



HAL
open science

A Theoretical and Numerical Study of Certain Dynamical Models of Synaptic Plasticity

David Conal Higgins

► **To cite this version:**

David Conal Higgins. A Theoretical and Numerical Study of Certain Dynamical Models of Synaptic Plasticity. Neurobiology. Ecole Normale Supérieure de Paris - ENS Paris, 2014. English. NNT : 158 . tel-01052580


HAL Id: tel-01052580

<https://theses.hal.science/tel-01052580v1>

Submitted on 28 Jul 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Thèse de Doctorat

En vue de l'obtention du grade de

DOCTEUR DE L'ÉCOLE NORMALE SUPÉRIEURE

École Doctorale
Cerveau, Cognition, Comportement

Discipline ou spécialité:
Neuroscience Computationnelle

Présenté et soutenue par:
David HIGGINS

le 16/7/2014

Titre:
A Theoretical and Numerical Study
of Certain Dynamical Models of Synaptic Plasticity

Unité de recherche UMR 8197

Thèse dirigée par Nicolas BRUNEL et Boris BARBOUR

Membres du jury Walter SENN Dominique DEBANNE
 Mark van ROSSUM Jean-Pierre NADAL
 Nicolas BRUNEL Boris BARBOUR

Numéro identifiant de la Thèse:

*Do mo thuismitheoirí
Tadhg agus Áine*

Abstract

Synaptic efficacy measures the ability of a presynaptic neuron to influence the membrane potential of a postsynaptic neuron. The process of changing synaptic efficacy, via plasticity, is thought to underlie learning and memory in the brain. Focusing on chemical synapses, we examine the abstract rules of synaptic plasticity which determine how changes in synaptic efficacy occur.

Beginning with an atypical, non-Hebbian synapse, the parallel fibre to Purkinje cell synapse, we develop a model which explains the burst frequency and length dependence of this particular synaptic plasticity rule. We present a model based on underlying calcium and NO pathways which accurately unifies much of the experimental literature. This model will be useful in future studies of synaptic plasticity for this synapse and its simplicity will allow for numerical studies involving large numbers of synapses in a network architecture.

We also examine a more typical plasticity rule for neocortical synaptic plasticity, developing analytical tools which accurately predict the behaviour of this synapse model under pre- and postsynaptic Poisson spiking. Building on this analysis we extend the theory to leaky integrate-and-fire (LIF) neurons in a network. We develop theoretical tools which can accurately describe the network response to both constant and transiently elevated noisy external inputs. Utilising these tools we examine the duration of synaptic memories under ongoing background (1/sec) spiking activity both in independent neurons and in a recurrent network. We find that lowering the extracellular calcium concentration extends memory time scales and that the further introduction of a bistability to the synaptic plasticity rule extends this memory time scale by several orders of magnitude. In all cases we provide theoretical predictions of memory time scales which match subsequent simulation comparisons.

Both sets of investigations reveal insights into the processes of learning and subsequent forgetting in the brain. Both models reveal the joint importance of burst frequency and relative spike timing in the induction of memory changes at the synaptic level. Adjustment of model parameters to more closely mimic *in vivo* conditions extends the retention time of memories, under ongoing activity, to biologically relevant time scales. Our work represents a coherent development right through from the biophysical processes of synaptic plasticity to the analytical mean-field level.

Résumé

L'efficacité synaptique quantifie la capacité d'un neurone présynaptique à influencer sur le potentiel de membrane d'un neurone postsynaptique. La plasticité synaptique, regroupant tous les processus d'altération de l'efficacité synaptique, est considérée comme le mécanisme fondamental de mémorisation et d'apprentissage dans le cerveau. Dans ce travail nous examinons des règles d'apprentissage formelles aux synapses chimiques et leurs conséquences sur la mémorisation de patrons appris.

La plasticité de la synapse entre les fibres parallèles du cervelet et la cellule de Purkinje n'est pas de type Hebbien et dépend de la fréquence et de la durée de l'activité présynaptique. Nous avons développé un modèle qui prend en compte les caractéristiques de cette règle d'apprentissage spécifique. Ce modèle, basé sur l'interaction entre des variables de signalisation par le calcium et l'oxyde nitrique (NO) reproduit un large corpus de données expérimentales publiées. La relative simplicité de ce modèle permettra son utilisation dans des simulations numériques faisant intervenir un grand nombre de synapses dans une architecture de réseau.

En parallèle nous avons étudié une règle de plasticité plus typique, telle qu'observée aux synapses corticales, et avons développé des outils analytiques prédisant le comportement de cette synapse modèle dans le contexte de régimes d'activité présynaptique et postsynaptique poissonniens. Nous étendons cette analyse formelle à un réseau de neurones « leaky integrate and fire » (LIF) et développons des outils théoriques qui décrivent la réponse du réseau à des entrées externes bruitées d'amplitude constante ou transitoirement augmentée. Nous utilisons ces outils pour mesurer la durée de rétention de mémoires synaptiques dans un régime de décharge de fond à 1/sec soit dans des neurones indépendants soit dans un réseau récurrent. Nous trouvons que l'abaissement de la concentration de calcium extracellulaire augmente les constantes de temps de rétention de la mémoire. L'introduction d'une bistabilité dans la règle d'apprentissage synaptique rallonge le temps de mémorisation de plusieurs ordres de grandeur. Dans tous les cas nous fournissons des prédictions théoriques sur les échelles de temps de rétention de mémoire qui s'accordent aux résultats de simulations numériques.

Les deux parties de cette étude traitent des processus régissant l'apprentissage et sa rétention dans les circuits cérébraux. Les deux modèles montrent l'importance de la fréquence de décharge et des corrélations temporelles entre potentiels d'action dans l'induction d'apprentissage au niveau synaptique. L'ajustement des variables du modèle pour mimer les conditions physiologiques *in vivo* permet d'allonger la rétention d'apprentissage, dans un réseau soumis à une décharge moyenne continue, sur des échelles de temps biologiquement significative. Notre travail présente une tentative d'unification entre les règles biophysiques détaillées régissant l'apprentissage et

une approche analytique en champ moyen.

Contents

Acknowledgements	xix
1 Introduction	1
1.1 Synaptic plasticity experiments: hippocampus and neocortex	5
1.2 Synaptic plasticity models: hippocampus and neocortex	11
1.3 Synaptic plasticity experiments: cerebellum	20
1.4 Summary	29
2 Parallel-fibre to purkinje cell plasticity	31
2.1 Model Description	32
2.2 Methods	35
2.2.1 Synaptic plasticity model	35
2.2.2 Simulating experimental protocols	38
2.2.3 Optimisation of the model fit	40
2.3 Results	50
2.3.1 Burst length dependence	50
2.3.2 Frequency dependence	53
2.3.3 Safo and Regehr	55
2.3.4 Model comparison	55
2.3.5 Prediction: PF \rightarrow LTD	57
2.4 Discussion	60
3 Maintenance of memory in a cortical network model	65
3.1 Introduction	67
3.2 Materials and Methods	69
3.2.1 Calcium-based plasticity model	71
3.2.2 Probability density function of the calcium concentration	73

3.2.3	Diffusion approximation for the synaptic efficacy with a flat potential	76
3.2.4	Kramers expected escape time from a double-well potential . .	77
3.2.5	Numerical methods: Event-based implementation	78
3.2.6	The network	82
3.2.7	Numerical methods: Network simulations	83
3.2.8	Computing analytically mean firing rates and E-E synaptic efficacy	84
3.3	Results	86
3.3.1	Memory behaviour for a synapse connecting two independent Poisson neurons	87
3.3.2	Memory decay for a bistable synapse	91
3.3.3	Steady-state behaviour of networks of LIF neurons with plastic synapses	93
3.3.4	Memory decay in a recurrent network of LIF neurons	97
3.3.5	Memory induction in a recurrent network of LIF neurons . . .	98
3.4	Discussion	101
4	Conclusion	107
4.1	Cerebellar plasticity model	108
4.2	Hippocampal and neocortical plasticity model	110
4.3	Outlook	113
A	Simulation parallelisation	117
B	Generation of random numbers in parallel	121
C	Shot noise	125
C.1	Mapping calcium dynamics to shot noise	126
C.2	Solving the shot noise equation	128

List of Figures

1.1	Neuronal diagram, including a chemical based synapse. Electrical impulses travel from the soma, down the axons, until they reach synaptic boutons where they induce release of neurotransmitter molecules from synaptic vesicles. These neurotransmitters traverse the synaptic junction and are taken up by postsynaptic receptors, inducing a change in postsynaptic membrane voltage. (From the NIH handbook on Alzheimers; downloaded from Wikipedia.)	2
1.2	Diagram of the processes typically involved in synaptic plasticity. Presynaptically synaptic efficacy can be modified by the regulation of synapsin (1) and RIM protein (2) phosphorylation. Postsynaptically, greater or lesser degrees of AMPARs may be endocytosed (3), their lateral transfer to the synapse is also activity dependent (4) and finally, their phosphorylation state influences synaptic efficacy (5). Synaptic adhesion molecules are also influenced by neuronal activity (6). (Taken from Ho et al. (2011))	4
1.3	Series of diagrams of a number of forms of LTP and LTD. (A) Postsynaptic NMDAR activation leads to insertion of AMPARs in the postsynaptic membrane. (B) Activation of cAMP dependent PKA modifies the function of Rab and RIM changing presynaptic glutamate release. (C) Postsynaptic NMDAR activation can also lead to internalisation of postsynaptic AMPARs. (D) Postsynaptic mGluR activation may also lead to NMDAR internalisation. (E) Retrograde endocannabinoid transmission can lead to decrease in neurotransmitter release. (Taken from Kauer and Malenka (2007))	5
1.4	Bliss and Lømo (1973). Early experimental demonstration of synaptic potentiation. The test pathway demonstrated a potentiation which lasted for hours (black circles) whereas the control pathway demonstrated no change (open circles).	6

1.5	Debanne et al. (1994). The first demonstration of a relative timing dependence between pre- and postsynaptic processes in order to produce potentiation/depression. (A) In a paired pre- post- protocol we observe LTP. (B) However, in an unpaired input pathway, which was stimulated 800ms following the paired input pathway to the same cell, LTD was observed. (C) The LTD was not dependent on a prior stimulation of a the paired pathway but only on the prior depolarisation of the postsynaptic cell.	7
1.6	Bi and Poo (1998). The spike pair based spike timing dependence plasticity timing window. A pre- before a post- spike leads to potentiation. Whereas post- before pre- leads to depression.	7
1.7	Sjöström et al. (2001). If pre- and post- relative timings are random with an average offset of 0ms, then for low frequencies LTD results, whereas for high frequencies LTP results.	9
1.8	Bienenstock et al. (1982). A simulation of a rate based model of synaptic plasticity, with a sliding threshold, is able to explain the development of orientation selectivity in the visual cortex of a cat. (A) Simulation of a normally reared cat develops binocular orientation selectivity. (B) Simulation of a dark reared cat develops a binocular randomly fluctuating response curve. (C) Binocular deprivation (suturing of the eyelids), leads to a disappearance of the response selectivity but the cell remains responsive. (D) Ocular dominance is obtained when one eye is sutured and the other allowed to remain open. (E) Uncorrelated rearing (where the eyes do not receive the same inputs) leads to both monocular and binocular selective states.	12
1.9	Billings and Van Rossum (2009). Development of synaptic weight distributions from two STDP models. (Left column) A non-weight dependent STDP rule (nSTDP) leads to a bimodal distribution in weights under Poisson neuron inputs and a postsynaptic LIF neuron. (Right column) A weight dependent STDP rule (wSTDP) leads to a unimodal distribution of weights, under the same conditions.	15
1.10	Shouval et al. (2002). A calcium-based plasticity rule is capable of reproducing both the postsynaptic depolarisation dependence (A) and frequency dependence (B) of STDP. This particular rule predicts LTD for both positive and negative relative spike timings and LTP only for positive offsets (C).	16

1.11	Graupner and Brunel (2012). This calcium-based plasticity rule can capture a plethora of different theoretical STDP curves depending on the relationship of the calcium influx parameters to the synaptic plasticity thresholds. (A) For different dt the amount of time above the thresholds varies. (B) The mean synaptic efficacy is a result of the proportion of time above the two plasticity thresholds. (C and D) The shape of the STDP curve varies as a function of the pre- and post-synaptic calcium influx parameters and the potentiation/depression plasticity thresholds. The plane is divided up into areas with similarly shaped STDP curves.	19
1.12	Sakurai (1990). Demonstration of calcium dependence of LTD. When the postsynaptic cell is filled with EGTA, LTD is blocked (closed circles).	22
1.13	Crepel and Jaillard (1991). Demonstration of LTD via pairing of parallel fibre stimuli with depolarisation induced Purkinje cell calcium spikes.	23
1.14	Chen and Thompson (1995). (A) For 100 parallel fibre and climbing fibre pairings optimal LTD is obtained for a 250ms delay between parallel fibre stimulus and climbing fibre stimulus. (B) The control pathway. (C) The difference between the test and control pathways.	24
1.15	Lev-Ram et al. (1995). Bath applied NOS inhibitor blocks LTD from PF stimulation (A) but not from uncaged NO in the Purkinje cell (B).	25
1.16	Lev-Ram et al. (2002). (Left) LTP requires 300 parallel fibre pulses. (Centre) Bath applied NOS inhibitor blocks the 1Hz LTP protocol (black circles, subsequently shown to be a postsynaptic protocol) but not the 4Hz protocol (clear circles, a presynaptic plasticity protocol). (Right) Uncaging of NO within the Purkinje cell results in LTP.	26
2.1	Schematic diagram of our understanding of protocols leading cerebellar plasticity. Increasing numbers of parallel fibre stimuli leads to greater release of NO. The coincidence of parallel fibres and climbing fibre leads to a greater increase in postsynaptic calcium. High NO and low calcium leads to LTP (left), whereas high calcium and even a small amount of NO lead to LTD (right).	32

2.2	C-N plane schema for dynamics of synaptic plasticity.	Three zones are defined, a lower-left no change area (white background), an upper-left LTP zone (pink) and an upper-right LTD zone (blue). Time spent in each zone contributes to a weighted change in synaptic efficacy.	33
2.3	Diagrammatic description of model implementation of cerebellar plasticity.	Climbing fibre stimulation leads to high postsynaptic calcium concentration. Parallel fibre stimulation, following a delay, leads to lower calcium and NO increases. Calcium influences a downstream C-pathway and NO an N-pathway. Combined, these pathways lead to synaptic plasticity.	36
2.4	Example of the difficulty of numerical optimisation in the case of a long narrow valley.	The steepest gradient in the displayed landscape is that between the facing walls of the ‘valley’. Due to issues of numerical accuracy it is highly unlikely that a steepest descent method will ever find the floor of the valley and subsequently find the much shallower descent direction along the valley floor. This leads to an oscillatory behaviour which will be very slow to descend towards the true solution. A conjugate gradient method will first follow the direction of steepest descent, to find a minimum in that direction. It then rotates in a conjugate direction and perform a new descent in that direction. In the illustrated example it can be expected to quickly follow the direction of the valley floor. If reduction in cost can no longer be obtained a new local gradient is calculated and the algorithm begins anew from the most recently found minimum point. Adapted from Martens (2010).	44

2.5 **Example of the Bidoret 5ms protocol, combining a Purkinje cell depolarisation with a 200Hz parallel fibre doublet of stimulation.** (A) Linear-N model, parallel fibre stimulation (top row) combined with PC depolarisation (row 2) leads to a crossing of the $C + N = 1$ threshold (red line, row 3), leading to a period of synaptic plasticity (shaded background). C remains greater than θ_D (blue line, row 4) for this period so LTD results. (B) NMDAR-based model, the first parallel fibre spike does not lead to an increase in the N variable. Weight change for a single stimulation pairing is depicted in row 5. In row 6 (bottom row) we compare the experimental results from Bidoret et al. (2009) with the results of our simulation at the end of the protocol. The colours in row 6 correspond to the colours in Fig. 2.7. The numbered points in rows 3 and 4 match similarly labelled points in Fig. 2.6. 51

2.6 **Examples of the behaviour of the Linear-N model under three different protocols.** In green, we see the 5xPF protocol with a 16.6/sec intra-burst frequency. In orange, the 5xPF protocol with a 200/sec intra-burst frequency. In cyan, the Purkinje cell depolarisation combined with a 2xPF doublet at 200/sec. The numbered labels illustrate similar points on the plane as those same labels in Figs. 2.5 and 2.8. 52

2.7 **Degree of LTP/LTD is dependent on both burst length and intra-burst frequency.** Solid lines depict NMDAR-based model results, dashed lines depict Linear-N model. Circles and Squares show experimental results with standard error bars. The two model results which showed zero change (crosses at 0 and 5ms) produce identical results for both model versions. 53

2.8	<p>Example of the Linear-N model where the parallel fibres are stimulated 5 times. (A) Stimulation at 200Hz leads to a crossing of the $C + N = 1$ threshold (row 3, red line), leading to a period of synaptic plasticity (shaded background). C remains less than θ_D (row 4, blue line) so LTP results. (B) Stimulation at 16.6Hz leads to no change, in both the synaptic plasticity model and in experiments. The numbered points on the C and $C + N$ curves match similar label points in Fig. 2.6. The synaptic efficacy response to a single repetition of the induction protocol are shown in Row 5. Experimental results are plotted in bottom row and compared with theoretical predictions of the Linear-N model (blue line) and the NMDAR-based model (cyan line).</p>	54
2.9	<p>Three examples of the Safo and Regehr (2008) protocol using Linear-N model. Climbing fibre to parallel fibre offset of (A) -200, (B) +80 and (C) +400ms. (Top row) Timing of 7 parallel fibre stimuli at 100/s. (Row 2) Timing of climbing fibre induced complex spike. (Row 3) $C + N$ in response to a single repetition of induction protocol. Red line depicts threshold of plasticity. Blue background shading represents periods of LTD, pink shading periods of LTP. (Row 4) Evolution in time of C variable in response to a single repetition of induction protocol. Blue line depicts threshold between LTP and LTD. The numbered points on the C and $C + N$ traces correspond to the similarly numbered points in Fig. 2.10. (Row 5) Evolution of synaptic efficacy variable in response to a single repetition of induction protocol. (Bottom row) Evolution of synaptic efficacy variable for full induction protocol.</p>	56
2.10	<p>Model fit to Safo and Regehr (2008). (A) C-N plane examples of the three protocols used in Fig. 2.9, using Linear-N model, the numbered points correspond to similarly numbered points in that figure. (B) Comparison of the two model predictions with the Safo and Regehr (2008) experimental results for different climbing fibre to parallel fibre offset times.</p>	57

2.11	Model C-N prediction for 40 parallel fibre stimuli at 200/s. (A) Linear-N model evolution of C and N for 40 parallel fibre stimuli. Due to the linear summation in the model increases and decreases follow a straight line to and from the origin. (B) NMDAR-based model evolution for same protocol. The first parallel fibre spike leads to no change in N so the up and down trajectories no longer overlap and the slope is different. In both cases the trajectory spends a considerable amount of time in the LTD region.	58
2.12	Example of the Linear-N model where the parallel fibres are stimulated 40 times. (A) High frequency parallel fibre stimulation (200/sec) leads to an increase in both <i>C</i> and <i>N</i> which leads considerably more LTD than LTP. The model correctly predicted that LTD would result from such a protocol, and indeed the NMDAR-based method correctly predicted the degree of LTD induced (cyan line, see blue line for Linear-N model prediction). (B) Example traces before (black) and after (red) induction protocol. (C) EPSC charge was reduced in all cells following the protocol.	59
2.13	Biochemical cascade leading to LTD postsynaptically at the parallel fibre to Purkinje cell synapse. Adapted from Kawaguchi and Hirano (2013).	62
2.14	Influence of different concentrations of CaMKII availability on the amount of un-phosphorylated AMPA receptors. Simulated result presented in (Kawaguchi and Hirano 2013) which shows the percentage of un-phosphorylated AMPA receptors remaining 30 mins after a conditioning stimulation in the presence of different amounts of available CaMKII.	63

3.1	<p>Dynamics of the synaptic plasticity model with the <i>in vitro</i> and <i>in vivo</i> parameter sets. (A) Pre- and postsynaptic spike trains generated as realisations of Poisson processes at 1/s. (B,C) The spike train in <i>A</i> induces large calcium transients (blue trace) with the <i>in vitro</i> parameter set ($C_{\text{pre}} = 0.562$ and $C_{\text{post}} = 1.240$; see Tab. 3.1). Whenever the calcium trace crosses the depression (cyan) or potentiation thresholds (orange), changes in the synaptic efficacy (green) are induced. (D,E) Same as in <i>B,C</i> but with calcium amplitudes corresponding to the <i>in vivo</i> case ($C_{\text{pre}} = 0.337$ and $C_{\text{post}} = 0.744$). The small calcium transients do not cross the depression/potentiation thresholds and no efficacy changes are observed. Note that the flat potential for ρ is used here and that noise is set to zero for clarity, $\sigma = 0$.</p>	71
3.2	<p>Example of 1/sec shot noise prediction of calcium concentration for <i>in vitro</i> parameter set The probability density, $P(c)$, is shown for individual Poisson processes for the calcium influx value of $C = 0.56$ (red) or $C = 1.23$ (green). When both processes combine at a single synapse we obtain the joint probability density function (blue). Plasticity thresholds are shown as vertical lines.</p>	74
3.3	<p>Possible potentiation and depression threshold crossing cases of the calcium trace (blue lines) between events at time t_i and t_{i+1}. The six possible cases are depicted with respect to the location of the potentiation, θ_P (orange dashed line), and the depression thresholds, θ_D (cyan dashed line).</p>	80
3.4	<p>Example of network firing in asynchronous irregular state. A sample of 1000 excitatory neurons from the network shows irregular spiking behaviour in the raster (top) and the averaged firing rate of all 8000 excitatory neurons is steady around 1/sec (bottom).</p>	83

3.5 **Memory decay for a single synapse with flat potential in the presence of uncorrelated pre- and postsynaptic Poisson firing.** **(A,B)** Temporal evolution of the mean synaptic efficacy in the presence of pre- and postsynaptic Poisson firing at $1/s$ for the *in vitro* (green in *A*) and the *in vivo* (light blue in *B*) parameter sets (mean shown for $N = 1000$ synapses). Blue and red lines show the mean dynamics as predicted by the Ornstein-Uhlenbeck theory. Grey lines show example traces of synaptic efficacy evolution in time. **(C)** Decay time constant as a function of the firing rate for *in vitro* and *in vivo* parameter sets. The blue and red lines show the calculated decay time constant, τ_{eff} , from the OU theory. The points show exponential decay times obtained by fitting single exponential decay functions to the mean synaptic dynamics as shown in *A* and *B* illustrating that the OU theory correctly describes the full model behaviour. The cyan and orange dotted lines illustrate the derived power law behaviour, $\tau_{\text{eff}} \sim 1/\nu^k$, between memory time scales and low firing rates (see text). The power reflects the number of spikes required to cross the plasticity thresholds, that is, $k = 1$ for *in vitro* (cyan dotted line) and $k = 2$ (orange dotted line) for *in vivo* case. **(D)** Asymptotic synaptic efficacy as a function of the firing rate for *in vitro* and *in vivo* parameter sets. The lines show the calculated asymptotic value, $\bar{\rho}$, from the truncated OU theory ($\rho \in [0, 1]$) for *in vitro* (blue line) and *in vivo* (red line) cases. The points show steady-state values obtained by fitting single exponential decay functions to the mean synaptic dynamics as shown in *A* and *B* (green: *in vitro*; light blue: *in vivo*). 89

3.6 **Memory decay for a bistable synapse in the presence of uncorrelated pre- and postsynaptic Poisson firing.** (A) Steady-states of synaptic efficacy as a function of firing rate for the *in vitro* (blue) and the *in vivo* (red) parameter sets. Stable states are shown by solid lines and unstable states by dotted lines. Synaptic efficacy is bistable at low rates ($< 0.04/s$ for *in vitro* and $< 1.3/s$ for *in vivo*) and monostable at high firing rates. The effective potential of synaptic efficacy is shown for three firing rates (0.1/s - magenta line; 1/s - orange line; 2/s - green line) and the *in vivo* parameter set in the inset (firing rates indicated by vertical lines). (B) Decay time constant as a function of the firing rate for the *in vitro* and the *in vivo* parameter sets. For the *in vivo* parameter set below $\sim 1.3/s$, the bistability greatly extends memory time scale compared to a synapse with flat potential (red line) and can be predicted using Kramers escape rate (magenta line). The vertical dashed line illustrates the frequency at the *in vivo* bifurcation point. For the *in vitro* parameter set, the bistability has no influence on decay time constants for firing rates above 0.1/s. The points show exponential decay times obtained by fitting single exponential decay functions to the mean synaptic dynamics. (C) Individual synaptic efficacy traces for the *in vivo* parameter set at 1/s pre- and postsynaptic firing. The synapses remain in the upper potential well for a long time and stochastically cross the potential barrier to the low efficacy state. (D) Averaged synaptic efficacy trace of many synapses for the *in-vivo* parameter set at 1/s. The bistability extends the memory time scale from hours for a flat potential to days. 92

3.7 *See caption overleaf.* 94

3.7 **Steady-state behaviour of a recurrent network with plastic synapses between excitatory neurons.** **(A)** Firing rate mean-field predictions compared with network simulation results. The mean-field theory predicted firing rate is higher (black line) than the actual firing rate of the excitatory neurons (green dots) in the recurrent network of 8000 exc. and 2000 inh. LIF neurons. Network simulation with fixed synapses yield a good match with the theory (blue dots). **(B)** Average synaptic weight prediction compared with asymptotic average synaptic weights in the network simulation. The observed average synaptic efficacy of excitatory to excitatory connections is smaller (mustard dots) than the theoretical prediction (black line). Even when using the asymptotic firing rate of the network in the calculations (green dots), the average synaptic efficacy is overestimated by the theory. **(C)** Mean and standard deviation of synaptic weights vs. firing rate for independent LIF neurons (magenta), networked LIF neurons (green) and LIF neurons in a network in which actual weights are held constant but we examine how their efficacy would have evolved in the presence of observed firing (blue dots). Asymptotic synaptic weights for LIF neurons underestimate the efficacy predicted by the theory (blue line). **(D)** Average synaptic weight vs. firing rate for independent LIFs with different reset potentials. The analytical prediction of the asymptotic synaptic weight based on Poisson firing is shown by the blue line (same as in *C*). The reset potential in the LIF model, V_{reset} , has a marked influence on the observed average synaptic efficacy. Depolarised/hyperpolarised reset potentials (e.g. $-55/-70$ mV, cyan/green dots) lead to an over/under-representation of short ISIs (left/right inset) compared to Poisson neurons (red line in insets). ISI histograms in inset are shown for LIF neurons firing at 1/s. 95

3.8 **Memory decay for a subset of potentiated synapses in a recurrent network with the *in vitro* parameter set.** (A) Temporal evolution of the average excitatory (red) and inhibitory (blue) firing rate. A network of 10,000 LIF neurons is initialised at the theoretically predicted steady-state and simulated for 20 min real time. (B) Temporal dynamics of synaptic efficacies in the network. The majority of synapses are initialised to the theoretically predicted asymptotic synaptic efficacy (mean: magenta; single synapse example: dark gray). A randomly selected subset of 5% are set to 1 at the beginning of the simulation (mean: green; single synapse example: light gray). (C) The exponential decay time constant of the potentiated synapses. The value obtained from fitting a single exponential to the mean decay (green dots) is well approximated by the analytically calculated decay time constant from the OU process (Eq. (3.16)). Introduction of a double-well potential does not modify the memory time constant for the *in vitro* parameter set (orange stars). The slight deviation of the decay time constants with respect to the OU theory, that is, the network decay time constants are slower, are due to the LIF firing statistics as can be seen from the comparison with independent LIF neurons (magenta dots). 98

3.9 **Memory decay for a subset of potentiated synapses in a recurrent network with the *in vivo* parameter set and double-well potential.** (A) Temporal evolution of the average excitatory (red) and inhibitory (blue) firing rate. A network of 10,000 LIF neurons is initialised at the theoretically predicted steady-state and simulated for 120 min real time. (B) Temporal dynamics of synaptic efficacies in the network. The average dynamics of the 95% initialised in the DOWN state (blue) and the 5% initialised in the UP state (red) is shown. The shaded grey region represents the range of values visited by synapses in the UP and in the DOWN state populations, indicating that no transition occurs. 99

3.10	Memory induction for a subset of synapses in a recurrent network with the <i>in vitro</i> parameter set. A network of 10,000 LIF neurons is initialised at the theoretically predicted steady-state and simulated for 20 min real time. (A) Temporal evolution of the average firing rate for a randomly selected population of 100 excitatory neurons (STIM pop, blue). These neurons receive a 3-sec increase in external stimulation after 200 secs, leading to a burst of approx. 50/sec activity. (B) Temporal evolution of the average firing rate of excitatory neurons (excluding the STIM pop., red) and inhibitory neurons (green). (C) Temporal evolution of mean synaptic efficacies for synapses, incorporating the flat potential, which have no connection with the high frequency stimulation population (cyan) and synapses which receive both pre- and postsynaptic stimulation from the high frequency stimulation population (magenta). Stimulation at 50/sec increases mean synaptic efficacy to $\bar{\rho}(50)$ before a return to 1/sec background firing leads to an exponential decay to the 1/sec stationary efficacy with a time scale as predicted for the <i>in vitro</i> parameter set.	100
3.11	Impulse response to a 3 sec high frequency stimulus. Mean synaptic response for an initial value $\rho \in \{0, 0.1, \dots, 1\}$ to a burst of length 3 seconds for frequencies from 1 to 50/sec pre- and post-synaptic rates (blue lines). The average synaptic response assuming a uniform prior distribution of synaptic efficacies (yellow line). Truncated Ornstein-Uhlenbeck predicted stationary synaptic efficacy, $\bar{\rho}$, at each frequency for the <i>in vitro</i> parameter set (black line). At low frequencies the response curves show little deflection, whereas at high frequencies stimulation leads to a decrease in the effective time constant and consequent convergence to the long run synaptic efficacy.	101

List of Tables

2.1	Parameter values used in Linear-N model. Fitted using the optimisation procedures described in Section 2.2.3. For the reasons outlined in the text the following parameters were forced to be equal in the fit: $\tau_c = \tau_n$, $D_c = D_n$, $C_{CS} = C_{\text{depol}}$	47
2.2	Parameter values used in NMDAR-based model. Fitted using the optimisation procedures described in Section 2.2.3. For the reasons outlined in the text the following parameters were forced to be equal in the fit: $\tau_c = \tau_n$, $D_c = D_n$, $C_{CS} = C_{\text{depol}}$	48
3.1	Parameters of the calcium-based synapse model. The <i>in vitro</i> values are obtained by fitting the model to cortical plasticity data (Graupner and Brunel 2012; Sjöström et al. 2001). <i>In vivo</i> calcium amplitudes are scaled from the <i>in vitro</i> values according to the change in extracellular calcium concentration.	73

Acknowledgements

Learning has been a lifelong feature of my life so it seems appropriate that I now study the processes underlying learning. My parents encouraged me from the youngest age by always providing an answer to my favourite question, "Why?" and its close cognate, "How?" Unlike many other children I was always encouraged and never dismissed when asking this question.

