average length of a 1-transition process is $\mathcal{N}_1$ spikes, obviously there are $n/\mathcal{N}_1$ 1-transitions in a chain of n spikes, on average. Each such 1-transition process induces, on average, a change in synaptic strength by an amount $\mathcal{T}_1$. The total change induced by n spikes is thus expected to be $\mathcal{T}_1 \times (n/\mathcal{N}_1)$, but this should just be approximately ΔS_n, the expected change induced by n spikes. When n is sufficiently large, the statistical error should become negligible, so we should have

$$\Delta S_n \sim n\frac{\mathcal{T}_1}{\mathcal{N}_1}, \tag{7.107}$$

which is, of course, Eq. (7.95).

It is a remarkable fact that the asymptotic solutions that were derived somewhat laboriously from the recurrence relations nevertheless can be extracted so quickly by carving up an arbitrary sequence of spikes into its natural articulation points, these being the times at which the synapse has returned to the OFF state. This is possible because of the underlying Markovian nature of 1-transition processes: once the synapse has returned to the OFF state, its subsequent evolution is independent of its history at that point. It is therefore natural to think in terms of a chain of 1-transition processes, the end point of each of which is the OFF state.

Of course, the recurrence relations and their exact solutions are primary, while the asymptotic solutions are only approximate and thus secondary. However, we can now show that the asymptotic solutions provide a generating function for ΔS_n that is, in fact, exact, for all n, not just large n.

7.7 Generating Function for ΔS_n

We shall now show that although the asymptotic solution for ΔS_n is approximate, it in fact contains all the information required to build the exact solution for ΔS_n, any n. Although this seems astonishing, we shall show that it is a natural consequence of the manner in which the asymptotic solution is obtained from the exact form of ΔS_n so constructed.

We consider again the non-resetting model, $r \to \infty$. Inspired by the earlier 1-transition analysis, and also the regularity in the terms for ΔS_n^+ for the 1-reset model, we explicitly calculate ΔS_n for small n by decomposing ΔS_n into terms representing the final potentiating or depressing 1-transition process in the sequence of n spikes. Lengthy calculation then reveals that the A_+-dependent

parts are given by

$$\begin{aligned}
\Delta S_2^+ &= \lambda'_p A_+ \left[F_0 \tilde{K}_1^+\right], \\
\Delta S_3^+ &= \lambda'_p A_+ \left[F_0 \tilde{K}_2^+ + (F_0+F_1)\tilde{K}_1^+\right], \\
\Delta S_4^+ &= \lambda'_p A_+ \left[F_0 \tilde{K}_3^+ + (F_0+F_1)\tilde{K}_2^+ + (F_0+F_1+F_2)\tilde{K}_1^+\right],
\end{aligned}$$

where

$$\begin{aligned}
F_0 &= 1, \\
F_1 &= 1-\Lambda_1, \\
F_2 &= 1-\Lambda_1-(\Lambda_2-\Lambda_1^2).
\end{aligned}$$

The expressions for the A_--dependent parts are given by the usual transformations of the A_+-dependent parts, namely $\lambda'_\pi \leftrightarrow \lambda'_p$ and $+ \leftrightarrow -$. Notice that F_0, F_1 and F_2 are invariant under these transformations, and so the ΔS_n^- have identical F_i coefficients to those for the ΔS_n^+.

In general, by this final 1-transition process decomposition, we know that we must be able to write

$$\Delta S_{n+1} = \sum_{i=1}^{n} G_{n-i}(\lambda'_p A_+ \tilde{K}_i^+ - \lambda'_\pi A_- \tilde{K}_i^-), \tag{7.108}$$

where the $\lambda'_p A_+ \tilde{K}_i^+$ term represents the final 1-transition process contribution from a final spike sequence $\pi^i p$, and $-\lambda'_\pi A_- \tilde{K}_i^-$ that from a final $p^i\pi$ sequence. They must have the same coefficient, G_{n-i}, because G_{n-i} essentially counts all the possible ways in which we arrive at the OFF state before the final 1-transition process occurs, and this must be independent of whether that final process is $\pi^i p$ or $p^i\pi$, by the Markovian property of 1-transitions. Moreover, although G_i will be a function of the $\tilde{K}_j^\pm$, it must be a symmetric function of them, so that $G_i \to G_i$ under $\lambda'_\pi \leftrightarrow \lambda'_p$ and $+ \leftrightarrow -$, because for any sequence Σ_i that arrives at the OFF state after i spikes, there is a complementary sequence $\bar{\Sigma}_i$ that also arrives at the OFF state after i spikes, where $\bar{\Sigma}_i$ is obtained from Σ_i by $\pi \leftrightarrow p$, i.e. replace every π spike by a p spike and every p spike by a π spike. So, the expansion in Eq. (7.108) is generic, exploiting the importance of the 1-transition processes. It remains to determine the coefficients G_i. Since $U_n = \Delta S_{n+1} - \Delta S_n$, we have the expansion

$$U_n = \sum_{i=1}^{n} F_{n-i}(\lambda'_p A_+ \tilde{K}_i^+ - \lambda'_\pi A_- \tilde{K}_i^-), \tag{7.109}$$

where $F_0 = G_0 = 1$ and $F_i = G_i - G_{i-1}$, $i > 0$, or $G_i = \sum_{j=0}^{i} F_j$. So, equivalently, we can determine instead the coefficients F_i rather than G_i.

Although we are now interested only in the non-resetting model, because we know how to extract the r-reset model from it, a determination of the coefficients F_i depends critically on the recurrence relations for the r-reset models. In fact, we can explicitly calculate the sequence of functions ΔS_n, $n = 2, 3, 4, \ldots$, for the non-resetting model in an incremental manner from the 1-, 2-, 3-, ..., reset models, in the following manner. Consider ΔS_2. The sequences that could possibly contribute to the function are just $\pi\pi$, πp, $p\pi$ and pp. The two sequences πp and $p\pi$ never probe any resetting properties, because we need a chain of at least two identical spikes for this. Although the sequences $\pi\pi$ and pp do probe 1-resetting properties, they give exactly zero contribution to the function. Hence, the function ΔS_2 is independent of r for $r \geq 1$, and we can therefore use the $r = 1$ recurrence relation, Eq. (7.68), with $n = 0$ and $\Delta S_0 = 0$ and $\Delta S_1 = 0$ to calculate it. Now consider ΔS_3. Eight sequences could possibly contribute to it. Only $\pi\pi\pi$ and ppp could probe 2-resetting properties, but, again, they give zero contribution. All the other sequences probe at most 1-resetting properties. Thus, ΔS_3 is independent of r for $r \geq 2$. We can therefore use the $r = 2$ form of the recurrence relation in Eq. (7.67), with $n = 0$ and the initial values $\Delta S_0 = 0$, $\Delta S_1 = 0$ and ΔS_2 calculated from the $r = 1$ recurrence relation, to calculate ΔS_3. Similarly, ΔS_4 is independent of r for $r \geq 3$, and we can use the $r = 3$ form of Eq. (7.67) with $n = 0$ to calculate it, using the forms of ΔS_i, $i \leq 3$, already established. In general, therefore, we see that ΔS_{m+1} is independent of r for $r \geq m$ because $m + 1$ spikes cannot probe the m-resetting properties with non-zero contributions, so we can use the $r = m$ form of Eq. (7.67) with $n = 0$ and all the previously established forms of ΔS_i, $i \leq m$, to calculate ΔS_{m+1}.

The important thing to notice about this procedure is that we always set $\Delta S_0 = 0$ and $\Delta S_1 = 0$. This means that the corresponding matrices multiplying the vectors $\mathbf{\Delta S}_0$ and $\mathbf{\Delta S}_1$ in Eq. (7.67) with $n = 0$ do not contribute to the calculation of ΔS_{m+1} for $r \geq m$. But it was the form of the matrix multiplying the $\mathbf{\Delta S}_1$ term (the $\mathbf{\Delta S}_{n+1}$ term for non-zero n) in particular that prevented us before from simply adding the two equations for ΔS_m^π and ΔS_m^p to obtain a scalar equation entirely in ΔS_m and not a matrix equation in $\mathbf{\Delta S}_m$. So, now with $\Delta S_0 = 0$ and $\Delta S_1 = 0$, we are free to add the equations, obtaining

$$\Delta S_{m+1} = (\lambda_p' A_+ Q_m^+ - \lambda_\pi' A_- Q_m^-) + \sum_{i=0}^{m-2} (\Lambda_i - \Lambda_{i+1}) \Delta S_{m-i}, \tag{7.110}$$

valid for $m \geq 1$ with $\sum_{i=0}^{m-2} \equiv 0$ for $m = 1$, and we have used $\Lambda_0 \equiv 1$. After a little re-arrangement, we have

$$\sum_{i=0}^{m-1} \Lambda_i U_{m-i} = \lambda_p' A_+ Q_m^+ - \lambda_\pi' A_- Q_m^-, \tag{7.111}$$

where we have used the definition $U_i = \Delta S_{i+1} - \Delta S_i$. This recurrence relation succinctly encodes the procedure described above, and thus determines the form of ΔS_{m+1} for an r-reset model in which $r \geq m$. Taking the limit $r \to \infty$ to give the non-resetting model, Eq. (7.111) is valid for any $m \geq 1$, and represents a simple recurrence relation to calculate the form of U_m, any m, in the non-resetting model.

Substituting the expansion in Eq. (7.109) into the recurrence relation in Eq. (7.111) and looking at, say, only the A_+-dependent piece, we have

$$\sum_{i=0}^{m-1} \Lambda_i \sum_{j=1}^{m-i} F_{m-i-j} \tilde{K}_j^+ = Q_m^+, \tag{7.112}$$

and we can re-arrange the double sum to give

$$\sum_{i=1}^{m} \tilde{K}_i^+ \sum_{j=0}^{m-i} \Lambda_j F_{m-i-j} = Q_m^+. \tag{7.113}$$

Since $Q_m^+ = \sum_{i=1}^m \tilde{K}_i^+$, we see that Eq. (7.109) is a solution of Eq. (7.111) provided that the F_i satisfy the recurrence relation

$$\sum_{i=0}^{n} F_i \Lambda_{n-i} = 1, \tag{7.114}$$

any n. With $F_0 = \Lambda_0 = 1$, we can compute the first few F_i's:

$$\begin{aligned}
F_0 &= 1, \\
F_1 &= 1 - \Lambda_1, \\
F_2 &= 1 - \Lambda_1 - (\Lambda_2 - \Lambda_1^2), \\
F_3 &= 1 - \Lambda_1 - (\Lambda_2 - \Lambda_1^2) - (\Lambda_3 - 2\Lambda_1\Lambda_2 + \Lambda_1^3), \\
F_4 &= 1 - \Lambda_1 - (\Lambda_2 - \Lambda_1^2) - (\Lambda_3 - 2\Lambda_1\Lambda_2 + \Lambda_1^3) \\
&\quad - (\Lambda_4 - 2\Lambda_1\Lambda_3 + 3\Lambda_1^2\Lambda_2 - \Lambda_2^2 + \Lambda_1^4),
\end{aligned}$$

the first three of which agree with those calculated directly earlier. We recognise the new terms at each iteration as very closely related to the ΔT_i's, computed explicitly before. Again, as expected, we see combinatorial results playing a critical role.

We now write $F_i = \sum_{j=0}^i C_j$, $C_0 \equiv 1$. By subtracting $\sum_{i=0}^{n+1} F_i \Lambda_{n+1-i}$ from $\sum_{i=0}^n F_i \Lambda_{n-i}$, we see that C_i must satisfy the recurrence relation

$$\sum_{i=0}^{n} C_i \Lambda_{n-i} = 0, \tag{7.115}$$

for $n \geq 1$, with $C_0 \equiv 1$. We then have

$$\begin{aligned} C_0 &= 1, \\ C_1 &= -\Lambda_1, \\ C_2 &= -(\Lambda_2 - \Lambda_1^2), \\ C_3 &= -(\Lambda_3 - 2\Lambda_1\Lambda_2 + \Lambda_1^3), \end{aligned}$$

etc., carving up the F_i's in the natural way according to the partition of the integer at each step in the C_i's. We have the standard combinatorial expansion in terms of integer partitions,

$$C_l = \sum_{k=1}^{l} \sum_{\{p_l \vdash k\}} \frac{k!}{a_1! \cdots a_l!} (-1)^k \Lambda_1^{a_1} \cdots \Lambda_l^{a_l}, \tag{7.116}$$

and, in fact, this is its own inverse, so that

$$\Lambda_l = \sum_{k=1}^{l} \sum_{\{p_l \vdash k\}} \frac{k!}{a_1! \cdots a_l!} (-1)^k C_1^{a_1} \cdots C_l^{a_l}. \tag{7.117}$$

It is straightforward to construct a generating function for the C_i's, i.e. the function $\sum_{i=0}^{\infty} C_i x^i$. Consider the product $(\sum_{i=0}^{\infty} C_i x^i) \times (\sum_{i=0}^{\infty} \Lambda_i x^i)$. The coefficient of the general x^n term in this product is just $\sum_{i=0}^{n} C_i \Lambda_{n-i}$. But, for $n \geq 1$, the recurrence relation for the C_i's tells us that $\sum_{i=0}^{n} C_i \Lambda_{n-i} \equiv 0$. The coefficient of the x^0 term is just $C_0 \Lambda_0$, which is 1. Hence, we have the result

$$\sum_{i=0}^{\infty} C_i x^i = \frac{1}{\sum_{i=0}^{\infty} \Lambda_i x^i}. \tag{7.118}$$

Equally, $1/(\sum_{i=0}^{\infty} C_i x^i)$ is the generating function for the Λ_i's, explaining the form of the inverse of Eq. (7.116), since we can replace $C_i \leftrightarrow \Lambda_i$ throughout by the reciprocity of their respective generating functions. Because the coefficient of the x^n term in the expansion $(1-x)^{-1} \sum_{i=0}^{\infty} C_i x^i$ is just $\sum_{j=0}^{n} C_j \equiv F_n$, we can see that the generating function for the F_i's is

$$\sum_{i=0}^{\infty} F_i x^i = \frac{1}{1-x} \frac{1}{\sum_{i=0}^{\infty} \Lambda_i x^i}, \tag{7.119}$$

and similarly, the generating function for the G_i's is

$$\sum_{i=0}^{\infty} G_i x^i = \frac{1}{(1-x)^2} \frac{1}{\sum_{i=0}^{\infty} \Lambda_i x^i}. \tag{7.120}$$

We can at last construct an explicit generating function for ΔS_n. By direct calculation, we find that

$$\sum_{i=0}^{\infty} \Delta S_i x^i = \frac{x}{(1-x)^2} \frac{1}{\sum_{i=0}^{\infty} \Lambda_i x^i} \sum_{i=1}^{\infty} x^i (\lambda_p' A_+ \tilde{K}_i^+ - \lambda_\pi' A_- \tilde{K}_i^-), \tag{7.121}$$

where we recognise this as essentially the generating function of the G_i's multiplied by that for $\lambda_p' A_+ \tilde{K}_i^+ - \lambda_\pi' A_- \tilde{K}_i^-$. We can directly evaluate the two sums $\sum_{i=0}^{\infty} \tilde{K}_i^{\pm} x^i$ that appear in these generating functions, but we choose not to because the resulting expressions make the reduction to the general r-reset model opaque.

We are now in a position to re-examine the form of the asymptotic limit of $\frac{1}{n}\Delta S_n$. We know that $G_i = \sum_{j=0}^{i} \sum_{k=0}^{j} C_k$, so that, explicitly,

$$G_i = \sum_{j=0}^{i} (i+1-j) C_j. \tag{7.122}$$

We then have

$$\begin{aligned} \frac{1}{n+1} G_{n-i} &= \frac{1}{n+1} \sum_{j=0}^{n-i} (n+1-i-j) C_j \\ &= \sum_{j=0}^{n} C_j - \frac{1}{n+1} \sum_{j=0}^{n-i} (i+j) C_j - \sum_{j=n-i+1}^{n} C_j. \end{aligned} \tag{7.123}$$

Under suitable assumptions about the large order behaviour of the C_i's, we then have the asymptotic behaviour

$$\frac{1}{n+1} G_{n-i} \sim \sum_{j=0}^{n} C_j \sim \sum_{j=0}^{\infty} C_j = \frac{1}{\sum_{j=0}^{\infty} \Lambda_j} \equiv \frac{1}{\mathcal{N}_1}. \tag{7.124}$$

For large enough n, therefore, the sum in the equation

$$\Delta S_{n+1} = \sum_{i=1}^{n} G_{n-i} (\lambda_p' A_+ \tilde{K}_i^+ - \lambda_\pi' A_- \tilde{K}_i^-)$$

basically factorises, so that

$$\begin{aligned}
\frac{1}{n+1}\Delta S_{n+1} &= \sum_{i=1}^{n} \frac{G_{n-i}}{n+1}(\lambda'_p A_+ \tilde{K}_i^+ - \lambda'_\pi A_- \tilde{K}_i^-) \\
&\sim \sum_{i=1}^{n} \frac{1}{\mathcal{N}_1}(\lambda'_p A_+ \tilde{K}_i^+ - \lambda'_\pi A_- \tilde{K}_i^-) \\
&= \frac{1}{\mathcal{N}_1}\sum_{i=1}^{n}(\lambda'_p A_+ \tilde{K}_i^+ - \lambda'_\pi A_- \tilde{K}_i^-) \\
&\sim \frac{1}{\mathcal{N}_1}\sum_{i=1}^{\infty}(\lambda'_p A_+ \tilde{K}_i^+ - \lambda'_\pi A_- \tilde{K}_i^-) \\
&= \frac{\mathcal{T}_1}{\mathcal{N}_1},
\end{aligned}$$

again giving the standard asymptotic behaviour $\frac{1}{n}\Delta S_n \sim \mathcal{T}_1/\mathcal{N}_1$. All this argument really does is to encode our expectation that, for n large enough, the number of ways of arriving back at OFF after $n-(i-1)$ and $n-i$ spikes tend to the same number, so the final 1-transition processes of length $i-1$ or i contribute equally to the overall sum. The expressions for $\mathcal{T}_1$ and $\mathcal{N}_1$ that appear in the asymptotic form contain all the information necessary to reconstruct the exact form of ΔS_n, precisely because of the manner in which the asymptotic limit is extracted from the exact, finite n result, and the individual components C_i can be extracted uniquely from $\mathcal{N}_1$ because the integer partitions can be uniquely identified with the integer i.

Summarising these various results, we have that

$$\begin{aligned}
U_n &= \sum_{i=1}^{n} F_{n-i}(\lambda'_p A_+ \tilde{K}_i^+ - \lambda'_\pi A_- \tilde{K}_i^-), \\
\Delta S_{n+1} &= \sum_{i=1}^{n} G_{n-i}(\lambda'_p A_+ \tilde{K}_i^+ - \lambda'_\pi A_- \tilde{K}_i^-),
\end{aligned}$$

where $G_i = \sum_{j=0}^{i} F_j$ and $F_i = \sum_{j=0}^{i} C_j$, and the C_i's are generated by the function $1/\sum_{i=0}^{\infty} \Lambda_i x^i$. By simple re-arrangements of the sums in these two representations of U_n and ΔS_{n+1}, we also have

$$\begin{aligned}
U_n &= \sum_{i=1}^{n} C_{n-i}(\lambda'_p A_+ Q_i^+ - \lambda'_\pi A_- Q_i^-), \\
\Delta S_{n+1} &= \sum_{i=1}^{n} F_{n-i}(\lambda'_p A_+ Q_i^+ - \lambda'_\pi A_- Q_i^-).
\end{aligned}$$

These equivalent expressions represent the exact solution of the non-resetting model. Of course, the r-reset model is obtained by the standard replacement

$K_{lr+s} \to K_r^l K_s$, so that

$$\sum_{i=1}^{\infty} \tilde{K}_i^{\pm} x^i \to \frac{\sum_{i=1}^{r} \tilde{K}_i^{\pm} x^i}{1 - \tilde{K}_r^{\pm} x^r} \tag{7.125}$$

in the various associated generating functions.

We conclude this section, finally, by indicating an alternative method of calculating the coefficients F_i. Consider a train of i spikes and consider the number of ways of arriving finally at the OFF state after the last, ith spike. After weighting the spike trains by their pdfs and individual spike probabilities, and summing over all possible spike trains, using the methods of section 7.4, we arrive at a recurrence relation of order $r+1$ for the r-reset model. This recurrence relation is identical to that for ΔS_n, except that it is homogeneous, i.e. the $\tilde{A}_r$ term is absent. Iterative solution of this homogeneous equation starting from $F_0 \equiv 1$ yields precisely the F_i coefficients appearing in the equation for U_n. Thus, the form for U_n represents a decomposition into all possible sequences of 1-transitions returning to the OFF state, giving the F_i coefficients, followed by a final 1-transition, giving the $\lambda'_p A_+ \tilde{K}_i^+ - \lambda'_\pi A_- \tilde{K}_i^-$ term, and summing over all possible lengths of the final 1-transition.

7.8 Analytical and Simulated Results

In this section, we turn to a comparison between the analytical results derived above and numerical results obtained by the computational simulation of the r-reset model. We are particularly concerned with establishing how rapidly $\frac{1}{n}\Delta S_n$ converges to the analytical asymptotic results; establishing that the non-resetting model's asymptotic result matches a large n simulation of the same model, confirming indeed that the $r \to \infty$ limit of the r-reset model's asymptotic result really is the asymptotic result of the non-resetting model; and establishing how rapidly the r-reset model tends to the non-resetting model as $r \to \infty$.

We run simulations of the stochastic model as in Chapter 6 (Appleby and Elliott, 2005), with the standard set of parameters used there:

$$\begin{aligned} A_+ &= 1.00, \\ A_- &= 0.95, \\ \tau_+ &= 13.3\,\text{ms}, \\ \tau_- &= 20.0\,\text{ms}, \\ n_+ &= 3, \\ n_- &= 3. \end{aligned}$$

These parameters were chosen because they reproduced a large variety of the experimental results available in the literature, including the replication of the spike triplet results (Froemke and Dan, 2002). These parameters also guarantee the presence of both a depressing and a potentiating regime in the function ΔS_2, so that depression and potentiation can both occur at the ensemble level. In most of the simulations below, each simulation is run over exactly 10^7 spikes. Thus, in obtaining the numerical form of $\frac{1}{n}\Delta S_n$, we perform $10^7/n$ separate runs over n spikes, and average over all these runs. Of course, as n increases, there are fewer runs over which to average, but there is more averaging within each sequence of n spikes. This means that the statistical error should be roughly equivalent for any value of n, this being the reason for keeping the total number of spikes fixed at 10^7. ΔS_n is, of course, a function of both λ_π and λ_p. Previously, we set

$$\lambda_p = \begin{cases} \lambda_\pi - \Theta & \text{for } \lambda_\pi \geq \Theta \\ 0 & \text{otherwise} \end{cases} , \tag{7.126}$$

as a simple model of postsynaptic firing with a threshold set at $\Theta = 5$ Hz. This gives a cross section through the λ_π–λ_p plane. We will also show contour plots of $\frac{1}{n}\Delta S_n$ in the λ_π–λ_p plane.

In Fig. 7.1, we investigate how rapidly $\frac{1}{n}\Delta S_n$ approaches the asymptotic limit derived above. We use the $r = 1$ reset model for convenience, since we also have very simple, exact solutions against which to compare the numerical results. We see very rapid convergence to the asymptotic limit, so that even $n = 8$ gives the limit to within a few percent. This means that after $n \approx 10$ spikes, we expect to see the scaling property observed earlier for the 1-reset model, so that doubling the spike train length merely doubles the expected change in synaptic efficacy, adding no new qualitative features to the dynamics.

The results for the non-resetting model are shown in Fig. 7.2. Here we do not plot the corresponding finite n analytical results, but we show the analytical asymptotic limit. This figure demonstrates that the $r \to \infty$ limit of the asymptotic result for the r-reset model is, indeed, the asymptotic result for the non-resetting model.

We address the question of how rapidly the r-reset model converges to the non-resetting model as $r \to \infty$ in Fig. 7.3. We see that even the 4-reset model is extremely close to the non-resetting model, and the 8-reset model is virtually indistinguishable from it. For $\lambda_\pi \approx \lambda_p$, we have that $\lambda'_\pi \approx 0.5$ and $\lambda'_p \approx 0.5$. Thus, the probability of a long sequence of identical spikes in this case is maximally suppressed, and the contributions of such sequences to ΔS_n are weighted by these small probabilities. However, to resolve the difference between the r-reset model and the $(r+1)$-reset model for large r, we require precisely such

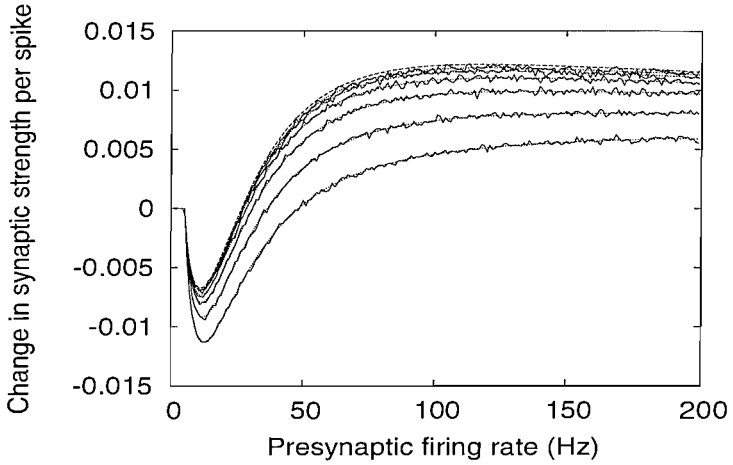

FIGURE 7.1: Analytical and numerical values of $\frac{1}{n}\Delta S_n$ plotted as a function of λ_π for the 1-reset model for various values of n. The solid lines show the numerical results for $n = 2, 4, 8, 16, 32$ and 64 spikes, while the corresponding dotted lines show the exact, analytical result. The $n = 2$ pair of lines corresponds to the bottom pair, $n = 4$ the next pair up, etc., the function $\frac{1}{n}\Delta S_n$ increasing as a function of n for fixed λ_π. The dashed line shows the exact, asymptotic limit, $n \to \infty$.

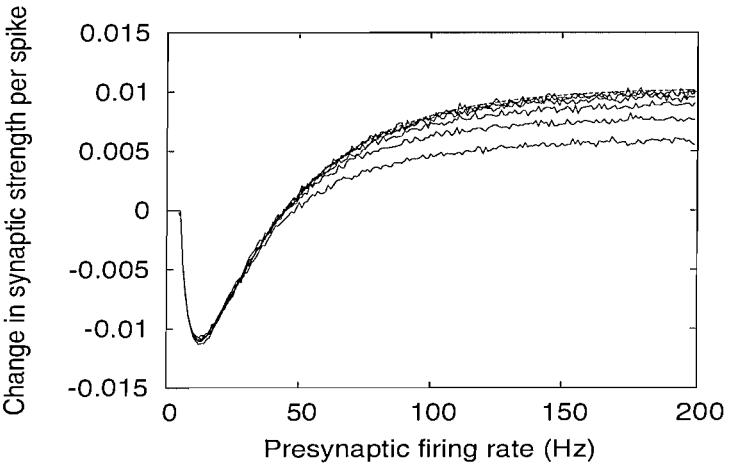

FIGURE 7.2: Numerical values of $\frac{1}{n}\Delta S_n$ plotted as a function of λ_π for the non-resetting model for various values of n. The solid lines again show the numerical results for $n = 2, 4, 8, 16, 32$ and 64 spikes, with the bottom line corresponding to $n = 2$, the next line up $n = 4$, etc., $\frac{1}{n}\Delta S_n$ exhibiting the same monotonicity as a function of n for fixed λ_π as in the 1-reset model. The dashed line shows the exact, analytical result for the asymptotic limit of the non-resetting model.

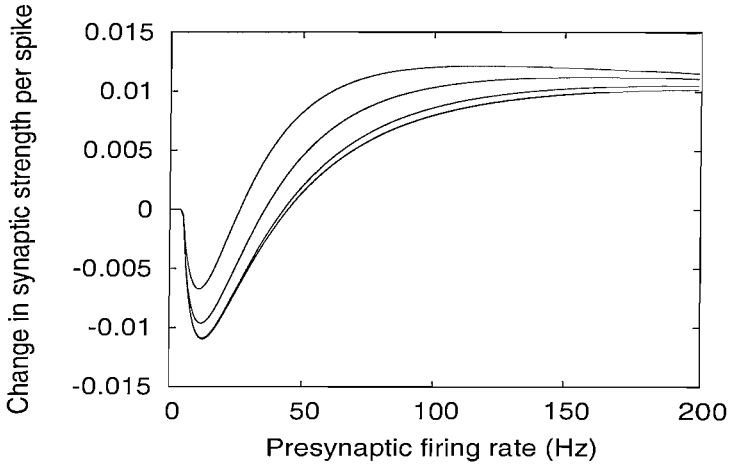

FIGURE 7.3: The analytical asymptotic limit of $\frac{1}{n}\Delta S_n$ for the r-reset model plotted as a function of λ_π for various values of r. Shown are results for $r = 1, 2, 4, 8, 16$, and the non-resetting model. The top curve corresponds to $r = 1$, the next one down $r = 2$, etc., $\frac{1}{n}\Delta S_n$ monotonically decreasing as a function of r for fixed λ_π. The curves for the 8-, 16- and non-resetting models are virtually identical, almost superimposed on top of each other.

long sequences of identical spikes. Hence, we should not be too surprised by this rapid convergence of the r-reset model to the non-resetting model.

The numerical simulations presented in Fig. 7.1 and Fig. 7.2 show data averaged over a large number of runs. However, we know that for n sufficiently large, ΔS_n scales linearly with n. This means that a single run over a typical spike train (without averaging over multiple runs) for a sufficiently large number of spikes should start to exhibit self-averaging properties, and the data should settle down to the exact, average result for ΔS_n. As the spike train length tends to infinity, the numerical result over this single spike train should converge to the analytical result for the average, asymptotic behaviour. A natural question is, therefore, how many spikes are required in such a train before the single-train form for the change in synaptic strength per spike begins to resemble the analytical result for $\frac{1}{n}\Delta S_n$, which is averaged over all spike trains of length n? In Fig. 7.4, we plot the raw data for the 1-reset model obtained from a single spike train containing 50,000 spikes, and, for comparison, we plot the exact, asymptotic limit of $\frac{1}{n}\Delta S_n$. We see that, although noisy, as expected, the single-train data follow the exact result. If the spike train length is increased to 100,000 spikes, the noise is much reduced, while for a spike train length of 10,000 spikes, although the trend in the mean is still clear, the noise begins to swamp the mean behaviour. We conclude that, for the parameter values used here, approximately 50,000 spikes are required for the self-averaging property of a long spike train to reduce the overall variance in the data significantly.

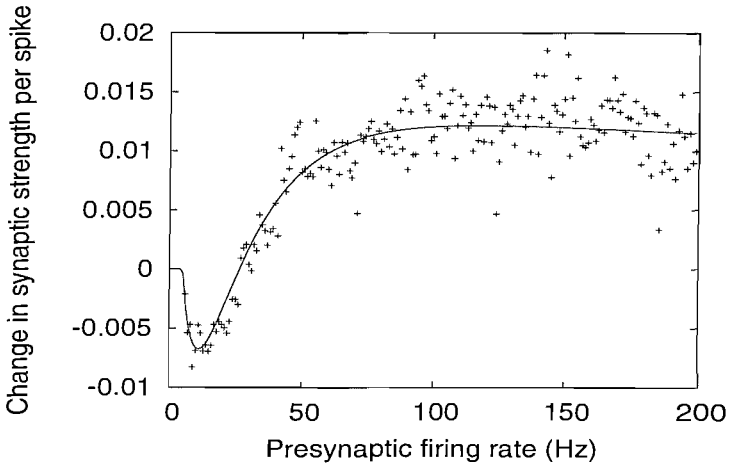

FIGURE 7.4: The change in synaptic strength per spike for a single, random spike train of 50,000 spikes for the 1-reset model plotted as a function of λ_π. The crosses indicate the change at each value of λ_π. For comparison, the solid line shows the exact, analytical result for the asymptotic limit of $\frac{1}{n}\Delta S_n$, which represents an average over all possible spike trains.

In Fig. 7.5, Fig. 7.6 and Fig. 7.7 we show contour plots for the analytical form of $\frac{1}{n}\Delta S_n$ for the 1-reset model for $n = 2$, $n = 3$ and the asymptotic limit $n \to \infty$, respectively. Contour plots for other values of n are all very similar to those for $n = 3$ and $n \to \infty$. These trends are also observed, for example, for the non-resetting model.

We see a significant difference between the $n = 2$ landscape and that for all other values of n. For the $n = 2$ rule, the landscape is symmetrical about the line $\lambda_\pi = \lambda_p$. This symmetry can be seen directly from an examination of the 2-spike learning rule derived in Chapter 6

$$\Delta S_2 = \lambda'_\pi \lambda'_p \left[A_+ K_1^+ - A_- K_1^-\right] \tag{7.127}$$

Where $\lambda'_\pi = \lambda_\pi/\beta$ and $\lambda'_p = \lambda_p/\beta$. We see that, in the 2-spike rule, the quantities λ_p and λ_π only ever appear as a sum or a product. The surface is therefore symmetric under the interchange of λ_π and λ_p, and hence is symmetric about the line $\lambda_\pi = \lambda_p$. In contrast, the $n = 3$ landscape in not symmetric. We can see this directly by examining the the 3-spike rule

$$\Delta S_3 = A_+ \left[2\lambda'_\pi \lambda'_p K_1^+ - \lambda'_\pi (\lambda'_p)^2 K_1^+ K_1^-\right] - A_- \left[2\lambda'_\pi \lambda'_p K_1^- + (\lambda'_\pi)^2 \lambda'_p K_1^+ K_1^-\right]. \tag{7.128}$$

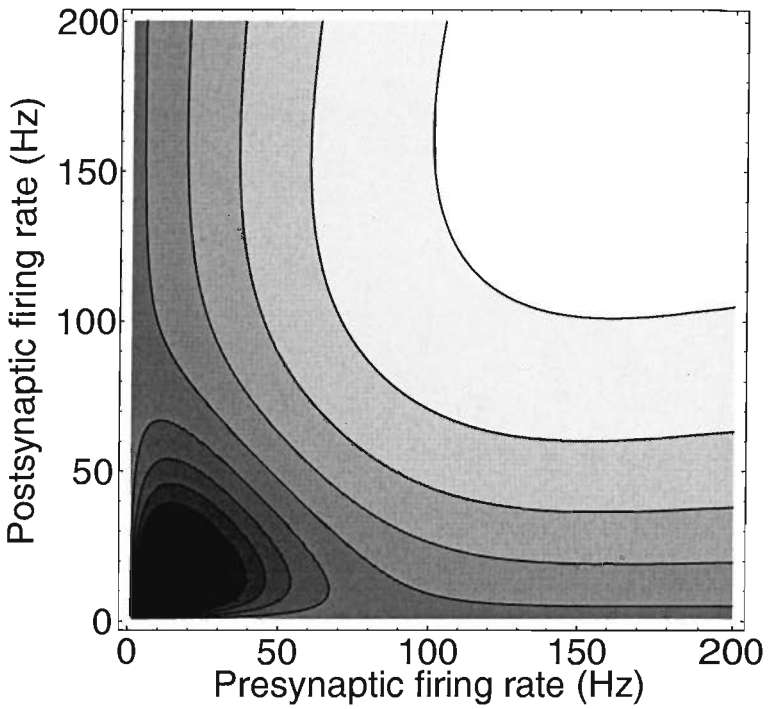

FIGURE 7.5: A contour plot of the function $\frac{1}{n}\Delta S_n$ for the 1-reset model, with $n = 2$, in the λ_π–λ_p plane. Black areas represent minimum values and white areas maximum values, and nine shades of grey interpolate between these extremes. The minimum value on this partial plane is -0.0122, while the maximum value is $+0.0059$.

Which contains terms which have unequal powers of λ_p and λ_π. Thus, the 3-spike rule is not invariant under the transformation $\lambda_p \Leftrightarrow \lambda_\pi$, and is therefore not symmetric about the line $\lambda_p = \lambda_\pi$.

We may offer an explanation of why the 2- and 3-spike rules, and therefore the learning surfaces they encode, are fundamentally different by examining directly the nature of the terms that make up the learning rules, ΔS_2 and ΔS_3. We do this explicitly, deconstructing the calculation of ΔS_n and examining the contribution and spike–probability weighting of each term. We may then identify the dominant terms and compare their contributions, explaining the resulting average learning behaviour as a sum these processes. Consider, first, the 2-spike rule. As discussed in Chapter 6, the rule is constructed from four terms, πp, $\pi\pi$, $p\pi$ and pp. Only πp and $p\pi$ contribute to plasticity. Examining these two terms more carefully we have that

$$\Delta S_{\pi p} = \lambda'_\pi \lambda'_p A_+ K_1^+, \tag{7.129}$$

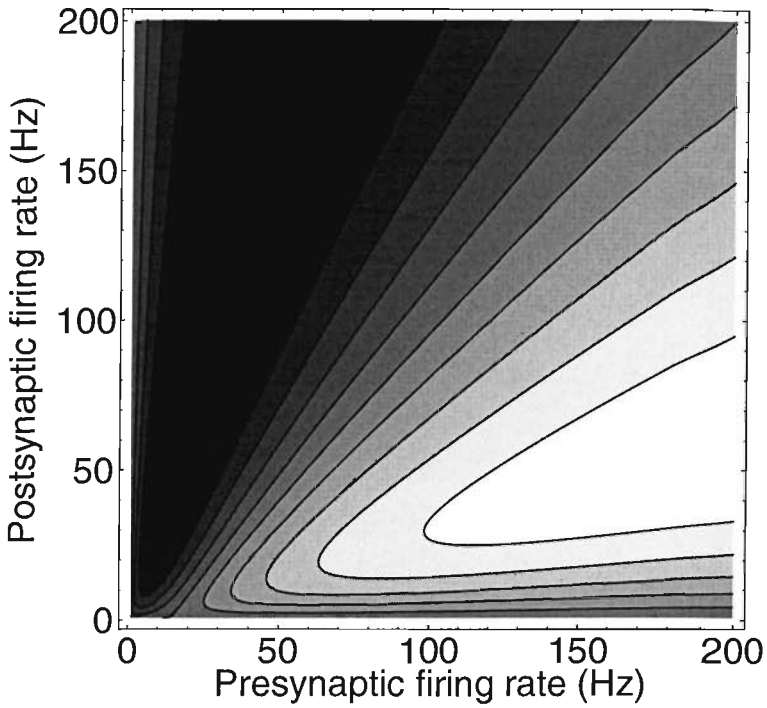

FIGURE 7.6: A contour plot of the function $\frac{1}{n}\Delta S_n$ for the 1-reset model, with $n = 3$, in the λ_π–λ_p plane. The minimum value on this partial plane is -0.0284, while the maximum value is $+0.0342$.

and

$$\Delta S_{p\pi} = \lambda'_\pi \lambda'_p A_- K_1^-. \tag{7.130}$$

We consider the learning surface away from the origin, so that at least one of λ_π or λ_p is large. Thus, as $\beta = \lambda_\pi + \lambda_p$, we have that $K_1^\pm(\beta) \to 1$. The two plasticity terms given above therefore reduce to

$$\Delta S_{\pi p} = \lambda'_\pi \lambda'_p A_+, \tag{7.131}$$

and

$$\Delta S_{p\pi} = -\lambda'_\pi \lambda'_p A_-. \tag{7.132}$$

The sum of these two processes is the overall learning rule, ΔS_2. When synaptic strengths become large, the postsynaptic firing rate, λ_p, also becomes large and $\lambda'_p \to 1$ while $\lambda'_\pi \to 0$. Alternatively, if presynaptic firing rates become large, we

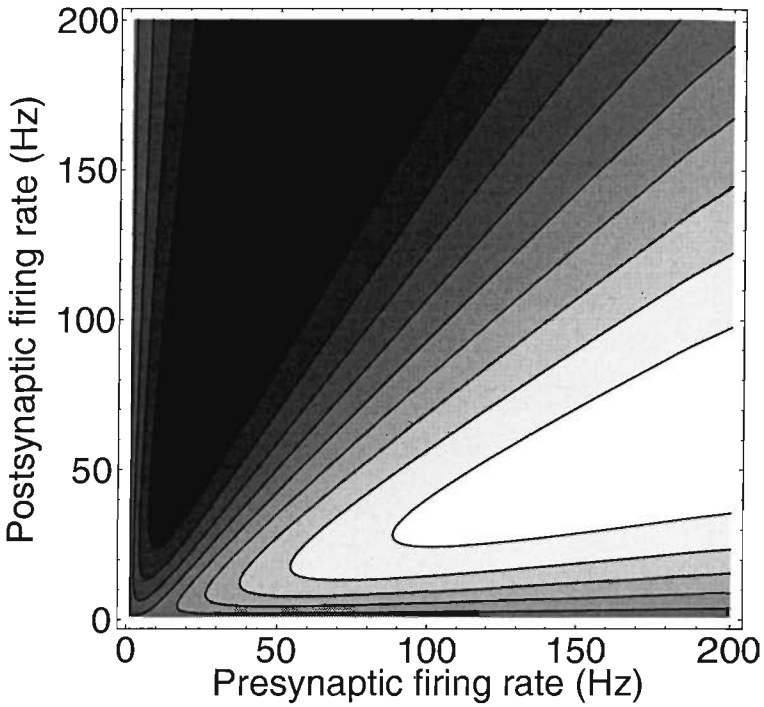

FIGURE 7.7: A contour plot of the function $\frac{1}{n}\Delta S_n$ for the 1-reset model, in the asymptotic limit $n \to \infty$, in the λ_π–λ_p plane. The minimum value on this partial plane is -0.1071, while the maximum value is $+0.1148$.

have that $\lambda'_\pi \to 1$ and $\lambda'_p \to 0$. In both cases, we see immediately that

$$\Delta S_2 \propto A_+ - A_-. \tag{7.133}$$

We previously assumed that $A_+ > A_-$, and thus, in the limit of high pre- or postsynaptic firing rates, the potentiation process πp dominates overall and the learning surface is positive.

For the resetting 3-spike rule, the contributing terms are $\pi\pi p$, $pp\pi$, $p\pi\pi$, πpp, $p\pi p$, and $\pi p\pi$. The first four are essentially spike-pair terms weighted by some spike probability. The final two terms are the "triplet" terms, and encode the multi-spike interactions that differentiate the multi-spike rules from the 2-spike rule. Examining these six terms more carefully, we find that

$$\Delta S_{\pi\pi p} = (\lambda'_\pi)^2 \lambda'_p A_+ K_1^+, \tag{7.134}$$

$$\Delta S_{pp\pi} = -\lambda'_\pi (\lambda'_p)^2 A_- K_1^-, \tag{7.135}$$

$$\Delta S_{p\pi\pi} = -(\lambda'_\pi)^2 \lambda'_p A_- K_1^-, \tag{7.136}$$

$$\Delta S_{\pi pp} = \lambda'_\pi (\lambda'_p)^2 A_+ K_1^+, \tag{7.137}$$

$$\Delta S_{p\pi p} = \lambda'_\pi (\lambda'_p)^2 (A_+ K_1^+ + A_- K_1^+ K_1^- - A_- K_1^-), \tag{7.138}$$

$$\Delta S_{\pi p\pi} = -(\lambda'_\pi)^2 \lambda'_p (-A_- K_1^- - A_+ K_1^+ K_1^- + A_+ K_1^+). \tag{7.139}$$

Collecting these terms, we have the A_+, potentiating terms

$$\Delta S_3^+ = A_+ \left[2\lambda'_\pi \lambda'_p K_1^+ - \lambda'_\pi (\lambda'_p)^2 K_1^+ K_1^-\right], \tag{7.140}$$

and the A_-, depressing terms

$$\Delta S_3^- = -A_- \left[2\lambda'_\pi \lambda'_p K_1^- + (\lambda'_\pi)^2 \lambda'_p K_1^+ K_1^-\right]. \tag{7.141}$$

The first term in both equations comes from the 3-spike interactions that reduce to terms resembling "pure" 2-spike terms. For example, for the A_+ equation $\pi\pi p$ and πpp both produce pure 2-spike terms. For the A_- equation it is $pp\pi$ and $p\pi\pi$. The second term in both equations stem from the irreducibly 3-spike interactions present in the interleaved spike trains $p\pi p$ and $\pi p\pi$. Consider, for example, the spike train $\pi p\pi$. This train gives rise to the switching process $\pi\gamma_+ p\pi$, where γ_+ represents a passive transition from POT to OFF, leading to a term of the form $A_-(1-K_1^+)K_1^-$. Thus, an irreducible 3-spike term $-A_- K_1^+ K_1^-$ is produced. An identical argument leads to the irreducible A_+ terms. Adding the A_+ and A_- terms together, we have

$$\Delta S_3 = A_+ \left[2\lambda'_\pi \lambda'_p K_1^+ - \lambda'_\pi (\lambda'_p)^2 K_1^+ K_1^-\right] - A_- \left[2\lambda'_\pi \lambda'_p K_1^- + (\lambda'_\pi)^2 \lambda'_p K_1^+ K_1^-\right]. \tag{7.142}$$

Which, on taking the same limit in β as for the 2-spike rule, reduces to

$$\Delta S_3 = A_+ \lambda'_\pi \lambda'_p \left[2 - \lambda'_p\right] - A_- \lambda'_\pi \lambda'_p \left[2 - \lambda'_\pi\right], \tag{7.143}$$

where the first and second terms in each square bracket are the pure 2-spike terms and irreducible 3-spike terms, respectively. Consider now the limit that $\lambda'_p \to 1$ and $\lambda_\pi \to 0$. We see immediately that

$$\Delta S_3 \propto A_+ - 2A_-, \tag{7.144}$$

which is negative for our standard parameter choices. In other words, for large input weights, the irreducible triplet term $-\lambda'_\pi(\lambda'_p)^2 K_1^+ K_1^-$ has the effect of reducing the magnitude of the A_+ potentiating term. In contrast, the corresponding irreducible triplet term $(\lambda'_\pi)^2\lambda'_p K_1^+ K_1^-$ has vanished and does not affect the A_- depressing term. We therefore experience depression overall, and the learning surface is negative.

In the limit where λ'_π is large, we have that $\lambda'_p \to 0$ and $\lambda'_\pi \to 1$, giving

$$\Delta S_3 \propto 2A_+ - A_-, \tag{7.145}$$

which is positive as $A_+ - A_- > 0$. Thus, for high presynaptic firing rates we experience potentiation overall, and the learning surface is positive.

Thus, we see that this process of decomposing the 3-spike rule into its constituents allows us to identify the exact processes which lead to the asymmetry of the 3-spike learning surface. These processes are the irreducible 3-spike terms arising from $\pi p\pi$ and $p\pi p$. This result contrasts with other models of STDP, where the spike trian $\pi p\pi$ would simply be a linear sum of two spike pairs which would approximately cancel (Song et al., 2000).

We find a similar asymmetric learning surface in all $n \geq 3$ rules. Such rules, therefore, will not necessarily suffer from runaway learning, but may display more interesting dynamics. These simple observations show that, although there tends to be an emphasis in the literature on the form of the 2-spike rule, multi-spike rules, at least in our model, are qualitatively different in character. We shall explore these differences in Chapter 8, showing that they have important consequences for the presence and nature of competitive dynamics in the switch model.

7.9 Discussion

In this Chapter we have extended our earlier, ensemble-based, stochastic model of STDP in order to derive exact, analytical results for general n-spike interactions for $n > 2$. By extending the model to include stochastic process resetting, we essentially constructed a mathematical ladder that we could ascend in order

to arrive at exact results for the non-resetting model. Once there, we can extract the general r-reset model from the non-resetting model. As suggested earlier, the most biologically plausible forms of these various resetting models are, arguably, the 1-reset model and the non-resetting model. Other finite r, $r \neq 1$, forms require us to postulate some kind of spike-counting machinery at individual synapses that, in the spirit of reducing the computational burden placed on any single synapse as much as possible, we consider to be less likely.

Although the general results for $\frac{1}{n}\Delta S_n$ for all but the 1-reset model are somewhat cumbersome (we regard the roots of the quartic equation associated with the 2-reset model as cumbersome, even though the n-spike interaction function is available explicitly in terms of these roots), in fact we find rapid convergence of $\frac{1}{n}\Delta S_n$ to its asymptotic, large n form, so that even by approximately 10 spikes, $\frac{1}{n}\Delta S_n$ is within just a few percent of the limiting result. Furthermore, in terms of the general form of the function ΔS_n in the λ_π–λ_p plane, we find no qualitative changes in the structure of the landscape for any $n \geq 3$. These two facts lead us to suppose that any properties exhibited by the asymptotic form of the n-spike function will also be exhibited by all the n-spike functions for $n \geq 3$. Of course, we expect quantitative differences, but the overall dynamics, as determined by the overall structure of the function, should be similar.

The n-spike interaction functions we have derived are unconditional expectation values that represent, in essence, rate-based learning rules in which the underlying spiking processes have been integrated out and summed over. We are thus led to ask whether a purely spike-based system experiencing a specific, typical pattern of pre- and postsynaptic spikes would exhibit a change in synaptic strength consistent with the rate-based, n-spike interaction function? In other words, what are the conditions that lead the spike-based behaviour to match the rate-based results derived above? The answer to this question, as shown in figure 7.4, is that n must be sufficiently large that enough self-averaging within a long spike sequence can take place, reducing the overall standard deviation in the data to well below the mean, so that the mean behaviour is not swamped by noise. We saw that "sufficiently large", for our parameter choice in the 1-reset model, means about 50,000 spikes. At a rate of around 10 Hz, this corresponds to 5000 seconds, or a little over an hour's worth of synaptic activity. At a rate of around 100 Hz, this is reduced to 500 seconds, or a few minutes' worth. These results are consistent with our general expectation that, if we consider synaptic plasticity over the time course of at least a few minutes, even though the neuron computes at the level of individual spikes, it behaves as though it were following an abstracted, rate-based rule.

In Chapter 6 we considered purely numerical simulations of n-spike interactions

for $n > 2$. There, we looked specifically at the cross section of the λ_π–λ_p landscape corresponding to Eq. (7.126) and found little quantitative difference between the 2-spike and the general n-spike results, all forms exhibiting a BCM-like profile with depression at low presynaptic rates switching over to potentiation for higher presynaptic rates. This observation reflects the highly specific cut through the λ_π–λ_p plane. Other cuts reveal profound differences between the 2-spike and the general n-spike results, as seen in figures 7.5, 7.6 and 7.7. The particular model of postsynaptic firing in Eq. (7.126) might be expected to be fairly reasonable (modulo issues of saturation, etc.) for a target cell innervated by precisely one afferent, so that the postsynaptic response follows closely the presynaptic input. However, for a target cell innervated by multiple afferents, it is entirely possible for there to be large differences between some of the presynaptic rates and the postsynaptic rate, because the postsynaptic rate will reflect all the inputs. We are then led to examine the whole λ_π–λ_p plane, which reveals these differences. Indeed, it is precisely these differences between the 2-spike and the general n-spike interaction functions, observed initially numerically, that motivated our derivation of the above analytical results for the n-spike function. We shall explore the learning dynamics of the different, n-spike rules we have derived in Chapter 8.

7.10 Conclusion

In conclusion, we have generalised the synaptic switch rule of STDP presented in Chapter 6. We now distinguish two forms of the model, based on the response of the process governing the stochastic transitions of the switch to multiple spikes of the same character. In the resetting model, multiple spikes can reset the stochastic process, while in the non-resetting model they cannot. We made this generalisation in order to extend our earlier analysis, in which spike interactions were limited to at most two spikes, to the case where spike interactions were not limited in any way. We thus derive exact, analytical expressions for the n-spike interaction function arising under our synaptic switch rule.

Even in retrospect, it is astonishing that many of these results can be derived exactly. This is partly due to the use of a stochastic model with intrinsic Markovian properties, and partly because the model is actually quite simple in its overall structure.

We note that the 2-spike rule is of qualitatively different character compared to all other multi-spike rules. This result is especially interesting given the essentially 2-spike nature of other models of STDP. In Chapter 8, we shall focus

on the profound differences between the 2-spike interaction function and higher-order interaction functions, and the consequences for the presence of stable, competitive dynamics in our model.

7.11 Appendix: Roots of Characteristic Equation

We return to a consideration of the roots θ_i, $i = 1, \ldots, 2r$, of Eq. (7.82),

$$\theta^{2r} + \sum_{j=1}^{r-1} \Lambda_j \theta^{2r-j} - \sum_{j=1}^{r-1} (\tilde{K}_j^+ \tilde{K}_r^- + \tilde{K}_j^- \tilde{K}_r^+) \theta^{r-j} - \tilde{K}_r^+ \tilde{K}_r^- = 0,$$

for specific cases in which the roots can be computed exactly. We know that for $r = 1$ and $r = 2$ the roots are available in closed form, but we seek cases for which exact expressions are available for the general r case.

A little re-arrangement of Eq. (7.82) gives

$$\left(\theta^{2r} - \tilde{K}_r^+ \tilde{K}_r^-\right) + \sum_{j=1}^{r-1} \theta^{r-j} \left[\tilde{K}_j^+ \left(\theta^r - \tilde{K}_r^-\right) + \tilde{K}_j^- \left(\theta^r - \tilde{K}_r^+\right)\right] = 0. \quad (7.146)$$

We see that if $\tilde{K}_r^+ = \tilde{K}_r^-$, then an overall factor of $\theta^r - \tilde{K}_r^+$ can be pulled out, yielding the r roots $\omega(\tilde{K}_r^+)^{1/r}$, where ω is an rth root of unity, $\omega^r = 1$. In general, however, we still have the roots of the remaining rth degree factor to extract. Moreover, $\tilde{K}_r^+ = \tilde{K}_r^-$ is very rarely satisfied, except for very particular choices of the rates λ_π and λ_p and the parameters $\tau_\pm$ and $n_\pm$.

A much less constraining case in which all the roots can be computed exactly is that for which $\tilde{K}_i^\pm = (\tilde{K}_1^\pm)^i$. This corresponds, for example, to the 1-reset model or to a general r-reset model in which $n_\pm = 1$. Defining $R^2 = \tilde{K}_1^+ \tilde{K}_1^-$, we then have

$$\left(\theta^{2r} - R^{2r}\right) + \sum_{j=1}^{r-1} \left((\tilde{K}_1^+)^j + (\tilde{K}_1^-)^j\right) \theta^j \left(\theta^{2(r-j)} - R^{2(r-j)}\right) = 0. \quad (7.147)$$

We see that a common overall factor of $\theta^2 - R^2$ can be pulled out. The solutions $\theta = \pm R$, of course, are already known for the $r = 1$ model, so we expect to reproduce these. We define the remaining $2(r-1)$ degree polynomial whose roots have yet to be determined to be $R^{2(r-1)} f(\theta)$, so that

$$R^{2(r-1)} f(\theta) = \sum_{i=0}^{r-1} \theta^{2i} R^{2(r-1-i)} + \sum_{j=1}^{r-1} \left((\tilde{K}_1^+)^j + (\tilde{K}_1^-)^j\right) \theta^j \sum_{i=0}^{r-j-1} \theta^{2i} R^{2(r-1-i-j)}. \quad (7.148)$$

Defining $\theta_\pm = \theta/\tilde{K}_1^\pm$, dividing through by $R^{2(r-1)}$, writing $\theta^2/R^2 = \theta_+\theta_-$ and $(\tilde{K}_1^\pm)^j\theta^j/R^{2j} = \theta_\mp^j$, we then have

$$f(\theta) = \sum_{i=0}^{r-1}\theta_+^i\theta_-^i + \sum_{j=1}^{r-1}\left(\theta_+^j + \theta_-^j\right)\sum_{i=0}^{r-j-1}\theta_+^i\theta_-^i. \tag{7.149}$$

Re-arranging the double sums, we have

$$\begin{aligned}\sum_{j=1}^{r-1}\sum_{i=0}^{r-j-1}\theta_+^{i+j}\theta_-^i &= \sum_{i=1}^{r-1}\theta_+^i\sum_{j=0}^{i-1}\theta_-^j \\ &= \sum_{0\le j<i\le r-1}\theta_+^i\theta_-^j.\end{aligned} \tag{7.150}$$

Similarly,

$$\sum_{j=1}^{r-1}\sum_{i=0}^{r-j-1}\theta_+^i\theta_-^{i+j} = \sum_{0\le i<j\le r-1}\theta_+^i\theta_-^j. \tag{7.151}$$

By observing that

$$\sum_{i=0}^{r-1}\theta_+^i\theta_-^i \equiv \sum_{0\le i=j\le r-1}\theta_+^i\theta_-^j, \tag{7.152}$$

we see that $f(\theta)$ factorises into separate θ_+ and θ_- pieces, so that

$$f(\theta) = \left(\sum_{i=0}^{r-1}\theta_+^i\right)\left(\sum_{j=0}^{r-1}\theta_-^j\right), \tag{7.153}$$

or

$$f(\theta) = \frac{1-\theta_+^r}{1-\theta_+}\frac{1-\theta_-^r}{1-\theta_-}. \tag{7.154}$$

Hence, $f(\theta)$ has the $2(r-1)$ roots $\theta = \tilde{K}_1^\pm\omega$, $\omega^r = 1$, $\omega \neq 1$. Since $\tilde{K}_1^\pm < 1$ for finite β, all $2r$ roots of the characteristic equation therefore have moduli strictly less than unity. This factorisation of the K^+ and K^- components does not appear to carry over to the completely general form of Eq. (7.82).

The cases $n_\pm = 1$ and $r = 1$ are not the only cases for which $\tilde{K}_i^\pm = (\tilde{K}_1^\pm)^i$. In the limit $n_\pm \to \infty$, we have that $K_i^\pm \to 1$. Hence, for example, $\tilde{K}_i^+ \to {\lambda'_\pi}^i$, so $\tilde{K}_i^+ = (\tilde{K}_1^+)^i$ in this limit too. For $\lambda'_\pi = 1$, we then have $\tilde{K}_1^+ = 1$, so, in this limit, the roots $\tilde{K}_1^+\omega$ approach the unit circle. However, they are not repeated because when $\lambda'_\pi = 1$ we have $\lambda'_p = 0$, so the other roots collapse around zero. Another case is that in which $\beta \to \infty$, for which, again, $K_i^\pm \to 1$. This case is valid for any $n_\pm$, and so, in the limit of large β, $2(r-1)$ of the roots of Eq. (7.82) will approach multiples of the rth roots of unity. Again, in the directions $\lambda'_\pi = 1$ and $\lambda'_p = 1$, these roots reach the unit circle, but are never repeated. All the cases for which we are able to solve Eq. (7.82) exactly (except for the general $r = 2$ solutions) appear to fall into subsets of all the above cases.

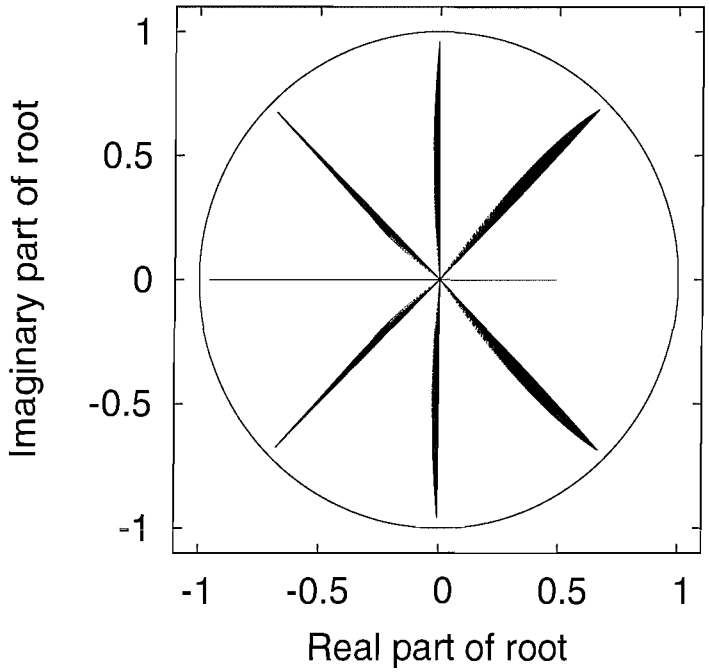

FIGURE 7.8: The distribution of roots of the characteristic equation of the 8-reset model with parameters $n_\pm = 3$, $\tau_+ = 13.3$ ms, $\tau_- = 20.0$ ms, plotted in the complex plane. The bounding unit circle $|z| = 1$ is also shown.

We know that $\tilde{K}_i \leq \tilde{K}_1^i$, and so this suggests that we should write $\tilde{K}_i = \tilde{K}_1^i - \epsilon_i$, $\epsilon_i \geq 0$ and then expand about the roots $\tilde{K}_1^\pm \omega$ obtained above. The validity of such solutions depends on the various ϵ_i's being small. For i large enough, this is always true (except in the limits considered above). For small i, however, ϵ_i's being small will depend on parameter choices. Nevertheless, if we expect the roots of the characteristic equation to be (at least piecewise) continuous functions of the parameters, then these arguments suggest that in the general case, the roots should roughly speaking cluster around the roots of unity, with some possibly large perturbations. In figure 7.8, we show the distribution of roots for parameter choices giving rise to the earlier figures, $n_\pm = 3$, $\tau_- = 20.0$ ms, $\tau_+ = 13.3$ ms, with $\lambda_\pi, \lambda_p \in (0, 200]$ Hz. We set $r = 8$ as an example value. We see that, indeed, the roots cluster around the eighth roots of unity. For all choices of (finite) $n_\pm$ and $\tau_\pm$ examined, all plots are qualitatively similar to the $n_\pm = 1$ case, but with perturbations around the known roots of this case.

Chapter 8

Stable, Competitive Dynamics in the Switch Rule

The work in this Chapter is based primarily on the paper "Stable Competitive Dynamics Emerge from Multi-Spike Interactions in a Stochastic Model of Spike Timing Dependent Plasticity " Appleby and Elliott (2005) (submitted).

8.1 Summary of learning rules

We make use of several of the n–spike learning rules derived in Chapter 7. For convenience, we reproduce these rules here. Defining

$$K_l^{\pm}(\beta) = \left(\frac{\tau_{\pm}\beta}{1+\tau_{\pm}\beta}\right)^l \sum_{i=0}^{n_{\pm}-1} \binom{i+l-1}{l-1} \frac{1}{(1+\tau_{\pm}\beta)^i}, \tag{8.1}$$

and

$$Q_r^{\pm} = \sum_{j=1}^{r} \tilde{K}_j^{\pm}, \tag{8.2}$$

we have the 2–spike result

$$\Delta S_2 = A_+ \tilde{K}_1^+ - A_- \tilde{K}_1^-, \tag{8.3}$$

and the asymptotic, r–reset, n–spike result

$$\frac{1}{n}\Delta S_n^r \sim \frac{\lambda_p' A_+ (1-\tilde{K}_r^-)Q_r^+ - \lambda_\pi' A_- (1-\tilde{K}_r^+)Q_r^-}{(1-\tilde{K}_r^+)(1-\tilde{K}_r^-) + (1-\tilde{K}_r^-)Q_r^+ + (1-\tilde{K}_r^+)Q_r^-}. \tag{8.4}$$

In particular, we are interested in the resetting ($r = 1$) and non-resetting ($r \to \infty$) forms of this asymptotic, n–spike result

$$\frac{1}{n}\Delta S_n^1 \sim \frac{\lambda_p' A_+(1-\tilde{K}_1^-)Q_1^+ - \lambda_\pi' A_-(1-\tilde{K}_1^+)Q_1^-}{1-\tilde{K}_1^+\tilde{K}_1^-}. \tag{8.5}$$

$$\frac{1}{n}\Delta S_n^\infty \sim \frac{\lambda_\pi' A_+ K_1^+(\lambda_p) - \lambda_p' A_- K_1^-(\lambda_\pi)}{1+\frac{\lambda_p}{\lambda_\pi}K_1^-(\lambda_\pi)+\frac{\lambda_\pi}{\lambda_p}K_1^+(\lambda_p)}. \tag{8.6}$$

8.2 The 2-Spike Rule and Beyond

We first examine the 2-spike rule and show that it leads, without further modification, to pathological learning dynamics. We then move beyond the 2-spike rule and perform an initial analysis of the multi-spike rules, indicating the reasons why their dynamics differ from the 2-spike rule's dynamics.

8.2.1 Failure of the 2-Spike Rule

To study 2-spike interactions, we consider a system of m afferents innervating a single target cell, and explore their behaviours in the context of the averaged, 2-spike learning rule, ΔS_2. We label the m afferents with indices such as i and j, so that $i, j \in \{1, \ldots, m\}$. Let afferent i support l_i synapses of strength $s_{i\alpha}$, $\alpha \in \{1, \ldots, l_i\}$. Let afferent i fire at rate λ_{π_i}, and the target cell fire at rate λ_p. We then define

$$\frac{ds_{i\alpha}}{dt} = \Delta S_2(\lambda_{\pi_i}, \lambda_p), \tag{8.7}$$

where we have noted the explicit dependence of ΔS_2 on the pre- and postsynaptic firing rates. Since ΔS_2 is independent of the synapse label α, all $s_{i\alpha}$ experience the same change at each time step: all of afferent i's synapses experience identical pre- and postsynaptic firing rates. Hence, we may consider the evolution of either the total synaptic strength supported by afferent i, $s_i^T \equiv \sum_\alpha s_{i\alpha}$, or the average synaptic strength, $s_i^A \equiv \frac{1}{l_i}\sum_\alpha s_{i\alpha}$, evolving according to $ds_i^T/dt = l_i\,\Delta S_2(\lambda_{\pi_i}, \lambda_p)$ or $ds_i^A/dt = \Delta S_2(\lambda_{\pi_i}, \lambda_p)$, respectively. We set the postsynaptic firing rate λ_p to be the standard linear sum of the presynaptic firing rates weighted by the synaptic strengths. Thus, if we consider the total strengths s_i^T, we set $\lambda_p = \sum_i s_i^T \lambda_{\pi_i}$, while if we consider the average strengths s_i^A, we set $\lambda_p = \sum_i l_i s_i^A \lambda_{\pi_i}$. Since we study competitive dynamics, we do not consider scenarios in which some afferents at least initially enjoy an advantage over other afferents, and in particular, we do not consider scenarios in which the numbers of synapses supported by a group of afferents differ significantly. Hence, for convenience we

take l_i to be independent of i: all afferents support the same number of synapses. We may then dispense with the factors of l_i since they can be absorbed into a re-definition of time (for the total strength) or a re-definition of afferent rates, which is equivalent to a re-definition of time (for the average strength). We therefore consider one common quantity s_i, which can be thought of as either total or average synaptic strength, evolving according to the equation

$$\frac{ds_i}{dt} = \Delta S_2(\lambda_{\pi_i}, \lambda_p), \tag{8.8}$$

for which we set $\lambda_p = \sum_i s_i \lambda_{\pi_i}$. As we consider only excitatory synapses, we insist that $s_i \geq 0$. Hence, if Eq. (8.8) tries to drive s_i negative, we truncate the evolution and set s_i to zero. The synapse is not frozen there: it can re-grow to non-zero strength.