My grandparents have served as particular role models in my life. They placed a strong emphasis on learning and scholarship, while also providing living demonstrations of contented industry and happiness.

I have had the good fortune to have had many excellent teachers in my life. I had three teachers during my thesis. Nicolas gave me the greatest investment of his time, for which I am eternally grateful, and taught me his dedication to rigour and fascination with his subject. Boris taught me how to live and work in science. And Stéphane taught me his passion and attention to detail. Outside of these official supervisors I have had the enormous good fortune to work in collaboration with Michael and Mariano each of whom have given me large amounts of their time and infused my work with their energies.

The work presented in this thesis has spanned two continents and no less than four labs. In each lab I have found friends and colleagues who encouraged my interest in neuroscience while providing a supportive social environment which made both the work and the frequent moves a great deal easier. Hongi, Laura, Alessandro, Omri, Claudia, Kasia, Guy, Elisa, Alexis, Sukbin, Loreen, in particular, thank you!

Friends and family are the substance of life. I am extremely grateful to the friends I have made throughout the world, who have accepted to remain supportive and in contact despite my frequent voyages. The past few years would not have been possible without the support of Tim, Anne, Garry and Nadine. Too many others to name have been there for me at just the right time. In particular I want to thank Shaynoor who ferried me to the doctor, when I had pneumonia, during the coldest winter in recorded history in Chicago, while teaching me about life and science in my healthier moments.

Introduction

The brain is generally accepted as the seat of consciousness in the body, and according to the *neuronal doctrine* information is processed in the brain via neuronal activity. Neurons are a series of specialised types of cells which tend to have an elongated morphology, determined by their type, and which have very different electro-chemical properties at different locations along their cell bodies. Neuronal membranes are lipid structures which contain active and passive structures which maintain a potential differential between the interior of the cell and the extracellular fluid. In this thesis we are particularly interested in the synapses through which neurons are joined. Synapses typically fall into two classes, electrical and chemical, electrical synapses involve a porous connection between two neurons through which charge can flow, chemical synapses involve the release of a chemical *neurotransmitter* from one of the cells which then interacts with receptors on the other cell resulting typically in a change in membrane potential.

Due to the ability of neurons to ‘compute’ a location for the storage of memories has long been sought, one long standing hypothesis is that memories might be stored in the connections between them. That is, in the ability of one neuron to influence another through their mutual synaptic connection. The first demonstration of synaptic learning, now called synaptic plasticity, was in *aplysia*, where the slug was able to learn a conditioned response through changing the strength of one particular synapse (Castellucci et al. 1970). Subsequent experimental studies led to the discoveries of long-term potentiation (LTP) and long-term depression (LTD), the ability to modify the influence of a presynaptic spike on the postsynaptic membrane potential, making it either stronger or weaker. Finally, a short-term form of plasticity was also discovered, in which the discrete nature of vesicle release and replenishment gives rise to




Figure 1.1. Neuronal diagram, including a chemical based synapse. Electrical impulses travel from the soma, down the axons, until they reach synaptic boutons where they induce release of neurotransmitter molecules from synaptic vesicles. These neurotransmitters traverse the synaptic junction and are taken up by postsynaptic receptors, inducing a change in postsynaptic membrane voltage. (From the NIH handbook on Alzheimers; downloaded from Wikipedia.)

changing influence of the presynaptic spike on the postsynaptic membrane potential over time periods on the order of the length of a burst. Since the first experimental discoveries much theoretical work has been performed predicting the ability of the brain to carry out computations utilising plastic synapses. Such theoretical studies have led to the exploration of questions of: memory capacity, given a certain number of neurons and a connectivity pattern how many discrete bits of information can be stored; memory representation, is a memory stored in a single synapse, in a series of synapses joining neurons linked in series, or more diffusely throughout a neuronal network; and memory duration, surely if we have the capacity to learn a memory at a given synapse then it is also possible to forget such a memory, either deliberately or due to other random events. The mechanisms of memory encoding or storage will be the topic of Chapter 2, whereas the issue of memory erasure will be the topic of Chapter 3.

An electrical impulse, called an action potential, typically initiates close to the soma and travels down the axon until it reaches the synapses. In the case of chemical synapses this charge induces a release of chemical neurotransmitter from presynaptic vesicles into the synaptic cleft (see Fig. 1.1). This neurotransmitter (most often glutamate or GABA) then diffuses through the cleft and is taken up by receptors, typically chemically keyed to the particular neurotransmitter, both pre- and more importantly postsynaptically. The uptake of the neurotransmitter postsynaptically induces a change in the postsynaptic membrane voltage, which may in turn open voltage-gated ion channels inducing an even greater change in voltage, and this postsynaptic potential goes on to influence the activity of the postsynaptic cell. There are a number of cellular methods by which the ability of the presynaptic cell to influence the postsynaptic membrane voltage may be altered, these are summarised in Figs. 1.2 & 1.3. Presynaptically either the probability that a vesicle will release its neurotransmitter payload (an inherently stochastic process) may be increased or the density of vesicles may increase. Postsynaptically, the expression of receptors such as AMPA receptors may be increased by phosphorylation of CaMKII, this increase in AMPA receptors means that more glutamate can be taken up and more trans-membrane channels can be opened, leading to a greater change in membrane potential.

In this thesis, we will concentrate on biochemical synapses as these have been studied for much longer than electrical synapses and lead to very different properties in the network. We will look at one relatively atypical synapse, the parallel fibre to Purkinje cell synapse in the cerebellum, which does not display the form of learning known as Hebbian learning typically found in other synapses, but rather displays a completely different form of associative learning between a set of ongoing inputs and




Figure 1.2. Diagram of the processes typically involved in synaptic plasticity. Presynaptically synaptic efficacy can be modified by the regulation of synapsin (1) and RIM protein (2) phosphorylation. Postsynaptically, greater or lesser degrees of AMPARs may be endocytosed (3), their lateral transfer to the synapse is also activity dependent (4) and finally, their phosphorylation state influences synaptic efficacy (5). Synaptic adhesion molecules are also influenced by neuronal activity (6). (Taken from Ho et al. (2011))

a motor error signal. We will also analyse a more typical, cortical and hippocampal, learning rule and look at issues of how memories might be maintained in a typical synapse under background firing, which is seen throughout the cortex.

In the rest of this chapter, we will review the experimental evidence for synaptic plasticity, we will then examine the development of models of synaptic plasticity, finally we will review the efforts to date to explain experimentally observed synaptic plasticity at the parallel fibre to Purkinje cell synapse.




Figure 1.3. Series of diagrams of a number of forms of LTP and LTD. (A) Postsynaptic NMDAR activation leads to insertion of AMPARs in the postsynaptic membrane. (B) Activation of cAMP dependent PKA modifies the function of Rab and RIM changing presynaptic glutamate release. (C) Postsynaptic NMDAR activation can also lead to internalisation of postsynaptic AMPARs. (D) Postsynaptic mGluR activation may also lead to NMDAR internalisation. (E) Retrograde endocannabinoid transmission can lead to decrease in neurotransmitter release. (Taken from Kauer and Malenka (2007))

1.1 Synaptic plasticity experiments: hippocampus and neocortex

Initial experiments demonstrating synaptic plasticity were performed by Bliss and Lømo (1973), in which they demonstrated, for the first time, the ability to change the

excitatory postsynaptic potential (EPSP) in a population, by extracellular stimulation of a bundle of presynaptic fibres. The effect was demonstrated in the hippocampus of an anaesthetised rabbit and in a subsequent paper for an unanaesthetised animal (Bliss and Gardner-Medwin 1973). They also demonstrated an increase in the excitability of the postsynaptic population, an effect which is related to but distinct from synaptic potentiation (Daoudal and Debanne 2003). They called the increase in synaptic transmission potentiation. The first demonstration of the opposite effect to potentiation, called depression, was by Lynch et al. (1977) where they demonstrated in hippocampal slice that the potentiation of one set of synapses led to the depression of other synapses terminating on the same set of postsynaptic cells.




Figure 1.4. Bliss and Lømo (1973). Early experimental demonstration of synaptic potentiation. The test pathway demonstrated a potentiation which lasted for hours (black circles) whereas the control pathway demonstrated no change (open circles).

Dudek and Bear (1992) demonstrated that presynaptic activity which does not provoke a postsynaptic spike can lead to synaptic depression. This work illustrated a clear frequency dependence of the effect and also a reversibility via what had previously been seen purely as potentiation protocols, suggesting that potentiation and depression are opposing rather than independent effects. In a follow-up paper,

Dudek and Bear (1993) confirmed that LTP is the reversal of LTD and vice-versa.




Figure 1.5. Debanne et al. (1994). The first demonstration of a relative timing dependence between pre- and postsynaptic processes in order to produce potentiation/depression. (A) In a paired pre- post- protocol we observe LTP. (B) However, in an unpaired input pathway, which was stimulated 800ms following the paired input pathway to the same cell, LTD was observed. (C) The LTD was not dependent on a prior stimulation of a the paired pathway but only on the prior depolarisation of the postsynaptic cell.




Figure 1.6. Bi and Poo (1998). The spike pair based spike timing dependence plasticity timing window. A pre- before a post- spike leads to potentiation. Whereas post- before pre- leads to depression.

In the 1990s there quickly followed a succession of papers, moving from previously exclusively extracellular techniques to intracellular recording, first, and subsequently also stimulation, demonstrating a precise timing dependence to synaptic plasticity.

An early paper (Debanne et al. 1994) demonstrated that pairing a postsynaptic depolarisation with a subsequent presynaptic stimulus leads to LTD (see Fig. 1.5), a result which had long been predicted as a theoretical corollary of Hebb's potentiation rule (Hebb 1949). In the same paper they also confirmed the dependence of LTD on postsynaptic NMDAR activation, suggesting a dependence on postsynaptic calcium influx via voltage-gated calcium channels. The same group followed-up, (Debanne et al. 1996) with further confirmation that, in fact, for single postsynaptic action potentials the precise timing of presynaptic relative to postsynaptic action potentials is extremely important for synaptic plasticity, whereas in the case of bursts the frequencies of spiking are more important. Using dual whole-cell recordings (previous work used extracellular presynaptic stimulation) Markram et al. (1997) demonstrated that presynaptic stimulation alone, which did not produce a postsynaptic action potential, does not lead to changes in synaptic efficacy. Instead there is a requirement for both pre- and postsynaptic activity in order to induce synaptic changes. This work implies a requirement for a backpropagating action potential in the postsynaptic cell for such a mechanism to work. They also further confirmed the pre- before post-paradigm, implied by Hebbian theory, leading to LTP and post- before pre- leading to LTD. Finally, Bi and Poo (1998) in hippocampal cell cultures thoroughly investigated the relative pre- post- timing dependence of synaptic plasticity demonstrating that synaptic changes can be induced by repeated single activation, at a low frequency, of both pre- and postsynaptic action potentials with only the relative timing between the potentials dictating the sign and magnitude of the changes (see Fig. 1.6). Closely separated action potentials were shown to lead to larger changes, while sufficiently largely separated action potentials lead to no synaptic change. Following the Hebb paradigm, pre- before post- synaptic potentials were shown to lead to LTP whereas post- before presynaptic potentials led to LTD.

Markram et al. (1997) reported that at protocol repetition frequencies lower than approx. 10Hz zero synaptic change occurred. Exploring this Sjöström et al. (2001) found that in fact, while a typically LTP pairing protocol leads to zero change at these frequencies, an LTD protocol still leads to depression even at low frequencies. Through a series of experiments they further demonstrated that this frequency dependence, on LTP, can be overcome by generating a larger postsynaptic EPSP, either through activation of more synaptic inputs or via a stronger stimulation, or alternatively via experimental postsynaptic depolarisation. In (Lisman and Spruston 2005) it was postulated that this is likely due to the requirement for a 'build-up' in synaptic current charge, due to repetition, in order to fully release the postsynaptic magnesium block, on the NMDA receptors, allowing for a subsequent opening of the voltage activated calcium channels. They reasoned that the heretofore accepted

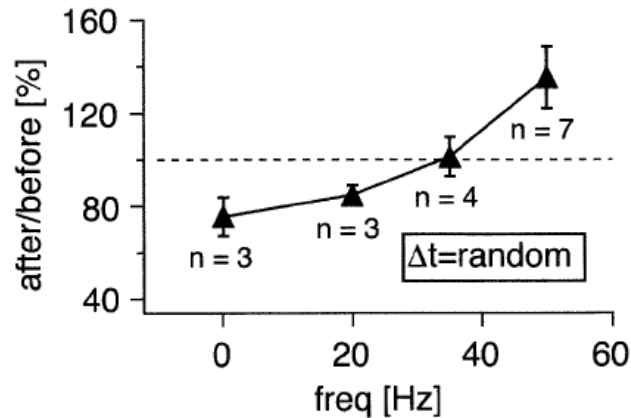


Figure 1.7. Sjöström et al. (2001). If pre- and post- relative timings are random with an average offset of 0ms, then for low frequencies LTD results, whereas for high frequencies LTP results.

influence of the backpropagating action potential may not in fact be an enabler of synaptic plasticity due to its attenuation at higher frequencies. Citing Golding et al. (2002) as evidence, they argue that it is a build-up of postsynaptic dendritic depolarisation which is a necessary signal for synaptic plasticity, rather than an actual postsynaptic action potential.

A number of groups have looked at the effects of higher numbers of pre- and postsynaptic spikes, than that originally examined in the basic STDP experiments, on synaptic plasticity. In (Froemke and Dan 2002) the authors explore a paradigm of a single presynaptic spike sandwiched closely in time between two postsynaptic spikes as well as the inverse setup; two presynaptic spikes surrounding a single postsynaptic spike. They find that the net synaptic plasticity cannot be explained as any kind of linear sum of the independent LTP and LTD effects, which can be expected by a single pre- spike followed by a post- spike and vice-versa. They propose a model, based on their observation that the time difference between the first pre- spike and the first (or only) post- spike is the greatest predictor of synaptic plasticity outcome, which requires that a presynaptic spike suppresses the influence of subsequent presynaptic spikes occurring within a short time window. In (Froemke et al. 2006) they extended this work, examining the effects of a burst of length 5 in combination with an isolated pre- or postsynaptic spike, and subsequently with a burst of equal length. By modifying their original model somewhat they find that suppression of the influence of non-initial spikes in a burst on the plasticity rule appears to explain

the data. Taking a similar approach, this time using cultured hippocampal neurons, Wang et al. (2005) showed a definite asymmetry in the LTP and LTD learning rules. Using a triplet protocol they demonstrated that a typically potentiating spike pair followed by a typically depressing pairing (in this case the postsynaptic spike is shared between both pairings) leads to a cancellation of the two processes, however reversing the protocol order leads to potentiation; suggesting that LTD can cancel LTP but LTP can completely overrule LTD. They then proceed to demonstrate the same effect for a quadruplet of sufficiently closely occurring spikes. Wittenberg and Wang (2006) further explore higher dimensional effects on the CA1 plasticity rule, at low extracellular calcium concentrations. At low frequency stimulation they demonstrate a complete absence of LTP effects, whereas LTD is produced for any closely paired pre- and postsynaptic spikes. Repeating this experiment at 5Hz leads to some cells potentiating whereas others depress for similar relative spike timings, however the addition of a second postsynaptic spike, along with the 5Hz tetanic frequency finally restores LTP for anti-causally oriented spike pairs. Finally, they demonstrate that LTP has a lower repetition dependence than LTD in order to produce effects of a similar magnitude.

Indications that synaptic plasticity saturates were first evident in (Bliss and Lømo 1973) where, furthermore, we see clear evidence that LTP may be an all-or-nothing process. Many subsequent experiments (cf. Debanne et al. (1994); Sjöström et al. (2001)) demonstrated that the two reversible processes, LTP and LTD, appear to have two distinct stable states, at least for the one hour duration following a plasticity induction experiment. In (Petersen et al. 1998) they demonstrated that potentiation of CA3 to CA1 synapses appears to be an all-or-none process, leading to saturation once the binary step-up in synaptic efficacy has occurred. As part of this demonstration it is important to note that different individual synapses have different individual efficacies and different thresholds to potentiation. This means that in a cooperative, distributed, synaptic system we will see synaptic behaviour which often resembles an analogue distribution of synaptic efficacies. The all-or-none nature of bidirectional plasticity was further explored in a pair of papers (O'Connor et al. 2005a; O'Connor et al. 2005b). In the *Journal of Neurophysiology* paper they examine the kinase and phosphatase pathways required for LTP and LTD. In the *PNAS* paper they clearly demonstrate that plasticity events are discrete, with only two states for a given synapse. Furthermore, the individual synaptic states are heterogenous again demonstrating a potential method of graded memory storage via different synapses.

The detailed biomechanical machinery responsible for synaptic plasticity involves a dizzying array of chemicals and pathways. An early substance which was found

to be vital for synaptic plasticity was calcium. Malenka et al. (1988) demonstrated, via a photo-induced increase in postsynaptic calcium concentration, that elevated calcium is a necessary and sufficient condition for postsynaptic LTP. This result had already been indicated in (Lynch et al. 1983) where intracellular EGTA was shown to block LTP. In practice, LTP activation is believed to be via the postsynaptic NMDA receptors, as blocking their activity via APV prevents LTP (Collingridge et al. 1983). In fact, NMDA receptors are believed to be critically implicated in LTP, requiring the coincidence of a presynaptic spike, inducing glutamate release, and postsynaptic depolarisation, releasing the magnesium block in the postsynaptic NMDA receptor, for induction (Bliss and Collingridge 1993). Calcium dependence of LTD was subsequently demonstrated by Neveu and Zucker (1996), again via photolysis of caged calcium compounds. In a follow-up paper, Yang et al. (1999) demonstrated that while LTP was sensitive to brief but high increases in calcium, LTD is induced via longer lasting intermediate levels of calcium increase. From the timing of synaptic plasticity experiments, the fact that elevated calcium concentrations decay much faster than the ultimate induction processes, it is clear that the ultimate actors of synaptic plasticity are downstream of any calcium activation and probably involve either, calcium dependent proteases, protein kinase C (PKC) or calcium-calmodulin dependent protein kinase II (CaMKII).

1.2 Synaptic plasticity models: hippocampus and neocortex

The original model of synaptic plasticity appears, somewhat informally, in Hebb's seminal work (Hebb 1949). It expresses a form of coincidence detection, whereby synapses joining neurons which 'fire together' either come into existence or become stronger; the details are left for the reader to imagine. There is no explicit definition of what happens to existing synapses joining neurons which are not firing together. However, over time it has come to be assumed that synapses which connect neurons which are not 'firing together' probably need to decay in strength and ultimately to disappear. This is in order to avoid a system where synaptic efficacies only grow, leading to increased firing rates and further synaptic strength increases.

In (Sejnowski 1977) the problem of interpreting Hebb's rule was approached using a covariance model. In this it was proposed that the influence of presynaptic activity on postsynaptic membrane potential is governed by the way in which the two membrane potentials or spiking processes covary. Kohonen (1982) proposed a firing rate based approach which saturates in order to avoid unrealistic growth of the




Figure 1.8. Bienenstock et al. (1982). A simulation of a rate based model of synaptic plasticity, with a sliding threshold, is able to explain the development of orientation selectivity in the visual cortex of a cat. (A) Simulation of a normally reared cat develops binocular orientation selectivity. (B) Simulation of a dark reared cat develops a binocular randomly fluctuating response curve. (C) Binocular deprivation (suturing of the eyelids), leads to a disappearance of the response selectivity but the cell remains responsive. (D) Ocular dominance is obtained when one eye is sutured and the other allowed to remain open. (E) Uncorrelated rearing (where the eyes do not receive the same inputs) leads to both monocular and binocular selective states.

synaptic efficacy. A power series expansion of this model was taken by Oja (1982), giving a learning rule which learns due to Hebbian type coincident activity but which has a normalising leak term. This model was then shown to implement a principal component analyser when embedded in a network. An alternative rate based plasticity rule was presented in (Bienenstock et al. 1982) where they showed that the introduction of a sliding threshold to such a rule allows them to account for such developmental traits as orientation tuning curves and ocular dominance in the visual cortex (see Fig. 1.8).

The Hopfield (1982) model of neuronal networks combines binary threshold neuronal units with a network structure which resembles spin-glass models studied in statistical physics, where synapses are updated in response to input patterns, representing memories to be recalled. This model proved tractable to mathematical analysis leading to a calculation for memory capacity in (Amit et al. 1985). They showed that, for a low synaptic noise, memories can be stored with perfect recall until the memories represent more than a critical proportion of the size of the network, after which a catastrophic failure of the memory process may occur. Nadal et al. (1986) showed that by modifying the Hopfield (1982) learning rule new memories can continue to be stored, indefinitely, at the cost of erasure of older memories. Finally, Tsodyks and Feigl'man (1988) proposed a modified learning rule for the Hopfield (1982) network, which featured greatly enhanced storage capacity for correlated patterns, when activity levels are low.

Rate based model development was driven by the experimental state-of-the-art at the time. Experimental results typically involved external stimulation of bundles of axonal fibres and recordings of population dynamics (cf. Bliss and Lømo (1973)). With the development of reliable intracellular recording techniques, the 1990s saw a switch to precise spike timing driven experiments and a commensurate switch in the modelling community. A forerunner however was Gerstner et al. (1993) where they showed that a Hebbian rate-based model, especially in a system with a reasonable variance in transmission delays, was not able to capture temporally coded information, whereas a spike-timing based rule can do so to a very high precision. This was followed up in (Gerstner et al. 1996) where the model was applied to explain the ability of barn owls to localise sound with a very high temporal precision despite apparently longer time constants in their auditory system than the latency between the two input channels. In (Song et al. 2000) they demonstrated that a simple spike-timing dependent plasticity (STDP) rule in a recurrent network naturally leads to selective weight potentiation. Synapses which are most involved in postsynaptic spiking tend to get potentiated whereas other synapses slowly get depressed. This demonstrates not only a clear motivation for the existence of STDP

in synapses but a potential method of homeostasis of incoming weight to a single neuron, without the requirement for a global signal. In order to achieve a weight distribution more closely matching that seen in experiments Van Rossum et al. (2000) introduced a weight dependence on the potentiation term, meaning that stronger weights get potentiated less than weaker weights. This leads to a unimodal distribution but requires an extra ‘leak’ term in order to maintain a target postsynaptic firing rate. The same authors analysed this model further in (Van Rossum and Turrigiano 2001) with the slight alteration, based on experimental observations, where there is no longer a weight dependence on depression, they demonstrate that such a rule leads to a stable distribution of synaptic weights under random background firing and that it will selectively learn for correlated inputs. Examples of the distributions obtained from additive and multiplicative STDP rules can be seen in Fig. 1.9, taken from Billings and Van Rossum (2009), a paper which will have relevance to our work in Chapter 3 as it pertains to memory maintenance and decay in these two models of STDP.

As experiments developed allowing the investigation of precise spike-timing effects on synaptic plasticity, it also became possible to more accurately explore the previously described phenomenon of short-term synaptic plasticity (Zucker 1989). This is a, typically, presynaptic effect, whereby there is a delay in replenishing presynaptic vesicles thus leading to a decrease in the observed effectiveness of the synapse in inducing an excitatory EPSP. The typical model of this process is the depletion-renewal model of Tsodyks and Markram (1997). Short-term facilitation was added to this model in (Tsodyks et al. 1998). Short-term plasticity was developed into a theory of a synaptic mechanism/trace of working memory storage in (Mongillo et al. 2008). In this thesis, we will focus only on models and mechanisms of long-term plasticity, their short-term counterparts are mentioned here for completeness and also to illustrate a point at which the literature separated into short-term and long-term components.