We may obtain a qualitative understanding of the dynamics induced by Eq. (8.8). Consider a scenario in which at least one afferent, say afferent i, has a large synaptic strength, s_i. The postsynaptic firing rate λ_p will thus typically be large, and so the variables $\beta_i = \lambda_{\pi_i} + \lambda_p$ will typically be large. In this limit, $K_1^{\pm}(\beta_i) \approx 1$, so the 2-spike learning rule reduces to $ds_j/dt \propto A_+ - A_- > 0$, for any afferent j. Thus, we see that once one afferent becomes strong, it induces all afferents to become strong, and this induces a positive feedback mechanism, in which all afferents reinforce each other. So, unless the first afferent that becomes strong is silenced with low λ_{π_i} for a sustained period, so that λ_p does not become large, all afferents' strengths will escape to infinity. Consider now a scenario in which all synaptic strengths s_i are small. Here we may write $\beta_i \approx \lambda_{\pi_i}$ since λ_p is small, and the 2-spike rule becomes $ds_j/dt \propto A_+ K_1^+(\lambda_{\pi_j}) - A_- K_1^-(\lambda_{\pi_j})$. Writing this out according to the definitions of $K_1^{\pm}$, we have two terms, a negative term that goes like $\gamma - 1$ and a positive term that grows with λ_{π_j}. However, only for presynaptic firing rates in excess of at least 100 Hz does ds_j/dt become overall positive. Thus, if the s_i are all small, then $ds_i/dt < 0$ for all but high, sustained firing rates, and so the s_i become even smaller. These weak afferents become trapped in the depressing phase and fall to zero strength. These simple limiting scenario arguments therefore suggest that the phase space for the 2-spike rule is partitioned into two regimes, with all afferents either pulled towards zero on average or pushed towards infinity on average. Although these arguments are not rigorous proofs, their conclusion is confirmed by a full fixed point analysis of the 2-spike rule. Indeed, we see that the fixed-point structure of the 2-spike rule is, in fact, identical to that of the early form of the BCM-rule which did not include a sliding threshold (Cooper et al., 1979). In both cases, the system in governed by an unstable fixed-plane, and any perturbations about this plane grow without bound. We may therefore equate the unmodified, 2-spike rule directly to the fixed-threshold form of the BCM-rule as these two models are, at least at this

averaged, rate-based level, identical. This allows us to directly deploy any of the analytical results derived elsewhere for the BCM-rule (Castellani et al., 1999). Moreover, the reciprocal relationship will also be true. That is, given the above analysis we may state immediately that, if postsynaptic spiking is assumed to occur according to a Poisson process with parameter given by the linear sum of input rates weighted by their connection strengths, the fixed-threshold BCM rule will also produce runaway learning in the same way as the unmodified 2-spike rule.

We therefore see that despite having a BCM-like form, the learning behaviour of the rate-based, 2-spike rule is, in fact, pathological and always leads to the afferents either all dying or all escaping to infinity. To demonstrate that the spike-based system also exhibits the behaviour characteristic of the rate-based system, in Fig. 8.1 we show a spike-based simulation of four afferents innervating a single target cell. Spike trains are truncated at two spikes: there is no interaction between successive pairs of spikes. We see that the spike-based simulation exhibits the same two regimes discussed above for the rate-based system. This behaviour is, as argued above, independent of the number of afferents because the governing factor in this behaviour is the postsynaptic firing rate, λ_p, which is common to all afferents synapsing on the target cell. Once λ_p begins to move towards or away from zero, uncontrolled learning ensues and the afferents either all die or all escape to infinity. The instability inherent in the 2-spike learning rule, whether rate- or spike-based, thus shows that it is unable to support the stable, competitive dynamics characteristic of, for example, ODC formation.

This situation appears rather unpromising. For both experimental and theoretical reasons, we require that a BCM-like form emerges from the 2-spike learning rule on average, yet this requirement leads directly to these pathological learning behaviours. Without further modification designed specifically to prevent runaway learning, such as hard upper bounds on synaptic strengths, the rule will always lead to uncontrolled learning. One possible remedy is to allow the threshold between the potentiating and depressing regimes, which is a function of various easily modifiable parameters, to depend on the recent time average of postsynaptic firing in a manner similar to the BCM rule Bienenstock et al. (1982). In effect, this "couples" potentiation and depression together, in the sense that the dependence of the plasticity threshold on the recent time average of postsynaptic firing allows the history of potentiation and depression to influence later plasticity events. The result of this coupling in the BCM rule is to stabilise the learning behaviour and prevent runaway learning.

We see a similar result when we modify the 2-spike rule to incorporate a sliding threshold. In the BCM model, the sliding threshold Θ_M is explicitly set as a function of the recent time-averaged postsynaptic firing rate, $\bar{\lambda}_p$. In our 2-spike

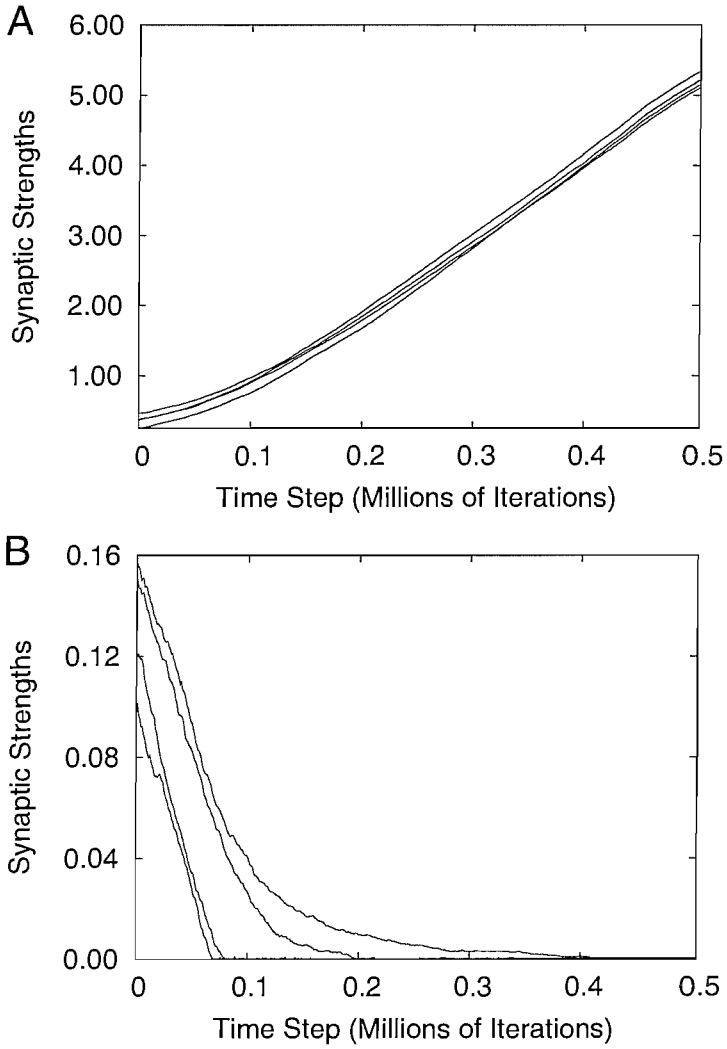

FIGURE 8.1: A spike-based simulation of the 2-spike rule for four afferents innervating one target cell. The dynamics are partitioned into two distinct regimes, determined by the postsynaptic firing rate. **(A)** When initial synaptic strengths are large, the afferents drive the postsynaptic cell to a high firing rate and runaway learning ensues. **(B)** When initial synaptic strengths are small, the postsynaptic firing rate is low and all four afferents fall to zero strength. In both cases, we simulate four afferents with 10 synapses each, innervating a single target cell then evolve the system according to the switch rule. Pre- and postsynaptic spiking occurred according to Poisson statistics, with the postsynaptic firing rate given by the simple linear sum of input rates weighted by connection strengths. One iteration translates to one second of simulated time.

rule, the threshold is dynamically determined by the solution of

$$A_+K_1^+(\beta) = A_-K_1^-(\beta). \tag{8.9}$$

With the six parameters $A_\pm$, $n_\pm$ and $\tau_\pm$ held fixed, this gives a value for β. However, for a given value of β (and fixed parameters $A_\pm$, $n_\pm$ and τ_-), we can instead regard this equation as determining a value of τ_+. We denote this value of τ_+ by $\tau_M = f(\beta)$, where f is the function that gives the solution of Eq. (8.9). Since it is unrealistic for the value of τ_+ to depend on the instantaneous value of β, we instead make it depend on the recent time average similarly to the BCM rule, so that $\tau_M = f(\bar{\beta})$. Thus, for a given value of $\bar{\beta}$ our preferred value of τ_+ should be set so that $\tau_+ = \tau_M$. Such a value would place the threshold at exactly the right location, putting some afferents in the depressing region and the others in the potentiating region. Thus, if we make τ_+ dynamically evolve towards $\tau_M(\bar{\beta})$, we will achieve a sliding threshold sensitive to the recent time average of postsynaptic firing, $\bar{\lambda}_p$, through $\bar{\beta} = \bar{\lambda}_\pi + \bar{\lambda}_p$. Setting

$$\frac{d\tau_+}{dt} = \epsilon\left[\tau_M(\bar{\beta}) - \tau_+\right] \tag{8.10}$$

represents one simple way in which to realise such sliding, where ϵ is some small inverse time constant for the relaxation of τ_+ to the (changing) value of $\tau_M(\bar{\beta})$. Implementing this sliding threshold in the 2-spike rule, we replicate the dynamics of the BCM rule. Synaptic strengths are therefore stabilised, and uncontrolled learning prevented. The behaviour of a set of four afferents operating under this modified, rate-based 2-spike rule is shown in Fig. 8.2. We have chosen initial conditions so that all synapses would grow without bound under the unmodified 2-spike rule. For the modified rule, we see that γ therefore decreases very rapidly (within the first few thousand time steps), moving all synapses into the depressing regime. As they depress, γ slowly increases. Eventually three of the synapses hit zero strength. The surviving synapse initially continues in the downward direction, but γ increases rapidly, moving the synapse into the potentiating regime. The synapse grows and stabilises, and γ remains approximately constant. We therefore see, as expected, that the introduction of a sliding threshold stabilises the learning behaviour.

8.2.2 Beyond Two Spikes

Although coupling potentiation and depression in the manner described above for the modified 2-spike rule stabilises the learning dynamics exhibited by it, doing so forces us to make assumptions concerning the exact dependence of our model's parameters on the recent firing history. To avoid such somewhat

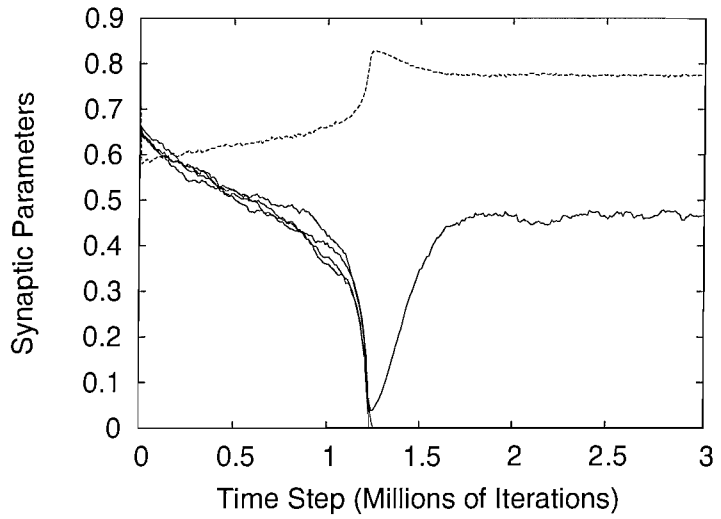

FIGURE 8.2: The evolution of four afferents innervating a single target cell under the modified 2-spike, rate-based rule, which allows the threshold between depression and potentiation to slide as a function of the time-averaged postsynaptic firing rate. The four solid lines show the strengths of the four afferents, while the dashed line shows the value of γ. The introduction of the sliding threshold has stabilised the dynamics and generated competition. We simulate four afferents with 10 synapses each, innervating a single target cell then evolve the system according to the 2-spike, rate-based learning rule. Pre- and postsynaptic spiking occurred according to Poisson statistics, with the postsynaptic firing rate given by the simple linear sum of input rates weighted by connection strengths. One iteration translates to one second of simulated time.

ad hoc complications, we seek to determine whether higher-order spike interactions can instead exhibit the required stable, competitive dynamics. One reason why higher-order spike interactions might achieve this is that they automatically provide the coupling between potentiation and depression that leads to stable behaviour under the modified 2-spike rule. This coupling takes place in the sense that once the synapse enters the POT (DEP) state, it cannot subsequently enter the DEP (POT) state without having first returned to the OFF state. That is, once a synapse has entered the potentiating mode, that synapse is prevented from entering the depressing mode without having first been deactivated to the resting state. For this coupling to be expressed, however, we require a minimum of three spikes. Consider, for example, the 2-spike train πp. Once the first presynaptic spike has elevated the synapse into the POT state, the paucity of spikes prevents the synapse from subsequently visiting the opposite half of the switch and undergoing depression. There is therefore never any coupling under 2-spike trains. This is not the case for higher-order spike trains. Consider, for example, the spike train $\pi p \pi$. As before, the first presynaptic spike elevates the synapse into the POT state. If a transition back to OFF does not occur before the p spike arrives, then a potentiation event occurs as usual, and the last π event is of

no importance. If, however, a transition back to OFF occurs before the second p spike arrives, then that p spike will cause the synapse to move to the DEP state, enabling the synapse possibly to undergo a depression event (assuming the final π spike arrives in good time). A similar argument applies for any number of spikes greater than two. Hence, if the switch is unified, all multi-spike rules couple potentiation and depression together, but the 2-spike rule does not. Given that the coupling of potentiation and depression endows the BCM rule with a structure that supports stable, competitive dynamics, it is natural to ask whether the presence of this coupling in the multi-spike rules alters the learning dynamics compared to the 2-spike rule.

In studying the dynamics of the multi-spike interactions we again consider a system of m afferents innervating a single target cell, and explore their behaviour in the context of the averaged, n-spike learning rules. The same arguments that led to Eq. (8.8) now lead to

$$\frac{ds_i}{dt} = \Delta S_n(\lambda_{\pi_i}, \lambda_p) \tag{8.11}$$

as the synaptic strength modification equation corresponding to the n-spike rule. We continue to truncate s_i at zero if it is driven negative.

A numerical exploration of the 3-spike, rate-based rule (either resetting or non-resetting) shows that, indeed, its learning dynamics differ significantly from those of the 2-spike rule. Under the 3-spike rule, the uncontrolled learning behaviour seen for the (unmodified) 2-spike rule is absent for a broad range of parameters. Afferents compete for control of the target cell, with stable segregation robustly occurring. Fig. 8.3 confirms that these observations are also true for the spike-based version of the 3-spike rule. Stable, segregated fixed points exist under the 3-spike rule, leading to the same competitive dynamics exhibited by the BCM rule. Thus, the ability of the 3-spike rule to probe the coupling of potentiation and depression under a unified switch, however modest, dramatically alters the dynamical landscape compared to that of the 2-spike rule, which cannot probe coupling. We observe a similar result for any multi-spike rule. Attempts to modify the 2-spike rule, such as introducing a sliding threshold, while adequate in terms of stabilising the learning dynamics, are therefore unnecessary: all that is required is to extend our consideration of spike interactions to three or more spikes with no *ad hoc* modifications of the learning rules.

The success of the multi-spike rules depends critically on the unification of the switch mechanism and thus the coupling of potentiation and depression. We can see this explicitly by examining the multi-spike learning rules for two ununified switches, so that we consider the potentiating and depressing lobes of the switch separately. We can reduce the unified switch to two separate potentiating and

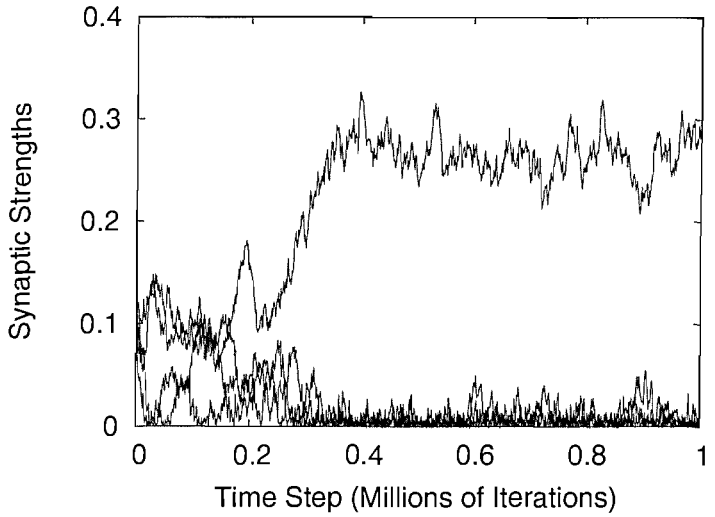

FIGURE 8.3: A spike-based simulation of the non-resetting 3-spike rule for four afferents innervating one target cell, with an explicit simulation of postsynaptic spiking via an integrate and fire neuron. As for the Poisson-based simulation, the rule is competitive, with stable segregation of the afferents being observed. One iteration translates to one second of simulated time.

depressing switches by setting $\tau_- = 0$ and $\tau_+ = 0$, respectively. Setting $\tau_- = 0$, for example, means that the depressing lobe of the unified switch is effectively unavailable, because any transition to the DEP state results in an instantaneous stochastic decay back to OFF with no depression. Similarly with $\tau_+ = 0$, the potentiating lobe is unavailable. Adding the two rules derived by separately setting $\tau_+ = 0$ and $\tau_- = 0$ gives the plasticity rule corresponding to the operation of two ununified, separate switches. Consider, for example, the resetting ∞-spike rule. Setting $\tau_- = 0$ gives

$$\Delta\hat{S}_\infty^{POT} = \lambda_p' A_+ \tilde{K}_1^+(\beta), \tag{8.12}$$

and setting $\tau_+ = 0$ gives

$$\Delta\hat{S}_\infty^{DEP} = -\lambda_\pi' A_- \tilde{K}_1^-(\beta), \tag{8.13}$$

so that the overall rule is

$$\Delta\hat{S}_\infty^{POT} + \Delta\hat{S}_\infty^{DEP} = \lambda_p' A_+ \tilde{K}_1^+(\beta) - \lambda_\pi' A_- \tilde{K}_1^-(\beta) \equiv \Delta S_2. \tag{8.14}$$

The resetting, ∞-spike rule therefore reduces to the 2-spike rule when the switch is split into two halves. Repeating this manipulation for any resetting, multi-spike rule produces the same result, so that $\Delta S_n^{POT} + \Delta S_n^{DEP} = \Delta S_2 \ \forall n$. Thus, the unification of the separate switches, which was proposed initially on the

grounds of simplicity, has unexpected consequences for the dynamics of the model. Examining the non-resetting, ∞-spike rule under an identical manipulation, we find that the potentiating half (when $\tau_- = 0$) gives

$$\Delta\hat{S}_\infty^{POT} = \frac{\lambda'_\pi A_+ K_1^+(\lambda_p)}{1 + \frac{\lambda_\pi}{\lambda_p} K_1^+(\lambda_p)} \tag{8.15}$$

and the depressing half (when $\tau_+ = 0$) gives

$$\Delta\hat{S}_\infty^{DEP} = -\frac{\lambda'_p A_- K_1^-(\lambda_\pi)}{1 + \frac{\lambda_p}{\lambda_\pi} K_1^-(\lambda_\pi)} \tag{8.16}$$

so that the overall rule is

$$\Delta\hat{S}_\infty^{POT} + \Delta\hat{S}_\infty^{DEP} = \frac{\lambda'_\pi A_+ K_1^+(\lambda_p)}{1 + \frac{\lambda_\pi}{\lambda_p} K_1^+(\lambda_p)} - \frac{\lambda'_p A_- K_1^-(\lambda_\pi)}{1 + \frac{\lambda_p}{\lambda_\pi} K_1^-(\lambda_\pi)} \tag{8.17}$$

Although this non-resetting rule has not reduced to the 2-spike rule, a similar analysis of the cases of large and small s_i (or large and small λ_p) as performed above for the 2-spike rule reveals identical conclusions, so that afferents either all escape to infinity or all die at zero. We therefore conclude that the presence of higher-order spike interactions under a unified, 3-state synaptic switch differentiates the 2-spike and multi-spike rules by allowing a probing of the coupling between potentiation and depression in the unified model. These higher-order interactions are responsible for giving rise to the stable, competitive dynamics that we observe for the multi-spike rules.

Examining the large β limit of the (unified) multi-spike rules reveals that their asymptotic behaviour differs significantly from that of the 2-spike rule. For large β, $K_l^\pm(\beta) \to 1$, so $\tilde{K}_l^+(\beta) \to \lambda_\pi'^{\,l}$ and $\tilde{K}_l^-(\beta) \to \lambda_p'^{\,l}$. Because $\lambda'_\pi + \lambda'_p = 1$, we introduce a new variable, $x \in [0, 1]$, such that

$$\lambda'_\pi = \frac{1}{2}(1 - x), \tag{8.18}$$

$$\lambda'_p = \frac{1}{2}(1 + x). \tag{8.19}$$

It is easy to see that x is the tangent of the angle between the vector $(\lambda_\pi, \lambda_p)^T$, the superscript T denoting the transpose, and the line $\lambda_\pi = \lambda_p$. If θ is the standard angle in a polar co-ordinate system ($\lambda_\pi = r\cos\theta$, $\lambda_p = r\sin\theta$), then x is just

$$x = \tan(\theta - \pi/4). \tag{8.20}$$

In this large β limit, we re-write the 2- and multi-spike rules in terms of the variable x and find that

$$\Delta S_2 \rightarrow \frac{1}{4}(1-x^2)(A_+ - A_-), \tag{8.21}$$

$$\Delta S_3 \rightarrow \frac{1}{8}(1-x^2)\left[A_+(3-x) - A_-(3+x)\right], \tag{8.22}$$

$$\Delta S_4 \rightarrow \frac{1}{16}(1-x^2)\left[A_+(9-4x-x^2) - A_-(9+4x-x^2)\right], \tag{8.23}$$

$$\Delta \hat{S}_\infty \rightarrow \frac{1}{2}\left(\frac{1-x^2}{3+x^2}\right)\left[A_+(1-x) - A_-(1+x)\right]. \tag{8.24}$$

Because the $\beta \rightarrow \infty$ limits of the resetting and non-resetting models are identical, the results above are independent of the form of the model, although the limit of the ∞-spike resetting rule is much easier to extract. We see that while the 2-spike rule, ΔS_2, is symmetric about $x = 0$, and thus symmetrical about the line $\lambda_\pi = \lambda_p$, the other rules exhibit an asymmetry about $x = 0$ due to the presence of odd powers of x. This property is true for all the multi-spike rules, not just ΔS_3, ΔS_4 and $\Delta \hat{S}_\infty$ given above. For $x^2 < 1$, the right-hand-side of Eq. (8.21) is strictly positive, because $A_+ > A_-$. Hence, the 2-spike rule always potentiates in the large β limit, as we saw earlier. This behaviour is not the case for the multi-spike rules. Consider the cases $x \approx +1$ and $x \approx -1$ in the above. Then we find

$$\Delta S_3 \propto \begin{cases} A_+ - 2A_- & \text{for } x \approx +1 \\ 2A_+ - A_- > 0 & \text{for } x \approx -1 \end{cases} \tag{8.25}$$

$$\Delta S_4 \propto \begin{cases} A_+ - 3A_- & \text{for } x \approx +1 \\ 3A_+ - A_- > 0 & \text{for } x \approx -1 \end{cases} \tag{8.26}$$

$$\Delta \hat{S}_\infty \propto \begin{cases} -A_- < 0 & \text{for } x \approx +1 \\ +A_+ > 0 & \text{for } x \approx -1 \end{cases} \tag{8.27}$$

Indeed, by using the general form of the $(n+1)$-spike rule Appleby and Elliott (2005), we find

$$\Delta S_{n+1} \propto \begin{cases} A_+ - nA_- & \text{for } x \approx +1 \\ nA_+ - A_- > 0 & \text{for } x \approx -1 \end{cases} \tag{8.28}$$

The multi-spike rules therefore always potentiate in the $x \approx -1$ or large λ_π, small λ_p direction. For the 3-spike rule, if $A_+ < 2A_-$, then it, and all higher-order rules, depress in the $x \approx +1$ or small λ_π, large λ_p direction. If, however, $A_+ > 2A_-$, then the 3-spike rule potentiates in all directions and will exhibit runaway learning just like the 2-spike rule. Although it may be the case that $A_+ > 2A_-$, it may not be the case that $A_+ > 3A_-$. Here, although the 3-spike rule fails, the 4-spike rule will depress in the $x \approx +1$ direction. Indeed, by looking at the behaviour of the general $(n+1)$-spike rule in Eq. (8.28), we

see that provided $A_- \neq 0$, there always exists a value of n above which the multi-spike rules will start to depress in the $x \approx +1$ direction. In particular, for $n > A_+/A_-$, the $(n+1)$-spike rule will always potentiate in the $x \approx -1$ direction and depress in the $x \approx +1$ direction. With such "mixed" dynamics at large β, large λ_p will not induce runaway learning. Since $A_+/A_- > 1$, the 2-spike rule can never achieve this. Hence, the 2-spike rule is irredeemably pathological in its learning behaviour due to a symmetry that is absent in all the multi-spike learning rules, and although the multi-spike rules can fail in the same way as the 2-spike rule, this is parameter-dependent (unlike the 2-spike rule), and we are guaranteed (for $A_- \neq 0$) that there exists an n above which the n-spike rules will not fail.

Plotting the ΔS_n in the λ_π–λ_p plane illustrates these results. For the 2-spike rule (Fig. 8.4), we see a depressing well around the origin, and a potentiating regime away from it. The 2-spike surface is symmetrical about the line $\lambda_\pi = \lambda_p$. Depression therefore always occurs at low β, and potentiation always occurs at high β. For the non-resetting forms of the 3- (Fig. 8.5), 4- (Fig. 8.6), and ∞-spike (Fig. 8.7) rules, the symmetry about the line $\lambda_\pi = \lambda_p$ is absent, and it is possible to induce either potentiation or depression at high β depending on the value of θ (or x), for the parameter choice used here. Thus, the values of λ_π and λ_p together determine whether potentiation or depression occurs.

8.3 Fixed Point Analysis of a Rate-Based Rule

Having derived various multi-spike, rate-based rules from our switch model and performed an initial study of the differences between the 2- and multi-spike rules, we may proceed to develop a deeper analytical understanding of the learning dynamics exhibited by the multi-spike rules. In particular, we continue to study Eq. (8.11) by performing a fixed point analysis of the ∞-spike learning rule.

We assume that all m afferents have the same mean firing rate, $\mu > 0$, and variance, $\sigma^2 > 0$. The firing rate of each afferent therefore fluctuates about a common mean. The fluctuations distinguish the afferents (unless they are perfectly correlated), while preventing any afferent from enjoying an overall advantage. We set $\lambda_{\pi_i} = \mu(1 + \alpha_i)$, where α_i is some small perturbation about the mean, $|\alpha_i| \ll 1$, and take the mean of α_i to be zero, $\langle \alpha_i \rangle = 0$, so that $\langle \lambda_{\pi_i} \rangle = \mu$, as required. As we will average over the ensemble of activity patterns, we must obtain an expression for $\langle \alpha_i \, \alpha_j \rangle$. The variance in the activity of afferent i is

$$\sigma^2 = \langle \lambda_{\pi_i}^2 \rangle - \langle \lambda_{\pi_i} \rangle^2 = \mu^2 \langle \alpha_i^2 \rangle \tag{8.29}$$

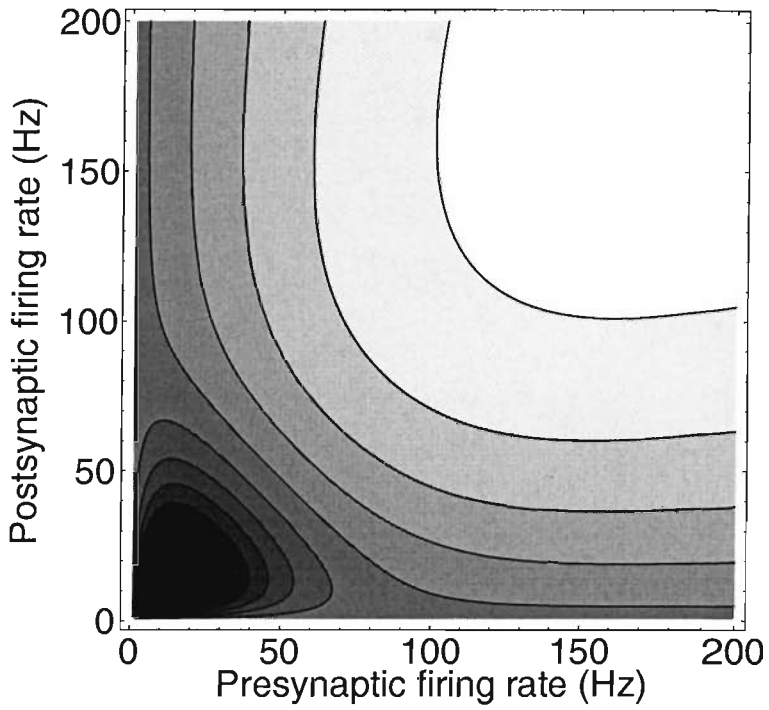

FIGURE 8.4: A contour plot of ΔS_2 in the λ_π–λ_p plane. Black areas represent minimum values and white areas maximum values, and nine shades of grey interpolate between these extremes. The minimum value on this partial plane is -0.0244, while the maximum value is $+0.0118$.

so that $\langle \alpha_i^2 \rangle = \hat{\sigma}^2$, where $\hat{\sigma} = \sigma/\mu$, with $\hat{\sigma}^2 \ll 1$. Assuming for convenience that the afferents' activities are uncorrelated, so that their covariance $\mathrm{Cov}(\alpha_i, \alpha_j) = 0$ for $i \neq j$, we then have

$$\langle \alpha_i \, \alpha_j \rangle = \hat{\sigma}^2 \delta_{ij}, \tag{8.30}$$

where δ_{ij} is the Kronecker delta, equal to one if $i = j$ and zero otherwise. Defining the vectors $\mathrm{s} = (s_1, \ldots, s_m)^T$ and $\alpha = (\alpha_1, \ldots, \alpha_m)^T$, we then have

$$\langle \alpha \cdot \mathrm{s} \rangle = 0, \tag{8.31}$$

$$\langle \alpha_i (\alpha \cdot \mathrm{s}) \rangle = \hat{\sigma}^2 s_i, \tag{8.32}$$

$$\langle (\alpha \cdot \mathrm{s})^2 \rangle = \hat{\sigma}^2 |\mathrm{s}|^2. \tag{8.33}$$

As the perturbations are small, $\hat{\sigma}^2 \ll 1$, we may expand any n-spike rule in α_i and then average over the ensemble of afferent activity patterns using the three equations above. This expansion must be to second order in α_i, as the mean of α_i is zero. We thus arrive at a set of equations describing the evolution of afferents governed by the n-spike switch rule when the activities of the afferents fluctuate about some common mean firing rate. We may then extract the fixed point structure that characterises the dynamics of this system. Because $s_i \geq 0$,

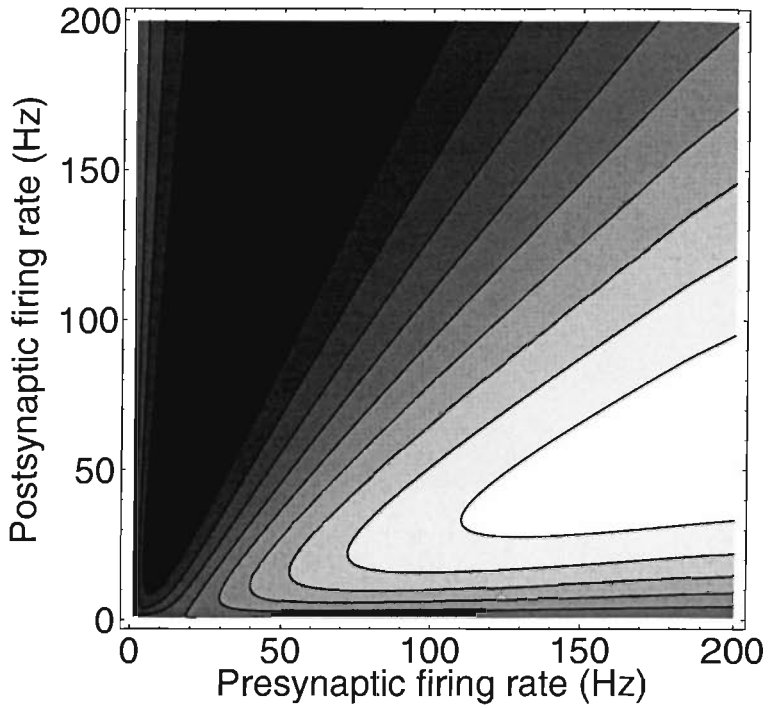

FIGURE 8.5: A contour plot of ΔS_3^{NR} for the non-resetting model in the λ_π–λ_p plane. The minimum value on this partial plane is -0.0826, while the maximum value is $+0.1007$.

any fixed points must lie in the non-negative hyperquadrant of the vector space defined by $\mathbf{s}$ in order to be accessible to the afferents.

We now explore the fixed point structure of a simplified form of the ∞-spike rule, the simplification merely allowing a less messy analytical characterisation of the locations and stabilities of the various fixed points. We then examine the full, unsimplified rule. Although analytical results can still be obtained for the full rule, they are messy, cumbersome and rather opaque, so we do not reproduce them here. Nevertheless, both the simplified and full forms of the model exhibit qualitatively identical dynamics.