Synaptic memory processes involve discrete processes, such as an increase in vesicle density presynaptically or in AMPA receptors postsynaptically (Lisman 1985). This has led to a number of theories of synaptic plasticity which involve discrete synaptic weights, or at least the use of discrete stable states connected by a continuous potential well. In (Amit and Fusi 1994) learning is modelled by a Markov process, representing the probability of transition between stable states, giving rise to a *palimpsest* memory process capable of continually learning new memories while decay of old memories is dependent on subsequent patterns presented. Fusi et al. (2005) advanced the concept of bistability in synapses to involve multiple internal (meta-)states in a synapse, which may refer to amenability towards plasticity in ei-

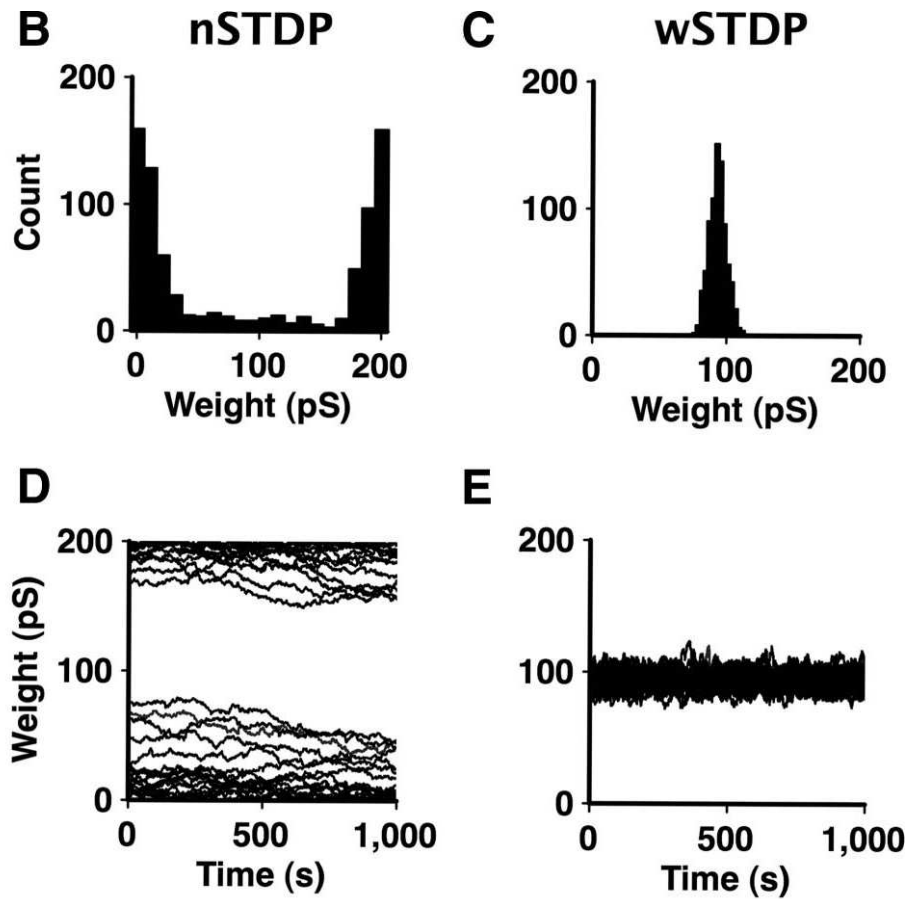


Figure 1.9. Billings and Van Rossum (2009). Development of synaptic weight distributions from two STDP models. (Left column) A non-weight dependent STDP rule (nSTDP) leads to a bimodal distribution in weights under Poisson neuron inputs and a postsynaptic LIF neuron. (Right column) A weight dependent STDP rule (wSTDP) leads to a unimodal distribution of weights, under the same conditions.

ther direction, with two or more discrete states visible to the network. Versions of this model have been found to have properties which contribute to both high memory capacity in a network and long memory time scales (Roxin and Fusi 2013). While not strictly discrete, the strong bistability in the Tag-Trigger-Consolidate model of Clopath et al. (2008) can be interpreted as a binary synapse and the underlying mechanism demonstrates an attractive link between the abstract model of Fusi et al. (2005) and the underlying biology involving early- and late-phase plasticity driven

by protein processes.

When typical synaptic plasticity models are implemented naively they may give rise to pathological behaviour. That is, an increase in synaptic strength will lead to increased postsynaptic spiking/excitability which tends in turn to lead to further increases in synaptic strength. A number of solutions to this problem have been proposed and fall generally under the term homeostatic plasticity or regulation (Turrigiano and Nelson 2000; Turrigiano and Nelson 2004). In brief, this means that there is some kind of normalisation process, which operates across all, or a subset, of the synapses inputting to an individual postsynaptic neuron, generally at a much slower time scale than that of spike-timing dependent plasticity, which leads to a moderation in postsynaptic spiking activity. Such an idea is implicitly present in the sliding threshold, at the individual synapse level, for the BCM rule (Bienenstock et al. 1982). Toyozumi et al. (2005) developed a synaptic plasticity rule based on the principle of maximisation of information between spiking neurons, a process which requires that homeostasis maintains the postsynaptic firing rate as close to its mean as possible, finding that such a rule has the same properties as the BCM rule in the absence of refractoriness. The weight dependent potentiation in the Van Rossum et al. (2000) model can be seen as an early attempt to incorporate such a process into a spike-timing plasticity model. This model does not necessarily prevent run-away excitation but it does at least greatly reduce the speed of its activation allowing for network processes to intervene. The Tag-Trigger-Consolidate model of Clopath et al. (2008) contains a more explicit implementation of homeostatic synaptic plasticity via a process called cross-tagging, a process global to the postsynaptic neuron although it could also be implemented locally to a dendritic branch.




Figure 1.10. Shouval et al. (2002). A calcium-based plasticity rule is capable of reproducing both the postsynaptic depolarisation dependence (A) and frequency dependence (B) of STDP. This particular rule predicts LTD for both positive and negative relative spike timings and LTP only for positive offsets (C).

The initial spike-timing dependent plasticity models were purely phenomenological models, eventually there was a move to base these models more closely on underlying physiological processes. As calcium had been shown to be necessary and sufficient for LTP (Malenka et al. 1988) and LTD (Neveu and Zucker 1996) it was natural for one of the first such models to use calcium as a central element (Shouval et al. 2002). In their model (see Fig. 1.10), Shouval et al. implement calcium influx postsynaptically via NMDA receptors, they do this by assuming that back-propagating action potentials induce long lasting after-depolarisations which release the magnesium blocks which had been preventing calcium entry via the NMDA receptors. This model can, at least qualitatively, fit the postsynaptic depolarisation dependence (Cummings et al. 1996), the frequency dependence (Bliss and Lømo 1973; Dudek and Bear 1992) and precise spike-timing based experiments (Bi and Poo 1998; Markram et al. 1997) in the early literature. It does however predict a late period of LTD plasticity, where a postsynaptic spike follows a presynaptic spike by approx. 100ms, over which there is much debate as to the real existence. The advantage of such a model is that it allows us to make precise predictions about experimentally verifiable variables in order to understand if our suspicions about the processes underlying synaptic plasticity are valid. A more detailed calcium dependent model was presented in (Rubin et al. 2005), in which they explicitly attempt to model the pathways regulating CaMKII in a single spine using a series of differential equations. Two of the equations filter the calcium concentration leading to LTP and LTD of the synapse at different calcium concentrations, with a strong duration dependence. They propose that a system of calcium detectors, which include a duration dependence for LTD, must be combined with a veto process on depression in order to be sufficiently robust to fit all experimental data. Such a veto might be compared with the competition of kinases and phosphatases in the CaMKII phosphorylation-dephosphorylation cascades.

In recent times a number of models have been devised which attempt to go beyond basic pair-wise spike-timing dependent plasticity, to capture higher dimensional features in spike patterns. An early such model was presented in (Sjöström et al. 2001) where they incorporated firing rate, depolarisation and relative spike timing into the synaptic plasticity rule. Fitting the data presented in the same paper they found their best model was one in which a spike which participates in an LTP process should not contribute to a separately calculated LTD process. An alternative approach was presented in (Froemke and Dan 2002) and more thoroughly developed in (Froemke et al. 2006) and was based upon experimental results which appeared in the same papers. In this model the underlying idea is that subsequent spikes in a burst have a lesser effect on the overall plasticity change. A model very much in

keeping with the original idea of pair-wise plasticity was the model of Pfister and Gerstner (2006) which moves to incorporating the three most recent spikes and their relative timings in order to predict plasticity changes. The emergence of this model was clearly influenced by the emergence of triplet based experiments such as Wang et al. (2005). An attractive feature of this model is that it can be directly mapped to a BCM type rule if spiking is assumed to be a Poisson process pre- and postsynaptically (Pfister and Gerstner 2006). A more biologically inspired model was presented by Clopath et al. (2010), in which thresholded low-pass filtered and instantaneous traces of postsynaptic membrane voltage are multiplied by learning constants upon presynaptic spiking. This brings the postsynaptic voltage dependence of synaptic plasticity (Sjöström et al. 2001) to the fore in the model and combines it with traces which can variously be explained as NMDAR activation, endocannabinoid release and postsynaptic calcium concentration. The model we will concentrate mostly upon, in this work, is the calcium-based plasticity rule of Graupner and Brunel (2012), which attempts to combine the direct biological relevance of a calcium-based rule with certain features, such as thresholding of the calcium trace, which capture higher order features of spike trains. This rule is based entirely upon duration above the relevant calcium thresholds in order to dictate plasticity outcomes.

Most of the work cited above developed theories of synaptic plasticity distributions which were then tested only in synapses surrounded by Poisson processes, or in feedforward networks of leaky integrate-and-fire neurons. Extension to simulations involving recurrent networks of spiking neurons has proved difficult for a number reasons, firstly, simulation of plasticity involves an increase in the number of equations from order N to order N^2 , secondly, the dynamics of synaptic plasticity are considerably slower than those of neuronal dynamics. This means that in order to simulate plasticity in neuronal networks it is necessary to run simulations which are not only considerably larger but must also run for much longer, than simulations involving fixed synaptic weights. Song et al. (2000) incorporated a Hebbian style plasticity rule into a recurrent network of leaky integrate-and-fire neurons by biasing the net effect of the rule towards depression, while setting the network activity into a balanced regime thus producing a surplus of presynaptic spikes which do not provoke a postsynaptic spike hence giving a surplus of potentiation. The network inputs and the learning rule parameters were calibrated such that the surplus of potentiation from the balanced network regime was perfectly complemented by the surplus of depression for closely correlated spikes from the learning rule. Mongillo et al. (2005) implemented a spiking network simulation with a two state synaptic plasticity model in order to explore effects of working memory experiments on synaptic plasticity, and vice-versa. Utilising the Poisson theory of neuronal spiking Gilson et al. (2010), one of

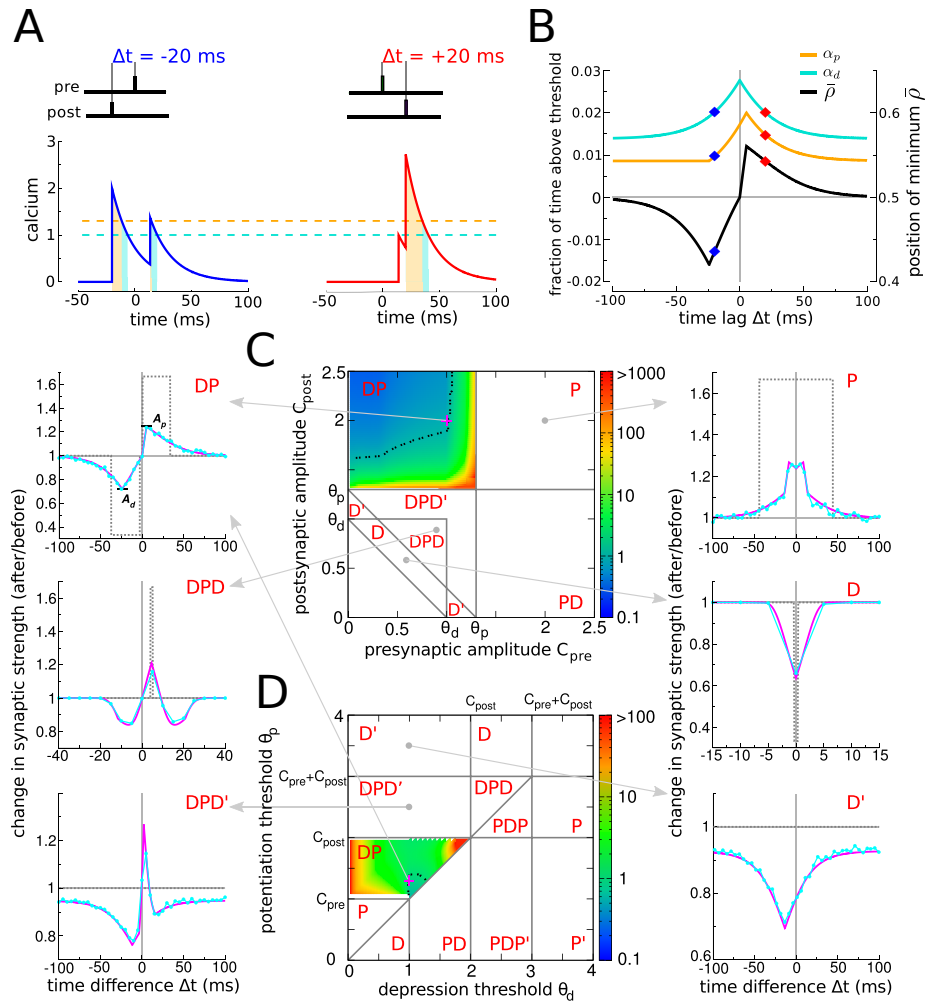


Figure 1.11. Graupner and Brunel (2012). This calcium-based plasticity rule can capture a plethora of different theoretical STDP curves depending on the relationship of the calcium influx parameters to the synaptic plasticity thresholds. (A) For different Δt the amount of time above the thresholds varies. (B) The mean synaptic efficacy is a result of the proportion of time above the two plasticity thresholds. (C and D) The shape of the STDP curve varies as a function of the pre- and postsynaptic calcium influx parameters and the potentiation/depression plasticity thresholds. The plane is divided up into areas with similarly shaped STDP curves.

a series of papers) use a mean-field theory approach to setup a network of recurrently

connected Poisson neurons (Kempster et al. 1999), for which they then make theoretical predictions, with accompanying simulated demonstrations. In particular, they observe the emergence of functional connectivity subpopulations, which correspond to correlated inputs, via symmetry breaking. Synaptic plasticity is not just reserved for excitatory connections, Vogels et al. (2011) study the impact of inhibitory-to-excitatory synaptic plasticity on network dynamics and learning. They demonstrate that inhibitory-to-excitatory synaptic plasticity is capable of reducing the firing rate of cell assemblies thus making them indistinguishable from background network activity, despite their still carrying an encoded excitatory memory trace, and restoring balanced network activity. Finally, Zenke et al. (2013) take a particular Hebbian based learning rule and demonstrate that in a recurrent network this particular rule requires a homeostatic process which operates on the order of seconds to minutes in order to avoid a run-away feedback leading to an explosion in firing rates. As will be shown later however, this behaviour is not a characteristic of all synaptic plasticity rules.

1.3 Synaptic plasticity experiments: cerebellum

Synaptic plasticity at the parallel fibre to Purkinje cell synapse in the cerebellum appears to be fundamentally different from that found in most other brain areas, in particular in the much studied areas of the hippocampus and the neocortex. The Purkinje cell is a particularly large cell with a large planar dendritic tree (Eccles et al. 1967). Parallel fibres typically run orthogonally to this dendritic tree, where they make multiple synaptic contacts. A single climbing fibre comes from the inferior olive and wraps itself around the Purkinje cell proximally to the soma, making strong synaptic contacts which greatly influence the activity of the Purkinje cell. In addition, some inhibitory interneurons contact the Purkinje cell, typically proximally, one class of which has as input the same Golgi cells whose parallel fibres project to the molecular layer leading to a form of feedforward inhibition. When the climbing fibre spikes this leads to a large calcium spike in the Purkinje cell. In contrast, the numerous parallel fibre contacts (of the order of 100,000 per Purkinje cell) have very little effect, on an individual basis, on the Purkinje cell depolarisation. Their effects are greatly attenuated in the large dendritic structure before they reach the soma. Finally, the Purkinje cell fires a complex spike somewhat regularly at a frequency of approximately 1/s and normal spikes at 30-40/s in vivo.

Plasticity at the parallel fibre to Purkinje cell synapse was originally postulated via theoretical motivations (Albus 1971; Marr 1969) before it was identified ex-

perimentally (Ito et al. 1982). Based on the extremely regular architecture of the cerebellum, described as far back as Cajal (1911), and thoroughly described in (Eccles et al. 1967), Marr predicted the importance of the cerebellar cortex in motor learning. In addition, he specified that plasticity at the parallel fibre to Purkinje cell synapse was essential for his model of cerebellar function to be valid. Working apparently in parallel and making use of the same basic approach of the inherent stereotyped structure of the cerebellar cortex, Albus (1971) developed a theory based on the Purkinje cell as a *perceptron* to explain cerebellar function. This theory also led to the implication of the primacy of parallel fibre to Purkinje cell plasticity in cerebellar function. In addition, Albus suggested that plasticity may exist at some of the inhibitory cell synapses but that this was not strictly necessary for the theory to be valid. In Albus' theory the climbing fibre is seen as a stimulus in the classical learning sense (Pavlov), meanwhile the parallel fibres should elicit a pause in Purkinje cell activity (a response), thus the climbing fibre signal should teach the parallel fibre synapses when to depress their synaptic weights. This weight change is opposite in sign to that predicted by Marr but both theories emphasise the location of plasticity to this particular synapse. The first experiments showing long-term depression (LTD) at the parallel fibre to Purkinje cell synapse, and confirming the theoretical prediction of Albus, were presented in (Ito et al. 1982).

Ito et al. (1982) demonstrated that conjunctive stimulation of a vestibular nerve and the inferior olive leads to long lasting (at least one hour) depression of the excitatory effect alone, of the vestibular stimulation, on the Purkinje cell. They further localised the site of plasticity to the Purkinje cell by comparing the application of glutamate locally to the Purkinje cell in conjunction with inferior olive stimulation alone to obtain the same results. This effect was further characterised by Ekerot and Kano (1989), where they found that there was an optimal relative timing between parallel fibre and climbing fibre stimulation in order to induce LTD. They also found an increase in the degree of depression induced by coincident stimulation at 4Hz, postulating that this was due to a saturation of calcium processes in the dendrites. All of this work was produced by extracellular stimulation in decerebrate rabbits.

The first demonstration of *in vitro* parallel fibre to Purkinje cell plasticity was by Sakurai (1987) where, using parasagittal slices of guinea pig brain and extracellular stimulation of the parallel fibres and climbing fibre, they demonstrated a clear long-lasting depression of the postsynaptic potential induced in the Purkinje cell by parallel fibre stimulation. This move to *in vitro* experimentation greatly enhanced the ability of experimentalists to localise the sites of stimulation and consequent plasticity. Furthermore, it allowed for the application of picrotoxin removing the effects of feedforward inhibition. In the same paper, Sakurai also demonstrated a form of

potentiation which decayed after approximately 20 minutes. This potentiation was a forerunner to subsequent discoveries which would uncover multiple forms of LTP in this synapse. Crepel and Jaillard (1991) utilised depolarisation of the Purkinje cell in place of climbing fibre activation to obtain LTD, demonstrating a requirement for calcium spikes in the Purkinje cell, rather than the smaller sodium spikes obtained at a lesser depolarisation. They also found a certain degree of potentiation in their synapses which they could not disentangle, by protocol, from the depressed synapses. Karachot et al. (1994) using a sagittally cut rat cerebellar cortex, demonstrated that the optimal protocol for LTD induction is 300 repetitions of a 1Hz stimulation of the climbing fibre and the parallel fibres, with the climbing fibre ideally following the parallel fibre stimulation.




Figure 1.12. Sakurai (1990). Demonstration of calcium dependence of LTD. When the postsynaptic cell is filled with EGTA, LTD is blocked (closed circles).

The influx of calcium into Purkinje cells was long postulated to be a signal for synaptic plasticity at the parallel fibre synapse (Ito et al. 1982). This was confirmed experimentally in (Sakurai 1990) where, by filling their pipette with a solution containing EGTA, a calcium chelator, they blocked the normal induction of LTD and indeed observed LTP. Shortly thereafter Konnerth et al. (1992), combining imaging techniques with electrophysiology, demonstrated that parallel fibre to Purkinje cell depression is coincident with a brief rise in calcium concentration typically induced by climbing fibre activation. They further repeated this result replacing climbing fibre stimulation with sub spiking-threshold Purkinje cell depolarisation. An analysis of the localisation and time course of the intracellular calcium concentration increase

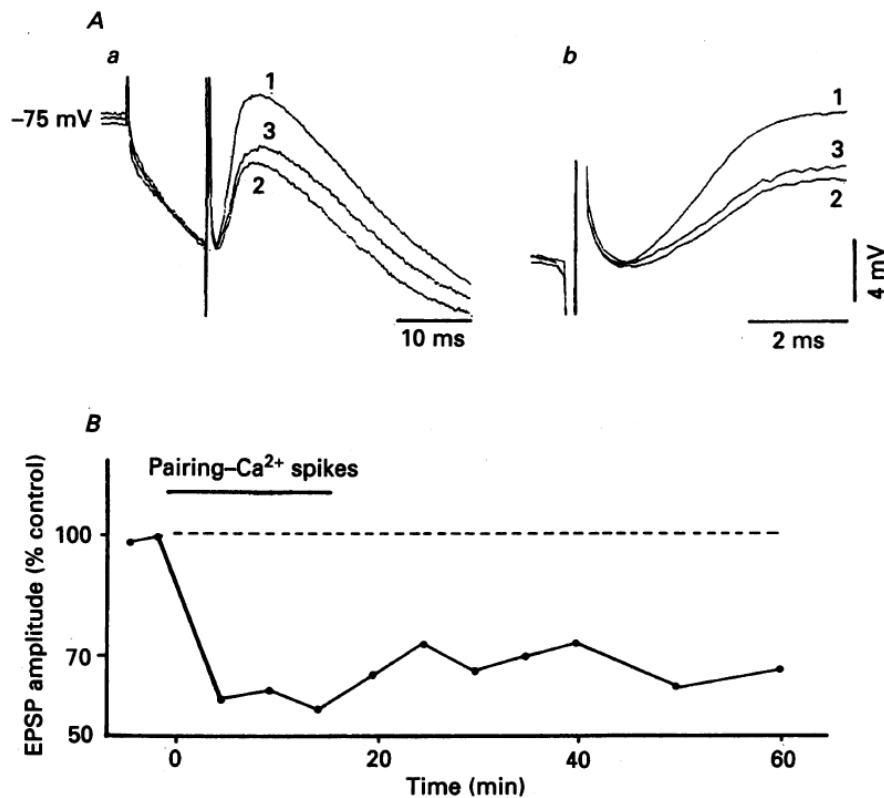


Figure 1.13. Crepel and Jaillard (1991). Demonstration of LTD via pairing of parallel fibre stimuli with depolarisation induced Purkinje cell calcium spikes.

(Eilers et al. 1995) showed that the maximal dendritic calcium concentration occurs much later (35ms) than in the soma.

The relative timing of the climbing fibre and parallel fibre stimuli was explored in (Chen and Thompson 1995). There they showed that in a typical setup, optimal LTD is obtained by stimulating the parallel fibres 250ms before the climbing fibre stimulus. When the blockage of inhibitory interneurons, via picrotoxin, is removed the sharpness of the timing dependency is particularly accentuated, whereas in the presence of picrotoxin it is possible to obtain some degree of LTD for a wide range of offsets. They also looked at the number of stimulus repetitions at a given strength required to obtain LTD, finding that for weak parallel fibre stimuli there is a very strong repetition dependence of LTD. For high numbers of repetitions (600), it is again possible to induce LTD for all the examined inter stimulus intervals. This shows a very strong associative element of parallel fibre to Purkinje cell plasticity. Later,




Figure 1.14. Chen and Thompson (1995). (A) For 100 parallel fibre and climbing fibre pairings optimal LTD is obtained for a 250ms delay between parallel fibre stimulus and climbing fibre stimulus. (B) The control pathway. (C) The difference between the test and control pathways.

using calcium imaging techniques Wang et al. (2000) explored the issue of relative parallel fibre activity to climbing fibre induced Purkinje cell spiking in order to induce LTD. They observed that maximal peak dendritic spine calcium concentration was attained when the climbing fibre activation followed the beginning of the parallel fibre burst by 64ms and, when looking at the integral of the calcium response, maximal concentration was obtained with a 94ms delay.

Lev-Ram et al. (1995) combined the calcium and timing based results and in addition demonstrated a further dependence on presynaptically released NO on postsynaptic plasticity. In this paper, using photolytic induction of NO release internally to the Purkinje cell combined with membrane depolarisation, they demonstrated LTD which could not be blocked by a bath applied NOS inhibitor indicating the target

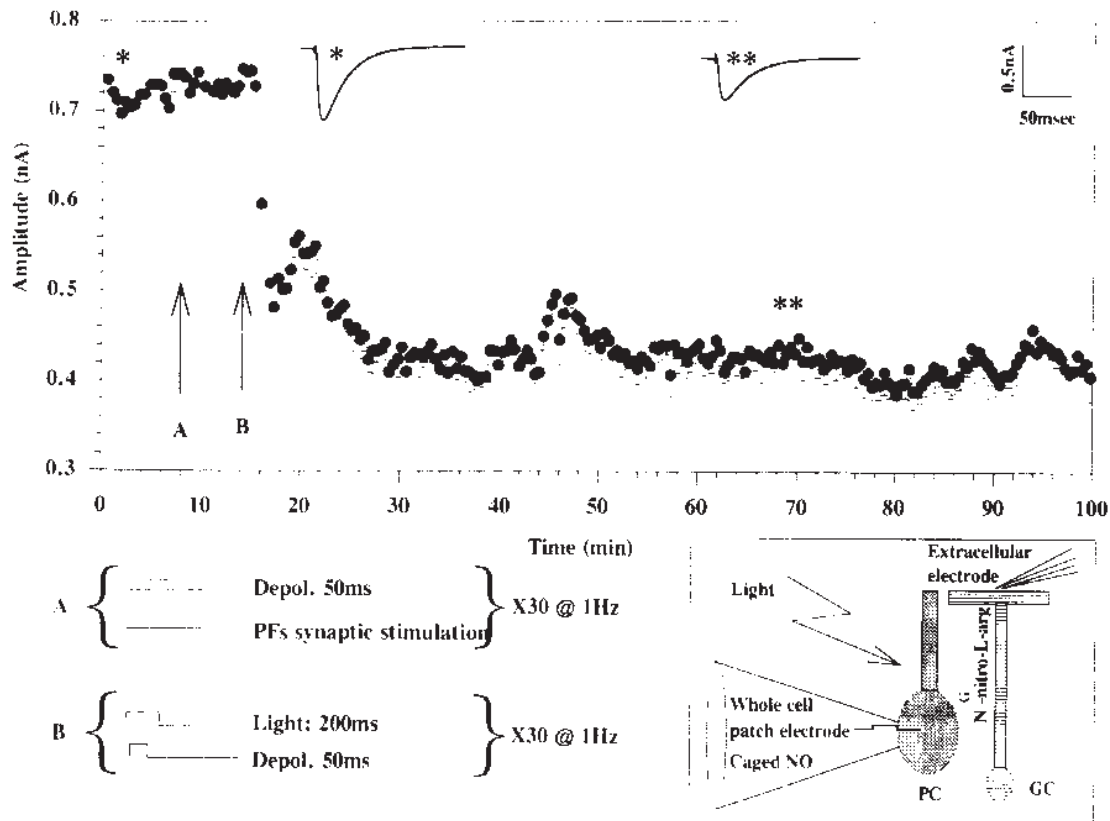


Figure 1.15. Lev-Ram et al. (1995). Bath applied NOS inhibitor blocks LTD from PF stimulation (A) but not from uncaged NO in the Purkinje cell (B).

was within the Purkinje cell (see Fig. 1.15). Furthermore, this effect could be blocked by intracellular BAPTA thus demonstrating the dual dependence on both NO and calcium. They also demonstrated the relative timing dependence, showing that the Purkinje cell depolarisation must precede or be coincident with the NO release in order for LTD to occur. Finally, using a typical LTD protocol with a bath applied NO scavenger and in the absence of photolytic NO release they saw no change in synaptic efficacy, demonstrating that the source of NO is not in the Purkinje cell, but likely presynaptic.

A postsynaptic form of LTP, with the ability to reverse the above postsynaptically expressed LTD, was discovered in (Lev-Ram et al. 2002). They revealed that not only was this form of LTP dependent on a presynaptic NO signal, but furthermore, it is enhanced by a reduction in postsynaptic intracellular calcium concentrations.




Figure 1.16. Lev-Ram et al. (2002). (Left) LTP requires 300 parallel fibre pulses. (Centre) Bath applied NOS inhibitor blocks the 1Hz LTP protocol (black circles, subsequently shown to be a postsynaptic protocol) but not the 4Hz protocol (clear circles, a presynaptic plasticity protocol). (Right) Uncaging of NO within the Purkinje cell results in LTP.