8.3.1 Simplified ∞-Spike Rule

We consider a simplified form of the full ∞-spike non-resetting rule given in Eqn. 8.6, that excludes the denominator. This exclusion simplifies the expansion, making the resulting expressions more transparent. The price for this transparency is that the location and stability of any fixed point in this simplified model will be slightly different from that of the full model. To zeroth order in $\hat{\sigma}^2$, the simplified model and full model are, however, identical in their

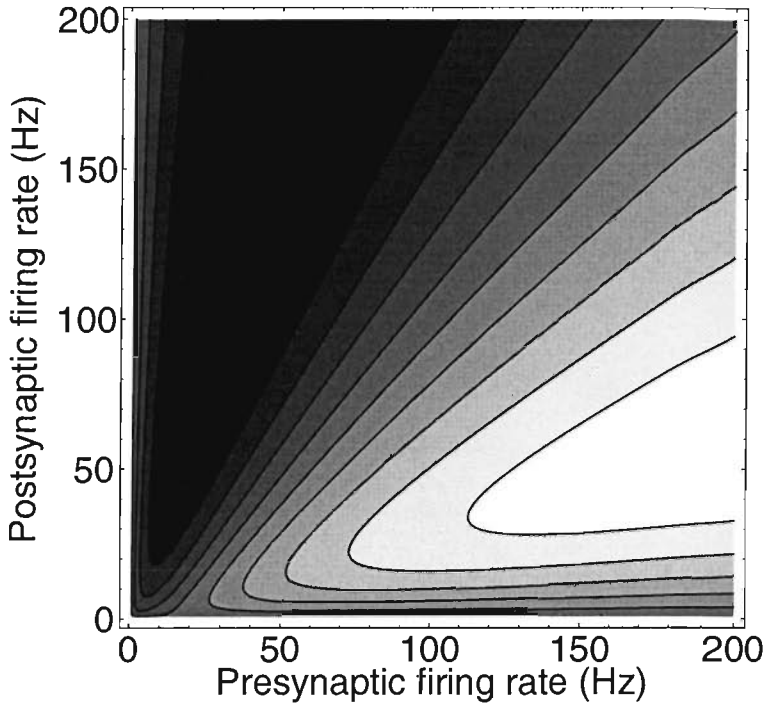

FIGURE 8.6: A contour plot of ΔS_4^{NR} for the non-resetting model in the λ_π–λ_p plane. The minimum value on this partial plane is -0.1686, while the maximum value is $+0.1926$.

fixed point structure. For reasons of analytical tractability, we set $n_\pm = 1$. The simplified ∞-spike, non-resetting rule is then

$$\Delta\hat{S}_\infty^{Sim}(\lambda_\pi, \lambda_p) = \lambda'_\pi A_+ K_1^+(\lambda_p) - \lambda'_p A_- K_1^-(\lambda_\pi), \tag{8.34}$$

and we set

$$\frac{ds_i}{dt} = \Delta\hat{S}_\infty^{Sim}(\lambda_{\pi_i}, \lambda_p). \tag{8.35}$$

We denote the sum of the synaptic weights as $s_+ = \sum_i s_i$ and we define $t_+ = 1 + s_+$. Expanding the simplified rule to second order in α_i and averaging over the ensemble of afferent activity patterns as set out above yields the averaged form of the rule. Dropping the angle brackets around the s_i for notational convenience, after lengthy algebra we obtain

$$\frac{ds_i}{dt} = A_-\tau_-\mu\frac{1}{t_+}\left(s_+ N_0 + \hat{\sigma}^2 N_1^i\right), \tag{8.36}$$

where

$$N_0 = \frac{\gamma}{1 + \mu\tau_+ s_+} - \frac{1}{1 + \mu\tau_-}, \tag{8.37}$$

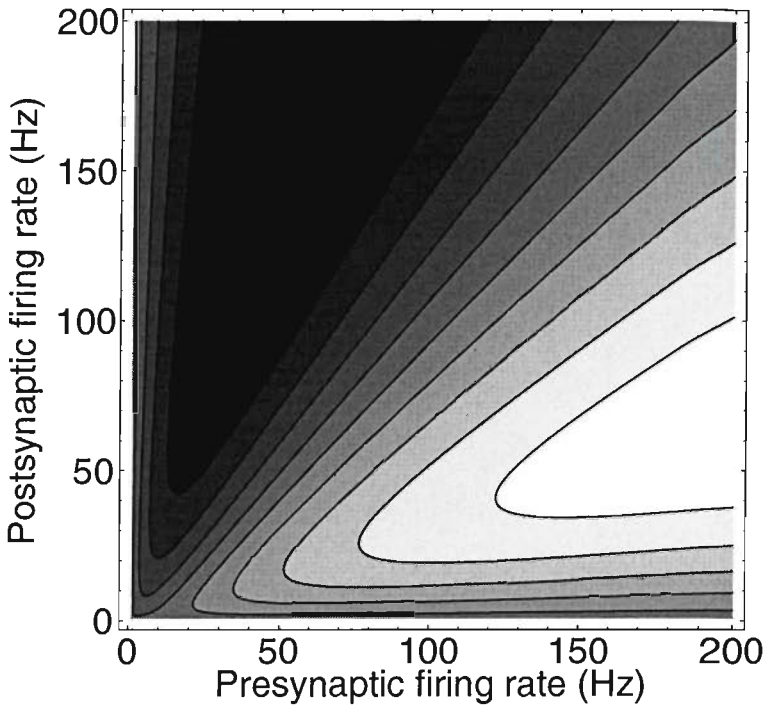

FIGURE 8.7: A contour plot of $\Delta\hat{S}_{\infty}^{NR}$ for the non-resetting model in the λ_{π}–λ_{p} plane. The minimum value on this partial plane is -0.1010, while the maximum value is $+0.1036$.

and

$$N_1^i = \frac{\gamma}{1+\mu\tau_+ s_+} X_i - \frac{1}{1+\mu\tau_-} Y_i, \tag{8.38}$$

where

$$X_i = \frac{1}{1+\mu\tau_+ s_+}\left(s_i - \frac{s_i+|\mathbf{s}|^2}{t_+}\right) + s_+ \frac{1+2s_i+|\mathbf{s}|^2}{t_+^2} - s_+\frac{1+s_i}{t_+} - \frac{\mu\tau_+|\mathbf{s}|^2}{(1+\mu\tau_+ s_+)^2}, \tag{8.39}$$

and

$$Y_i = \frac{1}{1+\mu\tau_-}\left(s_i - s_+\frac{1+s_i}{t_+}\right) + s_+ \frac{1+2s_i+|\mathbf{s}|^2}{t_+^2} - \frac{s_i+|\mathbf{s}|^2}{t_+} - \frac{\mu\tau_- s_+}{(1+\mu\tau_-)^2}, \tag{8.40}$$

At a fixed point, we require that $d\mathbf{s}/dt = \mathbf{0}$. Solving this equation exactly for the location of all the fixed points is usually difficult if not impossible, so we proceed by finding an approximation to zeroth-order in $\hat{\sigma}^2$, for which all the fixed points can be located, then calculate first-order corrections in $\hat{\sigma}^2$. We therefore write a fixed point as

$$\mathbf{s}_{FP} = \mathbf{x} + \hat{\sigma}^2\mathbf{y}, \tag{8.41}$$

where $\mathbf{x}$ is the zeroth-order approximation to the location of the fixed point,

$\mathbf{s}_{FP}$, and $\mathbf{y}$ the first-order correction. We define $x_+ = \sum_i x_i$ and $y_+ = \sum_i y_i$, so that $s_+ = x_+ + \hat{\sigma}^2 y_+$.

8.3.1.1 Zeroth-Order Solutions and Behaviour

To zeroth order in $\hat{\sigma}^2$, Eq. (8.36) becomes

$$\frac{ds_i}{dt} = A_-\tau_-\mu\frac{s_+}{t_+}\left(\frac{\gamma}{1+\mu\tau_+ s_+} - \frac{1}{1+\mu\tau_-}\right), \tag{8.42}$$

which depends only on s_+ and not the individual components s_i. By inspection we see that there are two fixed hyperplanes. One of these corresponds to the hyperplane $x_+ = 0$. All points except $\mathbf{x} = \mathbf{0}$ in this hyperplane have at least one negative component of $\mathbf{x}$, and all these points are therefore forbidden. Hence, the only fixed point on this hyperplane of fixed points accessible to the afferent weight vector is the origin. We refer to this point throughout as the zero fixed point. The other fixed hyperplane arises from the solution of

$$\frac{\gamma}{1+\mu\tau_+ x_+} = \frac{1}{1+\mu\tau_-}, \tag{8.43}$$

or

$$x_+ = \frac{\gamma(1+\mu\tau_-) - 1}{\mu\tau_+}. \tag{8.44}$$

So that at least some points in the hyperplane $s_+ = x_+$ have all non-negative components, we require $x_+ > 0$, or

$$\gamma > \gamma_0 \equiv \frac{1}{1+\mu\tau_-}. \tag{8.45}$$

We refer to the hyperplane $s_+ = x_+ > 0$ as the non-zero fixed hyperplane. If $\gamma = \gamma_0$, then the non-zero hyperplane becomes co-incident with the hyperplane $s_+ = 0$, and the only permitted fixed point that exists for $\gamma = \gamma_0$ is the origin, $\mathbf{x} = \mathbf{0}$.

To determine the stability of these fixed points we examine the behaviour of the system under small perturbations about them. We denote a perturbation by $\delta = (\delta_1, \ldots, \delta_m)^T$ and write $\mathbf{s} = \mathbf{s}_{FP} + \delta$. We define $\delta_+ = \sum_i \delta_i$.

Expanding and linearising Eq. (8.42) about the zero fixed point, $\mathbf{s}_{FP} = \mathbf{0}$, we have

$$\frac{d\delta_i}{dt} = A_-\tau_-\mu(\gamma - \gamma_0)\delta_+, \tag{8.46}$$

The m eigenvalues of the associated matrix characterising these linearised dynamics are therefore easily seen to be

$$\lambda_1 = mA_-\tau_-\mu(\gamma - \gamma_0), \tag{8.47}$$

and

$$\lambda_i = 0, \quad \forall i > 1. \tag{8.48}$$

The $m - 1$ repeated zero eigenvalues indicate that there is no flow along the associated eigenvectors. The flow towards or away from the origin will therefore only occur parallel to the eigenvector associated with λ_1, which is $(1, \ldots, 1)^T$. The stability of the zero fixed point is determined by the sign of λ_1. For $\gamma < \gamma_0$, the origin is stable, while for $\gamma > \gamma_0$, it is unstable. Note that when $\gamma > \gamma_0$, the non-zero fixed hyperplane $s_+ = x_+ > 0$ intersects the positive hyperquadrant, so the origin becomes unstable precisely when this hyperplane moves into the positive hyperquadrant.

Now expanding and linearising Eq. (8.42) about any point in the non-zero fixed hyperplane, we obtain

$$\frac{d\delta_i}{dt} = -A_-\tau_-\tau_+\mu^2\frac{x_+}{1+x_+}\gamma^{-1}\frac{1}{(1+\mu\tau_-)^2}\delta_+, \tag{8.49}$$

with associated eigenvalues

$$\lambda_1 = -mA_-\tau_-\tau_+\mu^2\frac{x_+}{1+x_+}\gamma^{-1}\frac{1}{(1+\mu\tau_-)^2}, \tag{8.50}$$

and

$$\lambda_i = 0, \quad \forall i > 1. \tag{8.51}$$

The flow towards the non-zero fixed hyperplane is parallel to $(1, \ldots, 1)^T$. As for the zero fixed point, the sign of λ_1 determines the stability of the non-zero fixed hyperplane. For $x_+ > 0$, i.e. when the non-zero fixed hyperplane intersects the positive hyperquadrant, the non-zero fixed hyperplane is stable. Note that $x_+ > 0$ requires $\gamma > \gamma_0$, so when the non-zero fixed hyperplane is stable, the origin is unstable.

The zeroth order dynamics are therefore uniquely determined by the sign of the quantity $\gamma - \gamma_0$. We briefly summarise the dynamics in each of the two possible regimes.

When $\gamma < \gamma_0$, a single fixed point $\mathbf{s}_{FP} = \mathbf{0}$ is permitted for the afferent weight vector. This point is stable, and the afferent weight vector will initially flow towards the hyperplane $s_+ = 0$ parallel to the vector $(1, \ldots, 1)^T$. When it hits a hyperplane defined by $s_i = 0$, for some i, it is prevented from crossing it because s_i is always truncated at zero. The weight vector therefore remains in this $s_i = 0$ hyperplane and flows towards the origin. It may hit other $s_j = 0$, $j \neq i$, hyperplanes as it further evolves, and again will be constrained to remain in them. The weight vector will therefore always arrive at the origin, regardless of the initial conditions.

When $\gamma > \gamma_0$, the origin is an unstable fixed point and there exists a hyperplane of stable, non-zero fixed points in the positive hyperquadrant. The afferent weight vector initially flows parallel to the vector $(1, \ldots, 1)^T$ towards the hyperplane $s_+ = x_+ > 0$, either from above or from below it. For some initial conditions, the weight vector will directly hit the non-zero hyperplane and stop evolving. For other initial conditions, the weight vector will hit a hyperplane $s_i = 0$, for some i, first, and then flow in this hyperplane until it reaches the intersection with the $s_+ = x_+ > 0$ hyperplane, and then stop evolving. Regardless of the initial conditions, the weight vector will always arrive at some point on the non-zero fixed hyperplane.

Of course, overall these dynamics are rather uninteresting, precisely because the zeroth-order solutions do not discriminate between afferents, since all afferents fire with a common rate μ. Nevertheless, the zeroth-order solutions are the foundation on which the first-order corrections are determined, and this is why we have laboured the analysis of the zeroth-order case somewhat. The first order corrections do permit a discrimination between afferents based on the fluctuations in their firing rates, and so we expect to find a more compelling set of dynamics at first order. We now turn to this case.

8.3.1.2 First-Order Corrections and Behaviour

We now examine the full form of Eq. (8.36), including the first-order corrections. For the first-order dynamics, we find that several different classes of fixed point exist.

By simple inspection of the terms in N_1^i, the first-order correction in Eq. (8.36), we see that the system still possesses a fixed point at $\mathbf{s} = \mathbf{0}$, and thus there are obviously no first order corrections to its location. However, there are corrections to the eigenvalues of the stability matrix. Expanding and linearising Eq. (8.36) about the zero fixed point as usual, we now find

$$\frac{d\delta_i}{dt} = A_-\tau_-\mu\left[\left(\gamma - \frac{1}{1+\mu\tau_-} - \hat{\sigma}^2\frac{\mu^2\tau_-^2}{(1+\mu\tau_-)^3}\right)\delta_+ + \hat{\sigma}^2\frac{\mu\tau_-}{(1+\mu\tau_-)^2}\delta_i\right], \quad (8.52)$$

with associated eigenvalues

$$\lambda_1 = mA_-\tau_-\mu\left[\gamma - \frac{1}{1+\mu\tau_-} - \hat{\sigma}^2\frac{\mu\tau_-}{(1+\mu\tau_-)^2}\left(\frac{\mu\tau_-}{1+\mu\tau_-} - \frac{1}{m}\right)\right] \quad (8.53)$$

and

$$\lambda_i = A_-\tau_-\mu\frac{\mu\tau_-}{(1+\mu\tau_-)^2}, \quad \forall i > 1. \quad (8.54)$$

The eigenvalues λ_i, $i > 1$, are always positive, so the first-order corrections have made the zero fixed point always unstable. The precise classification of the zero fixed point depends on the sign of λ_1. When $\lambda_1 < 0$, the origin is a saddle node and when $\lambda_1 > 0$, it is a repeller. The transition of the zero fixed point from a saddle to a repeller occurs at the value of γ given by

$$\gamma = \gamma_2 \equiv \frac{1}{1+\mu\tau_-}\left[1 + \hat{\sigma}^2 \frac{\mu\tau_-}{1+\mu\tau_-}\left(\frac{\mu\tau_-}{1+\mu\tau_-} - \frac{1}{m}\right)\right]. \tag{8.55}$$

Unlike the zero fixed point, the first order corrections do affect the location of the non-zero fixed points. Indeed, the corrections destroy the entire non-zero hyperplane of fixed points, leaving one real fixed point and m quasi-fixed points, together with a set of other fixed points that are essentially uninteresting because their existence merely renders the global fixed point structure consistent. Unless $\mathbf{s} = \mathbf{0}$, the first-order corrections in the expression for N_1^i break the symmetry between the afferents that is present at zeroth order. It is precisely this symmetry that endows the zeroth-order system with an entire hyperplane of fixed points. At first order, this symmetry is absent, and the hyperplane collapses into a set of isolated quasi-fixed points and real fixed points. We first define what we mean by quasi-fixed points and examine their stability, then examine the other fixed points.

In models of synaptic competition that possess a fixed point structure, there are usually fixed points in which all but one s_i are zero. Such fixed points correspond to segregated states, since at each fixed point, only one afferent innervates the target cell. With m afferents, there are m such segregated fixed points. Of course, the defining condition of a fixed point is that $d\mathbf{s}/dt = \mathbf{0}$ when the derivatives are evaluated at the fixed point. Suppose, however, that we have a point for which $s_i \neq 0$ and $s_j = 0$ $\forall j \neq i$ for which $ds_i/dt = 0$ and $ds_j/dt < 0$, $\forall j \neq i$, when the derivatives are evaluated at this point. Such a point is not strictly a fixed point. However, if a model's dynamics include truncation of s_j as it tries to pass through zero into a region of negative s_j, then s_j will be returned to zero. Such a point would therefore appear to be a fixed point, since the dynamics could evolve the weight vector to this point and then it would remain there. We refer to such points as quasi-fixed points. If $ds_j/dt < 0$ for those components that are zero, and if the non-zero s_i direction is stable, then we refer to the quasi-fixed point as stable; otherwise we refer to the quasi-fixed point as unstable.

Eq. (8.36) possesses m such points corresponding to segregated quasi-fixed points. Consider the point $\mathbf{s}_{QFP} = (0, \ldots, 0, s_+, 0, \ldots, 0)^T$, where only the ith component is non-zero, and where we write the usual expansion, $s_+ = x_+ + \hat{\sigma}^2 y_+$. The component s_i must therefore have zero derivative at $\mathbf{s}_{QFP}$, and this requirement determines the value of s_+. Solving the equation $ds_i/dt|_{\mathbf{s}_{QFP}} = 0$ we of course

obtain the zeroth order solution

$$x_+ = \frac{\gamma(1+\mu\tau_-)-1}{\mu\tau_+}, \tag{8.56}$$

which is Eq. (8.44), reflecting the fact that all non-zero solutions at zeroth order must live in the non-zero fixed hyperplane, and, after a little algebra, the first-order correction is found to be

$$y_+ = \frac{1}{\mu\tau_+}\frac{1-\gamma}{\gamma}\frac{1-\gamma\mu\tau_-}{1+\mu\tau_-}. \tag{8.57}$$

We require that $s_+ = x_+ + \hat{\sigma}^2 y_+ > 0$ for this point to be in the non-negative hyperquadrant, so we will obtain a first-order correction to the bound on γ. At zeroth order, the condition that $x_+ > 0$ forces $\gamma > \gamma_0$. We now write $\gamma = \gamma_0 + \hat{\sigma}^2\hat{\gamma}$ and determine a condition on $\hat{\gamma}$ so that $s_+ > 0$. We find that

$$\hat{\gamma} > -\frac{\mu\tau_-}{(1+\mu\tau_-)^3}, \tag{8.58}$$

and so, for $s_+ > 0$ at these m possible segregated points, we need

$$\gamma > \gamma_1 \equiv \frac{1}{1+\mu\tau_-}\left[1 - \hat{\sigma}^2\frac{\mu\tau_-}{(1+\mu\tau_-)^2}\right]. \tag{8.59}$$

It is easy to see that $\gamma_2 \geq \gamma_1$ since $m \geq 1$. Thus, these m points become accessible to the weight vector before the origin turns from a saddle into a repeller. We now need to determine the sign of ds_j/dt, $j \neq i$, at $\mathbf{s}_{QFP}$ in order to determine whether these points are possibly stable. We find that for $j \neq i$,

$$\left.\frac{ds_j}{dt}\right|_{\mathbf{s}=\mathbf{s}_{QFP}} = A_-\tau_-\mu\,\hat{\sigma}^2\frac{x_+}{(1+x_+)^2}\frac{1}{\gamma}\frac{1}{1+\mu\tau_-}\left[x_+\left(\gamma - \frac{1}{1+\mu\tau_-}\right) - \gamma\frac{\mu\tau_-}{1+\mu\tau_-}\right]. \tag{8.60}$$

This equation is purely first order in $\hat{\sigma}^2$ and thus any resulting bound on γ derived from it will be purely zeroth order in $\hat{\sigma}^2$. To calculate a higher order correction to any resulting bound on γ, we would be compelled to extend our expansion out to order $\hat{\sigma}^4$. One consequence of this is that while at first order in $\hat{\sigma}^2$ the quantity x_+ is allowed to be negative, because the correction $\hat{\sigma}^2 y_+$ can pull the sum $s_+ = x_+ + \hat{\sigma}^2 y_+$ overall positive, nevertheless, in determining the sign of the right-hand-side of Eq. (8.60), we must take x_+ strictly positive, because the resulting bound on γ will be only zeroth order, for which we are only allowed to have $x_+ > 0$. Hence, the sign of the right-hand-side of Eq. (8.60) is determined by the terms in square brackets. For negative derivatives, we require

$$x_+\left(\gamma - \frac{1}{1+\mu\tau_-}\right) - \gamma\frac{\mu\tau_-}{1+\mu\tau_-} < 0. \tag{8.61}$$

Replacing x_+ by its expression in Eq. (8.56) and writing $\tau_+ = \gamma\tau_- A_-/A_+$ since

we regard τ_+ as a function of γ with τ_- and $A_\pm$ fixed, we then obtain the quadratic equation

$$\left[1 - \frac{A_-}{A_+}(\mu\tau_-\gamma_0)^2\right]\gamma^2 - 2\gamma_0\gamma + \gamma_0^2 < 0. \tag{8.62}$$

Solving this equation for γ gives bounds on γ that, after a little algebra, are

$$\frac{A_+(1+\mu\tau_-) - \mu\tau_-\sqrt{A_+A_-}}{A_+(1+2\mu\tau_-) + (A_+ - A_-)\mu^2\tau_-^2} < \gamma < \frac{A_+(1+\mu\tau_-) + \mu\tau_-\sqrt{A_+A_-}}{A_+(1+2\mu\tau_-) + (A_+ - A_-)\mu^2\tau_-^2}. \tag{8.63}$$

Notice that $\gamma = \gamma_0$ satisfies Eq. (8.62), so the relevant bound on γ is the upper bound. Defining

$$\gamma_3 = \frac{A_+(1+\mu\tau_-) + \mu\tau_-\sqrt{A_+A_-}}{A_+(1+2\mu\tau_-) + (A_+ - A_-)\mu^2\tau_-^2}, \tag{8.64}$$

we require that $\gamma < \gamma_3$ in order that the segregated quasi-fixed points $\mathbf{s}_{QFP}$ may be stable. For $\gamma > \gamma_3$, the derivatives ds_j/dt, $\forall j \neq i$ become positive, so the segregated quasi-fixed points are certainly unstable in this region. It remains to be assured that the non-zero, s_i direction is stable. At zeroth order, the various points $\mathbf{s}_{QFP}$ are part of the non-zero fixed hyperplane, which we know to be stable for $\gamma > \gamma_0$. Hence, the s_i direction is always stable, at zeroth order. Thus, the segregated quasi-fixed points are certainly stable for $\gamma_0 < \gamma < \gamma_3$, where for consistency the lower bound γ_0 is of the same order in $\hat{\sigma}^2$, namely zeroth order, as the upper bound γ_3. This leaves open the small, first-order-sized window $\gamma_1 < \gamma < \gamma_0$ in which the segregated quasi-fixed points exist in the non-negative hyperquadrant. In fact, the quasi-fixed points are stable in this small region too. Strictly speaking, it is inconsistent to write the condition on γ that guarantees the stability of the segregated quasi-fixed points in the form $\gamma_1 < \gamma < \gamma_3$, since the orders of the two bounds differ, but we shall do so anyway, in order to close the small $\gamma_1 < \gamma < \gamma_0$ window.

We may now turn to a real, unsegregated fixed point. This fixed point is defined by $\mathbf{s}_{FP}$ having entirely non-zero components. From the definition of a fixed-point, we require that

$$\left(s_+N_0 + \hat{\sigma}^2 N_1^i\right)|_{\mathbf{s}=\mathbf{s}_{FP}} = 0, \quad \forall i \tag{8.65}$$

we must have $N_1^i = N_1^j$, $\forall i \neq j$. The simplest solution of this equation is, of course, $s_{FP}^i = s_{FP}^j$, $\forall i \neq j$, so that all the components of $\mathbf{s}_{FP}$ are equal (and non-zero). Because all the components are equal, we refer to this point as the unsegregated fixed point. We write $\mathbf{s}_{FP} = \frac{1}{m}(s_+, \ldots, s_+)^T$ and again expand s_+ as $s_+ = x_+ + \hat{\sigma}^2 y_+$. The zeroth order solution must lie on the zeroth order

non-zero fixed hyperplane, so we again have

$$x_+ = \frac{\gamma(1+\mu\tau_-)-1}{\mu\tau_+}. \tag{8.66}$$

We find that the first-order correction is

$$y_+ = \frac{\gamma}{\mu\tau_+}\left[\frac{1}{1+\mu\tau_-}\left(\mu\tau_- - \frac{1}{m}\gamma^{-2}\mu\tau_+ x_+\right) - \frac{\mu\tau_-}{1+x_+}\left(1-\frac{1}{m}\right)\right]. \tag{8.67}$$

Again, we require that $s_+ = x_+ + \hat{\sigma}^2 y_+ > 0$ for this unsegregated fixed point to be accessible to the weight vector, so, expanding γ as $\gamma = \gamma_0 + \hat{\sigma}^2\hat{\gamma}$ as for the segregated quasi-fixed points, we find that γ must satisfy

$$\gamma > \frac{1}{1+\mu\tau_-}\left[1+\hat{\sigma}^2\frac{\mu\tau_-}{1+\mu\tau_-}\left(\frac{\mu\tau_-}{1+\mu\tau_-}-\frac{1}{m}\right)\right]. \tag{8.68}$$

The right-hand-side of this inequality is precisely γ_2 defined in Eq. (8.55), which determines the value of γ at which the zero fixed point turns from a saddle into a repeller. Thus, as the unsegregated fixed point passes through the origin into the positive hyperquadrant, the zero fixed point turns into a repeller. To determine the stability of the unsegregated fixed point we expand and linearise about it as usual, and after lengthy algebra we find that

$$\frac{d\delta_i}{dt} = \frac{A_-\tau_-\mu}{1+\mu\tau_-}\frac{1}{(1+x_+)^2}\frac{1}{\gamma}\left[\frac{x_+(1+x_+)}{1+\mu\tau_-}\left(-\mu\tau_+ + \hat{\sigma}^2 J\right)\delta_+ - \hat{\sigma}^2 K\,\delta_i\right], \tag{8.69}$$

where J is a long and unwieldy expression that we do not reproduce here, and K is given by

$$K = x_+\left(\gamma - \frac{1}{1+\mu\tau_-}\right) - \gamma\frac{\mu\tau_-}{1+\mu\tau_-}. \tag{8.70}$$

The eigenvalues of the associated matrix are then just

$$\lambda_1 = \frac{mA_-\tau_-\mu}{1+\mu\tau_-}\frac{1}{(1+x_+)^2}\frac{1}{\gamma}\left[\frac{x_+(1+x_+)}{1+\mu\tau_-}\left(-\mu\tau_+ + \hat{\sigma}^2 J\right) - \frac{1}{m}\hat{\sigma}^2 K\right], \tag{8.71}$$

and

$$\lambda_i = -\frac{A_-\tau_-\mu}{1+\mu\tau_-}\frac{1}{(1+x_+)^2}\frac{1}{\gamma}\hat{\sigma}^2 K, \quad \forall i > 1. \tag{8.72}$$

Although the expression is messy, it is easy to show that λ_1, associated with the eigenvector $(1,\ldots,1)^T$, changes sign at $\gamma = \gamma_2$, being positive for $\gamma < \gamma_2$ and negative for $\gamma > \gamma_2$. Hence, as the unsegregated fixed point moves into the positive hyperquadrant, the zero fixed point turns into a repeller because the direction corresponding to $(1,\ldots,1)^T$ becomes unstable and the unsegregated fixed point becomes stable precisely in this same direction. The sign of all the other eigenvalues associated with the unsegregated fixed point is determined solely by K. Stability in all the directions orthogonal to $(1,\ldots,1)^T$ thus requires

$K > 0$, or

$$x_+ \left(\gamma - \frac{1}{1 + \mu\tau_-} \right) - \gamma \frac{\mu\tau_-}{1 + \mu\tau_-} > 0. \tag{8.73}$$

This is identical to Eq. (8.61), determining the stabilities of the segregated quasi-fixed points, except that the inequality is opposite. Thus, we see immediately that we must have $\gamma > \gamma_3$ for the stability of the unsegregated fixed point, and for $\gamma_2 < \gamma < \gamma_3$, the unsegregated fixed point is an unstable saddle node. The unsegregated fixed point becomes stable precisely when the segregated quasi-fixed points become unstable, and the unsegregated fixed point becomes accessible to the weight vector precisely when the zero fixed point turns into a repeller. Thus, all three sets of fixed points are dynamically coupled in terms of their stabilities.

Because the local fixed point structure must be globally consistent, we can deduce that there must exist other fixed points for $m \geq 3$. For example, in the interval $\gamma \in (\gamma_1, \gamma_3)$, there must exist saddles that partition the afferent weight vector space into m regions, each region being defined by the requirement that an initial weight vector in the region always flows to the same quasi-fixed point. Consider, for example, three afferents labelled s_i, s_j and s_k. Examining the surface $s_k = 0$, we still see segregation of s_i and s_j on this surface, with flow similar to that shown in Fig. 8.8A. That is, there exists a line of fixed points that divide the surface $s_k = 0$ into two regions, according to whether the flow is to the $s_i \neq 0$ or to the $s_j \neq 0$ segregated state. There exist, therefore, points of the form $(s, s, 0)$, which represent a partially segregated state. However, we are only interested in such points if they are stable and, for uncorrelated afferents, these points are not stable but saddles.

The situation is not so clear for correlated afferents. If two afferents were perfectly correlated then these afferents would be indistinguishable. A 3-afferent system with two perfectly correlated and one uncorrelated afferent would therefore be identical to a (uncorrelated) 2-afferent system. Moving between the two descriptions would simply involve combining the strength of the two perfectly correlated afferents into one. Thus, the partially segregated fixed-point $(s, s, 0)$ would, in the case of two perfectly correlated afferents, be equivalent to the segregated fixed-point $(s, 0)$ in the 2-afferent case, and would therefore be stable. Thus, the system would evolve to a either a segregated or partially-segregated fixed-point and remain there. As reported for other models of competition (Elliott and Shadbolt, 1996) we expect that perfect correlation is not, in practice, required to observe this behaviour. We would instead expect to find critical correlation threshold above which afferents cannot be segregated but instead evolve to unsegregated fixed-points like $(s, s, 0)$.

However, as we consider only uncorrelated afferents, these partially-segregated fixed-points are always saddles. We do not explore these additional, essentially

uninteresting fixed points here because they merely render consistent the global fixed point structure that is determined by the stability of the segregated quasi-fixed points and the unsegregated fixed point. It is the stabilities of the segregated and unsegregated states in which we are principally interested here, since these are the states relevant for a putative model of synaptic competition.

We see that the first-order dynamics are again uniquely determined by the value of γ, but in contrast to the zeroth-order dynamics, which had only two distinct regimes, the first order system has four distinct regimes. We now summarise the dynamics in each of the regimes.

For $\gamma < \gamma_1$, there is only one fixed point, at the origin. Although a saddle node, it is stable in the $(1, \ldots, 1)^T$ direction. Hence, all flow initially moves parallel to this vector in the direction of the origin. If the weight vector moves sufficiently close to the origin, it will experience a repulsion in the directions orthogonal to $(1, \ldots, 1)^T$, but for $\gamma < \gamma_1$ this repulsion is never sufficiently strong to reverse the downward components of flow towards the origin. In all cases, the weight vector will eventually hit an $s_i = 0$, for some i, hyperplane and be trapped in it by the truncation procedure. Because there are still negative components of flow, the weight vector continues to move towards the origin, becoming trapped in further $s_j = 0$, $j \neq i$, hyperplanes. The weight vector thus is always driven towards the origin and ends up there, despite the origin's being a saddle node. The truncation at zero thus overrides this fixed point's instability.

For $\gamma_1 < \gamma < \gamma_2$, the zero fixed point is still a saddle node, and now there exists a set of m stable, segregated quasi-fixed points. In this regime of γ the repulsion away from the origin is sufficiently strong to reverse the negative components of flow. Thus, for a weight vector sufficiently close to the origin, it hits an $s_i = 0$, for some i, hyperplane and is turned away from the origin, so that it starts moving in the opposite direction. It moves up $s_i = 0$ hyperplanes until it reaches a stable, quasi-fixed point, and remains there. For weight vectors sufficiently distant from the origin, they flow towards the origin parallel to the direction $(1, \ldots, 1)^T$. These vectors never go sufficiently close to the origin to have their negative components of flow reversed. They thus hit $s_i = 0$ hyperplanes and move down towards the stable, quasi-fixed points, and stay there.

In the regime $\gamma_2 < \gamma < \gamma_3$, the zero fixed point has now turned into a repeller, so there are never any components of flow towards the origin in the neighbourhood of the origin. An unsegregated fixed point has become accessible, but it is a saddle node. The segregated quasi-fixed points remain stable. Hence, this regime is essentially identical to the regime in which $\gamma_1 < \gamma < \gamma_2$, except with local differences near the origin and around the now-accessible unsegregated fixed point. All flow therefore ends up at the stable, segregated quasi-fixed points.

The combined regime $\gamma_1 < \gamma < \gamma_3$ therefore supports stable, competitive dynamics, allowing afferents to segregate on the target cell in an activity-dependent manner, as required, for example, in a model of ODC formation.