An extracellular NOS inhibitor blocks this postsynaptic LTP but not a presynaptic form of LTP (see Fig. 1.16). Conversely, uncaging of NO within the Purkinje cell induces LTP. The discovery of parallel fibre to Purkinje cell LTP took a long time to discover, relative to the early discovery of LTD, perhaps because of the protocol required for its induction. Lev-Ram found that it takes a full 300 repetitions at 1Hz, but is confounded with presynaptic effects when the frequency is increased, by comparison the LTD protocol requires only 120 repetitions. The mutual reversibility of the LTP and LTD processes was more completely characterised in (Lev-Ram et al. 2003). There they demonstrated that while both postsynaptic LTP and LTD require NO, LTP requires a moderate to low calcium concentration in comparison with a high concentration for LTD.

The NO dependence of LTD was found to be associated with presynaptic NMDA receptors in (Casado et al. 2000). Driven by the observation that the application of NMDA to a slice leads to a depression in the EPSC, they showed that the application of NMDA must be coincident with ongoing parallel fibre spiking, presumably in order to remove the magnesium block, in order to observe this effect. Furthermore the transmission of the presynaptic uptake of NMDA to the postsynaptic domain is blocked by NOS inhibitor (note there are no postsynaptic NMDARs in adult rat Purkinje cells). In (Casado et al. 2002) they showed that it is possible to block postsynaptic LTD by the application of specific NMDA receptor antagonists. This further gives rise to the theory that it must take at least two parallel fibre spikes in close temporal proximity in order to induce LTD, as the first spike would be required

to release the magnesium block in the absence of experimental manipulation. In a later paper from the same group Bidoret et al. (2009) described the presynaptic NMDARs as a high-pass filter of parallel fibre spiking activity. Alternative viewpoints were provided in (Piochon et al. 2010) where they believe that NMDARs exist at the climbing fibre to Purkinje cell synapse and not on the parallel fibre terminals, and (Shin and Linden 2005) where they argue that the NMDARs exist on the stellate cell (interneuron) to Purkinje cell synapses instead and that glutamate spillover from the parallel fibres causes their activation.

Alongside the direct electrophysiological manipulative examination of synaptic plasticity at the parallel fibre to Purkinje cell synapse, there has been a parallel interest in pharmacologically identifying the pre- and postsynaptic receptors and mechanisms of plasticity at this synapse. LTD appears to be caused by internalisation of the GluR2-type of postsynaptic AMPA receptors, thus reducing their ability to pick-up glutamate from the synapse (Matsuda et al. 2000). This is driven by a protein kinase C cascade (PKC) phosphorylating the GluR2 receptor apparently preparing it for clathrin-mediated endocytosis (Wang and Linden 2000; Xia et al. 2000). In the case of LTP, Kakegawa and Yuzaki (2005) showed that this is also a GluR2 process, this time governed by NO-mediated N-ethylmaleimide and not via the more typical calcium-calmodulin protein kinase II (CaMKII) pathway. Belmeguenai and Hansel (2005) further showed that LTP is governed by a phosphatase pathway, dubbing plasticity in this synapse an, "inverse phosphatase/kinase switch," with respect to hippocampal plasticity. See Hansel (2005) and Jörntell and Hansel (2006) for reviews. Sarkisov and Wang (2008) identified the IP3 receptor, previously identified in (Khodakhah and Armstrong 1997) as an essential element in the induction of LTD, with the selection for the order of arrival of parallel fibre vs climbing fibre stimuli. They showed that the activation of the climbing fibre before the arrival of a parallel fibre burst leads to a much larger release of calcium than if the order is reversed.

The *in vitro* examination of parallel fibre to Purkinje cell plasticity is particularly difficult due to the anatomical arrangement of the Purkinje cell dendritic tree orthogonally to the parallel fibres. This makes it particularly difficult to decide how to cut the slice and where precisely to apply a stimulus. Marcaggi and Attwell (2007) showed that the localisation of the stimulus can greatly change the plasticity outcome. Stimulation in the molecular layer leads to a dense bundle of synaptic contacts being stimulated simultaneously, whereas stimulation in the granule layer should lead to a sparse density of contacts on an individual postsynaptic Purkinje cell. When normalising the stimulus strength to produce the same magnitude fast EPSC, Marcaggi and Attwell did not see any long-term plasticity when delivering

the sparser stimulation. This suggests that many of the protocols used heretofore rely on spillover between synapses or a higher stimulus strength leading to much the same end result. It should also be noted that the synapses formed by ascending granule cell axons directly onto the Purkinje cell were shown not to display synaptic plasticity (Sims and Hartell 2006), so it is only the synapses between the parallel fibres and the Purkinje cell which are involved in synaptic plasticity. In terms of stimulation strength Hartell (1996) showed a clear LTD via strong stimulation of the parallel fibres. This effect has been largely ignored in the literature until the results we will present in this thesis.

The first models of parallel fibre to Purkinje cell plasticity were the theoretical works of Marr (1969) and Albus (1971). The perceptron theory of Albus, in particular, became the accepted model for the next thirty years. In a first attempt to improve the explanation as to how activity at the synapse can lead to bidirectional plasticity Coesmans et al. (2004) sketched a calcium concentration rule, whereby high postsynaptic calcium concentration leads to LTD and lower concentration to LTP. This of course, was not necessarily a new idea but rather an assembly of ideas in the literature at the time, and was presented in a review paper without a theoretical implementation. Taking as inspiration the work of Bhalla and Iyengar (1999), modelling molecular signalling pathways explicitly, Kuroda et al. (2001) built and implemented a model of the phosphorylation of AMPA receptors to model LTD. They found that there appears to be a key feedback loop involving PKC and driven by an initial calcium concentration, which must be above a certain threshold concentration, which leads to the phosphorylation of AMPA receptors and a reduction in the number of non-phosphorylated AMPA receptors. In a related paper, Doi et al. (2005) modelled the dynamics of postsynaptic calcium production as induced by parallel fibre and climbing fibre activity. They found that another feedback loop, involving IP3, leads to an accurate representation of the relative spike timing via the calcium concentration, an effect subsequently demonstrated experimentally in (Sarkisov and Wang 2008). Expanding on the detail presented in these models, Kawaguchi and Hirano (2013) present another model of LTD looking at the effects of CaMKII downstream in the plasticity cascade. They demonstrate that, while LTD is typically dependent on both CaMKII and NO acting synergistically, in the absence of CaMKII NO is sufficient to drive LTD. To our knowledge nobody has yet attempted to model LTP at this synapse beyond the sketched calcium theory of Coesmans et al. (2004).

1.4 Summary

In this thesis we will examine synaptic plasticity from the single synapse right up to the generic large-scale network level. In Chapter 2 we will develop a model for the relatively atypical parallel fibre to Purkinje cell synapse. This model accurately describes synaptic plasticity outcomes for a series of experiments, from the literature, performed under comparable experimental conditions. We further test the model by making predictions which were subsequently tested experimentally. The ultimate goal of this model is to explain in the most reduced form possible the necessary factors for predicting the results of an experimental protocol. In addition, we believe that the model we develop is compact enough to allow it to be implemented on a large-scale thus allowing for predictions of the global learning behaviour of a module of the cerebellar cortex.

In Chapter 3 we switch to a more traditional synaptic plasticity paradigm, using the Graupner and Brunel (2012) model of synaptic plasticity and a parameter set fitted to a set of cortical slice experiments (Sjöström et al. 2001). This model has been shown to be able to reproduce the experimental results obtained cortical synapses. We will examine how the synapses might be expected to behave under ongoing background activity. This is especially relevant when we consider that *in vivo* cortical neurons are rarely silent. But if they are continuously active, and yet typical STDP rules are influenced by every spike, how is it that memories are not destroyed? In a thorough analysis of the synaptic plasticity model, we show first theoretically and then using numerical simulations that, under the conditions of reduced extracellular calcium concentration typically found *in vivo* combined with a bistability in the model rule, memories can in fact be expected to be sustained for exceedingly long durations even under ongoing random background activity.

Parallel-fibre to purkinje cell plasticity

The cerebellum is a particularly interesting structure to study. As one of the older parts of the brain it can be expected to have mechanisms which date back many millions of years and which may be found in species from other genetic phyla. The primary function of the cerebellum appears to be in the correction of motor commands to compensate for a changing reference frame (e.g. the Vestibulo-Ocular response (VOR)). A non-mammal with a similar structure is the mormyrid, with its ‘cerebellar-like’ structure, which appears to use this structure to correct for the impact of its own electric field on its electric field detection organ (Bell and Emde 1995).

The anatomical structure of the cerebellum follows a rigidly repetitive pattern of purkinje cells and their granule cell afferents. This repeated structure, exhibiting very little variation, suggests that some form of canonical computation is being performed massively in parallel. What exactly this computation is remains to be discovered, but it is clear that incoming connectivity to the cerebellum is highly locally targeted and that the computation appears to use plasticity at the parallel fibre to purkinje cell synapse to perform some form of highly precise synchronisation of inputs.

We have chosen to construct a model of synaptic plasticity at the parallel fibre to purkinje cell synapse as we believe that the time is ripe to resolve much of the confusion in the literature and thus lead to a greater understanding of cerebellar function. Experiments to date have shown much disagreement on the exact protocols required in order to induce long-term depression (LTD) or potentiation (LTP) at this synapse. Progression of experimental techniques is the most likely source of this problem, Marcaggi and Attwell (2007) showed that the exact locus of stimulation leads to large differences in the outcome of cerebellar plasticity experiments.

Furthermore, there has been great disagreement regarding the existence of and necessity for active presynaptic NMDA receptors in the literature, with suggestions that they may instead exist on neighbouring intraneuron to purkinje cell synapses (Shin and Linden 2005). An imaging study has recently shown the active existence of these receptors on the parallel fibre termini (*under submission*). Here we are able to study the potential impact of these receptors on NO production and the consequent downstream postsynaptic plasticity processes.

2.1 Model Description

Our development of a model for parallel fibre to purkinje cell synaptic plasticity was influenced, in part, by the observation that parallel fibre stimulation leading to LTP is blocked by extracellular NO buffers (Lev-Ram et al. 2002), whereas climbing fibre stimulation, coincident with parallel fibre stimulation, leads to larger postsynaptic calcium transients (Wang et al. 2000) and subsequent LTD. This understanding is summarised in Fig. 2.1, where we see five parallel fibre stimuli leading to a larger release of NO, whereas only two parallel fibre stimuli leads to less NO but when combined with a climbing fibre stimulus a larger increase in postsynaptic calcium results.

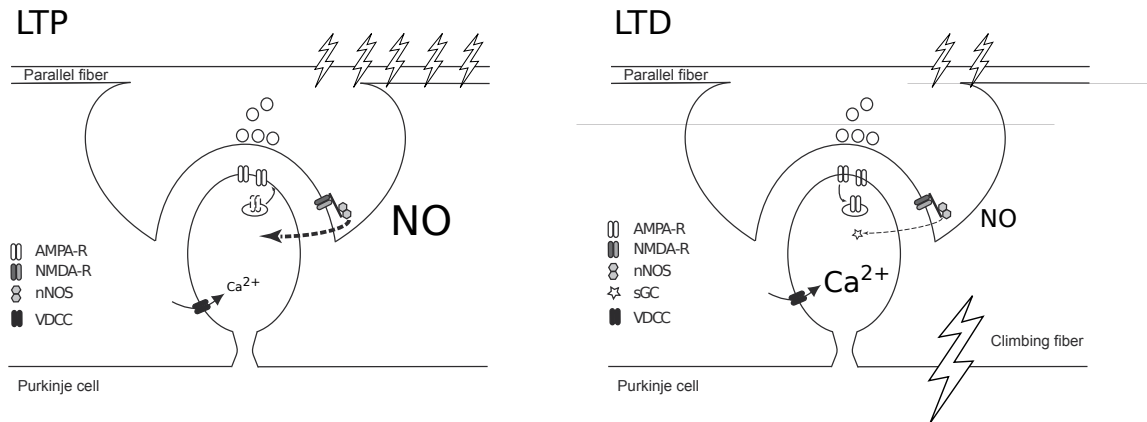


Figure 2.1. Schematic diagram of our understanding of protocols leading cerebellar plasticity. Increasing numbers of parallel fibre stimuli leads to greater release of NO. The coincidence of parallel fibres and climbing fibre leads to a greater increase in postsynaptic calcium. High NO and low calcium leads to LTP (left), whereas high calcium and even a small amount of NO lead to LTD (right).

Two controlling variables C and N are used to define a two-dimensional plane and the temporal variations of synaptic state can be represented as trajectories within this

plane (see Fig. 2.2). The first variable is a calcium-dependent cascade (C-pathway) which, in part enables plasticity, but more importantly, dictates the sign of plasticity. The second variable is a NO-dependent cascade (N-pathway) which, when combined with calcium, enables plasticity. The plane was divided into regions in which LTP, LTD or no plasticity occurs, reflecting the following properties of plasticity at the parallel fibre to purkinje cell synapse: (i) Except in the presence of extreme calcium elevations, NO is required for both LTP and LTD (Ito and Karachot 1990; Lev-Ram et al. 1995, 2002) (ii) The calcium elevation in protocols leading to LTP is believed to be lower than that in protocols leading to LTD (Coesmans et al. 2004); (iii) LTD can be induced by an increase in calcium alone (Finch and Augustine 1998; Miyata et al. 2000; Tanaka et al. 2007).




Figure 2.2. C-N plane schema for dynamics of synaptic plasticity. Three zones are defined, a lower-left no change area (white background), an upper-left LTP zone (pink) and an upper-right LTD zone (blue). Time spent in each zone contributes to a weighted change in synaptic efficacy.

When compared to LTD, long-term potentiation appears to depend on stronger

NO signalling (longer parallel fibre bursts) and weaker purkinje cell calcium signals (no climbing fibre activity) (Coemans et al. 2004; Lev-Ram et al. 2002). Our model does not rule out that LTP can be induced by strong NO signalling without any calcium rise, but experimental evidence of this point is unclear. The experiments that have come closest to testing this hypothesis involved application of an NO donor during normal test stimulation (which could still induce a small calcium signal) (Kakegawa and Yuzaki 2005). In these conditions, LTP was produced. For simplicity, we have therefore allowed a high N to induce LTP in the absence of a calcium rise.

We chose an implementation in which the degree of LTP and/or LTD is proportional to the time spent in the respective regions of the C-N plane. Thus, if the synapse crosses into the LTP region from the "no change" region, but does not reach the LTD region, only LTP occurs. Synapses visiting both LTP and LTD regions express a resultant change reflecting the difference of the time spent in the two regions, weighted by potentiation and depression rates. In simplified mathematical form then, the dynamics of synaptic efficacy can be written as a first-order differential equation

$$\frac{d\rho}{dt} = (1 - \rho)\gamma_P\Theta(\text{LTP}) - \rho\gamma_D\Theta(\text{LTD}) \quad (2.1)$$

where $\Theta(\text{LTP})$ represents presence in the LTP zone and $\Theta(\text{LTD})$ presence in the LTD zone, and $\gamma_{P/D}$ are the potentiation and depression learning rates respectively. The weight change is also dependent on the current synaptic efficacy, via $(1 - \rho)$ and ρ , which introduce bounds on the synaptic efficacy model variable at 1 and 0 respectively.

The boundaries of the LTP and LTD regions represent model nonlinearities that can be considered as corresponding to the activation of different mediators (i.e. kinases and phosphatases) involved in the signalling pathways leading to plasticity. The final result of a given parallel fibre and climbing fibre activity pattern will depend on the relative amount of time spent in the different plasticity regions (LTP or LTD). This integration in time may represent the dynamics of synthesis or accumulation of mediators in signalling pathways.

It is now generally agreed that optimal LTD induction occurs when parallel fibre activity precedes climbing fibre activity, by approximately 100ms (Ekerot and Kano 1989; Safo and Regehr 2008; Wang et al. 2000). To account for this, we introduced a delay between the parallel fibre action potential and the rise of the C signal. This delay could reflect lags introduced by downstream signalling pathways. Very little is known about the time course of the NO signal. For simplicity, we kept the same delay between the parallel fibre action potential and the rise of the N signal as for the C signal, but the model behaves similarly if the N signal is triggered by parallel fibre activation without delay.

The threshold on plasticity enforces the high-pass filter on parallel fibre activity seen in both LTP and LTD experiments, because high-frequency activity is necessary for effective summation to reach the threshold. The sign of plasticity is determined by the calcium level, which was previously depicted as an inverted BCM rule (Bienenstock et al. 1982; Coesmans et al. 2004). Both signals are typically required for both forms of plasticity (Lev-Ram et al. 1995, 1997, 2002).

2.2 Methods

We will now present the Methods used in the development and analysis of our model of parallel fibre to purkinje cell plasticity. We will begin by describing the model, which is dependent on underlying calcium and NO dependent cascade dynamics. We will then proceed to explain how the model can be applied to existing experimental data. Finally, we will describe a series of numerical optimisation techniques which we used to fit the parameters of the model to a set of experiments.

2.2.1 Synaptic plasticity model

We have developed a model of parallel fibre to purkinje cell plasticity which consists of a single dynamical equation describing synaptic efficacy

$$\frac{d\rho}{dt} = (1 - \rho)\gamma_P\Theta(C + N - 1)\Theta(\theta_D - C) - \rho\gamma_D\Theta(C + N - 1)\Theta(C - \theta_D) \quad (2.2)$$

which is dependent on two underlying dynamical equations, one describing calcium dependent C-pathway activation, C , and the other NO dependent N-pathway activation, N , acting in combination via Heaviside functions, $\Theta(x)$

$$\Theta(x) = \begin{cases} 1, & \text{if } x > 0 \\ 0, & \text{otherwise} \end{cases} \quad (2.3)$$

Two of the Heaviside functions institute a threshold dependence on plasticity, in which $C + N$ must be greater than 1 (i.e. $\Theta(C + N - 1)$) in order for plasticity to occur. The further two Heaviside functions indicate the sign of plasticity. In the case of LTP, $\Theta(\theta_D - C)$ means that potentiation only occurs when C is less than θ_D . For LTD, $\Theta(C - \theta_D)$ means that C must be greater than θ_D . Now, when C and N are in the LTP region of the C-N plane we see a weight dependent potentiation, which implements a soft bound at 1. Similarly, when C and N are in the LTD region we see a weight dependent depression, leading to a minimal synaptic efficacy at 0. γ_P and γ_D are the rates of learning for potentiation and depression respectively.

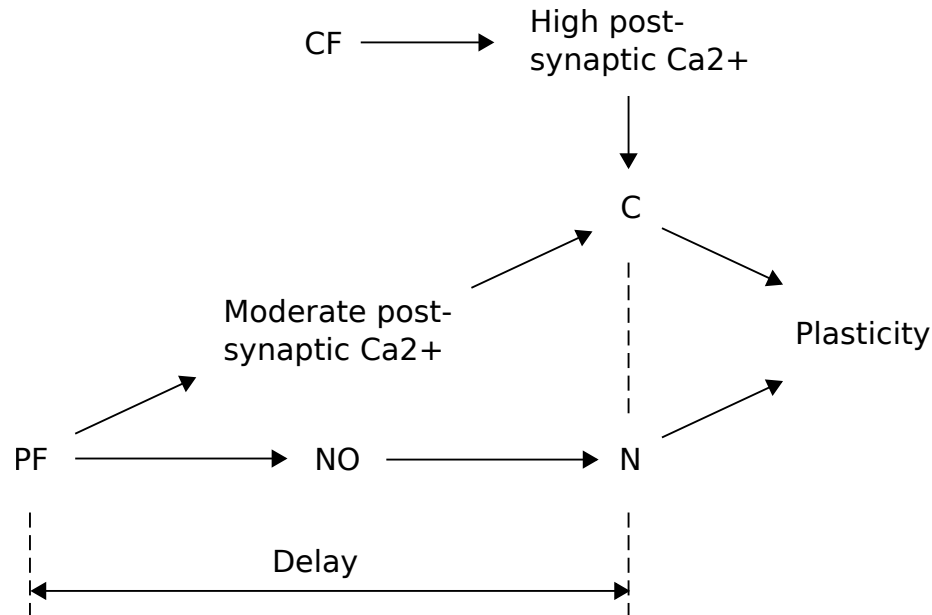


Figure 2.3. Diagrammatic description of model implementation of cerebellar plasticity. Climbing fibre stimulation leads to high postsynaptic calcium concentration. Parallel fibre stimulation, following a delay, leads to lower calcium and NO increases. Calcium influences a downstream C-pathway and NO an N-pathway. Combined, these pathways lead to synaptic plasticity.

2.2.1.1 Calcium dependent pathway dynamics

The calcium dependent pathway is activated by both pre- and postsynaptic spiking activity. When a parallel fibre spike (action potential) occurs, there is typically an influx of calcium into the purkinje cell via voltage gated ion channels. We model this as an increase in the C-pathway, of size C_{PF} , following a delay, D_c . The delay accounts for multiple delays in the system, the first of which is the opening of the voltage-gated ion channels, but most likely is dominated by the time required for the positive feed-back loop found in (Kuroda et al. 2001) to stabilise, leading to maximal LTD, and was fitted to the experimentally observed optimal delay as reported in (Safó and Regehr 2008). In the absence of other activity the C-pathway variable decays exponentially to 0 with a time constant τ_c .

Due to variations in the experimental techniques we have two methods of modelling postsynaptic activity effects on the C-pathway. In the case of a climbing fibre induced complex spike, we model C using an immediate increase of size C_{CS} . In this

case, the full calcium pathway dynamics are described by the equation

$$\tau_c \frac{dC}{dt} = -C + \tau_c C_{\text{PF}} \sum_j \delta(t - t_j - D_c) + \tau_c C_{\text{CS}} \sum_i \delta(t - t_i) \quad (2.4)$$

where $\delta(x) = 1$ when $x = 0$, t_j are the times of parallel fibre spikes and t_i the times of postsynaptic complex spikes.

Experimentally it is often easier to replace climbing fibre stimulation by purkinje cell depolarisation when inducing LTD. In reality this is likely to produce extremely nonlinear effects in postsynaptic calcium activity and their downstream effects on the C-pathway. In the absence of alternative data we have modelled this process as an immediate increase in C to a plateau level, C_{depol} , above which it is maintained for the duration of the depolarisation, with normal parallel fibre induced changes still fully operational above this level but with a floor C value at C_{depol} . This results in two equations, the first describing C dynamics in the absence of depolarisation, and at the moment of depolarisation

$$\tau_c \frac{dC}{dt} = -C + \tau_c C_{\text{PF}} \sum_j \delta(t - t_j - D_c) + \tau_c C_{\text{depol}} \sum_i \delta(t - t_i) \quad (2.5)$$

and the second, describing C dynamics during depolarisation

$$\tau_c \frac{dC}{dt} = -(C - C_{\text{depol}}) + \tau_c C_{\text{PF}} \sum_j \delta(t - t_j - D_c) \quad (2.6)$$

2.2.1.2 NO dependent pathway dynamics

The production of NO in the cerebellum has been shown to be external to the purkinje cell. Two main theories are that the NO is either: produced in the parallel fibre termini via an initial deactivation of the magnesium block on the presynaptic NMDA receptors and a subsequent calcium induced activation of presynaptic NOS (Casado et al. 2002); or that, the NO is produced via glutamate spillover activation of nearby interneuron AMPA receptors (Shin and Linden 2005) (the potential mechanics of which are less well understood). We have developed two models of NO production, the first a simple linear model (**Linear-N Model**) whereby the N-pathway increases linearly in response to parallel fibre activation and decays exponentially with a time constant τ_n via the first-order differential equation

$$\tau_n \frac{dN}{dt} = -N + \tau_n N_{\text{PF}} \sum_j \delta(t - t_j - D_n) \quad (2.7)$$

Following a delay, D_n , after a parallel fibre spike at time t_j , we see an increase of size N_{PF} in the N variable. This production process should not be considered a rejection of the NMDAR dependent model of NO production, but rather a linearisation of any potential NO production process which may have beneficial properties for analysis and simulation.

In order to model parallel fibre NMDAR dependent NO production more explicitly we have a second model (**NMDAR-based Model**) in which the activation level of the presynaptic NMDA receptors is described via

$$\frac{dV}{dt} = -\frac{V}{\tau_\nu} + \sum_j (V_{\text{max}} - V) V_{\text{PF}} \delta(t - t_j - D_\nu) \quad (2.8)$$

Here, a parallel fibre spike occurring at time t_j leads to an increase in NMDAR activation level, proportional to the current value V and by a maximal jump value V_{PF} , up to a maximal activation value V_{max} . In practice we have implemented $V_{\text{max}} = 1$ and $V_{\text{PF}} = 1$ meaning that a single spike, following a delay D_ν , saturates the NMDA receptors. This is a simplification but seems reasonable as in many experiments it is likely that a beam of parallel fibres are being stimulated leading to glutamate spillover and consequent receptor saturation. The delay, D_ν , accounts for the activation time constant of the receptor and results in a production of NO starting from the following spike.

NO production is multiplicatively dependent on the activation state of the NMDA receptors, via

$$\tau_n \frac{dN}{dt} = -N + \tau_n \sum_j N_{\text{PF}} V(t_j - D_n) \delta(t - t_j - D_n) \quad (2.9)$$

A parallel fibre spike at time t_j leads to an increase in the N-pathway variable, N , after a delay D_n , of magnitude $N_{\text{PF}} V(t_j - D_n)$. Note that the delay $D_\nu \ll D_n$, meaning presynaptic changes in NMDAR activation occur orders of magnitude faster than their influence on postsynaptic processes via NO production. N decays exponentially with a time constant τ_n in the absence of other activity.

2.2.2 Simulating experimental protocols

From the broad literature on synaptic plasticity experiments at the parallel fibre to purkinje cell synapse we selected a coherent subset, which were performed under similar experimental conditions with respect to animal type (rat), age (3 weeks), and experimental operating temperature (32° celcius). In general, we expect our model

to be expandable to account for experiments conducted under different conditions but, particularly in the case of temperature, we would expect to need to re-fit the parameters in order to fit such experiments correctly. The experiments which we fitted are some ‘in-house’ experiments which will be published alongside this model, the LTD experiments presented in (Bidoret et al. 2009), and a protocol presented in (Safo and Regehr 2008) which helped enormously to constrain the parameters. In total, we have 17 experimental data points produced under differing protocols, each of which has an associated experimental standard error bar to fit.

2.2.2.1 Fitted experiments

In order to constrain the effects of parallel fibre activity, in terms of both repetition number and frequency, on plasticity we were able to use the following ‘in-house’ experimental results. At 200/sec intra-burst stimulation it requires a burst of 5 spikes in order to reliably induce LTP, whereas a burst of 2 or 3 spikes leads to no discernible change from baseline synaptic efficacy. As the 2 spike protocol, as seen through our model, is a reduction of the 3 spike protocol, both of which leading to a "no change" result, we could only include one of these protocols in our parameter optimisation procedure (Section 2.2.3); since the 2 spike protocol showed much smaller standard error (5% vs 19%) we chose to be conservative and use it in our fitting procedure. As will be seen in Section 2.3.1 the resulting fits still showed no synaptic change for the 3 spike protocol. Dropping the intra-burst frequency, of a burst of 5 spikes, to 33.3/sec reduces the amount of potentiation but still results in a measurable increase. At 16.6/sec the synaptic change is close to zero. The precise data points from this protocol, along with their error bars, are plotted in blue in Fig. 2.7. This experimental protocol was performed with a burst repetition frequency of 1Hz, repeated 300 times.

We were further able to constrain the frequency dependence, this time of LTD, from the experiments of Bidoret et al. (2009). There they showed that combining a purkinje cell depolarisation of length 120ms, with a single parallel fibre spike, at the mid-point of the depolarisation, leads to zero synaptic change. Whereas, a parallel fibre doublet, the second spike of which occurs at the depolarisation mid-point, at 200/sec leads to significant LTD. They further examined the effects of performing the parallel fibre doublet at an internal frequency of 1000, 66.6, 33.3 and 16.6 / sec, giving us a total of 6 data points to fit. As with the LTP data points these experimental data points are reproduced, this time in red, in Fig. 2.7. This protocol, of purkinje cell depolarisation combined with parallel fibre stimulation, was repeated 120 times, at a frequency of once per second.