Finally, when $\gamma_3 < \gamma$, the segregated quasi-fixed points become unstable, and the unsegregated fixed point becomes an attractor. The origin remains a repeller. Hence, all flow ends up at the unsegregated fixed point.

The phase portraits for the two interesting, dynamically distinct γ regimes are shown in Fig. 8.8, for a 2-afferent system. We do not show the low γ regime, $\gamma < \gamma_1$, because the portrait is trivial, with all initial conditions flowing to the origin. Fig. 8.8A shows the regime in which γ takes an intermediate value, $\gamma_1 < \gamma < \gamma_3$, for which a set of stable, segregated quasi-fixed points exists. The system always evolves to one of these segregated points. For this intermediate value of γ, the unsegregated fixed point is unstable. When γ is too high, $\gamma_3 < \gamma$, shown in Fig. 8.8B, the segregated quasi-fixed points become unstable, and the unsegregated fixed point becomes stable, so afferent segregation on the target cell breaks down and the afferent weight vector always flows to the unsegregated fixed point.

8.3.2 Full ∞-Spike Rule

The above fixed point analysis for the simplified rule may be repeated for the full, non-resetting, ∞-spike rule, for the convenient parameter choice $n_\pm = 1$. We do not present the results of this analysis here because the expressions that arise from the full model are unwieldy and thus lack the transparency of those for the simplified model. The simplified model has the virtue, compared to the full model, that almost all the resulting expressions can be stated on one line, while those for the full model occupy several lines. Nevertheless, the full model possesses dynamics that are qualitatively identical to those of the simplified model discussed above.

In the full model, we observe similar critical values of γ at which new quasi- or real fixed points become accessible to the afferent weight vector, or at which the fixed points change their stability. In particular, corresponding to the values γ_1 and γ_2 for the simplified model, we have the values, say, $\overline{\gamma}_1$ and $\overline{\gamma}_2$ for the full model, where these differ from γ_1 and γ_2 by terms of order $\hat{\sigma}^2$. At $\gamma = \overline{\gamma}_1$, the segregated quasi-fixed points become available and are initially stable. At $\gamma = \overline{\gamma}_2$, the unsegregated fixed point moves into the non-negative hyperquadrant, being initially an unstable saddle, and the zero fixed point at the origin simultaneously changes from a saddle node into a repelling node. We still have $\overline{\gamma}_1 \leq \overline{\gamma}_2$ as in the simplified model.

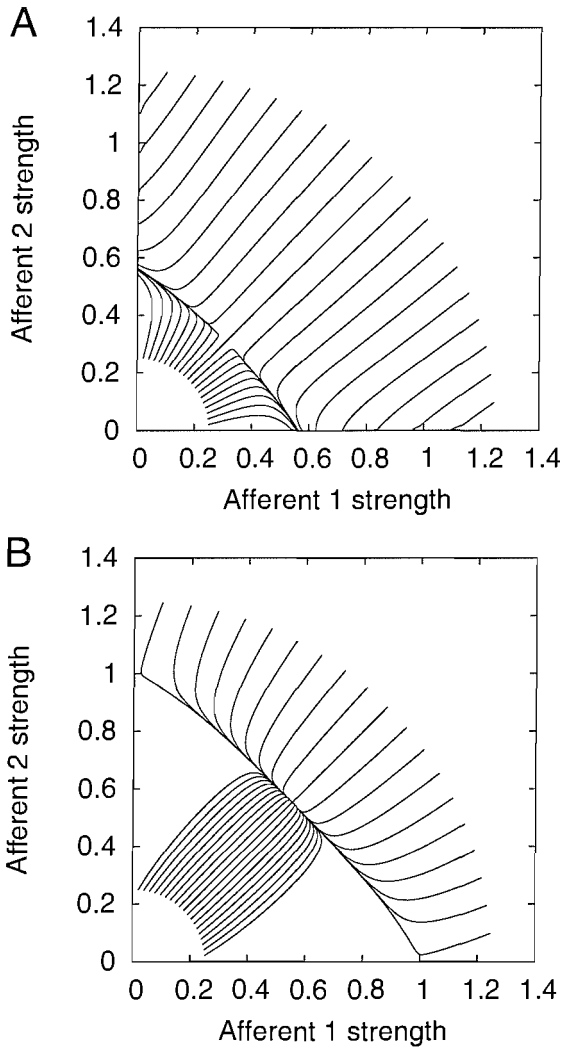

FIGURE 8.8: Phase portraits of the simplified, non-resetting ∞-spike rule, with $n_{\pm} = 1$, in a system with two afferents. **(A)** Evolution to the segregated quasi-fixed points with $\gamma = 0.65$. **(B)** Evolution to the unsegregated fixed point with $\gamma = 0.95$.

Corresponding to γ_3 in the simplified model, we also have a new value, say, $\overline{\gamma}_3$ in the full model. Although derived from first-order equations, γ_3 and $\overline{\gamma}_3$ are nonetheless purely zeroth-order in $\hat{\sigma}^2$. Hence, γ_3 and $\overline{\gamma}_3$ differ even at zeroth order. Despite this, the dynamics associated with the $\gamma = \overline{\gamma}_3$ transition in the full model are identical to those associated with the $\gamma = \gamma_3$ transition in the simplified model. At this point, the segregated quasi-fixed points become unstable and the unsegregated fixed point simultaneously becomes stable.

Although our analyses of the full and simplified models have been performed only for the case $n_{\pm} = 1$, for reasons of analytical simplicity, we can explore

numerically the impact of other values of $n_\pm$ on the full and simplified models. For $n_\pm = 3$, for example, we observe dynamics that are qualitatively similar to the $n_\pm = 1$ case, with the same three, essentially distinct parameter regimes in γ.

8.3.3 Beyond Small α_i

In the above fixed point analysis, we expanded, for example, Eq. (8.34) in the variables α_i, representing small fluctuations about a common mean afferent activity, μ, and then performed an ensemble average over these fluctuations. Although this permits us to make some progress in terms of understanding the fixed point dynamics of the model, it necessarily does not allow an examination of the large α_i regime, in which the fluctuations about the mean activity can be large. In principle, because we defined the size of these fluctuations with respect to the mean firing rate, so that $|\alpha_i| \leq 1$, we could continue the expansion to yet higher orders in $\hat{\sigma}^2$. Although possible, doing so would be tiresome and the resulting expressions an uncontrollable mess. It is therefore not clear that any additional analytical insight would be possible in the face of the growing complexity of the new terms.

We can, however, explore the fixed point structure for large α_i by means of numerical simulation. When we do so, we find essentially the same three regimes in γ for both the full and simplified models considered above. In particular, γ_1 and γ_2 still exist and define the points at which the segregated quasi-fixed points and the unsegregated fixed point, respectively, move into the non-negative hyperquadrant and thus become accessible to the weight vector. However, the γ_3 critical value in both models is somewhat modified. The large α_i fluctuations "split" this value of γ into two different values; call them, say, γ_3' and γ_3'', where $\gamma_3' < \gamma_3''$. At $\gamma = \gamma_3'$, the unsegregated fixed point becomes stable while the segregated quasi-fixed points remain stable. Only at $\gamma = \gamma_3''$ do the segregated quasi-fixed points become unstable. Thus, there is a narrow interval for γ, $\gamma \in (\gamma_3', \gamma_3'')$, of size of order $\hat{\sigma}^2$ or higher, in which the system may evolve either to a segregated quasi-fixed point or to the unsegregated fixed point. To which point the system evolves is determined by the initial conditions. Because the fixed point structure must be globally consistent, we deduce that there must exist in this narrow regime a new set of saddle fixed points that partition the space into a region containing the unsegregated fixed point, to which any initial point in this region will flow, and m regions containing the segregated quasi-fixed points. The existence of this narrow transition region in which both the segregated quasi-fixed points and the unsegregated fixed point are simultaneously stable is all that appears to distinguish the dynamics of both the full and simplified models in the

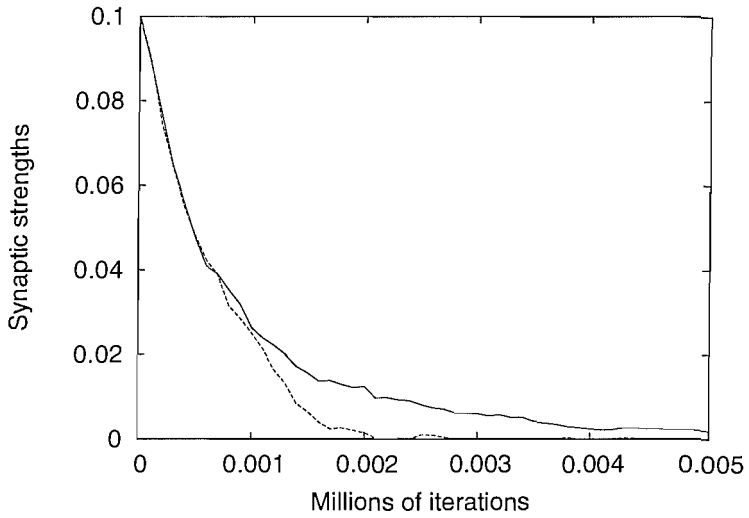

FIGURE 8.9: Spike-based simulation of two afferents, for large α, and $\gamma = 0.1$. The fixed-point at the origin is, effectively, stable and the afferents fall to zero. One iteration is equal to one second of simulated time.

small α_i, analytically-explored regime from the large α_i, numerically-explored regime. It is likely, in fact, that this transition region is present even for small α_i, but is so narrow as to be extremely difficult to observe numerically.

Fig. 8.9 shows the evolution of the spike-based, large α simulation, for $\gamma = 0.1$. At this low γ value, the fixed-point at the origin is, effectively, stable and the afferents fall to zero. Fig. 8.10 shows the medium γ regime, where the afferents engage in competitive interactions and evolve to a stable, segregated fixed-point. Fig. 8.11 shows the high γ regime where the only stable fixed-point is the unsegregated fixed-point, and the afferents evolve to a state of equal strength. Finally, Fig. 8.12 and Fig. 8.13 show the intermediate regime where stable fixed-points exist of both segregated and unsegregated character. The evolution of the system depends on the initial conditions, with the afferents segregating in one case and not in the other.

8.4 Computation in the Rate-Based Limit

To derive the n-spike, rate-based rules we integrated over the interspike intervals and averaged over all 2^n possible spike trains to compute an unconditional expectation value for the change in synaptic efficacy due to a typical n-spike train. The resulting rate-based rule is somewhat abstract, in the sense that the neuron does not really compute at the level of this rate-based rule, but continues to compute at the level of individual spikes. We may ask, however, whether there

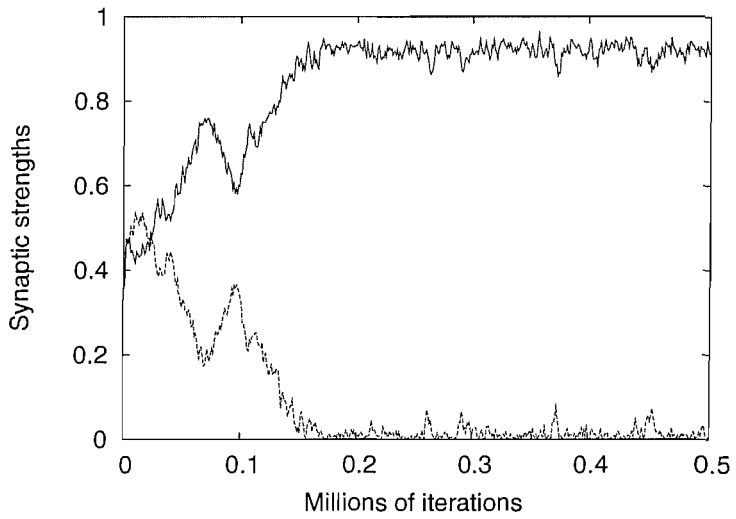

FIGURE 8.10: Spike-based simulation of two afferents, for large α, and $\gamma = 0.7$. The afferents engage in competitive interactions, and evolve to a stable, segregated fixed-point. One iteration is equal to one second of simulated time.

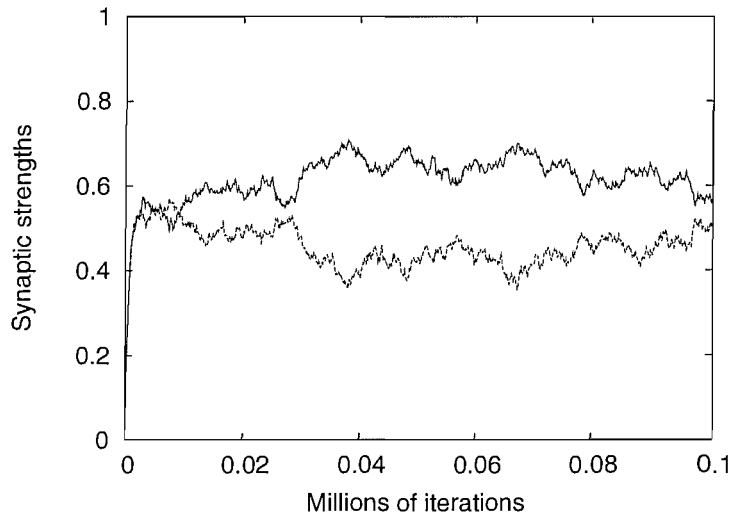

FIGURE 8.11: Spike-based simulation of two afferents, for large α, and $\gamma = 0.98$. The only stable fixed-point is the unsegregated fixed-point, and the afferents evolve to a state of equal strength. One iteration is equal to one second of simulated time.

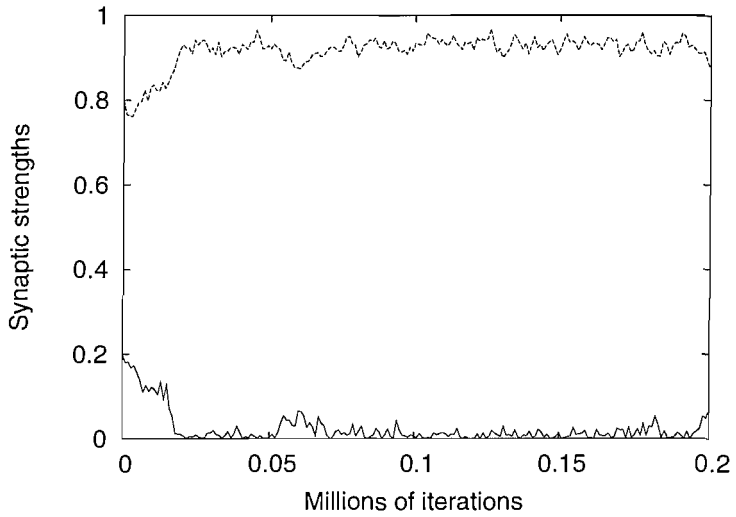

FIGURE 8.12: Spike-based simulation of two afferents, for large α, and $\gamma = 0.75$. Initial synaptic strengths are 0.8 and 0.2. Starting close to the segregated fixed-point causes the system to evolve to and remain there. One iteration is equal to one second of simulated time.

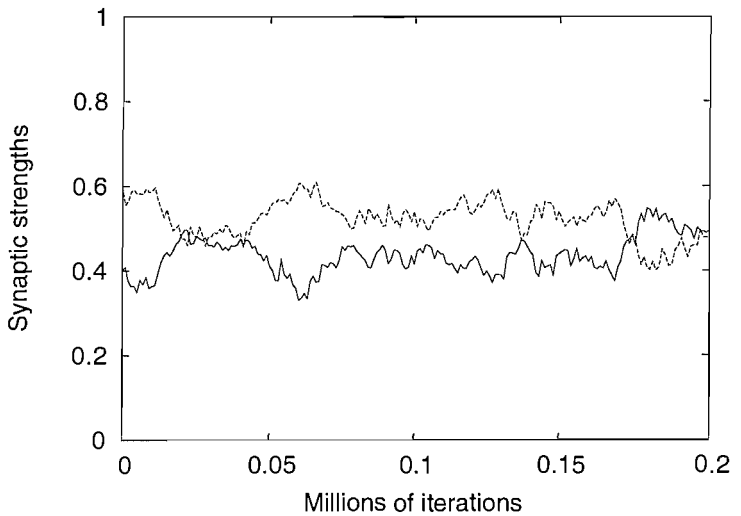

FIGURE 8.13: Spike-based simulation of two afferents, for large α, and $\gamma = 0.75$. Initial synaptic strengths are 0.6 and 0.4. Starting close to the unsegregated fixed-point causes the system to evolve to and remain there. One iteration is equal to one second of simulated time.

are any conditions under which, although the neuron is computing at the level of spikes, it nevertheless behaves as if it were following the rate-based, n-spike rule. If no such conditions exist, then our analysis is somewhat academic, since the derived rules would be merely mathematical abstractions to which the neuron's behaviour cannot ever approximate.

We can identify two limiting cases of interest. First, if the spike train is extremely long and not highly unusual, then we can think of the train as naturally decomposing into a set of shorter sub-trains. The neuron can then effectively average its behaviour over these sub-trains. If there are enough sub-trains, then sufficient averaging over them will occur in order for the neuron's mean dynamics to approximate the rate-based rules. However, the variance in this behaviour could be large, and these fluctuations could thus nevertheless prevent the emergence of stable, segregated states at the spike-based level. Second, we know that even for small n, $n \geq 3$, the rate-based rules exhibit stable, segregated states, and the above considerations for large n do not apply. Of course, if we present a neuron with a single instance of an n-spike train, n small, then we would not expect its synaptic strengths to evolve much during this short train. The neuron must therefore be presented with a long sequence of such n-spike trains, so that it can average over this sequence, and the trains must be sufficiently well separated that all the synapses return to the OFF state between trains. Again, the averaging will ensure that the mean behaviour is exhibited, but the fluctuations could destroy the stability. We therefore see that the key to stability is to ensure that the fluctuations are small.

Small fluctuations can be guaranteed provided that the neuron's dynamics are not dominated by the most recent spike train (or sub-train). If the most recent spike train essentially erases the neuron's state developed from exposure to earlier trains, then the neuron's behaviour will be dominated by train-to-train fluctuations. We therefore require that the change in synaptic strength induced by each spike train is small compared to the (non-zero) synaptic strengths. This can be achieved by setting $A_{\pm}$ sufficiently small, since these two parameters set the overall magnitude of plasticity.

Changing the magnitude of plasticity (the overall scale of $A_{\pm}$) is, of course, equivalent to changing the learning rate in a model's dynamics. The learning rate is essentially the step size in a numerical integration procedure. It is well-known that the stability of the numerical solutions of a set of differential equations depends critically on the step size, and that there usually exists a threshold above which the integration scheme fails to converge to the exact solutions, with chaotic instability or divergent behaviour ensuing. We should therefore not be too surprised that when the magnitude of plasticity is sufficiently small, the spike-based behaviour will be expected to converge to the rate-based behaviour.

Given the complexity of the switch model, however, determining the location of the threshold above which the neuron does not compute in the rate-based limit, and exhibits instead a strong dependence primarily on the last spike train, is a difficult matter. We therefore resort to a simple numerical search for the approximate location of this threshold.

To obtain a condition on the magnitude of plasticity below which the rate-based behaviour becomes dominant in a spike-based simulation, we determine when the spike-based system exhibits qualitatively the same fixed point structure known to exist in the rate-based system. We consider a system of two afferents for simplicity. A rate-based simulation of two afferents will stably segregate, with one afferent gaining complete control of the target cell, provided that $\gamma_1 < \gamma < \gamma_3$. Thus, we select a value of γ in this range and run a spike-based simulation for various values of the overall scale of $A_{\pm}$. In these simulations, in order to examine the n-spike rule, we present a series of n-spike trains to each of the afferents' synapses, and after every train force the synapses to return to the OFF state, which is equivalent to spacing the trains sufficiently far apart that they do not interact. Within each train, the afferents have their Poisson firing rates randomly fixed either "high" (75 Hz) or "low" (25 Hz). We perform an initial presentation of 2.5×10^7 spikes, partitioned into n-spike trains, in order to allow for sufficient time for the afferents to segregate on the target cell. At a typical, average rate of 50 Hz, this corresponds to approximately 6 days' worth of simulated synaptic activity, which is not too dissimilar to the typical time scale for developmental processes in the nervous system. After this initial period to allow time for segregation, we present another series of 2.5×10^7 spikes, again partitioned into n-spike trains, during which we probe the extent and stability of any segregation. After each train presentation, we calculate the segregation index, S_I, which we define to be

$$S_I = \frac{s_1 - s_2}{s_1 + s_2}. \tag{8.74}$$

If the afferents are well-segregated, then S_I will take values close to $+1$ or -1, depending on which afferent controls the target cell. If segregation is stable, then this index will not change much, except for small fluctuations. If the afferents are segregated, but not stably so, with control switching between the two afferents, then S_I will flip between $+1$ and -1. Averaged over sufficient trains, its value will be roughly zero. If the afferents are not segregated, but oscillate about a mean synaptic strength, then S_I will always be roughly zero, and of course its average will be roughly zero. Thus, for this probing phase, we determine $\langle S_I \rangle_P$, where $\langle \rangle_P$ denotes the average value of S_I during this second period. We then take the absolute value of this average, and average this value over 50 distinct runs for

each value of the overall scale of $A_{\pm}$. Thus, our final measure of segregation and stability is $\langle|\langle S_I\rangle_P|\rangle_R$, where $\langle\rangle_R$ denotes an average over runs.

In Fig. 8.14, we plot $\langle|\langle S_I\rangle_P|\rangle_R$ as a function of the overall scale of $A_{\pm}$ for 3-, 9- and 15-spike train simulations for the non-resetting model with $n_{\pm} = 3$ and $\gamma = 0.6$. We obtain qualitatively similar results for the resetting model and for different values of γ. Also shown in Fig. 8.14 is the fit of our raw data to logistic-like functions, $a - b\tanh(cx - d)$, where a, b, c and d are fitted parameters. That the fits to logistic-like functions match the raw data well indicates that the transition from stable segregation to unstable segregation is relatively sharp. For an overall scale of plasticity of approximately 10^{-3} or lower, depending on the number of spikes in the train, we observe robust and stable afferent segregation, while for a scale greater than this value, segregation is achieved, but is not stable, so that the afferents change their control of the target cell over time. For values of the overall scale very much greater than 10^{-3}, segregation is not achieved at all. We observe that as the train length increases, the mean segregation index increases for a fixed value of the overall plasticity magnitude. However, by about 10 to 20 spikes, asymptotic behaviour is reached, with no further increase in the index observed. This is in accord with our expectations, since the n-spike rate-based learning rules converge very rapidly as a function of n, with convergence achieved by $n \approx 10$ spikes Appleby and Elliott (2005).

The data in Fig. 8.14 are obtained by randomly fixing the afferents' Poisson firing rates within each spike train. For longer trains, therefore, each afferent fires for longer with the same firing rate. It could therefore be argued that the dependence of the magnitude of plasticity on the number of spikes in a train merely reflects this longer exposure to the same firing pattern. We can rule this out in two ways. First, instead of fixing each train's afferents' rates, we can instead fix each afferent's rate for a given period of time (a firing "epoch"), regardless of the number of spikes in each train. For a firing epoch length of 1000ms, for example, we obtain data that are essentially identical to those in Fig. 8.14 (data not shown). Second, we can, for example, consider 3-spike trains with fixed firing per train, but simply duplicate the firing patterns between consecutive trains, so that 2-train sequences of 3-spike trains have the same rates. Doing this, we produce data identical to the 3-spike data shown in Fig. 8.14 and not the 6-spike data with an enhancement in the segregation index (data not shown). Thus, longer exposure to the same activity patterns is not responsible for the trend exhibited in Fig. 8.14. Rather, the observed trend reflects the fact that as the number of spikes increases in a train, competition becomes stronger and stronger, as revealed by Eq. (8.28).

Although this numerical treatment is not exact, it is sufficient for our purposes to demonstrate that computation in the rate-based limit is realistically available to a

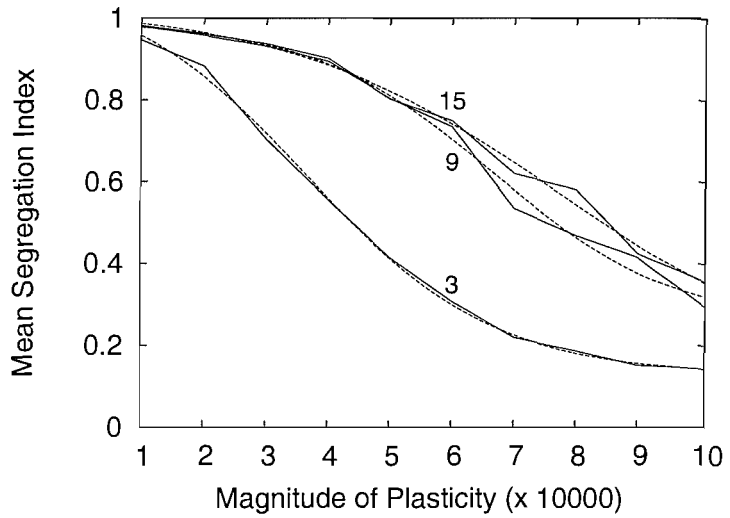

FIGURE 8.14: The dependence of the mean segregation index $\langle|\langle S_I\rangle_P|\rangle_R$ on the magnitude of plasticity, or the overall scale of $A_\pm$. Solid lines represent the numerically-obtained values of $\langle|\langle S_I\rangle_P|\rangle_R$ for the number of spikes in each train indicated by the attached number. The dashed lines show the fit of the raw data to logistic-like functions, as described in the main text.

real, spike-based system and to have an approximate idea of where that limit lies. Assuming that $A_\pm$ satisfy the numerically-obtained condition on the magnitude of plasticity below which the rate-based behaviour becomes dominant, we may dispense with spikes completely and work in the rate-based limit. The analysis of Section 4 was carried out in the rate-based limit and did not consider individual spikes. The stable, competitive dynamics that we observed will therefore be the dominant mode of computation provided that $A_\pm \sim 10^{-3}$ or lower.

8.5 Integrate and fire neuron

So far in our exploration of the switch rule we have assumed that postsynaptic firing is governed by a Poisson process with a firing rate given by the simple linear sum of input firing rates weighted by their connection strength. Postsynaptic spiking therefore occurs independently of presynaptic spiking. This was motivated partly by simplicity and the resulting analytical tractability, and partly by the desire to show that the dynamics of the model are not dependent upon the presence of the additional non-linearities, such as the hard resetting mechanism, present in an integrate-and-fire neuron. The assumption of Poisson postsynaptic firing is also, arguably, the most appropriate assumption to make when modelling cortical neurons that may possess many thousands of synaptic

inputs. In such cases the influence of any one individual input would be very small and ensemble input properties might reasonably be expected to dominate.

It is important, however, to show that this assumption is a realistic one, and that the direct correlations in pre- and postsynaptic firing that follow from an explicit model of postsynaptic spiking do not destroy any of our results. We therefore extend our simulations of the switch model to include an explicit integrate-and-fire postsynaptic neuron, similar in nature to those used in other theoretical studies of STDP (Song et al., 2000). If the assumption of Poisson spike trains is a reasonable one, we would expect to find that all the basic results presented in this Chapter are preserved. For example, we would expect simulation of 2-spike trains to display some form of pathological learning, and all higher-spike rules to display stable, segregation of afferents. However, given that the intergrate-and-fire neuron also contains various non-linearities such as a hard reset after each postsynaptic spike, we may expect to see some differences in the learning dynamics. We may indicate, however, which differences are due to the integrate-and-fire mechanism itself, and not due to the temporal correlations, by rerunning the simulations with Gaussian noise in the timing of postsynaptic spikes.

We implement an integrate-and-fire neuron identical in nature to that used in Song (2000) (see Chapter 6), with the exception that we use a more standard reset potential of -90mV. If we were to simulate a large number of afferents, then each individual afferent will exert a relatively small influence on postsynaptic spiking. Decreasing the number of afferents would increase the ability of individual afferents to directly influence the postsynaptic cell. As we are interested in showing whether the direct correlations between pre- and postsynaptic spiking due to the integrate-and-fire mechanism affects the learning behaviour, we deliberately simulate a low number of afferents and choose, as for the Poisson-based simulations, to simulate 4 afferents innervating a single target cell. We examine the 2-spike rule for low and high initial weights, with $\gamma < 1$, and the 3-spike and ∞-spike rules for the three γ regimes derived above. We define low γ to be 0.1, medium γ to be 0.7, and high γ to be 0.98. The behaviour of the model under these eight scenarios can be viewed as a test of whether the essential dynamics of the switch model are preserved.

Fig. 8.15 shows the simulation of 2-spike trains with an integrate-and-fire postsynaptic neuron, starting from all afferents at "normal" strength (0.15, roughly commensurate with that used in Song (2000)). We see that the dynamics are still dominated by runaway learning, in the sense that one afferent potentiates uncontrollably. However, we note that the potentiating afferents seems to suppress the other inputs, a behaviour that does not occur in the original Poisson-based simulations.

The integrate-and-fire neuron is known to introduce competition when combined with an overall depressive STDP window ($\gamma < 1$ in our notation) (Song et al., 2000). It has been argued elsewhere that this is because the integrate-and-fire mechanism allows one input to effectively control postsynaptic spiking, so that a spike arriving at this input always evokes a postsynaptic reply. Such an input is therefore always potentiated. Other inputs, on the other hand, hardly influence postsynaptic spiking at all, and, since they are uncorrelated with the dominant afferent, they will experience the "average" of the STDP window. As $A_+\tau_+ < A_-\tau_-$, so that $\gamma < 1$, these inputs are depressed. We note, however, that the introduction of Gaussian temporal noise to the timing of postsynaptic spikes does not destroy the result shown in Fig. 8.15. In other words, removing the temporal correlations does not restore the Poisson result. Thus, the competitive, "suppressive" interaction is, in fact, produced by some more complicated interaction between the STDP rule and the integrate-and-fire neuron, not simply from the introduction of temporal correlations as argued elsewhere (Song et al., 2000). In the 2-spike simulation of Fig. 8.15 this suppressive interaction prevents the potentiating afferent from pulling up the other inputs. In contrast, the Poisson-based simulations has neither temporal correlations nor the non-linearities of the integrate-and-fire mechanism, so this suppressive dynamic is absent. A high postsynaptic firing rate therefore produces, on average, potentiation for all inputs, which are consequently rescued from zero.

The case of runaway depression, where all afferents fall to zero, has no analogue in the integrate-and-fire case as when afferent weights become small, but non-zero, the postsynaptic neuron stops spiking altogether and afferents cease to be modified. We therefore do not show this case. We cannot try to remedy this by increasing synaptic strengths a little because, when the afferents become strong enough to evoke postsynaptic spiking, we immediately begin to experience the suppressive dynamics discussed above.

Fig. 8.16 shows the simulation of 2-spike trains with an integrate-and-fire postsynaptic neuron, starting from all afferents at "high" strengths (15, ten times that used above). The dynamics differ to that of the Poisson-based simulation, in that we do not see runaway potentiation, and also differ to the "normal" strength case above, in that we do not see competitive interactions. These differences are not, however, due to a genuine difference in the dynamics of the temporally correlated system compared to the Poisson-based or "normal" strength system, but is a result of the integrate-and-fire mechanism itself. When the postsynaptic firing rate gets very high, the details of the neuron model, such as the choice of reset potential, refractory periods and so on, begin to exert an important influence on the dynamics. When all afferents have high strengths, the postsynaptic neuron is spiking and being reset almost constantly (the injections to the excitatory

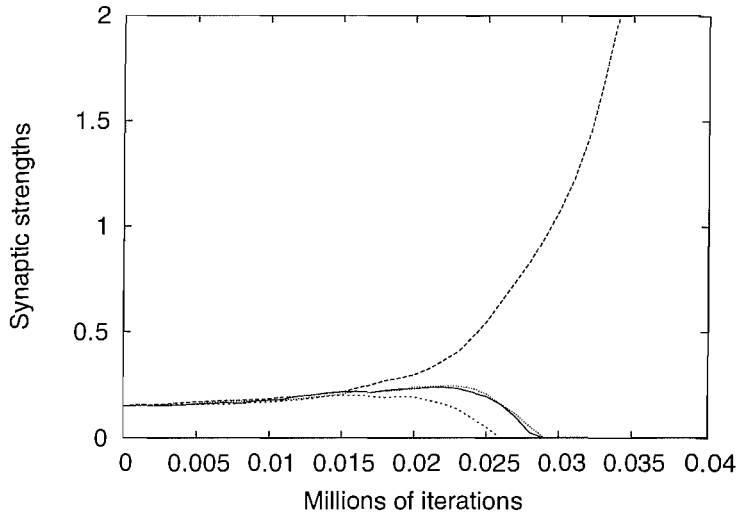

FIGURE 8.15: A spike-based simulation of the 2-spike rule for four afferents innervating one target cell, with an explicit simulation of postsynaptic spiking via an integrate and fire neuron. We start the afferents from normal strengths, as detailed in the text. As for the Poisson-based simulation, we see dynamics dominated by a form of runaway potentiation. However, we note that the integrate-and-fire neuron has introduced an additional suppressive interactions that forces all other afferents to zero. One iteration translates to one second of simulated time.

conductance persist for some time following a spike) and we begin to probe the machinery of the integrate-and-fire mechanism. In the 2-spike case this causes inputs to stabilise and prevents competition, but this result depends largely on the particular choice of implementation. A different choice of coupling pre- and postsynaptic spiking together may well have led to very different results. Again, we find that the introduction of Gaussian temporal noise to the timing of post-synaptic spikes does not destroy this result.

In summary, we see that simulations of 2-spike trains are still pathological in the integrate-and-fire model, and a form of runaway learning is still observed. We see, however, that the integrate-and-fire neuron has introduced an additional competitive mechanism in the form of a new, suppressive interaction. For normal synaptic strengths, this allows one afferent to potentiate away and suppress the others to zero. When high synaptic strengths, we see non-trivial behaviours stemming from the actual integrate-and-fire mechanism itself. This serves to stabilises the afferent in some manner, but we note that this would not necessarily occur under a different model of postsynaptic spiking. Repeating these simulations after introducing Gaussian noise into the timing of spikes, so that the actual temporal correlations become minimal, does not affect any of these results. We therefore see that these differences are due to the integrate-and-fire

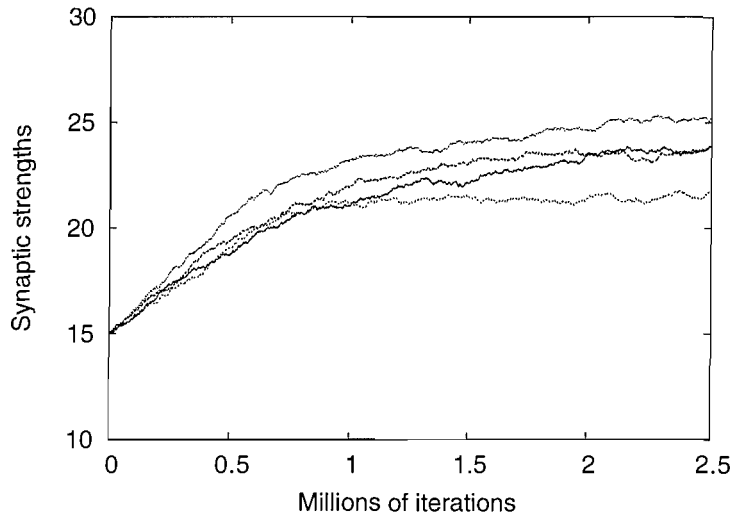

FIGURE 8.16: A spike-based simulation of the 2-spike rule for four afferents innervating one target cell, with an explicit simulation of postsynaptic spiking via an integrate and fire neuron. We start all afferents from high strengths. We see, in contrast to the Poisson-based simulation, that the afferents are stabilised and do not experience runaway potentiation. This arises due to non-trivial response of the spiking mechanism to a sustained level of high presynaptic input. One iteration translates to one second of simulated time.

neuron itself rather than the temporal correlations it introduces, and that the basic behaviour of the switch model is preserved.

Under the 3-spike rule, competition is already present in the switch model, arising from multi-spike interactions not present under the 2-spike rule. The suppressive dynamics introduced by the integrate-and-fire neuron might therefore be expected to play a less important role. The low and medium γ regimes for the non-resetting 3-spike rule are shown in Figs.8.17 and 8.18, respectively. For low γ we see that all the afferents fall towards zero, which we know acts as a stable fixed-point in this system. However, as for the 2-spike simulations, the afferents cannot ever reach zero due to the integrate-and-fire neuron shutting down, and instead they reach some lower bound on total synaptic strength then stop evolving. For medium γ, we see that, as for the Poisson-based simulations, the afferents engage in competitive interactions and reliably segregate, with all but one falling to zero. The system thus evolves to a stable, segregated fixed-point as expected.

Fig.8.19 shows the high γ regime. We expect to see the afferents evolve to the stable, unsegregated fixed-point. In other words, we expect to find that, in the high γ regime the switch rule is non-competitive. However, we find that this behaviour is destroyed, presumably because the competition which is now

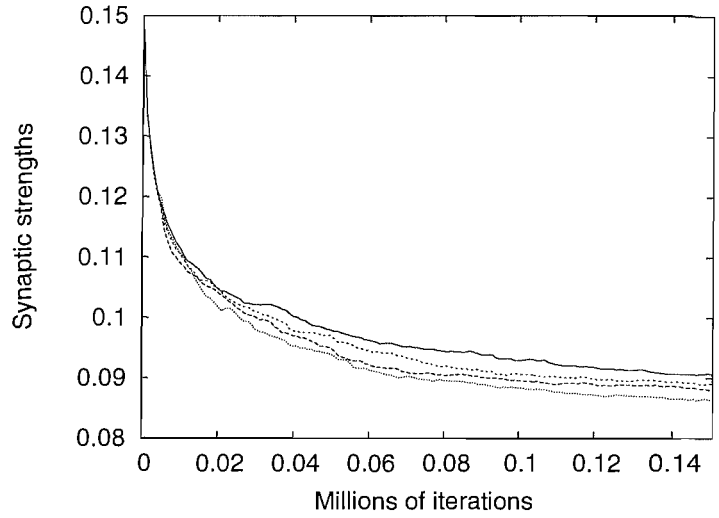

FIGURE 8.17: A spike-based simulation of the non-resetting 3-spike rule for low γ, with four afferents innervating one target cell and an explicit simulation of postsynaptic spiking via an integrate and fire neuron. As for the Poisson-based simulation, the afferents fall to zero. However, as the integrate-and-fire mechanism shuts down when afferent weights become small, the afferents stop evolving when total synaptic strength reaches some lower threshold. One iteration translates to one second of simulated time.

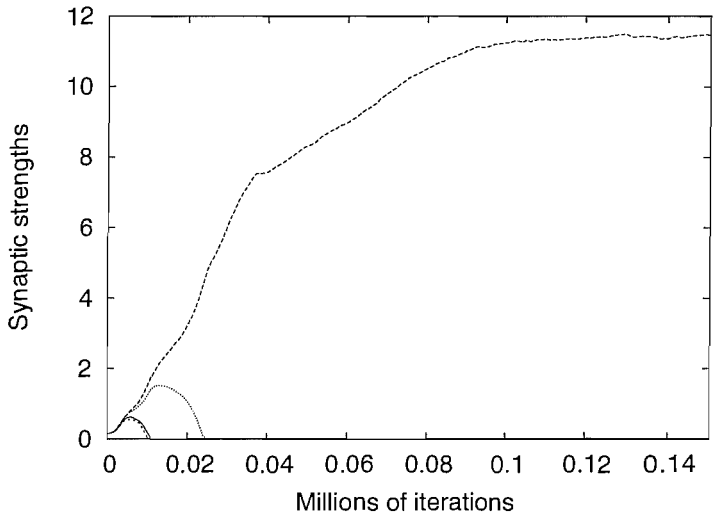

FIGURE 8.18: A spike-based simulation of the non-resetting 3-spike rule medium γ, with four afferents innervating one target cell and an explicit simulation of postsynaptic spiking via an integrate and fire neuron. As for the Poisson-based simulation, the rule is competitive, with stable segregation of the afferents being observed. One iteration translates to one second of simulated time.

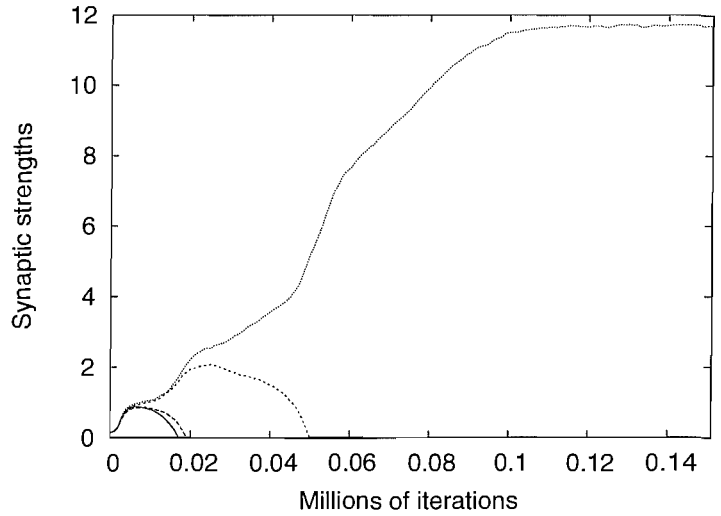

FIGURE 8.19: A spike-based simulation of the non-resetting 3-spike rule high γ, with four afferents innervating one target cell and an explicit simulation of postsynaptic spiking via an integrate and fire neuron. The suppressive dynamics of the integrate-and-fire neuron play a strong role and, unlike the Poisson-based simulation, we do not see the afferents evolve to the stable, unsegregated fixed-point. Instead, dynamics similar to the medium γ case are produced. One iteration translates to one second of simulated time.

absent from the switch is compensated for by the suppressive dynamics of the integrate-and-fire neuron, producing behaviour that resembles the medium γ regime. Again, we find that this behaviours is not destroyed by introducing Gaussian noise into the timing of postsynaptic spiking. It therefore arises due to the integrate-and-fire mechanism itself (specifically, the suppressive dynamics it introduces) rather than the temporal correlations in spike times.

We find that the observations made for the $n = 3$ case hold for all multi-spike simulations, including the ∞-spike rule. In particular, we find that, as for the Poisson-based case, we observe stable, competitive dynamics in the medium γ regime for all n-spike, $n > 2$ rules.

The use of an integrate-and-fire neuron, or any explicit model of postsynaptic spiking, introduces many additional features beyond the temporal correlations we seek to explore. These additional features can have strong implications for the dynamics observed. We have shown that these new dynamics are not due to the effect of temporal correlations on our model, but due to the integrate-and-fire neuron itself, by the observation that all of the differences are preserved even when we effectively remove the temporal correlations by introducing Gaussian noise into the postsynaptic spike times. A more detailed study, which would be a topic of future work, would seek to address this issue in a more controlled manner by introducing temporal correlations in spike timing without the need

to also introduce the spiking machinery associated with an integrate-and-fire mechanism.

8.6 Large-Scale Numerical Simulations

In Section 4 we established that the ∞-spike, non-resetting rule possesses a fixed point structure consistent with the presence of stable, competitive dynamics in the model. The presence of such dynamics has also been argued for in all multi-spike rules for an appropriate choice of parameters. In Section 5, we established that the rate-based multi-spike dynamics are actually available to neurons in the sense that, although computing at the level of single spikes, the neuron's dynamics will respect the fixed point structure of the rate-based rules and exhibit stable, competitive dynamics. We are therefore now in a position to consider a large-scale simulation of, for example, ODC development, in order to demonstrate that the results shown for a single target cell scale up without difficulty to multiple target cells. Moreover, while for convenience we considered above only uncorrelated afferent activity patterns, a simulation of ODC development will permit us to explore negatively- or positively-correlated afferent activity patterns too. For ODC development in the cat, which develops ODCs in the presence of presumably positively-correlated afferent activity after eye-opening, the case of positive correlations is more pertinent. This is also usually the much harder task: models that can segregate negatively-correlated afferents frequently cannot segregate positively-correlated afferents. For our switch model to be a candidate model of developmental synaptic plasticity, we must therefore show that it operates successfully in the face of positively-correlated afferent activity patterns.

We run ODC simulations according to existing, documented protocols Elliott and Shadbolt (1998). We use the rate-based, non-resetting ∞-spike rule with $n_{\pm} = 1$ as the synaptic modification rule. In brief, we simulate two patches of retinotopically equivalent lateral geniculate nucleus (LGN), each containing a square array of cells of size 13×13, with periodic boundary conditions imposed for convenience. The cortex is a 25×25 square array of cells, again with periodic boundary conditions. Each LGN cell arborises over a retinotopically appropriate patch of cortex, of size 7×7. LGN activity patterns are constructed according to the method of Goodhill (1993), taking the form of Gaussian correlated noise. The parameter $p \in [0, 1]$ determines the activity correlations between the two LGN patches, with $p = 0$ corresponding to perfectly anti-correlated patterns and $p = 1$ corresponding to perfectly correlated patterns. Cortical activity is just the standard linear sum of afferent input, but smeared by convolving cortical activity with a short-range Gaussian function. Such smearing could be achieved,

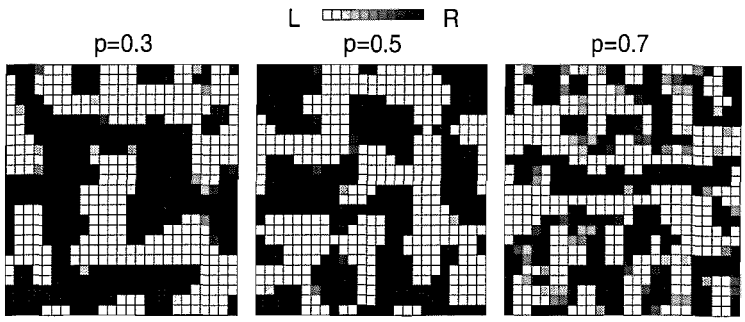

FIGURE 8.20: A simulation of ocular dominance column formation in the non-resetting, ∞-spike model with $n_{\pm} = 1$ for the three values of the inter-ocular correlation probability shown above each map. In these maps, each square represents a single cortical neuron. The shade of grey assigned indicates the relative degree of control by the two eyes. A white square indicates complete control by the left eye, while a black square indicates complete control by the right eye. Shade of grey, as shown in the key, interpolate between these two extremes.

for example, by lateral excitation. In this way nearby cortical cells fire similarly, which is necessary in order to develop structured ODCs rather than a pattern of salt-and-pepper segregation. It is well known that presynaptic constraints must typically be introduced in order that the pattern of ODCs exhibits some degree of regularity, with fairly constant stripe widths across the whole simulated patch of cortex. To achieve this, we introduce a lower bound on the total synaptic strength that an afferent may support. If it falls below 75% of the average total strength supported by all afferents, we simply freeze depression. We use this method because it is convenient and removes the need for us to calculate explicitly the expected total afferent strength in such a simulation, and then set a lower bound accordingly.

In Fig. 8.20 we show three ODC maps corresponding to three different values of p: $p = 0.3$ representing negatively-correlated activity patterns; $p = 0.5$ representing uncorrelated activity patterns; $p = 0.7$ representing positively-correlated activity patterns. We see clear patterns of ODCs in all cases, with well-segregated afferents, although with increasingly binocular boundaries between ODCs as p increases, as expected. We also observe a clear decrease in the widths of ODCs as the inter-ocular correlations increase. This phenomenon was first observed by Goodhill (1993) in simulation, and subsequent experimental results supported the possibility that the widths of ODCs are not fixed, but may be partially determined by visual experience Löwel (1994); Tieman and Tumosa (1997). Our results here show that our switch model of STDP scales up to a large-scale simulation without difficulty, and can comfortably segregate afferents whose activities are positively correlated.

8.7 Discussion

In this Chapter we have studied the synaptic dynamics induced by the 2- and multi-spike rules that are derived within the context of a stochastic model of STDP. The structure of the postulated unified 3-state synaptic switch, which was initially designed only to account for the basic phenomenology of the STDP learning rule seen in the context of interactions between two spikes, tightly constrains the form of the multi-spike interactions that exist in our model. One freedom that the model possesses is to allow stochastic process resetting in the *DEP* and *POT* states. Although resetting changes the precise form of the multi-spike interaction functions, the dynamics remain qualitatively unchanged.

We find that a broad parameter regime exists in which the rate-based, 2-spike learning rule is qualitatively BCM-like. This BCM-like form gives rise to instabilities in the 2-spike rule. Indeed, without further modification, the 2-spike learning rule is irredeemably pathological, sending all afferents either to zero strength or to unbounded growth. The 2-spike learning rule cannot therefore support the stable, competitive dynamics that are essential within the context of developmental synaptic plasticity. However, we find that all n-spike interaction functions, where $n > 2$, arising exhibit dynamics that differ significantly from those exhibited by the 2-spike interaction function. For a very broad range of parameters all these rules exhibit stable, competitive dynamics without any further modification.

Thus, extending our considerations from two spikes to three spikes immediately solves the inherent problems of the 2-spike rule. This result is quite remarkable. We argued that this is because in the unified switch rule, depression and potentiation are coupled in the sense that activation of the potentiation lobe of the switch precludes a simultaneous activation of the depression lobe of the switch, and vice-versa. By breaking the two lobes into two separate switches, so that they may be simultaneously active, we showed that competition immediately breaks downs, resulting in the runaway learning characteristic of the 2-spike rule. Critical to probing this coupling between potentiation and depression is the presence of three or more spikes; even in the unified switch, two spikes cannot ever probe this couplings, and so 2-spike interactions always induce uncoupled potentiation and depression dynamics. It therefore appears that this coupling, or interaction, between potentiation and depression processes at the level of a single neuron is vital to the presence of stable, competitive dynamics.

In the above analyses, we have truncated s_i at zero when it is driven negative. We saw, for example, that although the zero fixed point is a saddle for $\gamma < \gamma_2$ in the simplified ∞-spike resetting model, nevertheless, the truncation procedure dynamically converts it into a stable node when $\gamma < \gamma_1$. Moreover,

the truncation gave rise, partly, to the existence of quasi-fixed points. It might therefore appear that the truncation procedure is playing an important role in our model's dynamics, being partly responsible for the presence of stable, competitive dynamics in it. This is not, however, the case. We can see this in two different ways.

First, from a mathematical point of view, we set $ds_i/dt = \Delta S_n(\lambda_{\pi_i}, \lambda_p)$ and used this as our synaptic strength update rule. This is not the only choice available to us. In this formulation, s_i represents the average synaptic strength over all of afferent i's synapses on the target cell. Suppose, however, that we instead think of s_i as the (possibly scaled) number of synapses supported by afferent i, each of which experiences the single-synapse learning rule. Then the overall change in afferent i's synaptic strength will be the individual change multiplied by the number of synapses, s_i. Hence, on this view, we would instead write $ds_i/dt = s_i \Delta S_n(\lambda_{\pi_i}, \lambda_p)$. The presence of the extra factor of s_i in this rule immediately means that synapses cannot evolve below zero, and thus the quasi-fixed points are converted into real fixed points, with their location and stability completely unchanged. Furthermore, it is easy to see that the zero fixed point in this formulation is genuinely stable for $\gamma < \gamma_1$. We further note that truncation at zero is a hard non-linearity to which the afferents are completely insensitive when $s_i > 0$. Thus, their evolution to the vicinity of the segregated quasi-fixed points in the interval $\gamma \in (\gamma_1, \gamma_3)$ is entirely independent of truncation: it is a dynamical consequence of the overall structure of the model. The need for truncation, then, is merely a technical issue that does not fundamentally modify the model's prior dynamics.

Second, from a biological point of view, when a synapses reaches zero or close-to-zero synaptic strength, we expect it either to shut down and evolve no further, or to be retracted entirely. Biologically, an excitatory synapse of zero strength cannot turn into an inhibitory synapse. Our derivation of the learning rules, however, takes no account of this, for the purposes of tractability. However, the simplest remedy to this situation is merely to prevent a synapse from depressing when such depression would takes its synaptic strength negative. (Indeed, this it the presynaptic constraint that we used to model ODC development, in order not to obtain irregular patterns of afferent segregation, except that the lower limit on synaptic strength was set at some non-zero value.) This prevention of depression is tantamount to switching off the depressing lobe of the switch, perhaps by setting $A_- = 0$, when the synapse is too weak to permit further depression. In such a state, it could be potentiated, but not depressed. This, though, is essentially equivalent to truncation of strengths at zero. The key, competitive dynamics in the model are those that precisely allow a synapse to

move towards a quasi-fixed point near which truncation, or the closing down of depression, then becomes necessary.

We observe three essentially distinct parameter regimes in our model's dynamics, corresponding to different ranges for the parameter γ. For low γ, all afferent strengths fall to zero. For high γ, afferents evolve to an unsegregated fixed point in which all afferents equally control the target cell. Only for intermediate values of γ do we see the evolution of the system to stable, segregated final states. Considered within the context of a simulation of ODC formation, the low- and high-γ regimes would correspond to a breakdown of ODC formation, while the intermediate regime would correspond to the normal development of ODCs. Do these three regimes correspond to real experimental situations, or are they just mathematical fictions? Interestingly, the infusion of the neurotrophic factors brain-derived neurotrophic factor (BDNF) or NT-4/5 (Cabelli et al., 1995) or the blockade of their common, endogenous receptor trk-B (Cabelli et al., 1997) both result in the abolishing of ODC development. In the former case, autoradiographic labelling reveals a higher-than-normal labelling, while in the latter case, the labelling is lower than normal. One interpretation of these results is that BDNF or NT-4/5 infusion causes a growth of afferent axonal arbors, while removing available BDNF or NT-4/5 from afferents causes their axonal arbors to atrophy. A similar influence of neurotrophic factors on axonal branching is also observed in the frog retinotectal system (Cohen-Cory and Fraser, 1995). Low and high γ, which cause the breakdown of ODC development, could therefore correspond to these experimental regimes. Moreover, low γ corresponds to very weak synapses while high γ corresponds to strong synapses. Under an anatomical interpretation of synaptic strength, these would correspond to small and large axonal arbors, respectively. Intermediate values of γ then correspond to normal patterns of development. Given the capacity of neurotrophic factors rapidly to modulate synaptic transmission in the visual system (Akaneya et al., 1996; Carmignoto et al., 1997; Sala et al., 1998), we can thus imagine a scenario in which neurotrophic factors dynamically determine the parameter γ in our model of STDP.

We may offer an explanation of why the higher-spike rules emerging from the switch model are intrinsically stable, and why they do not suffer from runaway learning as with the 2-spike rule by examining the learning surface that each rule encodes. We demonstrated in Chapter 7 that, for a low numbers of spikes, we may examine the nature of the terms that make up the learning rule, ΔS_n explicitly, deconstructing the calculation of ΔS_n and examining the contribution and spike–probability weighting of each term. We find that the $n > 2$ learning rules, and their resulting learning surface, are asymmetric. This is because all higher-spike

rules contain asymmetric terms that cannot be reduced to sums of simple spike-pairs. For example, the 3-spike learning rule contains the processes $\pi p\pi$ and $p\pi p$ which give rise to the irreducible triplet terms $A_+K_1^+K_1^-$ and $-A_-K_1^+K_1^-$. Given that we understand the origins of the n-spike learning surfaces, we may explain the dynamics of any rule simply by examining the form of its learning surface at a point of interest.

For the 2-spike surface, when β is low, the afferents lie in a depressing regime around the origin. Afferents will therefore fall to zero. When β is high, on the other hand, the afferents always lie in a potentiating regime. Afferents governed by the $n = 2$ rule therefore will suffer from a form of runaway learning, with potentiated inputs driving λ_p (and therefore β) higher and higher and, as a result, experiencing more and more potentiation.

Examining the $n = 3$ surface we see that the symmetry present in the $n = 2$ case is absent. For the $n = 3$ rule, we may either have potentiation or depression in the high β regime, depending on the balance of pre- and postsynaptic firing rates. We will, therefore, not necessarily see runway potentiation at high synaptic strength. Consider, for example, the dynamics of the system along a line of constant λ_π. This line corresponds to a fixed presynaptic firing rate. Initially, synaptic strength are low. As the learning surface in this region is positive, synapses will be potentiated and the postsynaptic firing rate, λ_p, will increase. As the postsynaptic firing rate increases we pass into a region where the learning surface is negative. Synaptic strengths are therefore depressed and, unlike in the 2-spike case, the system is stabilised. We therefore do not experience runway potentiation.

Thus, we find stable, competitive dynamics emerge from our switch model that are able to segregate uncorrelated inputs. Although many models of competition perform similarly, it is relvent to ask if such a process is biologically useful, or even desirable. There exist several examples in biology of competitive interactions which occur between apparently uncorrelated inputs. For example, the development of ocular dominance columns in monkeys and ferrets is initiated prior to eye-opening, and therefore begins prior to the appearance of correlated firing patterns between the two eyes. Despite the lack of correlations between inputs, however, competitive interactions still occur and afferents are able to segregate into zones of ocular control in the primary visual cortex (Hubel and Wiesel, 1962). Likewise, the activity patterns of afferents during synaptic rearrangement in the development of the neuromuscular junction are also apparently unstructured, yet afferents reliably engage in competitive interactions to win control of muscle fibres (Sanes and Lichtman, 1999). From a theoretical point of view, it is well known that the overall correlation strength plays an important role in determining whether or not segregation occurs (von der Malsburg, 1973;

Kempter et al., 1981; Goodhill and Barrow, 1994; Miller and Mackay, 1994; Elliott and Shadbolt, 1996). These studies have shown that, if a model is unable to segregate uncorrelated inputs, it will not, in general, be able to segregate correlated inputs such as those present after eye–opening in the vertebrate visual system. Thus, segregation of uncorrelated inputs is necessary if a model of competitive plasticity is to explain various well characterised developmental processes.

8.8 Conclusions

In conclusion, we have extended our earlier analysis of a stochastic model of STDP (Appleby and Elliott, 2005) to include an examination of multi-spike interactions. We have found that a consideration of multi-spike interactions is sufficient to endow our model with a fixed point structure consistent with the presence of stable, competitive dynamics. In contrast, at least in our formulation, 2-spike interactions by themselves cannot give rise to competitive dynamics. Multi-spike interactions therefore appear critical to understanding the presence of stable, competitive dynamics under STDP.

Chapter 9

Discussion

STDP has attracted a great deal of interest from the theoretical community since it was first observed in the late 1990s (for review, see Roberts and Bell, 2002). This interest is due, in part, to the potential for new forms of computation arising from the dependence of the plasticity on the timing of pre- and postsynaptic action potentials, as well as the prospect that a spike-timing framework may be able to account for a range of different, apparently distinct, forms of synaptic plasticity.

STDP is apparently widespread, and a range of different STDP-like learning curves have been observed in experimental work Bell et al. (1997); Zhang et al. (1998); Bi and Poo (1998). The heterogeneity of STDP rules suggests that STDP could possibly underlie several distinct functions in the nervous system. The computational properties of STDP are therefore of potentially quite wide reaching consequence. In theoretical studies, a variety of phenomenological and biophysical models of STDP have been proposed. Phenomenological models typically approximate the two phases of the STDP curve by two exponential functions with different amplitudes, polarities and decay constants then apply the resulting rule directly to determine changes in synaptic strength. In conjunction with certain constraints, this method can give rise to stable distributions of synaptic efficacies with competitive dynamics either emerging directly or introduced by synaptic scaling (Song et al., 2000; van Rossum et al., 2000; Izhikevich and Desai, 2003). Biophysical models do not take the STDP curve over directly, but rather attempt to derive STDP from a more detailed, biophysically plausible analysis of the molecular machinery present at the synapse (Senn et al., 2000; Castellani et al., 2001; Karmarkar and Buonomano, 2002; Shouval et al., 2002; Saudargine and Porr, 2005). Typically, these models employ the idea that NMDA-receptor-dependent calcium dynamics serve as a molecular coincidence detector. Both phenomenological and biophysical approaches have had some success in reproducing the basic STDP (Bi and Poo, 1998) and multi-spike (Froemke and Dan,

2002) results, and connections to the BCM-rule have been made (Izhikevich and Desai, 2003).

In Chapter 6 we proposed a new model of synaptic plasticity. This model is based on the idea that the synapses making up a connection can exist, independently of the other synapses, in one of three functional states. Pre- and postsynaptic spiking causes the synapse to transition between these states, with certain transitions associated with changes in synaptic strength of fixed magnitude. The transitions that return the synapse to the resting state (which we labelled the OFF state) occur in a stochastic manner. This simple, 3-state, stochastic switch, when averaged across multiple synapse and multiple spike-pairings, gives rise to an overall change consistent with the experimentally observed STDP curve. The rule is robust under highly variable spike timings, and, as the changes in synaptic strength are of fixed magnitude, removes the need for precise coincidence detection machinery that is capable of translating spike timing differences into graded changes in synaptic strength. An explanation of spike triplet interactions also emerges as a natural consequence of the switch, with no need to introduce additional constraints.

We examined the effect of multi-spike interactions, and derived the various multi-spike, rate-based learning rules that emerge from the switch model in Chapter 7. We saw that a Bienenstock-Cooper-Munro (BCM)-like (Bienenstock et al., 1982), rate-based plasticity rule emerges directly from the switch model for all n-spike rules. In Chapter 8 we explored the learning dynamics of the multi-spike rules, and found that the 2-spike rule displays pathological learning behaviours. We also found that, by extending our consideration to more than two spikes, we could resolve these problems without the need to introduce additional modifications designed explicitly to ensure stability. We now discuss in more detail the various issues surrounding the switch model, and its relation to other models of STDP.

9.1 Interpreting Experimental Data

Any model must ultimately be compared to some form of experimental result. This comparison typically take place during the models initial formulation and also later, in the form of comparison of predictions to new experimental results.

If a model is to be compared to a particular experimental result, then that result must be carefully interpreted and critically evaluated. This need to critically evaluate the data is especially important in biology, as experimental preparations are often unavoidably rather different from the natural system the experimenter sets out to study. For example, as discussed in Chapter 2, acute cortical slices offer a way of gaining easy access to neuronal assemblies which are amenable

to a variety of manipulations. However, the vast loss of connectivity, trauma of slicing, silencing of background activity, and numerous other factors mean that the preparation is, in some senses, far removed from the "ideal observer" of cortical function that the experimenter seeks. Other, apparently uncontroversial decisions, such as the choice or balance of constituents of a buffer solution, can have unexpected consequences for neuronal function, and can possibly lead to misleading results.

In biology, then, just as in other sciences, experimental results must be carefully evaluated and taken within their experimental context. Consider, for example, the STDP results discussed in Chapter 2. Despite the variety of experimental techniques, methodologies, and stimulation protocols, some common issues arise in the majority of STDP experiments. In particular, in order to generate statistically meaningful changes in the measured postsynaptic response repeated pairing of some form of pre- and postsynaptic stimulation is required. It is also usually the case that multiple inputs are stimulated, both to increase the magnitude of response and to compensate for often low levels of synaptic connectivity (but see Petersen et al., (1998) and Wang et al., (2005)). Taken together with the general experimental issues already mentioned, it is clear that plasticity experiments typically produce a very particular kind of result - a population or ensemble measurement taken from neurons that have been considerably disrupted compared to their normal operating environment.

These limitations are, of course, to be expected in so technically challenging endeavor as measuring synaptic plasticity. A great deal of skill is required to produce usable results from even the simplest plasticity experiment. This does not mean, however, that these limitations may be ignored or the issues arising from them left unaddressed. Indeed, one of the principle motivations for pursuing a theoretical study is to interpret the data, draw attention to experimental issues, and explicitly examine any underlying assumptions of the theory. In contrast, accepting results at face value is likely to lead to models which are simply reparameterisations of the data, rather than those that can offer valuable insights into the phenomenon in question.

In theoretical work on neuronal plasticity, many of these issues are often ignored. Experimental results are often interpreted literally without a critique of the various issues surrounding them. Such a literal interpretation typically carries with it implicit assumptions that can, in fact, turn out to have important theoretical consequences. In the case of plasticity, it is typically assumed that experimentally observed plasticity rules, such as STDP, despite being measured at the ensemble and population level, are valid for each individual afferent (and, indeed, at each individual synapse) for every spike pair. The assumption that

synapses are capable of implementing STDP-like computation has become almost standard in the literature, and a number of modelling studies have been put forward on this basis (Song et al., 2000; van Rossum et al., 2000; Senn et al., 2000; Shouval et al., 2002; Karmarkar and Buonomano, 2002; Saudargine and Porr, 2005). These modelling studies have had some success in explaining various STDP and rate-based results (Izhikevich and Desai, 2003). It is interesting to note, however, that learning rules formulated on a literal interpretation of STDP data often struggle in even the early states, and especially when attempts are made to generalise the models to accommodate experimental results beyond those involved in the models initial formulation.

In our own studies, we find that a critical re-evaluation of STDP results allows an alternative, and equally valid, formulation of STDP in which apparently complicated spike-timing computation emerges due to the averaging of a much simpler synaptic rule. This synaptic rule explains a range of spike- and rate-based results, and greatly reduces the computational burden placed on individual synapses. The model also produces specific predictions regarding the response of inputs to multiple spike trains, which provides a specific focus for further experimental studies seeking to distinguish our proposed model from competing theories.

9.2 Emergent Computation

The phenomenological studies of STDP discussed in Chapter 4 typically set out to formulate a STDP-like learning rule then examine the computational consequences of that rule. Biophysical studies take a slightly less direct approach, typically proposing some biologically plausible mechanism that could underpin STDP rather than modelling the data directly. In both cases, a mechanism must exist at the synapse level that is capable of making graded changes based on spike-timing differences.

The computational properties of both types of learning rule are often heavily dependent on the details of its implementation. In both types of model, STDP is "built in" in the sense that they contain an explicit representation of the STDP curve. In the phenomenological models, it is immediately apparent that the two phase modification function used to adjust synaptic strength is this explicit representation (Song et al., 2000). In biophysical models, a representation of the STDP is often implemented in a less direct manner. For example, the NMDA-receptor model of Senn et al., (2000) uses a dynamic population of NMDA-receptors to record spike timing and make graded changes in synaptic strength. The population of NMDA-receptors, which decay back to some resting state, therefore act as a discretised counting device which together implement a STDP

curve. In the Shouval model (Shouval et al., 2002) a biphasic learning function is postulated which translates directly into a biphasic STDP curve.

Naturally, any model that contains an explicit representation of the STDP curve is guaranteed to explain the basic spike-pairing results of STDP. However, this does not mean that the model will generalise to accommodate other experimental results. The generalisation of existing models of STDP to accommodate, for example, spike-triplet or rate-based results is often only achieved through the introduction of additional non-linearities such as spike suppression (Froemke and Dan, 2002) or through the temporal limitation of spike interactions (Izhikevich and Desai, 2003). Although such modifications are often very successful (as would be expected given the empirical nature by which the modifications are made) this success does not generally extend beyond the immediate experimental result in question. For example, introducing spike suppression into simple models of STDP such as the Song model (Song et al., 2000) adequately accounts for spike-triplet results (Froemke and Dan, 2002) but, importantly, spike suppression does not offer an explanation of more general, multi-spike interactions. An explanation of multi-spike phenomena requires the introduction of additional, quite different modifications in the form of asymmetrical temporal limitation of spike interactions (Izhikevich and Desai, 2003).