Finally, Safo and Regehr (2008) performed a set of experiments in which they

combined a climbing fibre stimulus, calibrated to induce a purkinje cell complex spike, with seven parallel fibre stimuli with an intra-burst frequency of 100/sec. They systematically varied the timing offset between the climbing fibre stimulus and the fourth parallel fibre stimulus between -300ms and +500ms (-ve implies climbing fibre before parallel fibre), demonstrating for extremal values robust LTP whereas for more central values a large LTD, with a peak at parallel fibre spiking 80ms following the climbing fibre stimulus. The full set of data points from Safo and Regehr (2008) are reproduced in Fig. 2.10. This protocol was repeated just 30 times at a repeat frequency of 0.1/sec.

2.2.2.2 Mapping the model to experimental outcomes

In synaptic efficacy experiments it is common to measure synaptic efficacy, via an experimentally accessible parameter such as the amplitude of the postsynaptic EPSC induced by a single, low-strength presynaptic stimulus, both prior to and following the induction protocol. The post induction protocol synaptic efficacy is then divided by the pre induction value in order to produce a measure of change in synaptic efficacy, normalised at 1 for zero change.

In order to compare our synaptic efficacy variable, which has been bounded between 0 and 1, we apply a similar technique. We assume a uniform prior distribution of synaptic weights, thus leading to an average pre induction synaptic efficacy of 0.5. This allows us to begin all of our simulations from an initial value of 0.5 and to observe their outcomes without recourse to multiple initial start values. We then simulate the effects of a plasticity induction protocol on the synapse and record the final synaptic efficacy value divided by 0.5 (the average of the pre induction protocol synaptic efficacies) as the outcome of our simulated plasticity experiment. This value is then directly comparable with the experimental results without further transformation.

2.2.3 Optimisation of the model fit

The two models presented are entirely mechanistic and inspired by underlying biochemical processes, however it is difficult to directly relate their parameters to measurable biophysical properties. For this reason, we chose to fit our parameter values, using automated optimisation techniques, to the results from a coherent set of 17 experiments (explained in Section 2.2.2.1). These experiments were performed under almost identical experimental conditions and provide good coverage of the potential outcomes of synaptic plasticity experiments for this particular synapse. At a minimum, parameter optimisation techniques require the calculation of a cost function,

the difference between the simulation results and those observed in experiments. In some cases, access to the Jacobian or the Hessian of the cost function are also necessary.

We implemented three optimisation techniques in total. The Polak-Ribière Conjugate Gradient method (Polak and Ribiere 1969), the Levenberg-Marquardt method (Marquardt 1963) and the Nelder-Mead Simplex method (Nelder and Mead 1965). We did not have access to an analytical form for the shape of the cost landscape so we needed to resort to finite-difference methods in order to numerically calculate the gradient for the methods which required it.

The Polak-Ribière method works by computation of the gradient of the cost function with respect to the current position in parameter space, then performing a line search until it finds a local minimum, it then changes to a conjugate direction and repeats the line search until convergence. The Levenberg-Marquardt method is a specialist parameter optimisation technique for nonlinear data, it similarly requires knowledge of the local gradient of the cost function and performs an update based on this gradient. Finally, the Nelder-Mead Simplex method works by calculating the cost function of each node of a simplex in parameter space and performing transformations on the simplex which can generally be expected to improve the cost of the worst node of the simplex.

Our parameter optimisation search required a combination of all three methods. The Polak-Ribière method proved incapable of overcoming the non-linearities in our solution space, but sometimes performed much better local optimisation than the other two methods. The Levenberg-Marquardt method, despite its specialism in solving non-linear problems, still contains assumptions about the continuity of the gradient of the function being optimised, meaning it frequently got stuck in local minima. The Nelder-Mead Simplex method was the most successful method for our optimisation problem. It makes no assumptions about the cost landscape but tends to proceed in a region surrounding the current best point. However, it can be slow to converge to the base of a local minimum. Hence we frequently took the output of one procedure and used it as the input to one or both of the other procedures in order to further improve the fit.

2.2.3.1 χ^2 cost function

We defined two cost functions, both of which are based on the χ^2 function of the simulation fit to the data. The first cost function, $\phi(x)$, for a given parameter set x , was defined directly as the χ^2 value for the model fit to the 17 target experiments

using that parameter set. That is

$$\phi(x) = \sum_i \frac{(\text{Obs}_i - \text{Sim}_i(x))^2}{\sigma_i^2} \quad (2.10)$$

where Obs_i is the experimentally observed synaptic efficacy from experiment i and σ_i the associated standard error. $\text{Sim}_i(x)$ is the result of the simulation for experimental protocol i using parameter set x . This gives us a cost function whose value is 0 if it perfectly matches the experimental data and > 0 otherwise. This cost function was used with both the Polak-Ribière Conjugate Gradient method (Section 2.2.3.4) and the Nelder-Mead Simplex method (Section 2.2.3.6).

The Levenberg-Marquardt (LM) method (Section 2.2.3.5) calculates the gradient of a multidimensional cost function and then reduces the sum of the squares of this value. For consistency with the other methods we defined the cost function in this case to be

$$\Phi_i(x) = \frac{(\text{Obs}_i - \text{Sim}_i(x))}{\sigma_i} \quad (2.11)$$

where the subscript i denotes the experimental protocol. This cost function can take on both positive and negative values, which becomes important in the calculation of the search direction for the LM method. The value then reduced by the LM method is

$$\phi(x) = \sum_i \Phi_i^2(x) \quad (2.12)$$

giving rise to a χ^2 error function which is directly comparable with the cost functions used by the Polak-Ribière and Nelder-Mead Simplex methods.

2.2.3.2 Gradient calculation by finite-difference

Two of the optimisation methods used in this work required the calculation of the gradient of the cost function with respect to each of the model parameters. We calculated this simply by calculating, first the cost function, $f(x)$, at the point x , then modifying in turn each of the elements of x by δx_i and calculating the new cost function with respect to all of the experimental data points. Due to the differences in scale between the values of elements of the parameter vector, and indeed their completely nonlinear relative contributions to simulation results, we found it necessary to use a different δx_i value for different elements of the x vector. An optimal selection of δx_i was performed manually in advance of performance of the optimisation procedures. The gradient was then calculated as

$$\frac{\partial f(x)}{\partial x_i} = \frac{f(x + e_i \delta x_i) - f(x)}{\delta x_i} \quad (2.13)$$

where $e_i = (0, \dots, 0, 1, 0, \dots, 0)$ is the basis vector whose length is equal to the length of the vector x and whose element i is equal to 1, all other elements equal 0.

2.2.3.3 Simplification of the parameter space

In order to simplify the fitting of our models to the experimental data, and in the absence of experiments indicating direct biophysical constraints on our parameters, we reduced the number of parameters in our model by letting the delays in C and N increase following parallel fibre activity be equal, $D_c = D_n$. Subsequent examination showed that such a constraint was not necessary in the context of experiments under examination, letting $D_n = 0$ led to comparable results without modification of other parameters, but it greatly simplified the visual analysis of simulation results.

We further made the analysis of the model considerably easier by letting the exponential decays of C and N be equal, $\tau_c = \tau_n$. This means that for any given protocol, in the absence of increases, both C and N decay in a straight line towards the origin in C-N space. This allows for a more intuitive understanding of the effects of a given protocol when examined in C-N space.

A final simplification of the parameter search was found by setting $C_{\text{depol}} = C_{\text{CS}}$. Tests were run both incorporating and relaxing this restriction and showed very little improvement in the model fit through the separation of these parameters.

2.2.3.4 Polak-Ribière method

The Polak-Ribière method is a nonlinear *conjugate gradient* method. This means that it requires access to the local gradient of the cost function, which we calculate by numerically computing the gradient for a δx change in each of the elements of the parameter vector in turn (see Section 2.2.3.2). In order to seek a minimum, the gradient is first calculated, then a line search is performed downhill in the gradient, by incrementally increasing the step size, calculating the cost function (Eqn. (2.10)) at each step, until a minimum cost in that direction is reached. At this point a conjugate vector is calculated with respect to the original search direction and a new line search is performed. This method is greatly superior to a steepest descent approach due to its ability to navigate long narrow ‘valleys’ in the cost function, without bouncing from one side of the function to the other without ever descending in the numerically shallower descent but ultimately fruitful direction (see Fig. 2.4 for an example). Polak and Ribière contributed the iterative conjugate generation formula

$$\beta_n^{PR} = \frac{\Delta x_n^\top (\Delta x_n - \Delta x_{n-1})}{\Delta x_{n-1}^\top \Delta x_{n-1}} \quad (2.14)$$

giving a conjugate direction $s_n = \Delta x_n + \beta_n s_{n-1}$ where $s_0 = \Delta x_0$ is the downhill slope. If progress in the cost function stops then the gradient needs to be recalculated and the procedure can continue. If the equation was linear, as is the case for typical conjugate gradient methods, then we would expect convergence in less than N steps (the dimension of the vector space). There is an assumption in this method that, at least close to the solution, the cost function is quadratic in the parameters.




Figure 2.4. Example of the difficulty of numerical optimisation in the case of a long narrow valley. The steepest gradient in the displayed landscape is that between the facing walls of the ‘valley’. Due to issues of numerical accuracy it is highly unlikely that a steepest descent method will ever find the floor of the valley and subsequently find the much shallower descent direction along the valley floor. This leads to an oscillatory behaviour which will be very slow to descend towards the true solution. A conjugate gradient method will first follow the direction of steepest descent, to find a minimum in that direction. It then rotates in a conjugate direction and perform a new descent in that direction. In the illustrated example it can be expected to quickly follow the direction of the valley floor. If reduction in cost can no longer be obtained a new local gradient is calculated and the algorithm begins anew from the most recently found minimum point. Adapted from Martens (2010).

2.2.3.5 Levenberg-Marquardt method

The Levenberg-Marquardt method is a nonlinear least-squares minimisation technique which interpolates between the Gauss-Newton algorithm and the gradient descent method. The objective of this method is to minimise the sum of the squares function

$$S(x) = \sum_{i=1}^m [y_i - f_i(x)]^2 \quad (2.15)$$

where x is the parameter vector, y_i an element of the m experimental data points, and $f_i(x)$ is the evaluation of the simulation cost function defined in Eqn. (2.11). The update of the cost function is estimated by linearisation

$$f_i(x + \delta) \approx f_i(x) + J_i \delta \quad (2.16)$$

where

$$J = \frac{\partial f_i(x)}{\partial x} \quad (2.17)$$

is the Jacobian of the simulation cost function with respect to the current parameters. The key to solving the optimisation problem is to find a solution where $S(x) = 0$. Levenberg developed a solution, with a damping factor, via

$$(J^T J + \lambda I) \delta = J^T [y - f(x)] \quad (2.18)$$

solving for the required δ . An iterative update on the parameter function, x , is then performed following $x_{n+1} = x_n + \delta$. The cost and gradient functions are then updated and a new step δ calculated until convergence. Marquardt modified the damping factor to account for vastly different gradient values in different directions, a case we observe in our particular problem

$$(J^T J + \lambda \text{diag}(J^T J)) \delta = J^T [y - f(x)] \quad (2.19)$$

This avoids slow convergence when the gradient is small in certain directions.

2.2.3.6 Nelder-Mead Simplex method

Ultimately our model fit to the experimental data turned out to be a highly non-linear problem, with a non-convex solution space, therefore the method which served us the best for fitting the parameters to the data was the Nelder-Mead Simplex method (Nelder and Mead 1965). We defined the cost function, $\phi(x)$, for a given parameter set x , as the χ^2 value for the model fit to the 17 target experiments according to

Eqn. (2.10). We used an implementation of the algorithm as found in the GNU scientific library (GSL) (Galassi 2009) but will detail here the implementation of the algorithm.

The Nelder-Mead Simplex algorithm (not to be confused with the Simplex algorithm) requires $N + 1$ points in the N dimensional parameter space. In the case of the GSL library, we define an initial set of parameters, x , and then via the calculation of a set of orthonormal basis functions in the parameter space and an initial *step size*, which we define, for each direction in parameter space a further N points are automatically generated by performing a step of size *step size* in each direction of the orthonormal basis.

From this initialisation, the cost, $\phi(x)$, of each of the $N + 1$ vertices is calculated and the vertices are ordered according to their cost, with the cost of parameter set/vertex x_{N+1} being the highest. The barycentre of the N lowest cost vertices is then calculated and a number of operations are possible in order to improve the cost of the highest cost parameter set. These include: (i) reflection of the worst point through the barycentre; (ii) expansion of the reflected point; (iii) contraction of the reflected point; (iv) and ‘reduction’ of all but the best point, towards the best point. The processes of *reflection* and *expansion* allow for an exploration which may overcome local maxima, especially in a highly nonlinear space. *Contraction* and *reduction*, lead to the ultimate convergence of the algorithm to a simplex with internal area smaller than a user defined tolerance which can be seen as the ultimate convergence of the algorithm.

For our given problem set the Nelder-Mead Simplex algorithm is not guaranteed to find a global minimum. We approached the problem by exhaustive iteration of the method, using multiple restarts from different initial conditions. We also used an advanced form of the algorithm where the basis functions of the simplex are randomly oriented rather than being based on the fixed co-ordinate axes. We further generated new basis functions every 50 iterations of the algorithm, which turned out to be a very useful technique for escaping local minima close to the solution. Finally we fed the outputs of one solution as the initial starting point for a subsequent re-run. We were able to find two consistent sets of parameters from which it is unlikely that it is possible to find much improvement. It should be pointed out that the behaviour of the model variables C and N, as viewed in C-N space, combined with the full set of experimental protocols fitted (Section 2.2.2.1), and self-imposed restrictions on some of the parameters (Section 2.2.3.3), led to quite a restrictive solution space. This also allowed for the development of a high degree of ‘expert-knowledge’ of the influence of each of our parameters on the simulation outcomes, allowing us to hand-fit surprisingly successful initial conditions for our optimisation procedure for each

	Parameter	Value	Unit
C variable time constant	τ_c	182.936	ms
N variable time constant	τ_n	182.936	ms
C variable delay	D_c	78.9	ms
N variable delay	D_n	78.9	ms
Parallel fibre spike C influx	C_{PF}	0.081297	-
Complex spike C influx	C_{CS}	0.543789	-
Purkinje cell depolarisation C influx	C_{depol}	0.543789	-
Parallel fibre spike N influx	N_{PF}	0.2357006	-
LTP-LTD threshold in C	θ_D	0.5671626	-
Depression learning rate	γ_D	1.360871e-4	/ms
Potential learning rate	γ_P	3.466588e-5	/ms

Table 2.1. Parameter values used in Linear-N model. Fitted using the optimisation procedures described in Section 2.2.3. For the reasons outlined in the text the following parameters were forced to be equal in the fit: $\tau_c = \tau_n$, $D_c = D_n$, $C_{CS} = C_{depol}$.

of the models. That said, we did also try out many alternative solutions as starting points for our optimisation procedure.

2.2.3.7 Model selection

In order to measure the fit of a given model to the data we calculated χ^2 values for the ensemble of 17 fitted experiments, for the best fitting set of parameters for that model. The χ^2 value is defined as

$$\chi^2(x) = \sum_i \frac{(\text{Obs}_i - \text{Sim}_i(x))^2}{\sigma_i^2} \quad (2.20)$$

for a given parameter set x , where Obs_i is the experimentally observed synaptic efficacy from experiment i and σ_i the associated standard error. $\text{Sim}_i(x)$ is the result of the simulation for experimental protocol i using parameter set x .

A specific model, Model A, is said to be nested within a less restricted model, Model B, with more parameters and less degrees of freedom than Model A, if Model A can be derived from Model B by fixing at least one free parameter in Model B or by introducing other restrictions, e.g., by constraining a free parameter to equal one or more other parameters (Schermelleh-Engel et al. 2003). As the Linear-N model

	Parameter	Value	Unit
C variable time constant	τ_c	166.68557	ms
N variable time constant	τ_n	166.68557	ms
ν variable time constant	τ_ν	81.22673	ms
C variable delay	D_c	76.4	ms
N variable delay	D_n	76.4	ms
ν variable delay	D_ν	0.9	ms
Parallel fibre spike C influx	C_{PF}	0.0691449	-
Complex spike C influx	C_{CS}	0.783954	-
Purkinje cell depolarisation C influx	C_{depol}	0.783954	-
Parallel fibre spike N influx	N_{PF}	0.431781	-
LTP-LTD threshold in C	θ_D	0.52415	-
Depression learning rate	γ_D	9.1704836e-5	/ms
Potentiation learning rate	γ_P	2.8812565e-5	/ms

Table 2.2. Parameter values used in NMDAR-based model. Fitted using the optimisation procedures described in Section 2.2.3. For the reasons outlined in the text the following parameters were forced to be equal in the fit: $\tau_c = \tau_n$, $D_c = D_n$, $C_{CS} = C_{\text{depol}}$.

can be considered as *nested* within the NMDAR-based model it is appropriate to use a difference of χ^2 values to further compare the two models. The difference of χ^2 is defined as $\chi_{\text{diff}}^2 = \chi_s^2 - \chi_l^2$, where the subscript s denotes the smaller model, the Linear-N model in our case, and l the larger model, the NMDAR-based model. We further define the degrees of freedom of this distribution as $df_{\text{diff}} = df_s - df_l$. It is then possible, using χ^2 tables to decide whether the two models belong to the same distribution or not (Schermelleh-Engel et al. 2003; Steiger et al. 1985).

A potentially more robust comparison of the models can be performed using the Akaike Information Criterion (AIC) (Akaike 1974) and the related Bayes Information Criterion (BIC) (Schwarz 1978). These measures of model fit may be used exclusively for the comparison of models, and not for an absolute calculation of the ability of an individual model to fit data. More importantly, they do not require the assumption of *nestedness* used in the difference of χ^2 comparison. Both criteria measures use a maximum likelihood function for the estimate model combined with a penalty for the number of parameters in the model. They do not allow for significance testing. The Akaike Information Criterion is defined as

$$AIC = 2k - 2\ln(L) \quad (2.21)$$

where L is the maximum likelihood function for the estimate model, and k the number of parameters in the model. In our case, we can assume a χ^2 distribution on the error in the fit, due to the normality assumption on the underlying experimental data points, giving us $\ln(L) = G - \chi^2/2$, where G is a constant term whose value is particular to the underlying data being modelled and not to the particular model being used. So to a constant term we get

$$AIC = 2k + \chi^2 \quad (2.22)$$

The BIC is derived under the assumption that the distribution underlying the data follows an exponential distribution. It is then the marginal probability of seeing the observed data given the model proposed. Under similar arguments as those for the AIC regarding the distribution of the error, we define the BIC measure using our χ^2 values as

$$BIC = \chi^2 + k \ln(n) \quad (2.23)$$

where k is again the number of parameters, here n is the number of experimental values we are fitting (17 in our case).

Both the AIC and BIC allow us to directly compare our two models. A lower measure value implies a better model for the data, however there is no absolute scale allowing for definite categorisation of one model over the other and indeed both

methods suffer from the low number of experimental points fitted, although the BIC at least includes a penalty for this in the measure.

2.3 Results

Following selection of model parameter values (see Tab. 2.1 and 2.2) we were able to reproduce both the burst length and the frequency dependence of LTP and LTD protocols. In addition both models can correctly reproduce the timing offset curve, between climbing fibre activation and parallel fibre stimuli, demonstrated in (Safó and Regehr 2008). We compared the two versions of the model in order to examine how much we gain by the additional model parameters in the NMDAR-based model over the Linear-N model. Finally, we used the model to make an unintuitive experimental prediction, which was then checked *in vitro* and compared to the model results.

2.3.1 Burst length dependence

The first reported burst length dependence for LTD in the parallel fibre to Purkinje cell synapse was presented by Casado et al. (2002). They found that, when pairing a Purkinje cell depolarisation with parallel fibre stimulation, they needed a doublet of parallel fibre action potentials, at 200/sec, in order to induced LTD. Pairing of a single parallel fibre impulse with Purkinje cell depolarisation resulted in zero synaptic change. In Fig. 2.5 we demonstrate both versions of the model reproducing this 5ms inter-parallel fibre spike protocol and compare it with the experimental results reproducing this protocol in (Bidoret et al. 2009). We see clearly that for a two spike 5ms protocol interval we expect a significant synaptic depression in the postsynaptic EPSC, $34.7 \pm 6.2\%$ of it's pre induction protocol amplitude (n=9), the degree of which is accurately reproduced by the model. An example of this protocol in the C-N plane, following the Linear-N model implementation, is shown in Fig. 2.6, where we observe the high calcium induced via Purkinje cell depolarisation followed by increases in NO from the parallel fibres leads to entry into the LTD zone. By comparison, for a single spike protocol both model versions predict zero synaptic change, which matches the Bidoret et al. (2009) results (see Fig. 2.7, compare red square with overlapping orange cross).

Experimental demonstration that LTP is also burst length dependent has been more recent. In experimental results which will be published in a paper alongside this model (*under submission*) it has been shown that, under conditions of sparse

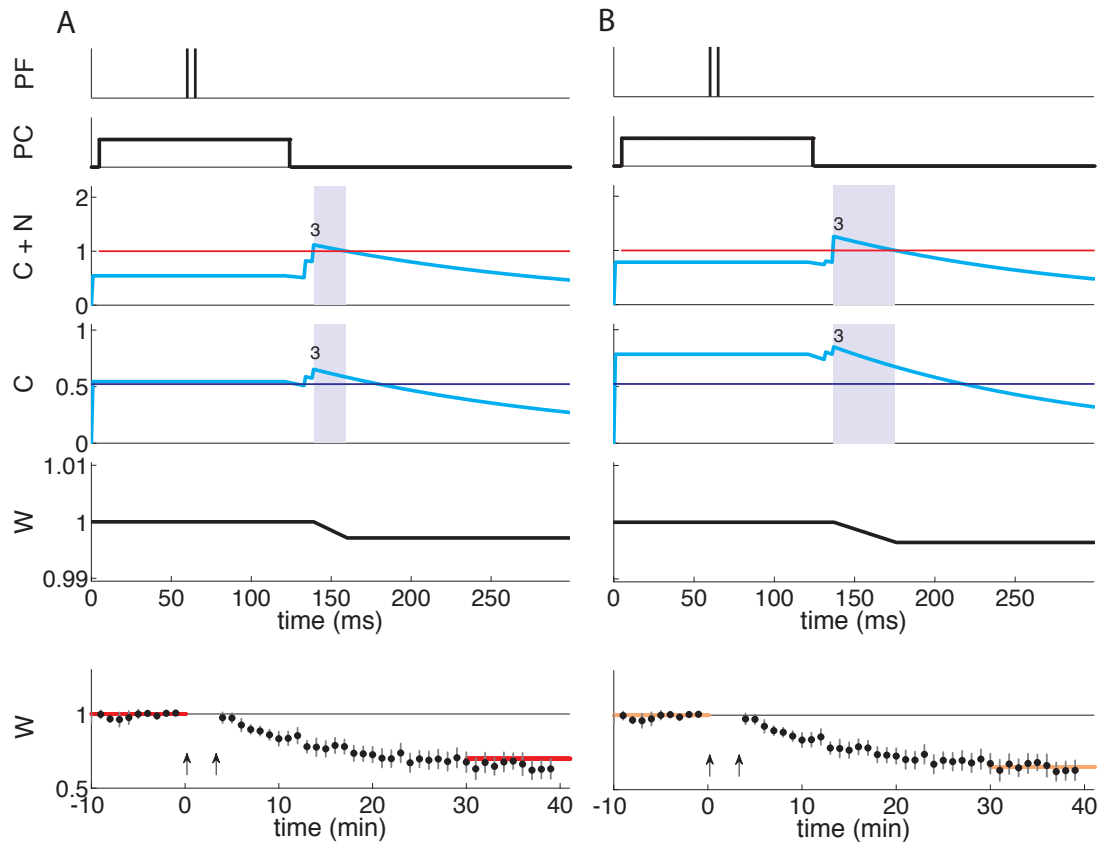


Figure 2.5. Example of the Bidoret 5ms protocol, combining a Purkinje cell depolarisation with a 200Hz parallel fibre doublet of stimulation.

(A) Linear-N model, parallel fibre stimulation (top row) combined with PC depolarisation (row 2) leads to a crossing of the $C + N = 1$ threshold (red line, row 3), leading to a period of synaptic plasticity (shaded background). C remains greater than θ_D (blue line, row 4) for this period so LTD results. (B) NMDAR-based model, the first parallel fibre spike does not lead to an increase in the N variable. Weight change for a single stimulation pairing is depicted in row 5. In row 6 (bottom row) we compare the experimental results from Bidoret et al. (2009) with the results of our simulation at the end of the protocol. The colours in row 6 correspond to the colours in Fig. 2.7. The numbered points in rows 3 and 4 match similarly labelled points in Fig. 2.6.

stimulation of parallel fibres, 5 parallel fibre stimuli at 200Hz are necessary in order to induce reliable LTP ($169 \pm 11\%$, $n=18$). If only two parallel fibre stimuli are

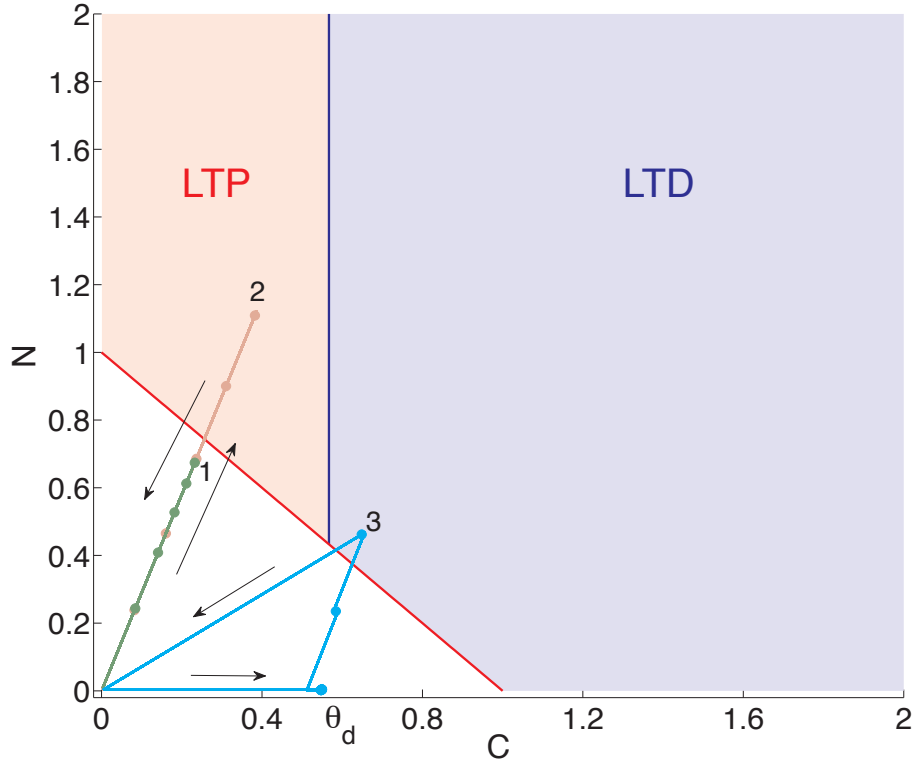


Figure 2.6. Examples of the behaviour of the Linear-N model under three different protocols. In green, we see the 5xPF protocol with a 16.6/sec intra-burst frequency. In orange, the 5xPF protocol with a 200/sec intra-burst frequency. In cyan, the Purkinje cell depolarisation combined with a 2xPF doublet at 200/sec. The numbered labels illustrate similar points on the plane as those same labels in Figs. 2.5 and 2.8.

provided then effectively no change occurs, $94 \pm 5\%$, $n=5$. And indeed for three parallel fibre stimuli at 200Hz the resulting synaptic efficacy is $104 \pm 19\%$, $n = 7$. As described in the Methods description of the experiments (Section 2.2.2.1), fitted using our parameter optimisation procedure, we prioritised fitting the two parallel fibre stimuli result. In fact, using the fitted parameter values (Tab. 2.1 & 2.2) we observe zero synaptic change using the two models for both the two and three spike protocols. In Fig. 2.7 we plot both protocols on top of one another using the cyan cross. In fact, when calculating our model χ^2 fits it appears that we have been over severely penalising ourselves. The tight error bars on the doublet protocol (Fig. 2.7, blue square) mean that our model can never get within one standard error of the




Figure 2.7. Degree of LTP/LTD is dependent on both burst length and intra-burst frequency. Solid lines depict NMDAR-based model results, dashed lines depict Linear-N model. Circles and Squares show experimental results with standard error bars. The two model results which showed zero change (crosses at 0 and 5ms) produce identical results for both model versions.

mean while still fitting the other protocols, but we can easily fit the triplet protocol, which has relatively large standard error ($\pm 19\%$), for a wide range of parameter values. In reality, we expect that the triplet protocol shows much wider variance as it is on the verge of plasticity, probably dependent on noise and initial synaptic efficacy, whereas the $94 \pm 5\%$ result for the doublet protocol likely represents zero real synaptic change, which is exactly what we obtain with the model.

2.3.2 Frequency dependence

Similarly to the parallel fibre burst length dependence it has been shown that the frequency of parallel fibre stimulation matters in producing both LTD and LTP. The frequency dependence of LTD was examined in detail in (Bidoret et al. 2009) where they combined a Purkinje cell depolarisation with a parallel fibre doublet with an intra-burst frequency of 1000, 200, 66.6, 33.3 and 16.6 Hz. Their results are summarised in the red circle data points in Fig. 2.7. Both our Linear-N and NMDAR-based models are able to reproduce these results (compare with the red dashed line and solid orange line, respectively, in Fig. 2.7).

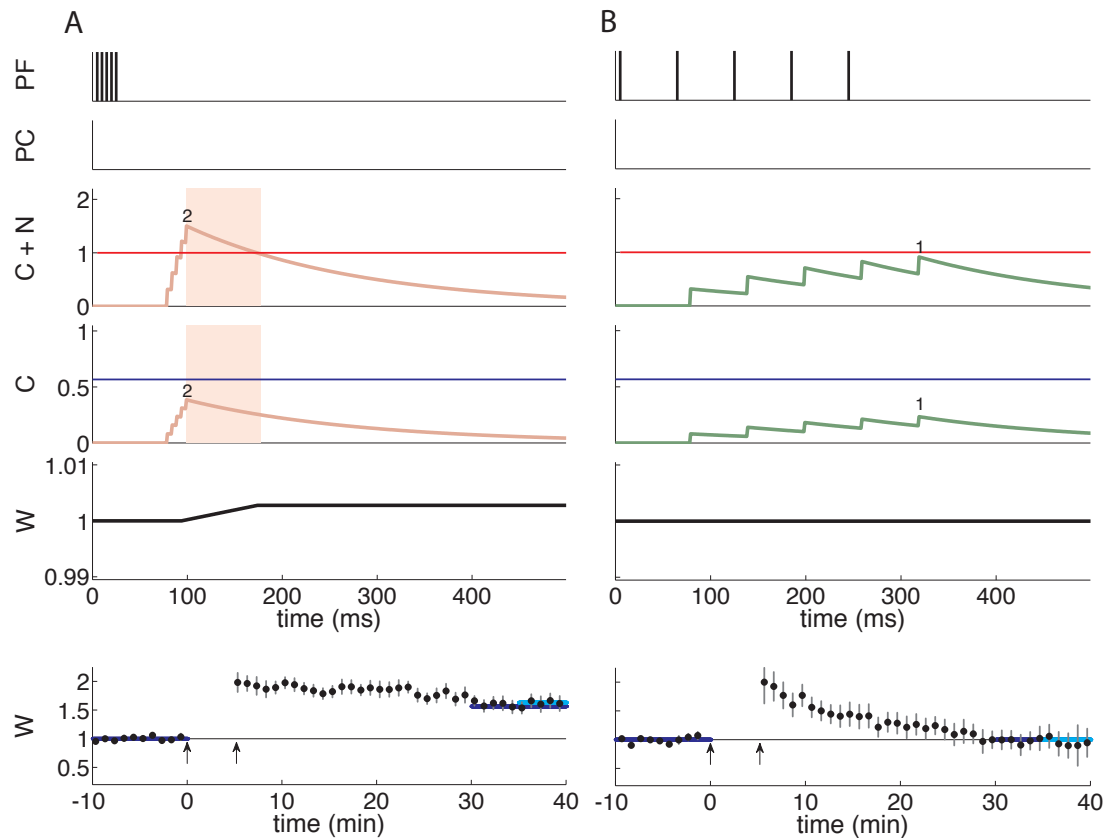


Figure 2.8. Example of the Linear-N model where the parallel fibres are stimulated 5 times. (A) Stimulation at 200Hz leads to a crossing of the $C + N = 1$ threshold (row 3, red line), leading to a period of synaptic plasticity (shaded background). C remains less than θ_D (row 4, blue line) so LTP results. (B) Stimulation at 16.6Hz leads to no change, in both the synaptic plasticity model and in experiments. The numbered points on the C and $C + N$ curves match similar label points in Fig. 2.6. The synaptic efficacy response to a single repetition of the induction protocol are shown in Row 5. Experimental results are plotted in bottom row and compared with theoretical predictions of the Linear-N model (blue line) and the NMDAR-based model (cyan line).

Until now experimental results showing the frequency dependence of LTP have been sorely lacking. Indeed many published protocols have assumed that a single parallel fibre impulse repeated at 1Hz is enough to induce LTP, or LTD when coincident with Purkinje cell depolarisation. The reason for this is most likely due to

the methods of experimentation. It has been shown (Marcaggi and Attwell 2007) that the positioning of the stimulating electrode will lead to a large difference in the density of parallel fibre spikes reaching a localised region on a Purkinje cell dendritic tree. Another likely factor is the temperature of experimentation, Isope et al. (2004) showed that at room temperature parallel fibre stimulation leads to a sustained parallel fibre influence on the postsynaptic cell. We have had access to recent experiments (*under submission*), which are summarised in Fig. 2.7 showing large scale LTP for a burst of five parallel fibre stimuli with an intra-burst frequency of 200Hz. The degree of LTP is reduced for an intra-burst frequency of 33.3Hz and is entirely absent at 16.6Hz. These experiments were included in the optimisation procedure fitting of our parameters and both of our models are able to fit these points, see the solid cyan (NMDAR-based model) and dashed blue (Linear-N model) lines in Fig. 2.7.

2.3.3 Safo and Regehr

The final set of experiments which we attempted to fit via our optimisation procedure was a set of experiments performed by Safo and Regehr (2008). This set of experiments was particularly restrictive on the parameter values used in our model due to the combination of a single climbing fibre induced complex spike with seven parallel fibre stimuli, with an intra-burst frequency of 100/sec. As seen in Fig. 2.9 for large time differences between the complex spike and the parallel fibre stimuli, the complex spike does not lead to any time spent above the plasticity threshold, hence the resulting synaptic change is driven entirely by the time spent in the LTP zone via the parallel fibre stimuli. As the climbing fibre and parallel fibre stimuli approach one another in time, considerably greater time is spent in the LTD zone leading to a dominance of LTD. Finally, a separation of timing in the opposite direction restores LTP. In Fig. 2.10 (left) we see the behaviour of these three protocols in the C-N plane, giving us an excellent example of how tightly constrained the model may in fact be. Finally, in Fig. 2.10 (right) we compare the results of both versions of the model, seeing a very close match with the experimental results from Safo and Regehr (2008).

2.3.4 Model comparison

We used a number of methods to compare the two model fits to the data. In an initial analysis we looked at χ^2 values for each of the model fits. In the case of the Linear-N model we see that $\chi^2 = 13.3212$, for a model with 8 parameters and 17 experimental

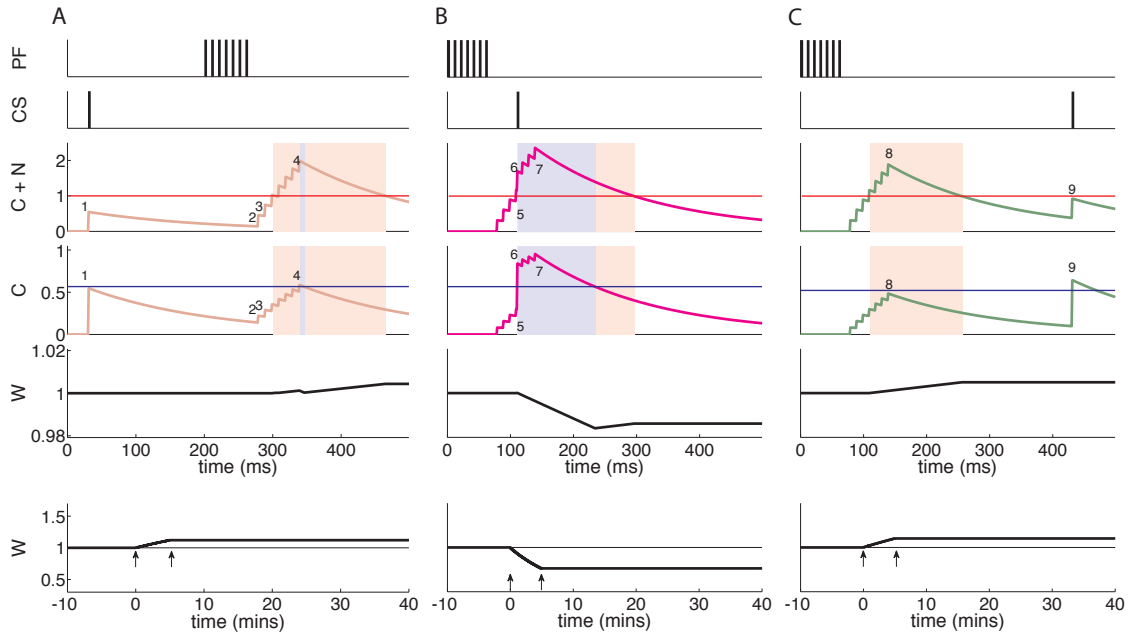


Figure 2.9. Three examples of the Safo and Regehr (2008) protocol using Linear-N model. Climbing fibre to parallel fibre offset of (A) -200, (B) +80 and (C) +400ms. (Top row) Timing of 7 parallel fibre stimuli at 100/s. (Row 2) Timing of climbing fibre induced complex spike. (Row 3) $C + N$ in response to a single repetition of induction protocol. Red line depicts threshold of plasticity. Blue background shading represents periods of LTD, pink shading periods of LTP. (Row 4) Evolution in time of C variable in response to a single repetition of induction protocol. Blue line depicts threshold between LTP and LTD. The numbered points on the C and $C + N$ traces correspond to the similarly numbered points in Fig. 2.10. (Row 5) Evolution of synaptic efficacy variable in response to a single repetition of induction protocol. (Bottom row) Evolution of synaptic efficacy variable for full induction protocol.

data points. By comparison the NMDAR-based model gives us $\chi^2 = 7.8359$ for a single additional parameter. We were able to compare the two model fits using a difference of χ^2 analysis, based on the assumption that the Linear-N model is *nested* within the NMDAR-based model (see Section 2.2.3.7). This led to a $\chi^2_{\text{diff}} = 5.4853$, with 1 degree of freedom. This gives us a p-value=0.025 that the models belong to different underlying distributions.

Two additional methods of model comparison are the AIC and BIC (for details see Section 2.2.3.7). In the case of the Linear-N model we get $AIC = 29.3212$

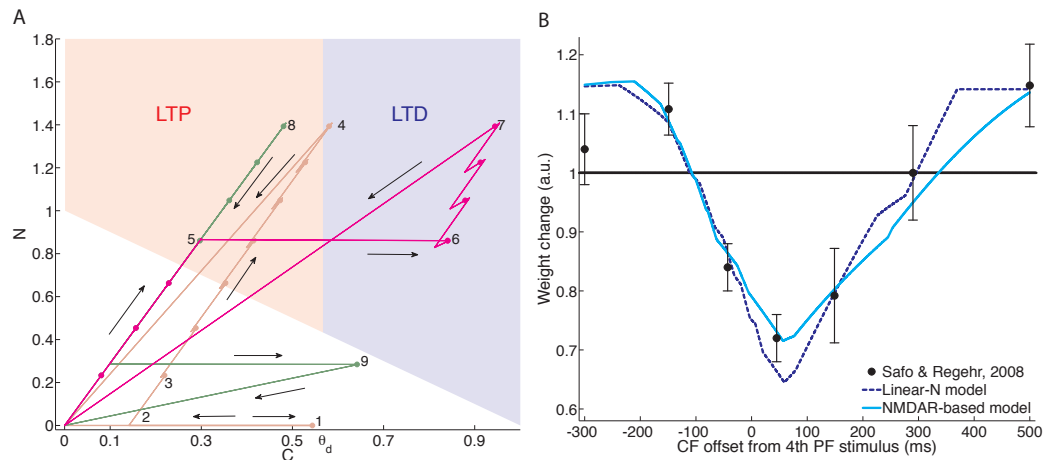


Figure 2.10. Model fit to Safo and Regehr (2008). (A) C-N plane examples of the three protocols used in Fig. 2.9, using Linear-N model, the numbered points correspond to similarly numbered points in that figure. (B) Comparison of the two model predictions with the Safo and Regehr (2008) experimental results for different climbing fibre to parallel fibre offset times.

and $BIC = 35.987$. This compares with $AIC = 25.8359$ and $BIC = 33.3348$ for the NMDAR-based model. In both cases the NMDAR-based model achieves a lower score than the Linear-N model, despite punitive terms in both methods for the increased number of parameters in the NMDAR-based model. This is somewhat suggestive that the NMDAR-based model is a better fit to the data, but does not mean that an entirely different model might not be significantly better.

2.3.5 Prediction: PF \rightarrow LTD

In order to test our model we looked to make an experimental prediction which could be tested by our collaborators. A number of ideas presented themselves in relation to the precise relative timing of the climbing fibre stimulus to the parallel fibre stimuli. Such an approach could be considered a refinement of the results of Safo and Regehr (2008), where they saw both LTP and LTD depending on the offset, we could have extended this to other numbers of parallel fibre spikes or frequencies of stimuli. An alternative idea was to explore the relative balance of LTP to LTD. Through the model it is possible to conceive a parallel fibre stimuli protocol which would have a high enough frequency to typically produce LTP but which would in fact lead to a balance of LTP and LTD, giving no change in synaptic efficacy. These

ideas were rejected as a simpler protocol presented itself which is a simplification of this second idea and which leads to a result which is completely unintuitive but which the model held to be true. We decided to examine the effects of a prolonged burst of parallel fibre activity and whether we could produce LTD via this method alone. It is appropriate to mention that Hartell (1996) previously reported that

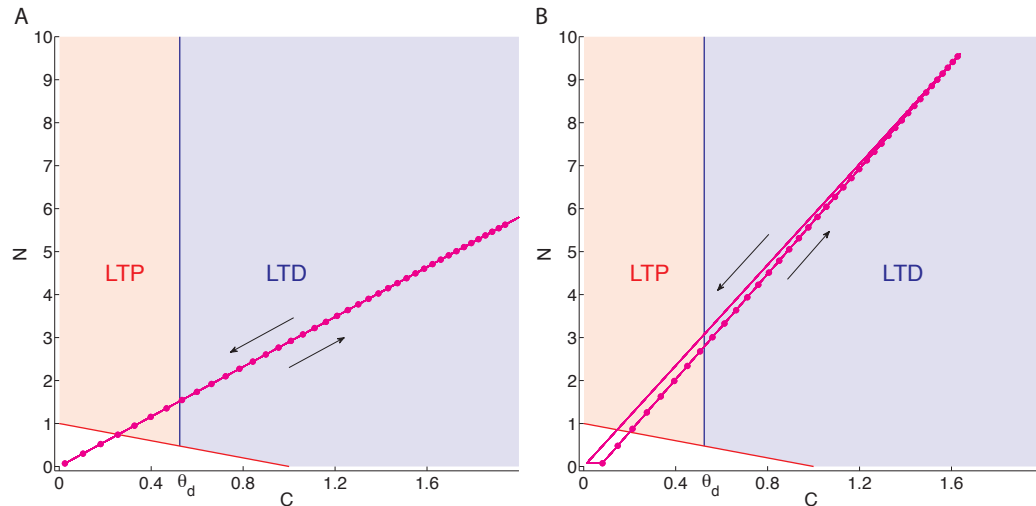


Figure 2.11. Model C-N prediction for 40 parallel fibre stimuli at 200/s. (A) Linear-N model evolution of C and N for 40 parallel fibre stimuli. Due to the linear summation in the model increases and decreases follow a straight line to and from the origin. (B) NMDAR-based model evolution for same protocol. The first parallel fibre spike leads to no change in N so the up and down trajectories no longer overlap and the slope is different. In both cases the trajectory spends a considerable amount of time in the LTD region.

strong stimulation of the parallel fibres alone led to LTD but this result relied on particularly strong stimulation, our model perhaps provides an explanation for why they obtained such a result but using weaker stimuli.

In Fig. 2.11 we see that the model predicts that, in the C-N plane, for high frequency stimulation if the burst is long enough, we will eventually reach the LTD zone. However, it is especially important to note that, at a lower frequency no matter how sustained the burst length, we will never reach LTD. In order to avoid wasting experimentalists time due to any minor fitting errors due to very few data points in the optimisation procedure, we suggested a sustained burst of 40 parallel fibre stimuli at 200Hz. Furthermore, we suggested that keeping the extracellular calcium concentration high (2.5mMol) might bias the experiment towards success,

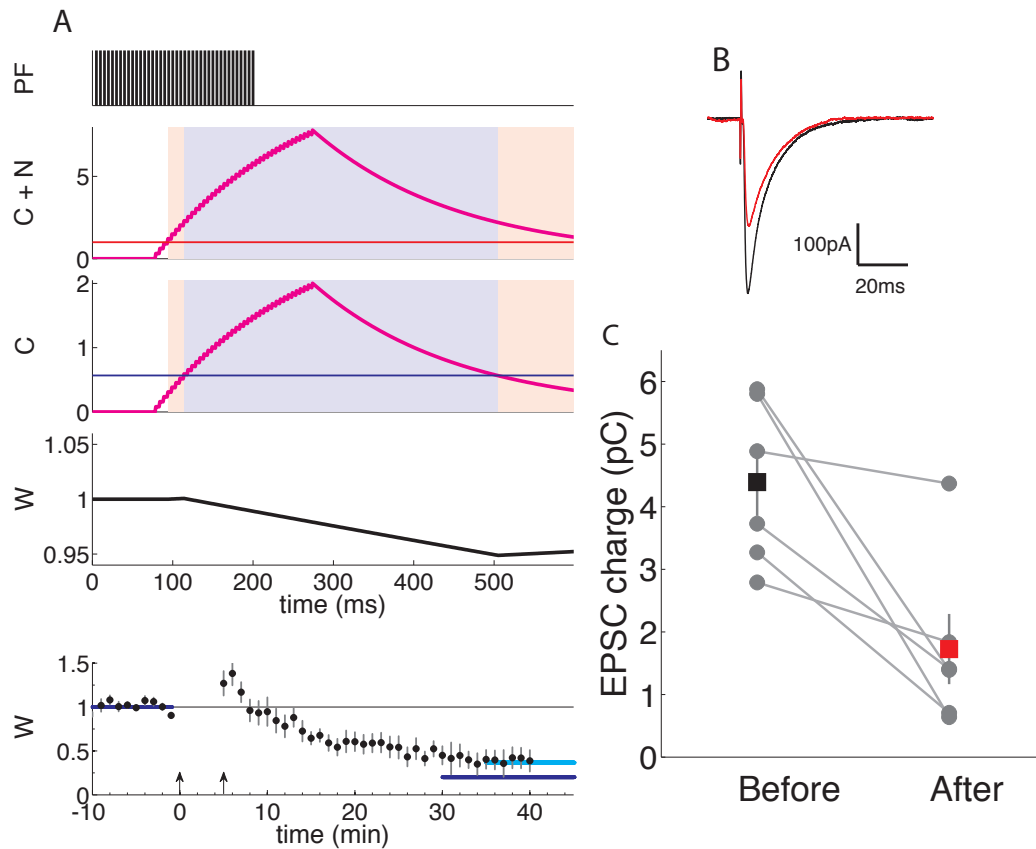


Figure 2.12. Example of the Linear-N model where the parallel fibres are stimulated 40 times. (A) High frequency parallel fibre stimulation (200/sec) leads to an increase in both C and N which leads considerably more LTD than LTP. The model correctly predicted that LTD would result from such a protocol, and indeed the NMDAR-based method correctly predicted the degree of LTD induced (cyan line, see blue line for Linear-N model prediction). (B) Example traces before (black) and after (red) induction protocol. (C) EPSC charge was reduced in all cells following the protocol.

as we did not want depletion of calcium to lead to reduced production of NO or influx of postsynaptic calcium. Happily, the experiment produced the predicted degree of synaptic depression. In Fig. 2.12 we compare the theoretical predictions with the results of the experiment. As with the fitted experiments both models correctly predict the sign of plasticity and both provide a reasonable estimate of the degree of synaptic change, the fit of the NMDAR-based model (cyan line) is however

considerably better.

2.4 Discussion

We were able to reproduce all the principal aspects of the experimentally determined plasticity rules using a simple model based on two variables, one driven by postsynaptic calcium and the other by NO. A plasticity threshold that depends on both calcium and NO, plus an LTP-LTD calcium threshold can explain the majority of the spike-pattern-based plasticity literature for the parallel fibre to purkinje cell synapse, at least when obtained under comparable experimental conditions. The threshold on plasticity enforces the high-pass filter on parallel fibre activity seen in both LTP and LTD experiments, because high-frequency activity is necessary for effective summation to reach the threshold. The sign of plasticity is determined by the calcium level, already depicted as an inverted BCM rule (Bienenstock et al. 1982; Coesmans et al. 2004). Both signals are required for both forms of plasticity (Lev-Ram et al. 1995, 1997, 2002). Our model can generate predictions for plasticity outcomes for arbitrary sequences of parallel fibre and climbing fibre activity.

We used the model to make and verify the prediction that bursts of parallel fibre activity could induce LTD if maintained for long enough. In addition to providing strong support for the model, this result reinforces our suggested mechanism according to which the calcium-driven signals from both parallel fibre and climbing fibre inputs are of a similar nature but different magnitude, implying a degree of equivalence. However, our model does not address the biochemical mechanisms underlying the plasticity thresholds. One can speculate the existence of qualitative differences in calcium activation between effectors for LTP and for LTD. These differences may rely on the distance between calcium sources and targets. Recent data showing different sensitivity of LTP and LTD to T-type calcium-channel blockade are supportive of this kind of a scheme (Ly et al. 2013).

For simplicity, our Linear-N model implements identical NO generation at all action potentials, including the first, yet is able to reproduce the requirement for multiple action potentials. The NMDAR-based version of the model which implements an absence of NO production for the first action potential of a burst behaves similarly, furthermore it is judged a better fit to the experimental data by the measures of χ^2 , AIC and BIC. A common general mechanism holds in all cases, which is that effective summation of decaying responses only occurs when several parallel fibre action potentials arrive within an interval that is short compared to their decay time. This holds whether the integration mechanism concerns presynaptic influx of

calcium through NMDARs, accumulation of mediators downstream of NO production, or both. Irrespective of the behaviour following the first spike of a burst, it is clear that a high-frequency burst is required for the induction of any plasticity and our results, when combined with associated imaging experiments (*under submission*), indicate that this reflects the necessity for NO production.

Previous attempts at modelling synaptic plasticity at this important synapse have confined themselves to modelling LTD (Doi et al. 2005; Kawaguchi and Hirano 2013; Kuroda et al. 2001). All three of these approaches have used detailed biochemical models to examine the cascade leading to the dephosphorylation of postsynaptic AMPA receptors and consequent depression of the ability of the postsynaptic Purkinje cell to be excited by presynaptic parallel fibre activity. Kuroda et al. (2001) identified a calcium dependent positive feedback mechanism (red arrows in Fig. 2.13), dependent on MAPK and PKC which leads to LTD. This feedback loop may appear to serve a similar role to our calcium threshold, θ_D , between LTP and LTD; above a certain calcium concentration, in their model, the feedback mechanism kicks in and LTD becomes a guaranteed result. Doi et al. (2005) examined the influence of the IP3 molecule, which appears earlier in the cascade, as a source of increased calcium postsynaptically. Similarly to our model, they induce small postsynaptic increases in calcium via parallel fibre stimulation, from the IP3 pathway, and larger increases via climbing fibre activation, from voltage-gated calcium channels. This pathway, in both their model and ours, appears to be vital in the determination of the optimal delay between climbing fibre and parallel fibre activity in order to induce LTD. Finally, Kawaguchi and Hirano (2013) extended the biochemical cascades to analyse the feedback loops involving CaMKII, which in high calcium disinhibits the NO pathway leading to a further enabling of the Kuroda feedback loop (red arrows in Fig. 2.13) and greater LTD. Clearly, their work suggests that while depression is possible in the absence of NO it is greatly aided by its presence. Figure 2.14 illustrates their results involving different amounts of CaMKII in the Purkinje cell. These results are immediately reminiscent of our C-N plot for LTP and LTD (Fig. 2.2) suggesting at least agreement in the determination of LTD between the two models.

It is tempting to speculate as to the source of LTP in the Purkinje cell. If the LTD cascades in Fig.2.13 are correct, this implies that there must be a phosphorylation pathway acting on the AMPA receptors perhaps in a continuous manner, whose set-point is adjusted by the CaMKII and NO loops from the LTD process. Perhaps the experiments demonstrating LTP in the presence of an NO-donor are in fact moving the set-point, via inhibition of the activity of PP-2A without activation of the calcium feedback loops, thus exposing the LTP process.