The construction of an explicit representation of an STDP curve designed specifically to reproduce spike-pairing data is not the only approach to modelling STDP. An alternative approach is to construct simpler, more general models of spike-based plasticity without building in a representation of the STDP curve so explicitly. The aim of such an approach would be to produce tractable models which, it is hoped, would yield valuable insights into the nature of spike-timing plasticity. In our switch model of synaptic plasticity, we proposed a simple rule where changes in synaptic strength are in jumps of fixed magnitude. We showed that we can derive an STDP-like plasticity rule due to the averaging of this simple rule over multiple synapse and multiple spike-pairings. Only at the synaptic or temporal ensemble level does the overall connection strength between the two neurons therefore respect a STDP-form of synaptic plasticity. Individual synapses obey a much simpler, double step-function response when presented with a single spike pair. The apparently complicated STDP rule is therefore, in our model, an emergent form of computation that is nowhere instantiated at any individual synapse.

This approach is analogous, for example, to the relationship between thermodynamics and statistical mechanics in physics. The gas laws, such as Boyle's law, are not followed by individual gas molecules, but rather emerge, statistically, from the underlying motions of molecules following Newton's laws. Temperature

and pressure are not intrinsic properties of individual gas molecules, but emerge as properties of the collective system.

The concept of emergent computation is a powerful one. By taking advantage of emergence wherever possible, a system may perform apparently complicated calculations with the minimum of effort. Consider the emergence of rate-based rules under our synaptic switch model. The rate-based rules emerge as a natural consequence of the simple, synaptic switch rule. There is no explicit representation of firing rates, with its attendant requirement that neurons perform some kind of recent time-averaging of their activity level. Nor is there some explicit learning function that requires the neurons to translate these recent time averages into changes in synaptic strength. From the point of view of the synapse, there is only a simple switch rule that governs its response on a spike-by-spike basis. The rate-based rule is emergent, appearing only at some abstract level, and is no where instantiated at any individual synapse.

Of course, biological machinery may be postulated that performs almost any required function, such an explicit representation of firing rates. However, under our approach we find that, by exploiting emergence, this machinery is simply unnecessary. In some sense, the rate-based behaviour in our model is "doubly" emergent, as it arises from the averaging of an STDP rule which is, itself, an emergent property of our model. This emergence is in contrast to simple phenomenological models of STDP, where additional non-linearities are required, in the form of temporal limitation of spike interactions, to ensure that the averaged learning behaviour is consistent with experimental data (Izhikevich and Desai, 2003).

Thus, despite being formulated as a model of spike-pair interactions, we see a natural explanation of a range of experimental results under our simple synaptic switch. We see an STDP-like plasticity rule appear at the level of temporal and spatial ensemble average over multiple synapses and spike-pairings. We see an explanation of spike-triplets emerge alongside these spike-pair results as a consequence of the structure of the switch. We also see that the averaged, rate-based behaviour of our learning rule is BCM-like in form, with a depressing regime followed by a transition to a potentiating one. All of these results arise from the intrinsic structure of the switch. By comparing our model to experimental data we may constrain the parameters of our model. At no stage are we required to modify our model in order to introduce a particular behaviour or prevent an undesirable one. This is a remarkable result, especially given that our emergent approach yields a model that is considerably simpler than many competing models of synaptic plasticity, and that greatly reduces the computational burden placed on the synapse.

9.3 Multi-Spike Interactions

Both phenomenological and biophysical models of STDP have generally been put forward as models of 2-spike interactions. That is, they are formulated to explain the results of spike-pairing experiments, such as those of Bi and Poo (1998). These 2-spike models are then generalised to accommodate additional experimental results, such as spike triplet data (Froemke and Dan, 2002) or rate-based results (Abraham et al., 2001). We note, however, that even at the early stages of formulation some of these 2-spike models begin to encounter difficulties. For example, the biphasic, exponential model of Song et al., (2000) in its simplest possible form always produces either runaway potentiation or runaway depression, depending on the choice of parameters. As runaway potentiation is undesirable in any model of plasticity, the authors choose to have depression dominate potentiation and impose a hard lower bound at zero to prevent synaptic strengths from becoming negative. In order to avoid the situation where all the inputs fall to zero and remain there, the authors are then obliged to introduce some non-linearity. In the original implementation this was in the form of a non-linear, integrate–and–fire target neuron. Unlike a Poisson model, the integrate-and-fire neuron responds "directly" to input spikes. Stronger inputs will therefore exert more of an influence than weaker inputs, and will tend to cause the postsynaptic cell to follow their spiking. Once an input becomes sufficiently potentiated it will therefore dominate postsynaptic firing in the sense that, when it fires, a postsynaptic action potential will invariably be triggered within a few milliseconds. The input will therefore be potentiated almost every time it fires, leading to runaway learning. This obliges a further non-linearity in the form of a hard upper bound on synaptic strengths. Thus, four of the five basic features of the Song formulation (the fifth being the exponential rule itself) are, in effect, forced upon the model in order to prevent undesirable learning behaviours. Each additional modification or non-linearity brings with it important consequences for the learning dynamics expressed, and the behaviour of the model is therefore heavily influenced by the details of its implementation. For example, the characteristic bimodal equilibrium distribution of input weights in the Song model is a direct consequence of the choice of non-linearities imposed to solve the various problems with runaway learning inherent in the biphasic, exponential learning rule.

In attempting to generalise such a model to accommodate spike-triplet results, we find that further modifications are required, such as the introduction of spike suppression (Froemke and Dan, 2002). Likewise, in order to explain various rate-based results (Abraham et al., 2001), even more non-linearities must be introduced, such as the temporal restriction of spike interactions (Izhikevich and Desai, 2003). At each stage, the model is found to be inadequate, and at each

stage additional modifications designed explicitly to rescue the model are introduced. All of these modifications may, of course, be justified in the sense that biological machinery plausibly exists capable of implementing the proposed non-linearities. However, following this approach, and introducing modifications in this ad hoc manner, produces models of substantial complexity which can be hard to understand. Indeed, almost all of the modifications to the Song model described above were introduced in an empirical manner. The justification for doing do was, in each case, little more than the observation that the particular modification produced the desired behaviour or incorporated a particular experimental result.

In more recent experimental work, further evidence has emerged that multi-spike interactions potentially play a very important role in STDP. In neonatal rat hippocampal cultures, Wang et al., (2005) use spike triplets and quadruplets to show that the evoked plasticity is inconsistent with the commonly adopted view that spike-pairs sum linearly. The authors suggest that more accurate, "second-order" STDP rules would consider these interactions, and find that they make significant contributions to plasticity.

In other models of STDP we see similar problems arise from a very early stage. A key aim of the biophysical model of Senn et al, (2000) was to capture the frequency dependence of STDP observed in L5 neocortical pyramidal cells (Markram et al., 1997). This frequency dependence is incorporated by means of explicit thresholds to prevent too much potentiation at low firing rates. However, when at attempt is later made to reproduce the low frequency spike-pairing results of experiments like Bi and Poo (1998), the authors find that the predicted change under their biophysical model of STDP the predicted change in around 2% of that seen experimentally. Thus, we again see that the choice of a non-linearity introduced explicitly to accommodate a particular experimental results has profound consequences for the overall learning dynamics of the rule. The introduction of an explicit threshold to explain the frequency dependence seen in one set of experimental results creates difficulty when it comes to explaining a different set of results. It is easy to imagine further modifications which would solve this problem as well, but there is no guarantee that the model would then explain any other experimental results. Thus, this approach of generalisation by modification typically only works on a case-by-case basis.

Under our synaptic switch rule we find that a range of spike- and rate-based results are accommodated. In particular, an explanation of spike-triplet results emerges as a natural consequence of the models structure. We are therefore not required to consider modifications designed to ensure compatibility with these results. In addition, we see a BCM-like, rate-based learning rule emerge when our synaptic switch rule is averaged over many spikes. Again, we are not required to

postulate any additional modifications or non-linearities to guarantee this result. In fact, the requirement for a BCM-like, rate-based learning rule to emerge on average from our synaptic switch model serves to neatly constrain the parameters of our model. Unlike other models of STDP, there is no need to introduce any additional non-linearities to achieve these results.

We note, however, that the 2-spike form of rate-based learning rule emerging from our synaptic switch model, despite possessing a BCM-like form, is irredeemably pathological in its learning behaviour. The 2-spike learning rule cannot therefore support the stable, competitive dynamics that are essential within the context of developmental synaptic plasticity. Thus, the presence of a BCM-like learning rule is not, in our case, sufficient to ensure the presence of stable, competitive dynamics. We may remedy this problem by introducing an explicit dependence of the threshold between potentiation and depression on the recent time average of postsynaptic firing rate, in a manner similar to the BCM-rule. Our model is then, at this level, equivalent to the BCM-rule and produces qualitatively similar dynamics. Given that other models of STDP can produce similar, BCM-like rate-based rules (Izhikevich and Desai, 2003; Shouval et al., 2002) we may expect that they too would display similar learning dynamics, provided that they were capable of supporting some form of sliding threshold. In all cases, however, we require the introduction of an explicit non-linearity designed to ensure stability.

However, as we saw in Chapter 8, extending our consideration to just three or more spikes creates a rather different dynamical landscape compared to that produced under 2-spike interactions only. The learning dynamics of all multi-spike rules, while possessing an overall BCM-like, rate-based form similar to that of the 2-spike rule, exhibit dynamics that differ significantly from those exhibited by the 2-spike interaction function. We found that the dynamics of all multi-spike rules are governed by the presence of stable, segregated fixed-points which exist for a broad range of parameters. The learning behaviour we see under multi-spike rules is therefore one of stable competition, where afferents vie for control of the postsynaptic target cell. Importantly, we do not require the introduction of a sliding threshold for these rules. Instead, the learning dynamics arise as a result of the intrinsic properties of the multi-spike rules.

Thus, by extending our analysis to consider to more than 2-spikes we find a novel method of introducing competition into our model that does not require modification of the underlying plasticity rule. We argued that the differences between the 2-spike and multi-spike interaction functions is that the multi-spike rules can probe the coupling of potentiation and depression that occur under our switch rule. This coupling occurs in the sense that, in the unified switch rule, activation of the potentiation lobe of the switch precludes a simultaneous

activation of the depression lobe of the switch, and vice-versa. By breaking the two lobes into two separate switches, so that they may be simultaneously active, we showed that competition immediately breaks downs, resulting in the runaway learning characteristic of the 2-spike rule. It therefore appears that this coupling between potentiation and depression processes at the level of a single neuron is vital to the presence of stable, competitive dynamics.

It is interesting to speculate as to the role of coupling in other models of neuronal plasticity. In the BCM model, for example, the sliding threshold between potentiation and depression is determined by the firing history of the cell, and this firing history is dependent on the strengthening and weakening of the synapses that the cell supports. Thus, a sliding threshold provides a means of coupling potentiation and depression events, albeit in a slightly different manner to our own model. Without this coupling, so that the threshold is fixed, stability and competition in the BCM model break down. It appears in general, therefore, that in order for a model of synaptic competition to operate successfully, the machinery of potentiation and the machinery of depression cannot be independent. Although the dependence of the two processes is rather indirect in the BCM model, via the sliding threshold, it is nevertheless critical to the model's successful operation. We therefore speculate that any model in which potentiation and depression are completely independent processes cannot be competitive.

The apparent ease with which all these results, from the explanation of spike-triplets to the emergence of a BCM-like learning rule and the emergence of stable, competitive dynamics, contrasts markedly with competing models of STDP which, at almost every stage, require modifications to be made and non-linearities to be introduced. The approach of introducing ad hoc modifications to rescue an otherwise ailing model is also unlikely to produce the kind of deep insights into STDP that the theoretical community seeks. Indeed, in other sciences, such as physics, the emphasis is on simple, tractable models upon which a deeper understanding of the phenomenon can be built. When a simple model is found to be inadequate in explaining some new experimental result then, rather than introducing modifications designed to prolong the models lifespan, the underlying assumptions are examined and new theories proposed which seek to reconcile the new data with established results. In the context of STDP, we propose that the difficulty in accommodating spike-triplet and rate-based experimental results into many existing theories of STDP arises due to the exclusively 2-spike basis of those theories.

As we saw in Chapter 4, restricting a theory to one of 2-spike interactions almost always obliges the introduction of additional nonlinearities to explain multi-spike phenomena. Although, naturally, any model of plasticity must be consistent with the basic STDP data, the consequences of such a model must also be understood

in terms of its response to natural spike trains. Certainly, multi-spike interactions are more "natural" in the sense that a neuron does not experience a series of isolated spike-pairs but rather continuous, overlapping trains of pre- and postsynaptic spikes. However, it is just such natural spike trains that create difficulties for many existing models of STDP. In contrast, under our synaptic switch rule we find an explanation of a range of spike- and rate-based results arising as a result of the models intrinsic structure. We find that the consideration of multiple-spike interactions provides a natural method of introducing stable, competitive dynamics. These dynamics are capable of explaining a range of developmental phenomena, such as the development of ODCs. These results are emergent in the sense that the switch rule is formulated as a model of spike-pairing interactions. At no stage do we introduce modifications designed to accommodate a particular experimental result. In addition, our model is considerably simpler than competing models of STDP, and dramatically reduces the computational burden place on the synapse.

9.4 Coincidence Detection

Common to both the phenomenological and biophysical approaches to modelling STDP, which we discussed in Chapter 4, is the view that STDP must be valid, at some level, at individual synapses for a single spike pair. Individual synapses in such models are required to possess some representation of spike-timing differences and adjust their strengths accordingly in a graded manner. Each synapse must therefore possess machinery capable of both resolving pre- and postsynaptic spike timing with millisecond accuracy and of translating the difference between these times into a graded changes in synaptic strength. That is, each synapse must perform as a millisecond resolution coincidence detector. This requirement places a considerable computational burden on the synapse.

Although mechanisms of coincidence detection have been suggested to exist at the synapse, such as the NMDA-receptor dependent calcium dynamics of the calcium control hypothesis discussed in Chapter 4, the conjecture that synapses are capable of precise coincidence detections is still far from proven.

Consider, for example, the calcium control hypothesis. It is true that the NMDA-receptor can operate in the manner of a coincidence detector as, following activation by presynaptic glutamate release, calcium entry through NMDA-receptor associated channels is only permitted once the characteristic voltage-dependent magnesium block is released. There is therefore a requirement for coincident presynaptic glutamate release and postsynaptic depolarisation. It is also true that the individual elements of the calcium control hypothesis have been shown

to be involved in plasticity, for example changes in the concentration of post-synaptic intracellular calcium concentration has been shown to induce up- or downregulation in synaptic strengths Yang et al. (1999). However, it has not been shown that NMDA-dependent calcium influx in the only relevant mechanism, and that a model of NMDA-receptor function is sufficient to explain STDP. Other processes may make equal, or even more, important contributions to plasticity at the synapse, and a range of secondary processes may also contribute in various ways.

In modelling studies, the calcium control hypothesis has been shown, under appropriate constraints, to account for various induction protocols of STDP (Senn et al., 2000; Shouval et al., 2002). Such models are, however, rather sensitive to parameters and often predict an additional LTD window at large pre-then-post spike timings which has rarely been reported (but see Nishiyama et al., 2002). A more complicated model, which relies on a second coincidence detector in addition to NMDA-receptors, has been proposed specifically to resolve the issue of the extra LTD phase at large spike timings (Karmarkar and Buonomano, 2002). More recently, a gradient-based rule has been proposed which achieves the same result (Saudargine and Porr, 2005). Although both these models eliminate the additional LTD window, it is difficult to view either approach as anything more than a specific remedy for a very particular problem with the calcium control hypothesis. In addition, biophysical based models are generally motivated by the view that they may offer a deeper understanding of the processes underlying STDP without the need to explicitly formulate, for example, two exponentials to account for the two halves of the STDP-curve. However, in the case of Karmarkar et al., (2002), the introduction of a second coincidence detector reduces the model to a reparamterisation of the double exponential STDP-curve, which is equivalent in some respects to the simple phenomenological models discussed elsewhere (Song et al., 2000; van Rossum et al., 2000). Thus, although the two coincidence detector model is largely successful in explaining STDP the increased complexity undermines the advantages of this approach, to explain STDP from simple biophysical arguments. In addition, the computational burden on the synapse increased even further, requiring it to possess two coincidence detectors rather than one. Although it is true that the synapse may possess such machinery, and that a faithful biophysical description of it is necessary to explain STDP, such conjectures are far from proven. In such a situation, it is the role of the theorist to examine the simplest theories first, and seek to explain a phenomenon without resorting to may be recognised later as unnecessary complexity.

That a biophysical model of the calcium control hypothesis can explain certain experimental results on STDP is by no means proof that the hypothesis is correct. Indeed, if NMDA-receptor-dependent calcium dynamics were the critical

process underlying STDP, it might be expected that a model of this process would be more successful in explaining STDP results. Throughout the literature, what is interpreted as considerable experimental support is often simply evidence that all the elements of a biophysical theory are somehow involved in plasticity. This is not the same as having direct experimental evidence that one particular pathway is responsible for STDP, and that a model of this particular pathway will offer a full explanation of the phenomenon. For example, a large number of signal transduction molecules have been implicated as having some role in synaptic plasticity, for example neurotrophins such as brain-derived neurotrophic factor (Zakharenko et al., 2003; Schinder and Poo, 2000), postsynaptically released endocannabinoids (Sjöström et al., 2003), arachidonic acid (Williams et al., 1989), or even nitric oxide and carbon monoxide (Zhuo et al., 1993). Determining which of these components are essential for, and which are involved more indirectly in, synaptic plasticity is often technically difficult. It has also been observed that some forms of NMDA-dependent LTP may be induced presynaptically (Zakharenko et al., 2003; Humeau et al., 2003), which is not consistent with the exclusively postsynaptic induction of the calcium control hypothesis.

Furthermore, for even the most basic predictions of the calcium control hypothesis to be consistent with experimental data, some careful parameter selection is required. For example, the binding constant of NMDA-receptors to glutamate is typically quite large. Under the calcium control hypothesis, this leads an unnaturally long potentiation tail on the STDP curve which is not seen experimentally (Bi and Poo, 1998; Zhang et al., 1998). In the absence of any understanding of why pre-post spike interactions are temporally restricted to timescales much shorter than the timescale governing glutamate binding to NMDA-receptors, the time constant for glutamate binding is simply reduced (Karmarkar and Buonomano, 2002). Likewise, in its basic form, the calcium control hypothesis predicts an LTD phase of a much shorter temporal duration than is observed experimentally (Bi and Poo, 1998). It was noted empirically that introducing a long after depolarisation lasting tens of milliseconds following a postsynaptic action potential extends the predicted LTD time window to a more appropriate size (Shouval et al., 2002; Karmarkar and Buonomano, 2002). A long after depolarisation was duly introduced, and subsequently elevated to the level of a basic assumption of the hypothesis (Shouval et al., 2002).

Moreover, it is not clear that calcium transients can offer the kind of accuracy demanded by the calcium control hypothesis. Natural variability in synaptic and dendritic transmission, as well as in calcium entry and reuptake, may wash out the precise temporal profile required by modelling studies (Shouval et al., 2002).

Thus, although attractive, the translation of the calcium control hypothesis into a coherent theoretical framework for synaptic plasticity is still incomplete. Direct experimental proof has not been produced, and modelling studies have been only partially successful in explaining STDP. Indeed, in its basic form, the calcium control hypothesis produces rather eccentric STDP curves, and various modifications must be made to ensure a temporal profile consistent with the majority of experimental work (Senn et al., 2000; Shouval et al., 2002; Karmarkar and Buonomano, 2002; Saudargine and Porr, 2005).

Of course, it is unreasonable to expect a biophysical model of STDP to explain every experimental result in its first formulation. It is reasonable to expect, however, that it explain a range of experimental data without the need to introduce modifications designed specifically to ensure consistently with those result. It is also desirable, from a theoretical point of view, that the model be as simple as possible, and only place as much computational burden on the synapse as is absolutely necessary.

In contrast to the other models of STDP, our switch mode requires only a simplified form of resolution coincidence. We only require that the synapse is capable of recording whether there has been a pre- or postsynaptic spike, and that this trace is destroyed in a stochastic manner. Importantly, we do not ask that the synapse records the times of pre- and postsynaptic events. We also do not require a time difference to be translated into some graded change in synaptic strength, as all changes occur in jumps of fixed magnitude. By eliminating the need for precise coincidence detection in this way, our switch model considerably reduces the computational burden on the synapse. From a biological point of view, the switch may be implemented in a variety of ways, and we do not commit to one particular implementation. We are therefore free to postulate many ways in which the machinery present at the synapse may contribute to plasticity.

In addition, our switch model of STDP successfully explains range of spike-pair, spike-triplet, and rate-based results, without the need for additional modifications or the introduction of additional non-linearities. Thus, we show that alternative interpretations of STDP may give rise to models which are more successful, while at the same time being much simpler, than competing models.

Of course, exploring the machinery of the synapse is a worthwhile activity in its own right. However, entering into the task with the goal of determining if the synapse can support the computation demanded by a particular hypothesis of synaptic plasticity is potentially misleading. Given the range of mechanisms in operation at the synapse, the ubiquity of signalling molecules, and issues such as the active properties of dendrites, it is likely that processes exist that may be construed as supporting almost any form of computation. As many of these

processes are only partially understood it is possible to imagine a whole range of computational activities that the synapse could engage in. In theoretical terms, therefore, it is more interesting to ask what is *necessary* to explain STDP, rather than ask what forms of computation the synapse can possibly support.

As we only require a minimal form of coincidence detection, we are freed from having to postulate some complicated mechanism as a substrate for our model. All that we require is that the synapse can records that pre- and postsynaptic spiking has occurred and then make changes in efficacy of fixed magnitude. As we discuss in Section 9.5, this may be implemented in a variety of ways. Thus, despite the considerable experimental effort that has been invested in looking for coincidence detection mechanism at the synapse, it is not clear that such mechanisms are even necessary to explain STDP.

9.5 Molecular Implementation

We have proposed that single synapses implement something akin to a three-state switch, with changes in synaptic strength occurring in discrete jumps of a fixed magnitude, and shown that the model can explain a variety of spike- and rate-based experimental results. The discrete nature of our model is in contrast to other theories of STDP, such as the calcium control hypothesis, which involve graded changes in synaptic strength and, as a result, are required to postulate coincidence detection machinery that is capable of translating differences in spike-timing into some graded change. This reduces the level of computation that individual synapses are required to perform, and opens up a range of possible molecular implementations.

Where might the machinery that implements this switch reside? The switch moves from the OFF state to the DEP state following a postsynaptic spike. All of a target's input synapses might thus be expected to move into the DEP state, and this suggests that the natural locus for the DEP state of the switch may be the postsynaptic aspect of the synapse. Similarly, when the switch moves from the OFF state into the POT state following a presynaptic spike, all of the afferent's output synapses might be expected to move into the POT state, and so this perhaps suggests a presynaptic locus for the POT state of the switch. Thus, the switch may actually be distributed across the entire synapse, rather than confined exclusively to the pre- or postsynaptic side of the synapse. Under such a distributed switch, an active transition back to the OFF state caused by a presynaptic spike may cause a postsynaptic change in the synapse, for example the removal of some neurotransmitter receptors from the postsynaptic membrane, leading to a decrease in synaptic strength. Similarly, an active transition to

the OFF state caused by a postsynaptic spike may change the presynaptic side of the synapse, perhaps enhancing neurotransmitter release, or even inducing terminal synaptic sprouting, so that synaptic strength is increased.

It is also relevant to discuss how our synaptic switch might be implemented. Given the range of signalling molecules available at the synapse, such as the neurotrophins (Zakharenko et al., 2003; Schinder and Poo, 2000), and endocannabinoids (Sjöström et al., 2003) mentioned in Section 9.4, there is plenty of scope for implementing our proposed synaptic switch at a molecular level. In the absence of any conclusive experimental data, we are therefore cautious in identifying one particular mechanism as a more likely than another. Indeed, it may be that the switch is implemented in different ways in different parts of the nervous system, or across different species, while still subserving the same broad function. All that we require is that synapses are capable of recording whether a pre- or postsynaptic spike has occurred, that this trace is destroyed in a stochastic manner, and that subsequent pre- or postsynaptic spikes which occur in a timely fashion are capable of triggering an up- or down-regulation in synaptic strength in jumps of fixed magnitude.

The requirement that synapses record that a pre- or postsynaptic spike has occurred requires a form of coincidence detection that is, in some respects, similar to the initial stages of the coincidence detection present in other models of STDP. We may therefore appeal to similar kids of mechanism as appealed to elsewhere. For example, postsynaptic NMDA-receptors, with their requirement for both presynaptic glutamate releases and concurrent postsynaptic depolarisation, would serve adequately in our model, as elsewhere, for detecting coincident pre- and postsynaptic spiking. However, we stress that although we require some level of coincidence detection in our model, it is minimal in the sense that we do not require the synapse to record the spike-timings with millisecond accuracy or that these spike-timings are then translated into some graded change in synaptic strength. That is, although we require some form of coincidence detection, it is, in some sense, equivalent to the "first step" of other models. We do not require the subsequent steps of comparing spike-timings or making graded changes based upon them. As a result, the demands we place upon the synapse, and the complexity of the molecular machinery we are required to postulate, are greatly reduced. We are therefore not committed to this particular form of postsynaptic NMDA-receptor-dependent coincidence detection, but may instead consider a range of possibilities.

The step-like changes in synaptic strength that occurs under our switch model could also be implemented in a number of ways. Switch-like activation of a process is not an uncommon phenomenon in biology. Indeed, CaMKII, which plays an key role in effecting the graded changes in synaptic strength required

by the calcium control hypothesis, has itself been shown to display a switch-like activation to calcium concentration levels (Bradshaw et al., 2003). This arises due to the steep dependence of the autophosphorylation of CaMKII on intracellular calcium concentration, an effect which is amplified by the competing dephosphorylation action of protein phosphatase 1 (PP1). The CaMKII-PP1 system may therefore function together as a simple molecular device that translates a calcium signal into all–or–none potentiation of individual synapses. That all–or–none potentiation or depression consistent with such a molecular switch can occur during plasticity episodes has recently been demonstrated in the hippocampus (Petersen et al., 1998; O'Connor et al., 2005). Although these results concerned putative individual afferents identified by means of a minimal stimulation protocol, they do provide evidence that all–or–none, binary changes such as the kind required by our switch model can occur in the nervous system. We also note that the observed differences in number of stimuli required to evoke potentiation and depression could easily be accommodated in our synaptic switch model through an appropriate choice of $\tau_{\pm}$. These results raise the question as to whether the apparently graded changes seen in other plasticity experiments, which have largely motivated the search for high-level coincidence detection at the synapse, are due instead to the recruitment of multiple inputs with different plasticity thresholds.

Thus, the implementation of our proposed synaptic switch may occur in a variety of ways. Given the ubiquity of both anterograde and retrograde messengers in the nervous system, we should be cautious in utilising any arguments about a specific locus or mechanism for the switch. The mathematical details of the switch do not commit us to any particular view concerning its exact locus, nor to the precise molecular machinery involved in its implementation. We note that several elements of existing theories of STDP, such as the NMDA-receptor/CaMKII pathways of the calcium control hypothesis, would adequately serve to underlie certain aspects of our model. However, we do not require the more complicated elements of this theory such as the translation of accurately recorded spike times into graded changes in synaptic strength. The level of coincidence detection we require is of a much reduced form, and we are therefore not required to commit to this particular mechanism but are free to imagine a variety of other implementations.

9.6 Experimental Tests

We may propose various experimental tests of our synaptic switch rule that would distinguish it from other competing models of STDP. These experimental

tests might come in two forms. Firstly, we could examine the key underlying assumptions of the model, such as the requirement that for the purposes of driving plasticity synapses operate as discrete switches, and propose direct experimental tests of theses assumptions. Secondly, perhaps more realistically, we may make predictions about the response of a system to various stimulation protocols, such as 3-, 4- or 5-spike trains. In addition to distinguishing our model of synaptic plasticity from competing models, these tests could also serve to distinguish the different forms of our model such as the resetting and non-resetting forms discussed in Chapter 7.

As discussed in Section 9.5, the mathematical formulation of our synaptic switch rule is broadly consistent with a range of possible biological implementations. Attempts to deconstruct the synapse and determine if the molecular machinery capable of supporting our switch exists at the synapse are therefore unlikely to provide clear evidence for or against our hypothesis. In terms of distinguishing between our synaptic switch rule and other competing models of STDP, the search for molecular evidence in support of one model over another is hindered by the fact that models of STDP often have certain common requirements, such as the need for some form of coincidence detection. Indeed, certain molecular pathways that have been identified as potential candidate mechanisms for other models of STDP, such as the well characterised NMDA-receptor-CaMKII pathway, would serve quite adequately to explain certain aspects of our switch rule. We therefore discuss experimental tests that focus on the computational properties of our plasticity rule rather than its underlying molecular implementation, as these tests are more likely to clearly distinguish our model from its competitors.

The most direct test of our synaptic switch hypothesis would be to examine changes occurring at individual synapses in response to controlled pre- and postsynaptic spiking. We predict that changes would occur in jumps of fixed magnitude, rather than in a graded fashion predicted by other formulations of STDP (Song et al., 2000; van Rossum et al., 2000; Shouval et al., 2002; Senn et al., 2000). We also predict that the timing window in which postsynaptic spiking must follow presynaptic spiking in order to induce these fixed jumps varies in a stochastic manner. Thus, we expect that repeated pre- and postsynaptic spike pairs of a fixed time difference will sometimes evoke potentiation of a fixed magnitude and sometimes not. This is contrast to other models of STDP which would predict that the spike pairings would consistently evoke some intermediate level of potentiation with every spike pair. In our model, only when these step-like changes are averaged over many individual spike pairings, or over may synapses, will the overall change resemble the apparently continuous changes observed in conventional STDP experiments (Bi and Poo, 1998). That all–or–none changes in synaptic efficacy can occur in a biological setting has recently been

demonstrated in the hippocampal formation (Petersen et al., 1998; O'Connor et al., 2005). Although these experiments were concerned with putative single inputs, identified by means of a minimal stimulation protocol, and a number of issues still remain to be clarified, they do at least show that step-like or binary changes can occur in the context of synaptic plasticity. We note that the kind of changes observed in these experiments, with their characteristic fixed-magnitude and rapid onset, cannot easily be accommodated into existing models of STDP which are formulated explicitly around the need to invoke graded changes in synaptic strength.

As the technical difficulties associated with recording from single synapses are formidable, it is prudent to discuss other, less direct tests of out synaptic switch hypothesis. We therefore propose some hypothetical experimental protocols which would, while still being within easy technical reach, would nevertheless effectively distinguish our model from competing models of STDP. We originally presented our switch model in Chapter 6 as a model of two-spike interactions, which guarantees that the model is consistent with basic spike-pairing STDP data (Bi and Poo, 1998). Examining the effect of spike-triplet interactions, such as those explored by Froemke and Dan (2002), we find that our synaptic switch model produces a natural explanation of spike-triplet results due to its intrinsic structure. No additional modifications or the introduction of additional non-linearities is required. This is in contrast to other models of STDP, such as the Song model (Song et al., 2000), which require the introduction of additional devices, such as spike suppression (Froemke and Dan, 2002), to explain the spike-triplet data.

Of course, we cannot really view this explanation of spike-triplets as a prediction because, despite our model being formulated as purely a model of spike-pair interactions and this emergence of an explanation of spike-triplet interactions for free, the spike-triple experiments were carried out prior to the publication of our model. We can, however, pursue this idea of multi-spike interactions further, and examine the predicted change in synaptic strength induced by more complicated spike trains. These predictions could be used to distinguish our synaptic switch rule from competing models of STDP, as well as differentiate the two forms of our model, resetting and non-resetting, that we presented in Chapter 7.

Consider a simple spike pairing, πp, where spikes occur at times t_0 and $t_0 + t_1$. Most models of STDP (Song et al., 2000; van Rossum et al., 2000; Senn et al., 2000; Shouval et al., 2002) give qualitatively similar predictions for the expected change in synaptic strength. In all cases, the spike pairing will induce some graded change in synaptic strength whose precise value is dependent on the time difference t_1. Under our switch model, we would observe a very similar

result. Of course, under our model the apparently graded change arises due to the temporal averaging of discrete changes in synaptic strength over multiple spike pairings. However, as multiple spike pairings would always be required in order to produced a statistically meaningful change in afferent strength (Bi and Poo, 1998), we cannot appeal to arguments about observing graded as opposed to discrete changes in synaptic strength. We can, however, consider the effect of introducing an additional π spike at time in front of the πp pair, to produce a spike-triplet $\pi\pi p$. In such a triplet there are two possible πp interactions. In the Song model, these spike interactions are summed linearly. The second πp interaction, with its shorter time difference and resulting large change in synaptic strength, would therefore make a larger contribution to plasticity than the first πp interaction. Under the spike suppression model, however, the initial presynaptic spike will suppress the influence of the second presynaptic spike. As a result, the total potentiation evoked will be somewhat less than predicted by the song model. The exact value of this potentiation will depend on the strength of the suppressive influence of the first π spike. Under both the Song model and the spike-suppression model, therefore, the addition of this extra presynaptic spike changes the expected plasticity level from that of the simple πp pair. Thus, in general, $\Delta S^{\pi\pi p} \neq \Delta S^{\pi p}$. This observation holds for all other models of STDP. For example, in the calcium-based model of Senn et al., (2000) , additional presynaptic spikes move more of the population of NMDA-receptors into their activated state. The potentiation induced by the final postsynaptic spike is therefore to a higher level than if only a single presynaptic spike had occurred.