Our model leads to a number of interesting questions. Is calcium, in fact, the

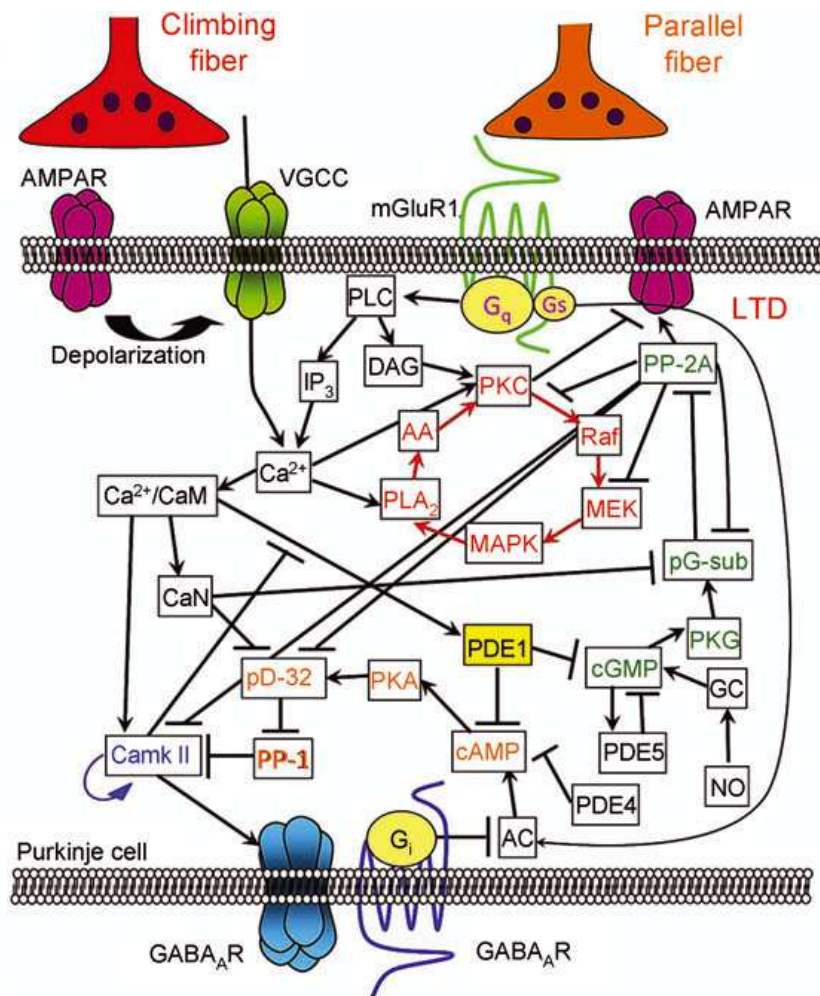


Figure 2.13. Biochemical cascade leading to LTD postsynaptically at the parallel fibre to Purkinje cell synapse. Adapted from Kawaguchi and Hirano (2013).

underlying decider of the sign of synaptic plasticity in this synapse? It certainly appears to be, certainly the other modelling attempts in this domain concur that high calcium is not just correlated with but actually necessary for LTD. A more interesting question is whether it really is the combination of NO and calcium which enables plasticity. Our model is based on this assumption, the experimental data is highly suggestive, but until the LTP pathway is correctly identified it remains impossible to give a definitive answer to this question. From the perspective of our model, it would

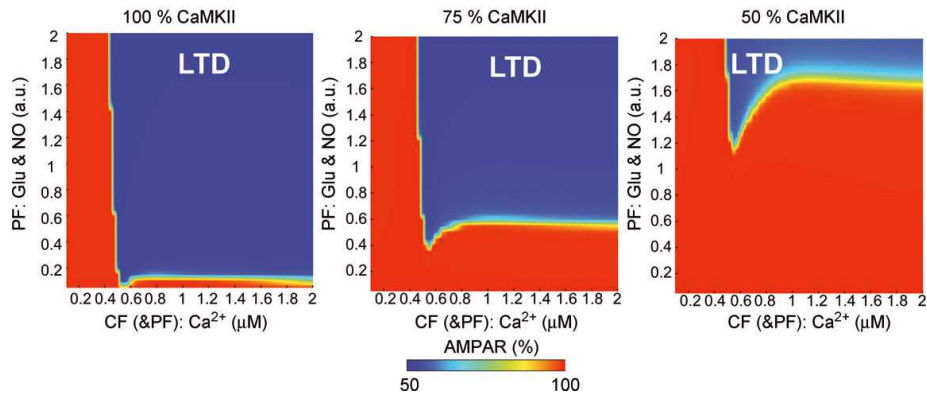


Figure 2.14. Influence of different concentrations of CaMKII availability on the amount of un-phosphorylated AMPA receptors. Simulated result presented in (Kawaguchi and Hirano 2013) which shows the percentage of un-phosphorylated AMPA receptors remaining 30 mins after a conditioning stimulation in the presence of different amounts of available CaMKII.

be interesting to discover whether the calcium threshold between LTP and LTD is in fact a straight line. Postsynaptically there are a number of feedback loops which lead to nonlinear release of calcium which might give a curved appearance to this line in the context of our linear calcium summation model. Ideally, we would like to use the model to make further experimentally testable predictions. The low hanging fruit of the parallel fibre based induction of LTD has been picked. Now, it is a question of whether it is more valuable to demonstrate mastery of the understanding of the balance of LTP to LTD via the system delays or whether it would be more interesting to go straight to the multi-synaptic or network level, making predictions which might be more easily experimentally verifiable, of population level behaviour.

Our parameter fit for the model was based on experiments performed in 2mMol extracellular calcium, with 1mMol Mg, at approximately 32° celsius. In-vivo, the extracellular calcium concentration is more likely to be 1.5mMol and 1.8mMol Mg (Silver and Erecińska 1990). How is this likely to change the behaviour of the synaptic plasticity? A simple assumption is that this will reduce available calcium for influx via voltage-gated calcium channels. Unfortunately it is also likely to reduce NO production and, particularly due to the increased magnesium ion concentration, this may occur in a highly nonlinear manner. In any case, we would expect much reduced synaptic plasticity, requiring much more sustained protocols, most likely with a much higher precision required on the timing coincidence, in order to produce

similar degrees of synaptic change. To date, we have performed some minor experiments in changing the model parameters by reducing the presynaptically induced calcium influx in a manner linearly proportional to the reduction in extracellular calcium concentration and we observe results with the same sign of plasticity as that seen in comparable experiments; a result which would not have held if we had not modified the simulation calcium parameter. But without knowing the related change in NO production, or in the absence of further experiments to induce constraints, we can make no clear estimate as to the magnitude of synaptic change. It would be extremely interesting to perform a number of electrophysiological experiments under these conditions in order to re-fit the model parameters. This would allow us to build simulations which might allow us to disentangle the behaviour of the parallel fibre to Purkinje cell learning rule under *in vivo* behavioural conditions.

Mechanistic plasticity models of the type we have proposed here are necessary to understand the potentially complex relationship between activity patterns and plasticity outcomes. The model of the Purkinje cell as a perceptron (Albus 1971; Marr 1969) has previously been used in order to predict the storage capacity of the parallel fibre to Purkinje cell synapses in the presence of learning of uncorrelated (Brunel et al. 2004) and correlated (Clopath et al. 2012) patterns of input. The perceptron model was further extended to an analogue form (Clopath and Brunel 2013) in order explain the distribution of synaptic weights seen at the parallel fibre to Purkinje cell synapse (Brunel et al. 2004). Our model is concise, easily implementable and accounts for a diverse array of experimental results. It would be interesting to use it to bridge the gap between these previous theoretical results and the underlying spike patterns necessary to induce particular distributions of synaptic efficacies and their influence on output spiking behaviour.

Ultimately, each Purkinje cell in the cerebellum receives contacts from on the order of 100,000 parallel fibre termini. If we are to one day understand the role of synaptic plasticity on the dynamics of Purkinje cell activity it will be through the use of simplified models such as ours. Our model is complex enough to capture the known experimental dynamics of plasticity while remaining simple enough to simulate on a large scale. We hope that this will help in deciphering the mechanisms of learning and memory in the cerebellum.

Maintenance of memory in a cortical network model

In the previous chapter we focused on a synapse with particular properties not normally seen in other synaptic connections throughout the brain. In addition to being dominated by its detection of an error signal, it was completely non-Hebbian. In this chapter we turn our analysis to a previously developed model of synaptic plasticity (Graupner and Brunel 2012) which has been shown to successfully fit data from synaptic plasticity experiments in hippocampal cultures (Bi and Poo 1998; Wang et al. 2005), hippocampal slices (Wittenberg and Wang 2006) and slices from the visual cortex (Sjöström et al. 2001). This model allowed us to perform an analysis which is generally applicable to many sites of synaptic plasticity throughout the brain, and particularly to those sites which are most associated with memory or learning (hippocampus, neocortex).

Much of our work here has been in the development of tools, which can then be used to analyse the particular model of synaptic plasticity. We have then applied these tools with the aim of predicting how long memories survive in such a system under ongoing background activity, but these tools are equally applicable in the prediction of the behaviour of a synapse and indeed the encompassing network under situations involving memory induction and recall, for which we provide an example implementation.

Following an Introduction (Section 3.1), which motivate the work, we present the Methods (Section 3.2) used in our analysis. Beginning with a description of the model dynamics and its dependence on an underlying calcium parameter, we develop a theory for the prediction of the activity of this calcium parameter under Poisson

firing. This theory in turn allows us to derive equations for the average behaviour of the synaptic plasticity rule, again under Poisson pre- and postsynaptic firing. We incorporate these equations into an Ornstein-Uhlenbeck process in order to develop tools for predicting the mean synaptic weights, the tendency of the mean to change over time, and the time scale of this change. Finally, we supplement the theoretical tools with an extension which accounts for a bistability induced by a double-well potential in the synaptic plasticity rule. Following the development of these analytically based tools we move on to advancing numerical methods, without which we could not have performed such detailed and long-run simulations as can be seen in the Results (Section 3.3). In particular, we have developed a numerically precise event-based implementation of the synaptic plasticity rule which allows us to avoid the overhead of updating every synapse equation on each time step. Similarly, we have greatly enhanced our ability to update neuronal dynamics equations by implementing their update using parallel computation technology. Finally, we introduce the reader to the analytical methods (Brunel 2000) used to predict the mean firing rate of a network, with random connectivity under a given set of inputs. We coupled these techniques to our model for the population dynamics of synaptic plasticity in order to choose parameters which set our network to a desired firing state.

In the Results section (3.3) we demonstrate the power and applicability of our developed tools, first in the case of the basic synaptic plasticity rule without a bistability term. This allows us to demonstrate a power law in the memory decay time scale for low firing rates, and the correctness of our predictions of average memory behaviour across a wide range of frequencies for Poisson processes. We then move on to examining the influence of the addition of a double-well potential to the synaptic plasticity rule and how it changes the memory time scales, again for uncorrelated Poisson processes. Finally, we use our theoretical tools in order to successfully incorporate the synaptic plasticity rule into a large-scale randomly connected network of leaky integrate-and-fire (LIF) neurons. With our theoretical predictions we demonstrate that it is possible to initialise such a network such that it will fire in a stable asynchronous regime and indeed we can then implant ‘memories’ in the network and observe them survive or decay under background network activity. Finally, we present a mechanism for memory induction via high frequency bursting in a recurrent network.

Most of the work which follows in this chapter will appear as a paper in *PLoS Computational Biology*. For coherence we have chosen to keep some aspects of the Introduction which follows, despite some overlap with the literature review in Chapter 1. We have changed somewhat the order of presentation of the contents with regards to the paper; adding explanatory text and enlarging on certain details. The

section on memory induction in a network is new and will be presented in a follow-up paper, here it is presented in a simplified format as a proof of concept.

3.1 Introduction

Many experiments have shown that long-lasting changes in synaptic efficacy can be induced by spiking activity of pre- and postsynaptic neurons (Bliss and Lømo 1973; Dudek and Bear 1993). In hippocampal and neocortical *in vitro* preparations, both long-term potentiation and depression can be induced by protocols in which pre- and postsynaptic neurons emit tens to hundreds of spikes in specific temporal patterns (Bi and Poo 1998; Campanac and Debanne 2008; Debanne et al. 1994; Froemke and Dan 2002; Markram et al. 1997; Sjöström et al. 2001; Wang et al. 2005; Wittenberg and Wang 2006). In those preparations, plasticity has been shown to depend both on relative timing of pre- and postsynaptic spikes ('spike timing dependent plasticity', or STDP), and the firing rates of pre- and postsynaptic neurons. These induced changes in the connectivity of a neural circuit have then been assumed to maintain or initiate a long-term memory trace of external inputs that triggered these synaptic changes (Hebb 1949). However, in order for this theory to be valid, the induced synaptic changes need to survive both activity triggered by other inputs, and the ongoing background activity that is pervasive in cortex (Burns and Webb 1976; Churchland et al. 2010). How changes in synaptic connectivity survive the continuous presentation of other inputs has been the subject of several studies (Amit et al. 1994; Fusi et al. 2005). Here, we focus most of our study on the decay of the synaptic memory trace due to background activity, using a theoretical approach.

Synaptic plasticity has been described using a multitude of different models and approaches (Albers et al. 2013; Amit and Fusi 1994; Bhalla and Iyengar 1999; Bienenstock et al. 1982; Clopath et al. 2010; El Boustani et al. 2012; Froemke and Dan 2002; Gerstner et al. 1996; Graupner and Brunel 2012; Kumar and Mehta 2011; Lisman 1985; Oja 1982; Pfister and Gerstner 2006; Sejnowski 1977; Shouval et al. 2002; Sjöström et al. 2001; Song et al. 2000; Yger and Harris 2013). In early plasticity models, synaptic changes were purely induced by the firing-rates of pre- and postsynaptic neurons (Bienenstock et al. 1982; Oja 1982; Sejnowski 1977). At the end of the nineties, theorists introduced purely spike-timing based models (Gerstner et al. 1996; Song et al. 2000). Finally, more recent models have been striving to capture a wide range of experimental data, and as a result capture both the spike-timing and firing rate dependence of synaptic plasticity (Albers et al. 2013; Clopath et al. 2010; El Boustani et al. 2012; Graupner and Brunel 2012; Kumar and Mehta 2011; Pfister

and Gerstner 2006; Shouval et al. 2002; Sjöström et al. 2001; Yger and Harris 2013). These models are natural candidates for studies of the stability of synaptic changes during ongoing activity. In our work, we choose the model of Graupner and Brunel (2012) for the following reasons: (i) the model includes the calcium concentration in the post-synaptic spine, which is known to be a crucial link between pre- and post-synaptic activity and long-term synaptic changes; (ii) the model exhibits bistability of synaptic changes accounting for experimental evidence suggesting that CA3-CA1 synapses in the hippocampus are bistable (O’Connor et al. 2005b; Petersen et al. 1998); (iii) the model is simple enough that the dynamics of the synaptic efficacy variable can be computed analytically.

Postsynaptic calcium entry has been identified to be a necessary (Ismailov et al. 2004; Mizuno et al. 2001; Nevian and Sakmann 2006) and sufficient (Malenka et al. 1988; Neveu and Zucker 1996; Yang et al. 1999) signal for the induction of synaptic plasticity (but see ref. Nabavi et al. 2013). However, most of the *in vitro* experiments evoking synaptic changes use elevated extracellular calcium concentrations, while *in vivo* calcium levels are estimated to be around 1.5 mM (Silver and Erecińska 1990). The impact of reduced calcium entry due to the lower extracellular calcium concentration *in vivo* on the time scale of synaptic decay has not been considered heretofore.

In the Results (Section 3.3), we study the persistence of synaptic changes, first in a synapse connecting a pair of independent Poisson neurons, and second in a large network of excitatory and inhibitory leaky integrate-and-fire (LIF) neurons. We show that in the absence of bistability, synaptic changes decay exponentially during ongoing activity and that the time constant exhibits a power-law like behaviour with respect to the present firing rate. We demonstrate that the reduced extracellular calcium concentration *in vivo* leads to several orders of magnitude longer memory time scales. The introduction of bistability in the synaptic plasticity rule further stabilises synaptic changes at low firing rates and extensively prolongs memory time scales when combined with the *in vivo* extracellular calcium conditions. Finally, we extend our results to a large recurrent network of LIF neurons, where we demonstrate network firing rate stability under synaptic plasticity, high-frequency stimulation induced implantation and decay of an implanted memory for *in vitro* parameters and long term memory maintenance for *in vivo* parameters.

3.2 Materials and Methods

We have examined the time scales of memory change under background cortical firing, in a calcium-based model of synaptic plasticity (Graupner and Brunel 2012). We start here by shortly introducing the model and we explain how we account for *in vitro* and *in vivo* conditions in the model.

In the model, the temporal evolution of the synaptic efficacy variable, $\rho(t) \in [0, 1]$, is described by

$$\tau \frac{d\rho}{dt} = -\frac{\partial U(\rho)}{\partial \rho} - \gamma_D \rho \Theta(c(t) - \theta_D) + \gamma_P (1 - \rho) \Theta(c(t) - \theta_P) + \sigma \sqrt{\tau} \sqrt{\Theta[c(t) - \theta_D] + \Theta[c(t) - \theta_P]} \eta(t), \quad (3.1)$$

where τ is the time constant of synaptic efficacy changes, and $c(t)$ is the calcium concentration. The dynamics of ρ depends on four terms:

1. The dynamics are governed by a potential $U(\rho)$ for low calcium concentrations ($c(t) < \theta_d$) since all other terms on the right-hand side of Eq. (3.1) are then zero. In the following we consider two possible scenarios for $U(\rho)$: (i) a flat potential, $U(\rho) = 0$ - in this case the synaptic efficacy variable stays constant in time in the absence of calcium transients. This means all possible values of $\rho \in [0, 1]$ are stable; (ii) a double well potential,

$$U(\rho) = \frac{1}{4} \rho^2 (1 - \rho)^2. \quad (3.2)$$

In this case, ρ evolves towards one of two possible stable fixed points (the minima of U), one at $\rho = 0$ - the DOWN state -, the other at $\rho = 1$ - the UP state -, depending on the initial condition. This corresponds to a bistable synapse.

2. For intermediate calcium concentrations ($\theta_p > c(t) > \theta_d$), the synapse is depressed, with a depression rate γ_d . This depression remains active for $c(t) > \theta_p$ but its magnitude is outweighed by potentiation.
3. For large calcium concentrations ($c(t) > \theta_p > \theta_d$), the synapse is potentiated, with a potentiation rate γ_p .
4. A noise is only active when calcium concentration is above the lowest plasticity threshold θ_D , and increases in magnitude when the upper plasticity threshold, θ_P , is also crossed. σ defines the amplitude of the noise, and $\eta(t)$ is a Gaussian white noise process with unit variance.

Changes in ρ are induced by increases in the postsynaptic calcium concentration, $c(t)$ (see Eq. (3.3)), evoked by pre- and postsynaptic spikes. The calcium concentration increases by an amount, $C_{\text{pre}}/C_{\text{post}}$, in response to pre-/postsynaptic spikes respectively, while it decays exponentially with a time constant τ_{Ca} in between spikes. Calcium transients induced by presynaptic activity are assumed to represent calcium influx through NMDA receptors, while calcium transients induced by postsynaptic spikes are assumed to represent activation of voltage-gated calcium channels (Sabatini et al. 2002) (see Graupner and Brunel 2012 for more details of the model).

This calcium-based model of synaptic plasticity has been used to successfully fit data from various experimental preparations (Graupner and Brunel 2012). Here, we use the data-set that best fits plasticity data obtained in visual cortex slices (Sjöström et al. 2001) - hereafter called the ‘*in vitro*’ parameter set. In this experiment, the extracellular calcium concentration was set to be 2.5 mM (Sjöström et al. 2001), which is significantly higher than the estimated *in vivo* concentration of about 1.5 mM (Silver and Erecińska 1990). Here we assume that a decrease in extracellular calcium concentration leads to a proportional decrease in the calcium influx into the postsynaptic spine. This assumption is quite different from the case in Chapter 2, where we could not predict the influence of reduced extracellular calcium concentration on the presynaptic production of NO as the actual production process would be nonlinearly dependent on calcium. In the current model both sources of calcium are believed to come from extracellular sources (at least to a first approximation) so decreasing available calcium will reduce the influx variables. Using this assumption, we can readily predict the effects of decreasing the extracellular calcium concentration on the plasticity rule in the calcium-based model by scaling the amplitudes of the pre- and postsynaptically evoked calcium transients according to the ratio of calcium concentrations, i.e. $1.5/2.5=0.6$. This leads to what we call the ‘*in vivo*’ parameter set. Values of all parameters for both conditions are indicated in Tab. 3.1.

The dynamics of the synaptic efficacy in response to calcium transients under the *in vitro* and the *in vivo* conditions are illustrated in Fig. 3.1. The synaptic efficacy is only modified when the calcium concentration increases above the depression threshold θ_D (Fig. 3.1,**B-E**). For the *in vitro* case, this happens whenever a postsynaptic spike occurs since $C_{\text{post}} > \theta_D$, but for the *in vivo* parameter set this happens much more rarely because of the smaller calcium amplitudes ($C_{\text{pre}}, C_{\text{post}} < \theta_D$ in the *in vivo* case; see Tab. 3.1). In the latter case, synaptic changes are only induced whenever subsequent spikes occur in short succession such that calcium accumulates and crosses the depression and/or potentiation threshold. Such events are rare at low firing rates (Fig. 3.1,**D-E**).




Figure 3.1. Dynamics of the synaptic plasticity model with the *in vitro* and *in vivo* parameter sets. (A) Pre- and postsynaptic spike trains generated as realisations of Poisson processes at 1/s. (B,C) The spike train in A induces large calcium transients (blue trace) with the *in vitro* parameter set ($C_{\text{pre}} = 0.562$ and $C_{\text{post}} = 1.240$; see Tab. 3.1). Whenever the calcium trace crosses the depression (cyan) or potentiation thresholds (orange), changes in the synaptic efficacy (green) are induced. (D,E) Same as in B,C but with calcium amplitudes corresponding to the *in vivo* case ($C_{\text{pre}} = 0.337$ and $C_{\text{post}} = 0.744$). The small calcium transients do not cross the depression/potentiation thresholds and no efficacy changes are observed. Note that the flat potential for ρ is used here and that noise is set to zero for clarity, $\sigma = 0$.

3.2.1 Calcium-based plasticity model

The temporal dynamics of the synaptic efficacy in the calcium-based model are given in Eq. (3.1). Changes in ρ are driven by the postsynaptic calcium concentration, c . The underlying basis for most of the analytical tools which we develop are based on the development of a probabilistic prediction for the mean activity of this calcium concentration, thus we present it separately here.

The calcium dynamics are modelled using instantaneous increases of size C_{pre} and C_{post} in response to pre- and postsynaptic spikes, respectively, followed by an

exponential decay

$$\frac{dc}{dt} = -\frac{c}{\tau_{Ca}} + C_{\text{pre}} \sum_i \delta(t - t_i - D) + C_{\text{post}} \sum_j \delta(t - t_j), \quad (3.3)$$

where τ_{Ca} is the calcium decay time constant, and C_{pre} , C_{post} the pre- and postsynaptically evoked calcium amplitudes. The sums go over all pre- and postsynaptic spikes occurring at times t_i and t_j , respectively. The time delay, D , between the presynaptic spike and the occurrence of the corresponding calcium transient accounts for the slow rise time of the NMDAR-mediated calcium influx.

In the supplementary material of their paper Graupner and Brunel (2012) also developed a more complicated calcium influx model. This model used a difference of exponentials to model the slow rising NMDA-mediated calcium increase as a result of a presynaptic spike, and a further difference of exponentials to account for the much faster voltage-dependent calcium channel (VDCC) calcium influx as a result of a postsynaptic spike. This VDCC term is further influenced nonlinearly by the NMDAR term leading to a much greater calcium influx when a postsynaptic spike follows a presynaptic spike. Graupner and Brunel (2012) show that this model for calcium dynamics coupled to their model of synaptic efficacy (Eq. (3.1)) can reproduce the standard STDP curve of Bi and Poo (1998, experimental data points shown in Fig. 1.6). They further motivate the utility of such an extension to their basic model (Eq. 3.3) in the modelling of the activity of certain pharmacologically interventions (e.g. NMDAR blocking). The two models of calcium dynamics presented in their paper are somewhat reminiscent of the two models of NO production which we presented in our parallel fibre to purkinje cell plasticity model in Chapter 2. For our analysis, we have chosen to implement the simplified calcium dynamics for a number of reasons (i) this is the version of the model which was successfully fitted to the greatest number of experiments, and for which we have access to the fit parameters (Graupner and Brunel 2012); (ii) our theoretical prediction of calcium concentration, which we will present in the next section, is more readily derived for the linear model; (iii) our event-based implementation for the evolution of synaptic efficacy is again easier in the linear case. This is an excellent example of a case where a more detailed model may be more biologically accessible but a linearisation of the model provides much greater scope for analysis without losing much in overall dynamics.

Parameter	In-vitro	In-vivo
C_{pre}	0.56175	0.33705
C_{post}	1.23964	0.74378
τ_{Ca} (ms)	22.6936	
θ_D	1	
θ_P	1.3	
γ_D	331.909	
γ_P	725.085	
σ	3.3501	
τ (sec)	346.3615	
ρ_*	0.5	
D (ms)	4.6098	

Table 3.1. Parameters of the calcium-based synapse model. The *in vitro* values are obtained by fitting the model to cortical plasticity data (Graupner and Brunel 2012; Sjöström et al. 2001). *In vivo* calcium amplitudes are scaled from the *in vitro* values according to the change in extracellular calcium concentration.

3.2.2 Probability density function of the calcium concentration

The probability density function (PDF) of the calcium concentration can be calculated if pre- and postsynaptic neurons fire as independent Poisson processes with rates ν_{pre} and ν_{post} respectively. In these conditions, the calcium concentration follows a shot noise process, where the effect of a single event (a ‘shot’) is the shape of a single calcium transient, $F(t) = a \exp(-t/\tau_{\text{Ca}})$, where a is either C_{pre} or C_{post} and τ_{Ca} the time constant of calcium decay. The general solution for a single shot noise process, with $a = 1$, was derived in (Gilbert and Pollak 1960). We were able to adapt their work to two independent processes, with $a \in \{C_{\text{pre}}, C_{\text{post}}\}$ and $\nu = \nu_{\text{pre}} = \nu_{\text{post}}$, deriving a probability density function given by the master equation

$$cP'(c) = (2\nu\tau_{\text{Ca}} - 1)P(c) - \nu\tau_{\text{Ca}}P(c - C_{\text{pre}}) - \nu\tau_{\text{Ca}}P(c - C_{\text{post}}) \quad (3.4)$$

where $P(c)$ is the probability density of calcium c . For a complete description of shot noise processes, including generalisation to the case $\nu_{\text{pre}} \neq \nu_{\text{post}}$, and how we applied them in the current work see Appendix C.

The probability density function, $P(c)$, allows us to calculate the fraction of time

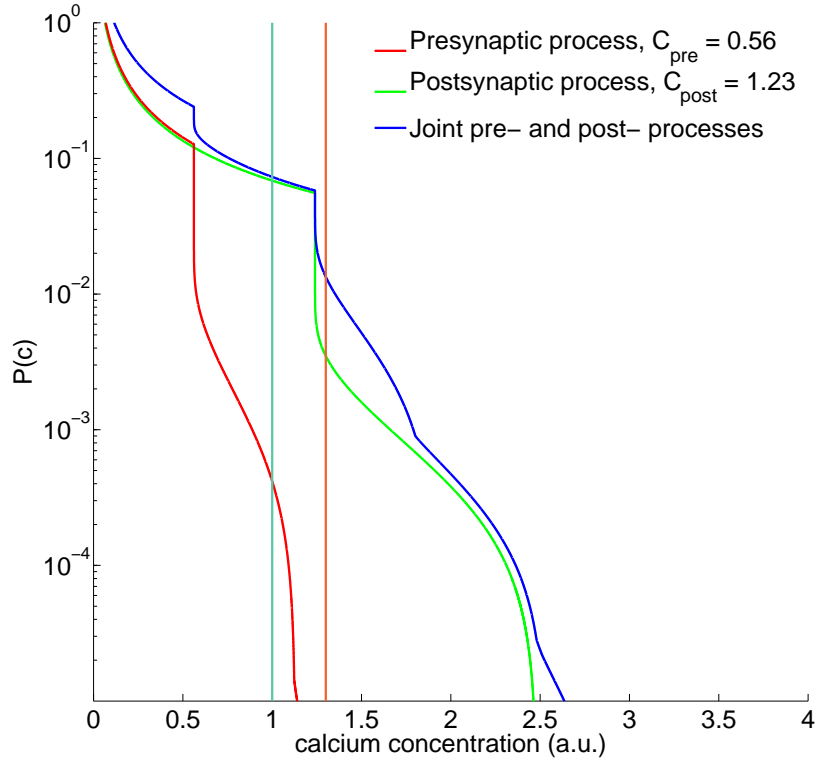


Figure 3.2. Example of 1/sec shot noise prediction of calcium concentration for *in vitro* parameter set The probability density, $P(c)$, is shown for individual Poisson processes for the calcium influx value of $C = 0.56$ (red) or $C = 1.23$ (green). When both processes combine at a single synapse we obtain the joint probability density function (blue). Plasticity thresholds are shown as vertical lines.

spent above the depression and potentiation thresholds according to

$$\begin{aligned}\alpha_D(\nu) &= 1 - \int_0^{\theta_D} P(c, \nu) dc \\ \alpha_P(\nu) &= 1 - \int_0^{\theta_P} P(c, \nu) dc\end{aligned}\tag{3.5}$$

That is, for a given pre- and postsynaptic Poisson firing rate, ν , the proportion of time that the calcium will spend above the potentiation/depression threshold is equal to 1 minus the area of the density function lower than the threshold. Recall that the area of a particular zone of a probability density function can be expected to equal the proportion of time spent in that zone, if the firing rate is held constant. These formulae will be useful to us in calculating the average long-run potentiation and depression rates.