In the resetting form of our switch model, it is not the case that additional presynaptic spikes modify the expected change in synaptic strength induced by the final p spike. In the spike train $\pi\pi p$, the initial π spike elevates the synapse to the POT state. By the time the second π spike arrives, the synapse will either still remain in the POT state or it will have spontaneously returned to OFF. If the synapse has returned to OFF then the second π spike elevates the synapse to the POT state. If, on the other hand, the synapse remained in the POT state, the second π spike resets the stochastic process governing the return to OFF. Thus, the stochastic process governing the transition to OFF is always restarted, regardless of whether the synapse returned to OFF before the arrival of the second spike or not. In effect, the initial presynaptic spike may be ignored, and only the final πp spike pair is relevant. This is a general feature of the resetting model, only the final spike in a long train of identical spikes is relevant for the purposes of inducing plasticity, and was discussed in Chapter 7. The expected potentiation under the resetting switch model for this $\pi\pi p$ train is therefore the same as that for spike pair πp provided the time difference of the last two spikes is the same. Thus, the spike trains πp, $\pi\pi p$, and $\pi\pi\pi p$ are,

in terms of evoking plasticity, all identical under our switch model provided the time difference of the last two spikes is the same. We may therefore write $\Delta S^{\pi^i p} = \Delta S^{\pi p}$, any i, provided that the time difference of the last two spikes is the same.

This property of the resetting form of the switch model, that, in a train of identical spikes, only the last spike is relevant with respect to inducing plasticity provides a convenient experimental focus for distinguishing our model from other models of STDP. In general, the observation that $\Delta S^{\pi^i p} = \Delta S^{\pi p}$, any i, does not hold for the non-resetting form of our model. We may still, however, construct various spike trains and compute the expected change in synaptic efficacy arising from the non-resetting forms of our switch mode to provide predictions for comparison to experimental results. It is easy to conceive of certain spike trains that, while perhaps slightly contrived, would produce different predictions from various competing models of STDP and from our own switch model.

Of course, this also means that we may use the hypothetical experiment discussed above to distinguish the resetting and non-resetting forms of our synaptic switch model. In the resetting form of the model, the spike train $\pi\pi...\pi p$ would induce the same plasticity as the spike pair πp, provided the timing difference of the last two spikes was the same. In the non-resetting model, this would not be the case, and the two spike trains would, in general, not evoke the same change.

We proposed in Chapter 7 that the resetting and non-resetting forms of the switch model are more biologically plausible than the general r-reset case. This is because, in the resetting and non-resetting models, additional pre- or post-synaptic spiking can either reset the stochastic process governing the return to OFF (the resetting model) or it cannot (the non-resetting model). Other, general r, cases require that some form of spike counting machinery exists at the synapse that, in the spirit of reducing the computational burden on the synapse as far as possible, we consider to be less likely. Nonetheless, for completeness we discuss an experimental test that would differentiate the different r-reset models from each other. This test would be very similar in character to the experiment already discussed. This time, spike trains of increasing length, say πp, $\pi\pi p$, $\pi\pi\pi p$, and so on, would be used. As the r-reset model resets after r identical spikes, we would expect that the $\pi^r p$ train would induce the same plasticity as a simple πp pairing, again provided that the time difference of the final two spikes was the same. All that would be required, therefore, is to determine r in an iterative manner.

In summary, direct experimental evidence of our proposed synaptic switch rule would be technically difficult. A search for molecular evidence would suffer from similar problems, and in any case, would be unlikely to distinguish clearly

between competing models of STDP. We therefore propose that a consideration of plasticity evoked by trains of multiple spikes could distinguish both the resetting and non-resetting forms of our switch model, and differentiate our model from competing models of STDP. This approach is also within easy technical reach, representing a simple extension of existing experimental plasticity protocols.

9.7 Comparison to Other Models of STDP

Although some existing models of STDP bear some superficial resemblance to our own switch model of synaptic plasticity, these similarities belie much deeper conceptual differences.

The biophysical model of Senn et al., (2000), which we discussed in Chapter 4, describes an NMDA-receptor-dependent plasticity rule in which the activation state of a fixed population of NMDA-receptors drives changes in synaptic strength. NMDA-receptors are postulated to transition from a resetting state to one of two activated states in response to pre- or postsynaptic spiking. Subsequent pre- or postsynaptic spiking is translated into a graded change in synaptic strength, with size and polarity determined by the relative populations of the three states. In the absence of further spiking, the population of NMDA-receptors in the activated states decay back to the resting state. Thus, the distribution of the population of NMDA-receptors between the three states is used to explicitly record spike timings, and these spike-timing are translated into graded changes in synaptic strengths. In effect, the NMDA-receptors act as a discrete-state counting device and encode an exponential, biphasic learning rule similar in character to that of various phenomenological models (Song et al., 2000; van Rossum et al., 2000).

Thus, in the Senn model, there is still the need for a comparatively high level of coincidence detection, and a translation of this coincidence detection into graded changes in synaptic strength. The population of NMDA-receptors, which have 3 distinct activation states, provide a mechanism for implementing this coincidence detection and the translation of spike times into graded changes. In our switch model we also utilise a discrete, 3-state mechanism. However, in the switch model this 3-state mechanism is in the form of a synaptic controller switch which governs plasticity at individual synapses. Transitions of the synaptic switch to different activation states refer not to some population of receptors present at the synapse, but to the whole synapse. Once pre- or postsynaptic spiking has activated the synaptic switch, the arrival of further pre- or postsynaptic spikes triggers step-changes in synaptic strength, provided those spikes arrive before the switch has deactivated. There is no decaying population with which spike times may be recorded, nor any translation of differences in spike timing into graded

changes in synaptic strength. Thus, unlike the Senn model, there is no explicit representation of the STDP curve in our switch rule, with the exponential STDP curve emerging instead due to averaging over multiple synapses and over multiple spike pairings.

Of course, in the Senn model, the 3 NMDA-receptor activation states are needed to encode the two phases of the STDP curve. In our switch model, we also seek to explain the bidirectional change characteristic of STDP and, therefore, our synaptic controller switch was also postulated to have 3 discrete states. However, this similarity is only superficial in the sense that, despite the Senn model and our synaptic switch model both containing some form of discrete, 3-state mechanisms, the two models apply these mechanisms in very different ways.

Binary synapses, whose strength is constrained to one of two values, have been studied in various contexts (Fusi, 2000; Fusi et al., 2002), and models based on such synapses have been shown to be capable of reproducing various rate- and spike-based results. In the Fusi model (Fusi et al., 2002), synapse are permitted to exist in a state of either lower high efficacy. Some internal variable determines when transitions are made between these two states. A threshold is chosen to control state transitions. When the internal variable is above the threshold, the synaptic efficacy is set to the high value, and when the internal variable is below the threshold the synaptic efficacy is set to the low value. Presynaptic spiking cause the internal variable to receive injections whose magnitude and sign are dependent on the level of postsynaptic depolarisation. In the absence of presynaptic spiking, the internal variable will drift to one of the limits, depending on its value compared to the threshold. Thus, the binary synapses are stabilised over time by the drift of the internal control variable. Correlations in pre- and postsynaptic spiking are then required to drive the variable across the threshold and cause the synapse to transition to the state of opposite strength. In contrast, in our model we refer to a discrete state switch that governs the plasticity of a continuous synaptic strength variable. Pre- and postsynaptic spiking induces changes in synaptic strength of a fixed amount according to the state of the synaptic switch. Although these changes are of fixed-magnitude, the synaptic efficacy is still continuous, and nowhere do we constrain its value to be one of a set of discrete values. Nor do we include a drift term to stabilise synapse strengths.

In some sense, a more accurate name for our switch model would be the "synaptic plasticity-switch" model, which highlights the fact that we consider continuous synapses that are governed by a discrete plasticity rule. Thus, although sharing some conceptual similarities, our model is quite distinct from existing studies of binary synapses. We consider synaptic strengths to be continuous variables,

and it is only the controlling switch that occupies discrete states. We note, however, that we do not rule out the possibility of binary-strength synapses. A variation of our switch model, where synaptic strengths are constrained to be either one of two states, can also be formulated. Provided that each input is comprised of a large enough number of these binary synapses, sufficient averaging will occur and the basic findings of our model will be preserved. Indeed, when the number of synapses becomes very large, this averaging becomes exact and the binary and continuous formulations of the model are identical. Thus, we do not commit, necessarily, to a continuous- or binary-synapse description, as our model is largely independent of this choice. Our choice to present a continuous formulation was originally motivated on the grounds of simplicity, but we note that a continuous formulation permits a more transparent demonstration of the stability of the multi-spike learning rules that emerge under our switch rule. These rules were discussed in Chapter 7 and Chapter 8, where we showed that stable, competitive dynamics emerge under all multi-spike rules, but not under the 2-spike rule. With binary synapses, a natural upper bound exists on synaptic strength and, as a result, total input strengths. Had we chosen to formulated our model in a binary manner, this natural hard upper bound may have concealed the intrinsic stability of multi-spike rules in our switch model.

9.8 Future Work

Our synaptic switch rule provides an explanation of a range of spike- and rate-based results. Compared to other models of STDP that attempt to explain similar results our synaptic switch rule explains the phenomena in a much more natural manner - relying on intrinsic, emergent properties rather than empirical modifications designed specifically to incorporate experimental results on a case–by–case basis. Several issues arise from our switch model of synaptic plasticity that we will explore in future work. We discuss these issues in turn.

9.8.1 Frequency-Dependence of STDP

We have deliberately presented our switch rule in the simplest form possible that is consistent with a variety of spike-timing data. This permits a degree of analytical understanding of the rule, enabling us to determine the properties of the model without resorting to a purely numerical approach. Although our switch rule accommodates a wide range of experimental results, we cannot expect our model to be comprehensive at this stage.

We note, in particular, that the frequency-dependence of the STDP observed in thick-tufted L5 neocortical pyramidal cells (Markram, 1996) is not captured

by our model in its present form. In this experiment, bursts of pairings were presented at approximately 1/2 second intervals. A burst consisted of five presynaptic spikes paired with five postsynaptic spikes. The postsynaptic spikes were induced by current injection approximately 5ms after each presynaptic spike. It was observed that the bursts of pairings only induced potentiation when the frequency of pairings within a burst was greater than around 10Hz. In other STDP experiments, such a protocol would be expected to induce potentiation even at burst rates as low as 1Hz (Bi and Poo, 1998). This apparent frequency-dependence could reflect important functional differences in the STDP rule in different areas of the brain.

9.8.2 Rate-based Computation

The rate-based rules that emerges from our switch model are, as we described them before (Appleby and Elliott, 2005), "doubly emergent", because they depend on the emergence of STDP at the synaptic or temporal ensemble level in our model, and then this STDP is itself averaged over time and over many spike patterns, giving rise to the rate-based behaviour.

Mathematically speaking, the STDP rule takes the form of a conditional expectation value for synaptic change, conditional, that is, on a given spike time difference. Second, to turn this conditional expectation value into an unconditional value, we must weight it by some probability density function for the spike time difference, and the probabilities of each possible spike pair. Thus, in order to obtain the BCM-like rule, we assumed that the spike time difference distribution originated from Poisson spike trains. Had we chosen to drive the afferents in some other, non-Poisson manner, a different rate-based synaptic plasticity rule may have emerged. In this sense, the BCM-rule emerges from an interaction between the STDP rule, which is itself emergent, and how we have decided to drive the afferents.

It would be interesting to explore more fully the kind of rate-based learning rules that can emerge from our synaptic switch rule under different afferent firing patterns. In particular, to explore whether the BCM-like learning rule we obtain for Poisson driven inputs is a general feature of the model. If different firing patterns were found to give rise to different learning rules, then this would suggest the possibility that apparently different rate-based rules may, in fact, merely reflect different choices available to the experimenter in how he decides to probe his experimental system or, indeed, that as synaptic patterns change and hence neuronal firing patterns change, the nervous system may slowly change its own rate-based learning rule.

9.8.3 Coupling of Potentiation and Depression

In our switch model we find that the presence of stable, competitive dynamics requires that the machinery of potentiation and the machinery of depression cannot be completely independent, but rather potentiation and depression must be coupled together. This coupling occurs in the sense that once the synapse has entered the POT state it cannot subsequently enter the DEP state without having first returned to the OFF state.

A similar coupling is present in other models of competitive plasticity, in particular the BCM model (Bienenstock et al., 1982) and the Song model (Song et al., 2000), albeit in a slightly different manner. In the BCM model, the sliding of the threshold between potentiation and depression provides the coupling. Again, this coupling is critical to the model's successful operation. Without the sliding threshold, the BCM rule breaks down with afferents either all falling to zero or experiencing runaway potentiation (Bienenstock et al., 1982). In the Song model, it is the non-linear integrate-and-fire target neuron that couples potentiation and depression together. With an integrate-and-fire neuron, inputs which are potentiated are more likely to experience further potentiation as they exert a greater influence on the postsynaptic spike timing. Without the nonlinearity in postsynaptic spike timing provided by the integrate-and-fire neuron competition in the model breaks down. If, for example, postsynaptic spiking was governed according to a Poisson process with rate given by the linear weighted sum of its inputs afferents in the Song model will invariably fall to zero or experience runaway potentiation, depending on the choice of parameters.

Thus, we see that in the Song, the BCM, and our own synaptic switch models, coupling of potentiation and depression is required to avoid pathological learning behaviours. In other models of STDP, despite coupling of potentiation and depression being only partially explored, it has been observed that some form of coupling could stabilise learning dynamics (Shouval et al., 2002). It is interesting to speculate about the role of coupling in a more general model of synaptic plasticity. In particular, whether a model in which potentiation and depression are completely independent processes can be competitive or not. It would be desirable to seek to prove this within the context of the most general class of model possible, although this is likely to be very difficult.

9.9 Closing Summary

In this thesis we have been concerned with activity-dependent neuronal plasticity in the nervous system. We motivated our study of plasticity with a discussion

of the often dramatic role plasticity processes can play throughout the life of an animal. We then presented a range of experimental results on plasticity, from early tetanus-based experiments to more recent spike-timing and single-input results. A critical reinterpretation of this experimental evidence led us to make several key observations. In particular, we noted that spike-timing results are often interpreted literally, despite the range of experimental issues that still surround these results. We discussed various rate- and timing-based models of neuronal plasticity and examined their learning properties. A detailed examination of one particular model of STDP, the Song model (Song et al., 2000), highlighted several important issues surrounding the interpretation of STDP results.

Drawing upon this work, we proposed a new model of plasticity, in the form of a synaptic switch rule. The rule governs changes at individual synapses and only when the rule is averaged over multiple synapses and multiple spike-pairs does an STDP-like rule emerge. The STDP-like rule is therefore an ensemble property of our model, which is nowhere instantiated at any individual synapse. In addition, we found that our switch rule explains a variety of spike- and rate-based plasticity results essentially for free. We also found that stable, competitive dynamics emerge naturally due to multi-spike interactions. At no stage were we required to introduce additional modifications to accommodate particular experimental results or avoid otherwise undesirable learning behaviours. Indeed, ensuring consistency with various experimental results serves to neatly constrain the parameters of our model in a concise manner. This is in contrast to many other models of STDP, which are often required to introduce additional modifications and non-linearities to explain experimental results on a case-by-case basis. Furthermore, our synaptic switch rule is considerably simpler than many competing models of STDP and places a much lower computational burden on individual synapses. We are therefore freed from the need to postulate precise coincidence detections mechanisms and, as a result, out synaptic switch rule is broadly consistent with a range of possible biological implementations.

Bibliography

W.C. Abraham, S.E. Mason-Parker, M.F. Bear, S. Webb, and W.P. Tate. Heterosynaptic metaplasticity in the hippocampus in vivo: A bcm-like modifiable threshold for ltp. *Proc. Natl. Acad. Sci.*, 98(19):10924–10929, 2001.

Y. Akaneya, T. Tsumoto, and H. Hatanaka. Brain-derived neurotrophic factor blocks long-term depression in the rat visual cortex. *J. Neurophysiol.*, 76: 4198–4201, 1996.

P.A. Appleby and T. Elliott. Synaptic and temporal ensemble interpretation of spike-timing-dependent plasticity. *Neural. Comp.*, 17:2316–2336, 2005.

W. Bair, C. Koch, W. Newsome, and K. Britten. Power spectrum analysis of bursting cells in area mt in the behaving monkey. *J. Neurosci.*, 14(5):2870–2892, 1994.

C.C. Bell, V.Z. Han, Y. Sugawara, and K. Grant. Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature*, pages 278–281, 1997.

G.Q. Bi and M.M. Poo. Synaptic modifications in cultured hippocampal neurons: Dependence on spike timings, synaptic strength, and postsynaptic cell type. *J. Neurosci.*, 18:10464–10472, 1998.

E.L. Bienenstock, L.N. Cooper, and P.W. Munro. Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *J. Neurosci.*, 2:32–48, 1982.

T.V.T. Bliss and A.R. Gardner-Medwin. Long-lasting potentiation of synaptic transmission in the dentate area of the unanaesthetized rabbit following stimulation of the perforant path. *J. Physiol.*, 232:357–374, 1973.

T.V.T. Bliss and T. Lømo. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.*, 232:331–356, 1973.

J.M. Bradshaw, Y. Kubota, T. Meyer, and H. Schulman. An ultrasensitive ca2+/calmodulin-dependent protein kinas ii-protein phosphatase 1 switch facilitates specificity in postsynaptic calcium signaling. *PNAS*, 100:10512–10517, 2003.

D.V. Buonomano. Distinct functional types of associative long-term potentiation in neocortical and hippocampal pyramidal neurons. *J. Neurosci.*, 19(16):6748–6754, 1999.

R. J. Cabelli, A. Hohn, and C. J. Shatz. Inhibition of ocular dominance column formation by infusion of nt-4/5 or bdnf. *Science*, 267:1662–1666, 1995.

R. J. Cabelli, D. L. Shelton, R. A. Segal, and C.J. Shatz. Blockade of endogenous ligands of trkb inhibits formation of ocular dominance columns. *Neuron*, 19: 63–76, 1997.

G. Carmignoto, T. Pizzorusso, S. Tia, and S. Vicini. Brain-derived neurotrophic factor and nerve growth factor potentiate excitatory synaptic transmission in the rat visual cortex. *J. Physiol.*, 498:153–164, 1997.

G.C. Castellani, N. Intrator, H. Shouval, and L.N. Cooper. Characterizing solutions of a bcm learning rule in a network of lateral interacting non-linear neurons. *Network*, 10:111–121, 1999.

G.C. Castellani, E.M. Quinlan, L.N. Cooper, and H.Z. Shouval. A biophysical model of bidirectional synaptic plasticity: Dependence on ampa and mnda receptors. *PNAS*, 98:12772–12777, 2001.

V. Castellucci, H. Pinsker, I. Kupfermann, and E.R. Kandel. Neuronal mechanisms of habituation and dishabituation of the gill-withdrawal reflex in aplysia. *Science*, 167:1745–1748, 1970.

V.F. Castellucci and E.R. Kandel. Presynantic facilitation as a mechanism for behavioral sensinzation in aplysia. *Science*, 194:1176–1178, 1976.

S. Cohen-Cory and S. E. Fraser. Effects of brain-derived neurotrophic factor on optic axon branching and remodelling *in vivo*. *Nature*, 378:192–196, 1995.

M. Constantine-Paton and M.I. Law. Eye-specific termination bands in tecta of three eyed frogs. *Science*, 202:639–641, 1978.

L.N. Cooper, F. Lieberman, and E. Oja. A theory for the acquisition and loss of neuron specifity in visual cortex. *Biol. Cybern.*, 33:9–28, 1979.

P. Dayan and L.F. Abbott. *Theoretical Neuroscience*. MIT Press, 2001.

D. Debanne, B.H. Gahwiler, and S.M. Thompson. Asynchronus pre- and post-synaptic activity induced associative long-term depression in area ca1 of the rat hippocampus in vitro. *Proc. Natl. Acad. Sci.*, 91(3):1148–1152, 1994.

D. Debanne, B.H. Gähwiler, and S.M. Thompson. Long-term synaptic plasticity between pairs of individual ca3 pyramidal cells in rat hippocampal slice cultures. *J. Physiol. (London)*, 507:237–247, 1998.

D. Debanne, B.H. Gähwiler, and S.M. Thompson. Heterogeneity of synaptic plasticity at unitary ca3-ca1 and ca3-ca3 connections in rat hippocampal slice cultures. *J. Neurosci.*, 19(24):10664–10671, 1999.

D. Debanne, B.H. Gawiler, and S.M. Thompson. Cooperative interactions in the induction of long-term potentiation and depression of synaptic excitation between hippocampal ca3-ca1 cell pairs in vitro. *Proc. Natl. Acad. Sci.*, 93: 11225–11230, 1996.

V. Doyere, B. Srebro, and S. Laroche. Heterosynaptic ltd and depotentiation in the medial perforant path of the dentate gyrus of the freely moving rat. *J. Neurophysiol.*, 77:571–578, 1997.

S.M. Dudek and M.F. Bear. Homosynaptic long-term depression in area ca1 of hippocampus and effects of n-methyl-d-aspartate receptor blockade. *Proc. Natl. Acad. Sci.*, 89:4363–4367, 1992.

S.M. Dudek and M.F. Bear. Bidirectional long-term modification of synaptic effectiveness in the adult and immature hippocampus. *J. Neurosci.*, 13:2910–2918, 1993.

V. Egger, D. Feldmeyer, and B. Sakmann. Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. *Nat. Neurosci.*, 2:1098–1105, 1999.

T. Elliott. An analysis of synaptic normalisation in a general class of hebbian models. *Neural Comput.*, 15:937–963, 2003.

T. Elliott and N. R. Shadbolt. A mathematical model of activity-dependent, anatomical segregation induced by competition for neurotrophic support. *Biol. Cybern.*, 75:463–470, 1996.

T. Elliott and N. R. Shadbolt. Competition for neurotrophic factors: Ocular dominance columns. *J. Neurosci.*, pages 5850–5858, 1998.

D.E. Feldman. Timing-based ltp and ltd at vertical inputs to layer ii/iii pyramidal cells in rat barrel cortex. *Neuron*, 27:45–56, 2000.

R.M. Fitzsimmonds, H.-J. Song, and M.-M. Poo. Propagation of activity-dependent synaptic depression in simple neural networks. *Nature*, 388:439–448, 1997.

R.C. Froemke and Y. Dan. Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature*, 416:433–438, 2002.

S. Fusi. Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. *Biol. Cybern.*, 87:459–470, 2000.

S. Fusi, M. Annunziato, D. Badoni, A. Salamon, and D. J. Amit. Spike-driven synaptic plasticity: Theory, simulation, vlsi implementation. *Neural Comput.*, 12:2227–2258, 2002.

C.F. Gerhald and P.O. Wheatley. *Applied numerical analysis.* Addison Wesley, 1994.

G. J. Goodhill. Topography and ocular dominance: A model exploring positive correlations. *Biol. Cybern.*, 69:109–118, 1993.

G. J. Goodhill and H. G. Barrow. The role of weight normalisation in competitive learning. *Neural Comput.*, 6:255–269, 1994.

B. Gustafsson, H. Wigström, W.C. Abraham, and Y-Y. Huang. Long-term potentiation in the hippocampus using depolarizing current pulses as the conditioning stimulus to single volley synaptic potentials. *J. Neurosci.*, 7:774–780, 1987.

D.O. Hebb. *The organisation of behavior: A neuropsychological theory.* Wiley, New York, 1949.

B. Hille. *Ionic channels of excitable membranes.* Sinauer, 1992.

L.L. Holland and J.J. Wagner. Primed facilitation of homosynaptic long-term depression and depotentiation in rat hippocampus. *J. Neurosci.*, 18(3):887–894, 1998.

D.H. Hubel and T.N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.*, 160:106–154, 1962.

D.H. Hubel and T.N. Wiesel. Binocular interaction in striate cortex of kittens reared with artificial squint. *J. Neurophysiol.*, 28:1041–1059, 1965.

D.H. Hubel and T.N. Wiesel. Anatomical demonstration of columns in the monkey striate cortex. *Nature*, 221:747–750, 1969.

D.H. Hubel and T.N. Wiesel. Functional architecture of the macaque monkey visual cortex. *Proceedings of the Royal Society of London Series B*, 198:1–59, 1977.

Y. Humeau, H. Shaban, S. Bissiere, and A. Lüthi. Presynaptic induction of heterosynaptic associative plasticity in the mammalian brain. *Nature*, 426: 841–884, 2003.

E.M. Izhikevich and N.S. Desai. Relating stdp to bcm. *Neural Comput.*, 15 (1511-1523), 2003.

U.R. Karmarkar and D.V. Buonomano. A model of spike-timing dependent plasticity: One or two coincidence detectors? *J. Neurophysiol.*, 88:507–513, 2002.

R. Kempter, W. Gerstner, and J. van Hemmen. Differential equations for the development of topographical nerve fibre projections. In *SIAM-AMS*, volume 13, 1981.

A. Kirkwood, M. Rioult, and M.F. Bear. Experience-dependent modification of synaptic plasticity in rat visual cortex. *Nature*, 381:526–528, 1996.

C. Koch and I. Segev. *Methods in neuronal modelling: From ions to networks.* MIT Press, 1998.

M.E. Larkum, J.J. Zhu, and B. Sakmann. Dendritic mechanisms underlying coupling of the dendritic with the axonal action potential initiation zone of the adult rat layer 5 pyramidal neurons. *J. Physiol.*, 533:447–707, 2001.

S. LeVay, M.P. Stryker, and C.J. Shatz. Ocular dominance columns and their development in layer IV of the cat's visual cortex: A quantitative study. *Journal of Comparative Neurology*, 179:223–244, 1978.

S. LeVay, T.N. Wiesel, and D.H. Hubel. The development of ocular dominance columns in normal and visually deprived monkeys. *Journal of Comparative Neurology*, 191:1–51, 1980.

S. Löwel. Ocular dominance column development: Strabismus changes the spacing of adjacent columns in cat visual cortex. *J. Neurosci.*, 14:7451–7468, 1994.

J.C. Magee and D. Johnston. A synaptically controlled, associative signal for hebbian plasticity in hippocampal neurons. *Science*, 275:209–213, 1997.

R.C. Malenka and R.A. Nicoll. Long-term-potentiation - a decade of progress? *Science*, 285:1870–1874, 1999.

H. Markram. Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature*, 382:807–810, 1996.

H. Markram, J. Lübke, M. Frotscher, and B. Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science*, 275:213–215, 1997.

S.J. Martin, P.D. Grimwood, and R.G. Morris. Synaptic plasticity and memory: an evaluation of the hypothesis. *Ann. Rev. Neurosci.*, 23:649–711, 2000.

M. Migaud, P. Charlesworth, M. Dempster, L.C. Webster, A.M. Watabe, M. Makhinson, Y. He, M. Ramsay, R.G.M. Morris, and J.H. Morrison. Enhanced long-term potentiation and impaired learning in mice with mutant postsynaptic density-95 protein. *Nature*, 396:433–439, 1998.

K.D. Miller and D.J.C. Mackay. The role of constraints in hebbian learning. *Neural Comput.*, 6:100–126, 1994.

R.G.M. Morris, E. Anderson, G. Lynch, and M. Baudry. Selective impairment of learning and blockade of long-term potentiation by an n-methyl-d-aspartate receptor antagonist, apv. *Nature*, 319:774–776, 1986.

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

C.M. Norris, D.L. Korol, and T.C. Foster. Increased susceptibility to induction of long-term depression and long-term potentiation reversal during aging. *J. Neurosci.*, 16:5382–5392, 1996.

D.H. O'Connor, G.M. Wittenberg, and S.-H. Wang. Graded bidirectional synaptic plasticity is composed of switch-like unitary events. *PNAS*, 102:9679–9684, 2005.

E. Oja. A simplified neuron model as a principal componant analyser. *J. Math. Biol.*, 16:267–273, 1982.

S.H.R. Oliet, R.C. Malenka, and R.A. Nicoll. Two distinct forms of long-term depression coexist in ca1 hippocampal pyramidal cells. *Neuron*, 18:969–982, 1997.

C.C.H. Petersen, R.C. Malenka, R.A. Nicoll, and J.J. Hopfield. All-or-none potentiation at ca3-ca1 synapses. *Proc. Natl. Acad. Sci.*, 95:4732–4737, 1998.

B.D. Philpot, J.S. Espinosa, and M.F. Bear. Evidence for altered nmda receptor function as a basis for metaplasticity in the visual cortex. *J. Neurosci.*, 23: 5583–5588, 2003.

H. Pinsker, I. Kupfermann, V. Castellucci, and E.R. Kandel. Habituation and dishabituation of the gillwithdrawal reflex in aplysia. *Science*, 167:1740–1742, 1970.

D. Purves and J.W. Lichtman. *Principles of Neural Development.* Sinauer, 1985.

H.O. Reiter and M.P. Stryker. Neural plasticity without postsynaptic action potentials: Less active inputs become dominant when kitten visual cortical cells are pharmacologically inhibited. *Proc. Natl. Acad. Sci.*, 85:3623–3627, 1988.

B.L. Sabatini and T.G. Oertner. The life cycle of ca2+ ions in dendritic spines. *Neuron*, 33:439–452, 2002.

R. Sala, A. Viegi, F. M. Rossi, T. Pizzorusso, G. Bonanno, M. Raiteri, and L. Maffei. Nerve growth factor and brain-derived neurotrophic factor increase neurotransmitter release in the rat visual cortex. *Europ. J. Neurosci.*, 10: 2185–2191, 1998.

J.R. Sanes and J.W. Lichtman. Development of the vertebrate neuromuscular junction. *Ann. Rev. Neurosci.*, 22:389–441, 1999.

A. Saudargine and Worgotter F. Porr, B. Local learning rules: predicted influence of dendritic location on synaptic modification in spike-timing-dependent plasticity. *Biol. Cybern.*, 92:128–138, 2005.

A.F. Schinder and M.-M. Poo. The neurotrophin hypothesis for synaptic plasticity. *Trends Neurosci.*, 23:639–645, 2000.

T.J. Sejnowski. Storing covariance with nonlinearly interacting neurons. *J. Math. Biol.*, 4:303–321, 1977.

D.K. Selig, R.A. Nicoll, and R.C. Malenka. Hippocampal long-term potentiation preserves the fidelity of postsynaptic responses to presynaptic bursts. *J. Neurosci.*, 19(4):1236–1246, 1999.

W. Senn, H. Markram, and M. Tsodyks. An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Comp.*, 13:35–67, 2000.

C.J. Shatz and M.P. Stryker. Ocular dominance in layer iv of the cat's cortex and the effects of monocular deprivation. *J. Physiol.*, 281:267–283, 1978.

H.Z. Shouval, M.F. Bear, and L.N. Cooper. A unified model of nmda receptor-dependent bidirectional synaptic plasticity. *PNAS*, 99:10831–10836, 2002.

P. J. Sjöström, G. G. Turrigiano, and S. B. Nelson. Rate, timing and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32:1149–1164, 2001.

P. J. Sjöström, G. G. Turrigiano, and S. B. Nelson. Neocortical ltd via coincident activation of presynaptic nmda and cannabinod receptors. *Neuron*, 39:641–654, 2003.

T.R. Soderling. Cam-kinases: modulators of synaptic plasticity. *Curr. Opin. Neurobiol.*, 10:375–380, 2000.

S. Song, K. Miller, and L.F. Abbott. Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.*, 3:919–926, 2000.

S.B. Tieman and N. Tumosa. Alternating monocular exposure increases the spacing of ocularity domains in area 17 of cats. *Visual Neurosci.*, 14:929–938, 1997.

A. Triller and H. Korn. Transmission at a central inhibitory synapse iii: Ultrastructre of physiologically identified and stained terminals. *J. Neurophysiol.*, 48:708–736, 1982.

T.W. Troyer and K.D. Miller. Physiological gain leads to high isi variability in a simple model of a cortical regular spiking cell. *Neural Comp.*, 9:971–983, 1997.

J.Z. Tsien. Linking hebb's coincidence-detection to memory formation. *Curr. Opin. Neurobiol.*, 10:266–273, 2000.

M.C.W. van Rossum, G.Q. Bi, and G.G. Turrigiano. Stable hebbian learning from spike timing-dependent plasticity. *J. Neurosci.*, 20:8812–8821, 2000.

C. von der Malsburg. Self organisation of orientation selective cells in the striate cortex. *Kybernetik*, 14:85–100, 1973.

J.J. Wagner and B.E. Alger. Gabaergic and developmental influences on homosynaptic ltd and depotentiation in rat hippocampus. *J. Neurosci.*, 15:1577–1586, 1995.

T.N. Wiesel and D.H. Hubel. Effects of visual deprivation on morphology and physiology of cells in the cat's lateral geniculate body. *J. Neurophysiol.*, 26: 978–993, 1963a.

T.N. Wiesel and D.H. Hubel. Single-cell responses in striate cortex of kittens deprived of vision in one eye. *J. Neurophysiol.*, 26:1003–1017, 1963b.

T.N. Wiesel and D.H. Hubel. Comparison of the effects of unilateral and bilateral eye closure on the cortical responses in kittens. *J. Neurophysiol.*, 28:1029–1040, 1965.

J.H. Williams, M.L. Errington, M.A. Lynch, and T.V.P. Bliss. Arachidonic acid induces a long-term activity-dependent enhancement of synaptic transmission in the hippocampus. *Nature*, 341:739–742, 1989.

S.-N. Yang, Y.-G. Tang, and R.S. Zucker. Selective induction of ltp and ltd by postsynaptic $[ca^{2+}]_i$ elevation. *J. Neurophysiol.*, 81:781–787, 1999.

X.-D. Yang and D.S. Faber. Initial synaptic efficacy influences induction and expression of long-term changes in transmission. *Proc. Natl. Acad. Sci.*, 88 (10):4299–4304, 1991.

S.S. Zakharenko, S.L. Patterson, I. Dragarsis, S.O. Zeitlin, S.A. Siegelbaum, E.R. Kandel, and A. Morozov. Presynaptic bdnf required for a presynaptic but not postsynaptic component of ltp at hippocampal ca1-ca3 synapses. *Neuron*, 39: 975–990, 2003.

L.I. Zhang, H.W. Tao, C.E. Holt, W.A. Harris, and M.M. Poo. A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, 395:37–44, 1998.

M. Zhuo, S.A. Small, E.R. Kandell, and R.D. Hawkins. Nitric oxide and carbon monoxide produce activity-dependent long-term synaptic enhancement in hippocampus. *Science*, 260:1946–1950, 1993.