In order to solve Eq. (3.4) we first present a simplified case where $C_{\text{pre}} = C_{\text{post}} = 1$. This case is useful as an example as it allows us to develop a completely analytical solution demonstrating the piecewise approach necessary for solving the equation. Furthermore, we will use this form later (Section 3.3.1) to derive a simple power law result which appears to carry over, even to the case $C_{\text{pre}} \neq C_{\text{post}}$. Solving then for $P(c)$ we get

$$P(c) = Ac^{2\nu\tau_{\text{Ca}}-1} \quad c \in [0, 1]; \quad (3.6)$$

$$= Ac^{2\nu\tau_{\text{Ca}}-1} \left(1 - 2\nu\tau_{\text{Ca}} \int_0^{c-1} \frac{z^{2\nu\tau_{\text{Ca}}-1}}{(z+1)^{2\nu\tau_{\text{Ca}}}} dz \right) \quad c \in [1, 2]; \quad (3.7)$$

$$= Ac^{2\nu\tau_{\text{Ca}}-1} \left(1 - (c-1)^{2\nu\tau_{\text{Ca}}} {}_2F_1(2\nu\tau_{\text{Ca}}, 2\nu\tau_{\text{Ca}}; 2\nu\tau_{\text{Ca}} + 1; 1 - c) \right) \quad c \in [1, 2]; \quad (3.8)$$

$$= c^{2\nu\tau_{\text{Ca}}-1} \left(A - 2\nu\tau_{\text{Ca}} \int_0^{c-1} \frac{P(z)}{(z+1)^{2\nu\tau_{\text{Ca}}}} dz \right) \quad c > 2, \quad (3.9)$$

where ${}_2F_1(a, b, c, z) = \sum_{n=0}^{\infty} \frac{(a)_n (b)_n}{(c)_n} \frac{z^n}{n!}$ is the ordinary hypergeometric function (Abramowitz and Stegun 1970),

$$A = \frac{\exp(-2\nu\tau_{\text{Ca}}\gamma)}{\Gamma(2\nu\tau_{\text{Ca}})}, \quad (3.10)$$

γ is Euler-Mascheroni constant, $\gamma \sim 0.577215665$, and Γ is the gamma function.

Fitting the calcium-based model to cortical plasticity data yields $C_{\text{pre}} < C_{\text{post}}$ (see Tab. 3.1). In this case, the initial parts of the solution to Eq. (3.4) read

$$P(c) = Bc^{2\nu\tau_{\text{Ca}}-1}, \quad c \in [0, C_{\text{pre}}] \quad (3.11)$$

$$= Bc^{2\nu\tau_{\text{Ca}}-1} \left(1 - \nu\tau_{\text{Ca}} \int_0^{c-C_{\text{pre}}} \frac{P(x - C_{\text{pre}})}{(x + C_{\text{pre}})^{2\nu\tau_{\text{Ca}}}} dx \right), \quad (3.12)$$

$$c \in [C_{\text{pre}}, \min(2C_{\text{pre}}, C_{\text{post}})]$$

$$= Bc^{2\nu\tau_{\text{Ca}}-1} \left(1 - \frac{1}{2} D^{2\nu\tau_{\text{Ca}}} {}_2F_1(2\nu\tau_{\text{Ca}}, 2\nu\tau_{\text{Ca}}; 2\nu\tau_{\text{Ca}} + 1; -D) \right), \quad (3.13)$$

$$c \in [C_{\text{pre}}, \min(2C_{\text{pre}}, C_{\text{post}})]$$

where $D = \left(\frac{c - C_{\text{pre}}}{C_{\text{pre}}}\right)$, and $B = A / (C_{\text{pre}} C_{\text{post}})^{\nu \tau_{\text{Ca}}}$ is a normalisation parameter which assures that $\int P(I) dI = 1$. This provides us with an analytical solution up to $c = \min(2C_{\text{pre}}, C_{\text{post}})$. While it is technically possible to continue this piecewise analytical approach for a few more intervals it becomes quickly unruly. Therefore we calculate the PDF for $c > \min(2C_{\text{pre}}, C_{\text{post}})$ by a numerical integration of Eq. (3.4) (see Eq. (C.30)).

3.2.3 Diffusion approximation for the synaptic efficacy with a flat potential

Performing a diffusion approximation of the synaptic efficacy ρ turns Eq. (3.1) into an Ornstein-Uhlenbeck process (Graupner and Brunel 2012). The requirements for this are that single calcium transients cause very little change in the synaptic efficacy. This is certainly the case for the chosen parameter set (Tab. 3.1), where $\tau_{\text{Ca}} \ll \tau$ means that many calcium transients must be integrated before a large change in synaptic efficacy will occur. The temporal evolution of ρ is then described by the Fokker-Planck equation

$$\tau \frac{d\rho}{dt} = \Gamma_P(1 - \rho) - \Gamma_D \rho + \sigma \sqrt{\tau} \sqrt{\alpha_P + \alpha_D} \eta(t), \quad (3.14)$$

for the case of a flat potential (*i.e.* $\partial U / \partial \rho = 0$). Here τ is the time constant of the synaptic efficacy equation (Eq. (3.1)), σ the magnitude of synaptic noise, $\eta(t)$ a gaussian noise process with zero mean and unitary standard deviation. We define $\Gamma_P(\nu) = \gamma_P \alpha_P$ and $\Gamma_D(\nu) = \gamma_D \alpha_D$ as the mean potentiation and depression rates, respectively, and the α_i are defined in Eq. (3.5).

The bounds at $\rho = 0$ and $\rho = 1$ lead to a truncated Ornstein-Uhlenbeck process, whose distribution is a truncated Gaussian, whose mean converges exponentially to

$$\bar{\rho}(\nu) = \frac{\Gamma_P(\nu)}{\Gamma_D(\nu) + \Gamma_P(\nu)} + \sigma_\rho \frac{G\left(\frac{-\bar{\rho}}{\sigma_\rho}\right) - G\left(\frac{1-\bar{\rho}}{\sigma_\rho}\right)}{H\left(\frac{-\bar{\rho}}{\sigma_\rho}\right) - H\left(\frac{1-\bar{\rho}}{\sigma_\rho}\right)} \quad (3.15)$$

where $\sigma_\rho^2 = \frac{\sigma^2(\alpha_P(\nu) + \alpha_D(\nu))}{2(\Gamma_P(\nu) + \Gamma_D(\nu))}$, $G(z) = \frac{1}{\sqrt{2\pi}} \exp\left(\frac{-z^2}{2}\right)$ is the Gaussian with zero mean and unit variance, and $H(z) = \frac{1}{2} \left(1 - \text{erf}\left(\frac{z}{\sqrt{2}}\right)\right)$ is the complementary cumulative density function of G . The time constant, τ_{eff} , of the exponential decay to $\bar{\rho}$ is defined by

$$\tau_{\text{eff}}(\nu) = \frac{\tau}{\Gamma_D(\nu) + \Gamma_P(\nu)}. \quad (3.16)$$

3.2.4 Kramers expected escape time from a double-well potential

In the case of a double-well potential, the diffusion approximation turns Eq. (3.1) into a Fokker-Planck equation

$$\tau \frac{\partial P}{\partial t} = \frac{(\alpha_D + \alpha_P)}{2} \sigma^2 \frac{\partial^2 P}{\partial \rho^2} + \frac{\partial}{\partial \rho} \left(\left((\Gamma_P + \Gamma_D) \rho + \Gamma_P + \frac{\partial U}{\partial \rho} \right) P \right). \quad (3.17)$$

This equation can be rewritten as

$$\tau \frac{\partial P}{\partial t} = \frac{1}{2} \sigma_{\text{eff}}^2 \frac{\partial^2 P}{\partial \rho^2} + \frac{\partial}{\partial \rho} \left(\left(\frac{\partial U_{\text{eff}}}{\partial \rho} \right) P \right). \quad (3.18)$$

where the *effective potential*, U_{eff} , is the sum of the ‘bare’ potential U and two quadratic terms proportional to the potentiation and depression rates, respectively (Eq. (3.20)), and σ_{eff}^2 is the amplitude of the effective noise

$$\sigma_{\text{eff}}^2(\nu) = \sigma^2(\alpha_D(\nu) + \alpha_P(\nu)). \quad (3.19)$$

The effect of background activity on the dynamics of ρ can be explained by the fact that it modifies the potential, $U(\rho)$, leading to an *effective potential*

$$U_{\text{eff}}(\nu, \rho) = \frac{1}{4} \rho^2 (1 - \rho)^2 + \frac{\gamma_D \alpha_D(\nu)}{2} \rho^2 + \frac{\gamma_P \alpha_P(\nu)}{2} (1 - \rho)^2. \quad (3.20)$$

In Eq. (3.20), the first term on the r.h.s. corresponds to the ‘bare’ double well potential $U(\rho)$ (Eq. (3.2)); the second term describes the effect of depression on the potential, that tends to strengthen the stability of the lower well (DOWN state) at $\rho \sim 0$, at the expense of the other well that tends to disappear when $\alpha_D(\nu)$ increases; finally, the last term describes the effect of potentiation, that shifts the minimum of the only remaining well towards higher values of ρ when $\alpha_P(\nu)$ increases.

When the effective potential is dominated by the double-well term (first term on the r.h.s. of Eq. (3.20)), the ‘escape’ rate from the UP state is driven by noise and can be estimated using Kramers theory (Gardiner 1986; Kramers 1940). The height of the potential barrier, $\Delta U_{\text{eff}} = U_{\text{eff}}(\rho_{\text{un}}) - U_{\text{eff}}(\rho_{\text{up}})$, determines the mean dwell time in the UP state, where ρ_{up} and ρ_{un} are the local minima and maxima of the effective potential and are obtained solving $\frac{\partial U_{\text{eff}}}{\partial \rho} = 0$. The mean first passage time for a synapse beginning in the UP state to the peak of the effective potential well, ρ_{up} , dominates the dwell time, as once a synapse crosses this threshold it is much more likely to fall quickly down into the DOWN state rather than to cross

back into the UP state. This allows us to calculate the expected escape time from the potential well

$$E(\tau_{\text{eff}}) = \frac{2\pi\tau}{\sqrt{|U''_{\text{eff}}(\rho_{\text{up}})U''_{\text{eff}}(\rho_{\text{un}})|}} \exp\left\{\frac{2\Delta U_{\text{eff}}}{\sigma_{\text{eff}}^2}\right\} \quad (3.21)$$

where

$$U''_{\text{eff}}(\rho) = \frac{1}{2} - 3\rho + 3\rho^2 + \Gamma_D - \Gamma_P. \quad (3.22)$$

3.2.5 Numerical methods: Event-based implementation

The temporal evolution of individual synaptic weights in the calcium-based model can be calculated in an event-based manner (as opposed to a finite difference method with a fixed time step Δt) in an analytically exact way. This greatly accelerates network simulations since the network variables are updated at the occurrence of spikes only. In the following we describe the practical implementation of the event-based update.

For the event-based update, three variables have to be stored per synapse: the calcium concentration, c , the synaptic efficacy variable, ρ , and the time of the last spike seen by the synapse, t . The synapse variables c and ρ must be updated on the occurrence of three events: (1) at the presynaptic spike when the postsynaptic voltage is increased, (2) with delay D after a presynaptic spike when the presynaptically evoked calcium increase occurs (see Eq. (3.3)), (3) and at the postsynaptic spike when the postsynaptic calcium increase occurs.

3.2.5.0.1 Calcium update The calcium concentration decays exponentially between events and is instantaneously increased by the amplitude C_{post} when a postsynaptic spike occurs. In the case of a presynaptic spike, the calcium increase C_{pre} occurs after the delay D (Eq. (3.3)). In consequence, we update the calcium concentration as a decay followed by a calcium concentration increase where the amplitude depends on the identity of the spike and the delay D (at time t_m for pre-synaptic spikes and t_n for post-synaptic spikes). The calcium update for time t_{i+1} (after the last update at t_i) at the three update events described above reads

$$c_{i+1} = c_i \exp(-(t_{i+1} - t_i)/\tau_{\text{Ca}}) + \begin{cases} 0 & (1) \\ C_{\text{pre}}\delta(t_{i+1} - D - t_m) & (2) \\ C_{\text{post}}\delta(t_{i+1} - t_n) & (3) \end{cases} \quad (3.23)$$

where (1), (2), (3) denote the events: (1) upon the occurrence of a presynaptic spike, (2) after a delay following a presynaptic spike, (3) upon the occurrence of a postsynaptic spike.

3.2.5.0.2 Synaptic efficacy update For the propagation of the synaptic efficacy variable, we distinguish between two different regimes, that is, stochastic and deterministic. When the calcium concentration at time t_i is lower than both thresholds, $c_i < \theta_D, \theta_P$, the dynamics of ρ are described deterministically. When the calcium concentration crosses either or both thresholds, the update is stochastic, for the duration of the suprathreshold period, and the dynamics of the mean and the standard deviation of the synaptic efficacy are described by the corresponding Ornstein-Uhlenbeck process. The different regimes may be updated sequentially in a piecewise manner, accounting for first suprathreshold and then subthreshold behaviour.

Here, we determine the possible potentiation/depression threshold crossings of the calcium trace between events i and $i + 1$ with the inter-event interval $\Delta t = t_{i+1} - t_i$. t_P is the time the calcium trace spends above the potentiation threshold within that interval, t_D is the time the calcium trace spends between the potentiation and the depression threshold (note this is different from α_D which represents time above the depression threshold, *including* time above the potentiation threshold), and t_0 the time the calcium trace spends below the depression threshold given by $t_0 = \Delta t - t_P - t_D$. Here, we consider $\theta_P > \theta_D$ only. The updates for the case $\theta_P < \theta_D$ are equivalent. c_{end} refers to the calcium concentration right before the update at event $i + 1$, that is $c_{\text{end}} = c_i \exp(-\Delta t / \tau_{\text{Ca}})$.

Between events at t_i and t_{i+1} , the following six possible threshold crossings can occur (see Fig. 3.3):

- | | | |
|-----|--|--|
| I | $c_i > \theta_P$ and $c_{\text{end}} > \theta_P$ | $t_P = \Delta t$ and $t_D = 0$ |
| II | $c_i > \theta_P$ and $\theta_D < c_{\text{end}} \leq \theta_P$ | $t_P = \tau_{\text{Ca}} \ln\left(\frac{c_i}{\theta_P}\right)$ and $t_D = \Delta t - \tau_{\text{Ca}} \ln\left(\frac{c_i}{\theta_P}\right)$ |
| III | $c_i > \theta_P$ and $c_{\text{end}} \leq \theta_D$ | $t_P = \tau_{\text{Ca}} \ln\left(\frac{c_i}{\theta_P}\right)$ and $t_D = \tau_{\text{Ca}} \ln\left(\frac{\theta_P}{\theta_D}\right)$ |
| IV | $\theta_D < c_i \leq \theta_P$ and $\theta_D < c_{\text{end}} \leq \theta_P$ | $t_P = 0$ and $t_D = \Delta t$ |
| V | $\theta_D < c_i \leq \theta_P$ and $c_{\text{end}} \leq \theta_D$ | $t_P = 0$ and $t_D = \tau_{\text{Ca}} \ln\left(\frac{c_i}{\theta_D}\right)$ |
| VI | $c_i \leq \theta_D$ | $t_P = 0$ and $t_D = 0$ |

The efficacy, ρ , is updated in a piece-wise fashion. Stochastic updates are per-




Figure 3.3. Possible potentiation and depression threshold crossing cases of the calcium trace (blue lines) between events at time t_i and t_{i+1} . The six possible cases are depicted with respect to the location of the potentiation, θ_P (orange dashed line), and the depression thresholds, θ_D (cyan dashed line).

formed when the calcium trace spent time above the potentiation threshold ($t_P > 0$, cases : I, II, III), or between the potentiation threshold and the depression threshold ($t_D > 0$, cases : II, III, IV, V; and for $\theta_P > \theta_D$). A deterministic update is performed if the calcium trace spent time below the depression threshold ($t_0 > 0$, cases : III, V, VI).

In case either the depression, or potentiation, or both thresholds are crossed in the interval $(t_i, t_{i+1}]$, the propagation of ρ is described by the temporal dynamics of the mean and the standard deviation of an Ornstein-Uhlenbeck (OU) process. This approximation is valid if: (i) single calcium transients induce small changes in the synaptic efficacy, and (ii) the depression and potentiation rates (γ_D and γ_P) are sufficiently large so that one can neglect $U(\rho)$ in Eq. (3.1) during synaptic stimulation, in the case of the double-well potential. This reduces Eq. (3.1) to an Ornstein-Uhlenbeck process, for which the potential of ρ during stimulation is quadratic with the minimum at $\bar{\rho}$. Both approximation conditions hold in our case as the effects of single calcium transients are indeed small, and $\gamma_D, \gamma_P \gg \max|\partial U(\rho)/\partial \rho|$ (see Tab. 3.1, the difference of magnitude between $\max|\partial U(\rho)/\partial \rho|$ and the influence of Γ_P, Γ_D is of order 10^4), which means we can ignore the potential well during the suprathreshold period.

If the calcium concentration is larger than either threshold, the update of ρ consists of the temporal evolution of the OU mean and a stochastic part drawn from a distribution given by the temporal evolution of the OU standard deviation. The updates for the stochastic cases are given by:

$$\begin{aligned}
t_P > 0 : \quad \rho(t_i + t_P) &= \frac{\gamma_P}{\gamma_P + \gamma_D} ((1 - \exp(-t_P(\gamma_P + \gamma_D)/\tau)) + \\
&\quad \rho(t_i) \exp(-t_P(\gamma_P + \gamma_D)/\tau)) \\
&\quad + \sigma z_P \sqrt{2} \sqrt{\frac{1 - \exp(-2(\gamma_P + \gamma_D)t_P/\tau)}{2(\gamma_P + \gamma_D)}}, \\
t_D > 0 : \quad \rho(t_i + t_P + t_D) &= \rho(t_i + t_P) \exp(-t_D\gamma_D/\tau) + \sigma z_D \sqrt{\frac{1 - \exp(-2\gamma_D t_D/\tau)}{2\gamma_D}},
\end{aligned}$$

where z_P and z_D are Gaussian random variables of unit variance and zero mean. Note that $t_D = 0$ in case I, and therefore only the first update is performed. Equivalently, $t_P = 0$ in cases IV as well as V and only the second update rule is performed.

When the calcium concentration, c_i , is smaller than the potentiation and the depression threshold ($t_0 = \Delta t - t_P - t_D > 0$, cases III, V and VI), the first two terms on the rhs and the noise term of Eq. (3.1) are zero. That reduces the rhs of Eq. (3.1) to $-\delta U/\delta \rho$ which can be integrated analytically for the flat- and the double well potential considered here. The update of ρ is therefore deterministic and depends on the choice of the potential:

1. flat potential

$$\rho(t_{i+1}) = \rho(t_i + t_P + t_D) \quad (3.24)$$

2. double-well potential

$$\rho(t_{i+1}) = \begin{cases} \frac{1}{2} - \frac{1}{2} \sqrt{1 + (\chi_0 e^{t_0/(2\tau)} - 1)^{-1}} & \text{if } \rho < 0.5 \\ \frac{1}{2} + \frac{1}{2} \sqrt{1 + (\chi_0 e^{t_0/(2\tau)} - 1)^{-1}} & \text{if } \rho \geq 0.5, \end{cases} \quad (3.25)$$

$$\text{with } \chi_0 = (\rho(t_i + t_P + t_D) - 1/2)^2 / (\rho(t_i + t_P + t_D)(\rho(t_i + t_P + t_D) - 1)).$$

By chaining the calculation of first $\rho(t_i + t_P)$, then $\rho(t_i + t_P + t_D)$, and finally $\rho(t_{i+1})$, each of which is dependent on the outcome of the previous update equation we are able to calculate the updated synaptic efficacy in an event-driven manner. This means that the order of 5 million synapses in our simulations are not slowing the progress of the simulation down.

3.2.6 The network

We implemented and simulated a recurrent network of 10,000 leaky integrate-and-fire (LIF) neurons, 8,000 of which are excitatory (E) neurons and 2,000 of which are inhibitory (I). Any two neurons have a spatially uniform probability of connection of 0.05. Autapses are specifically disallowed. Synapses between E neurons are plastic and their weight dynamics are described by the calcium-based plasticity model (Eq. 3.1, Graupner and Brunel (2012)). All other synapses have fixed strength $w_{\alpha\beta}$ ($\alpha, \beta \in \{E, I\}$). A presynaptic spike induces a voltage jump of size $w_{\alpha\beta}$ in the post-synaptic neuron.

The membrane potential of neuron i of population α evolve according to

$$\tau_m \frac{dV_{\alpha i}}{dt} = -(V_{\text{leak}} - V_{\alpha i}) + I_{\alpha i X} + \tau_m \sum_{\beta} \sum_j \sum_k w_{\alpha i \beta j} \delta(t - t_{\beta j k} - \tau_L) \quad (3.26)$$

where

$$I_{\alpha i X} = \mu_{\alpha X} + \sqrt{\tau_{m\alpha}} \sigma \eta_{\alpha i}(t) \quad (3.27)$$

is a common external drive to all neurons, comprising a constant input, $\mu_{\alpha X}$, and a white noise of amplitude $\sigma = 5\text{mV}$. $\eta_{\alpha i}(t)$ is a Gaussian white noise process with unit variance and zero mean, which is uncorrelated from neuron to neuron. In the absence of synaptic inputs each membrane potential decays exponentially to the leak potential, $V_{\text{leak}} = -70\text{mV}$, with a time constant $\tau_m = 20\text{ms}$. Spiking occurs when the voltage crosses a threshold, $V_{\text{thr}} = -50\text{mV}$, after which it is reset to the reset potential, $V_{\text{reset}} = -60\text{mV}$. During all of our simulations, we set the refractory period, during which the membrane potential is fixed at V_{reset} after spiking, to zero. The three sums in the r.h.s. of Eq. (3.26) go over the two populations $\beta \in \{E, I\}$, all presynaptic neurons j , and presynaptic spikes of neuron j in population β , that occur at times $t_{\beta j k}$. Each presynaptic spike of neuron j in population β causes a jump of amplitude $w_{\alpha i \beta j}$ in the voltage of neuron i after a delay τ_L . Here, the delay is chosen to be equal to the integration time step $dt = 0.01 \text{ ms}$ (see Section 3.2.7).

For all connections involving inhibition (i.e. all $(\alpha, \beta) \neq (E, E)$), the connectivity matrix is set as $w_{\alpha i \beta j} = c_{ij} w_{\alpha\beta}$ where c_{ij} are independent, identically distributed (i.i.d.) Bernoulli variables, $c_{ij} = 1$ with probability 0.05, 0 with probability 0.95, and the fixed synaptic weights are $w_{\text{IE}} = 0.1\text{mV}$, $w_{\text{II}} = -0.4\text{mV}$ and $w_{\text{EI}} = -0.4\text{mV}$. E-E synapses are given by $w_{\alpha i \beta j} = c_{ij} \rho_{ij} w_{\text{EE}}$ where c_{ij} are again i.i.d. Bernoulli variables, $c_{ij} = 1$ with probability 0.05, 0 with probability 0.95, ρ_{ij} obeys Eq. (3.1), and $w_{\text{EE}} = 0.2\text{mV}$. The average value of ρ is initially, and remains throughout our simulations, much smaller than 0.5, which means that with a 4 : 1 ratio in the E to I populations, for $\langle \rho \rangle \leq 0.5$ recurrent inhibition dominates excitation, leading to a stable asynchronous irregular state (see Fig. 3.4) (Brunel 2000).




Figure 3.4. Example of network firing in asynchronous irregular state. A sample of 1000 excitatory neurons from the network shows irregular spiking behaviour in the raster (top) and the averaged firing rate of all 8000 excitatory neurons is steady around 1/sec (bottom).

3.2.7 Numerical methods: Network simulations

We numerically simulated the recurrent network of LIF neurons using the forward Euler method with a time step of $dt = 0.01\text{ms}$. Synapses were updated using the event-based implementation described above (Section 3.2.5). The simulations were implemented in C and OpenCL and run on general-purpose GPUs. Parallel generation of random numbers, for the Gaussian noise in the LIF equations, was implemented using the Random123 library (Salmon et al. 2011). Without these innovations it would not have been possible to simulate large networks for such long time periods as will be shown in the Results (Section 3.3). For comparison, simulations which took two hours on our system took approximately one week using the Brian simulator (Goodman and Brette 2009) on a similarly high-performance calculation environment. For a more detailed introduction to our parallel simulations using

OpenCL see Appendix A. We provide a more detailed comparison of the alternative methods of random number generation in large-scale numerical simulations in Appendix B.

In most of our simulations it was sufficient to initialise all of the plastic synapses in the network to the mean value predicted by theory (we will explain this process in the next section, Section 3.2.8). However, in the case of our *in vivo* parameter set, with the double-well potential, the time scales of synaptic change were extremely long which meant that the simulation took longer to reach its long-run steady state, leaving us less time to test the effects of manipulations. To this end, we chose to initialise these simulations closer to their steady-state using the probability density function (PDF) of the effective potential well, U_{eff} (Eq. (3.20)), at the chosen network firing rate (1/s pre- and postsynaptic Poisson firing in our case). The PDF is the limiting distribution of the Kramers process defined by

$$P(\rho) = K \exp\left(\frac{-2U_{\text{eff}}(\nu, \rho)}{\sigma_\rho^2}\right) \quad (3.28)$$

where K is a normalisation constant such that $\int_0^1 = 1$, as we define our synaptic efficacies only on the interval $[0, 1]$. We first generate this PDF numerically and from it the cumulative distribution function (CDF). When initialising the network we generate a uniform random deviate between 0 and 1 for each synapse and by reverse look-up from the CDF we obtain a synaptic weight with value between 0 and 1, whose distribution follows that of the effective potential well Kramers process.

3.2.8 Computing analytically mean firing rates and E-E synaptic efficacy

In a network of excitatory and inhibitory LIF neurons receiving white noise inputs, the mean firing rates of excitatory and inhibitory neurons are given by Amit and Brunel (1997b); Brunel (2000)

$$\nu_E = \Phi(\mu_E, \sigma_E) \quad (3.29)$$

$$\nu_I = \Phi(\mu_I, \sigma_I) \quad (3.30)$$

where Φ is the standard LIF static transfer function (Amit and Brunel 1997b; Amit and Tsodyks 1991; Brunel 2000; Siegert 1951)

$$\frac{1}{\Phi(\mu, \sigma)} = \tau_{\text{rp}} + \tau_m \sqrt{\pi} \int_{\frac{V_{\text{reset}} - \mu}{\sigma}}^{\frac{V_{\text{thr}} - \mu}{\sigma}} e^{u^2} (1 + \text{erf}(u)) du, \quad (3.31)$$

where $\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$ is the error function, μ_α are the mean inputs to population $\alpha \in \{E, I\}$,

$$\mu_E = \mu_{EX} + C_{EE}\tau_{m,E}\nu_E\bar{\rho}(\nu_E)w_{EE} - C_{EI}\tau_{m,E}\nu_I w_{EI} \quad (3.32)$$

$$\mu_I = \mu_{IX} - C_{II}\tau_{m,I}\nu_I w_{II} + C_{IE}\tau_{m,I}\nu_E w_{IE} \quad (3.33)$$

and σ_α is the amplitude of the fluctuations in the inputs to population $\alpha \in \{E, I\}$,

$$\sigma_E = \sigma_{\text{ext}} + (C_{EE}\bar{\rho}(\nu_E)^2 w_{EE}^2 \tau_{m,E} \nu_E) + (C_{EI} w_{EI}^2 \tau_{m,E} \nu_I) \quad (3.34)$$

$$\sigma_I = \sigma_{\text{ext}} + (C_{IE} w_{IE}^2 \tau_{m,I} \nu_E) + (C_{II} w_{II}^2 \tau_{m,I} \nu_I) \quad (3.35)$$

The function $\Phi(\mu, \sigma)$ can be understood as 1 divided by the mean inter-spike interval. The first term on the RHS of Eq. (3.31), τ_{rp} , is the refractory period during which no spiking can occur. The second term represents the mean first-passage time from V_{reset} , the reset potential, to V_{thr} the spike generation threshold, under synaptic drive μ and a noise σ . The synaptic drive term $\mu_{E,I}$, Eqs. (3.32,3.33), breaks down into the constant external drive per population, $\mu_{EX,IX}$ and the membrane time constant of the target population $\tau_{m,\{E,I\}}$, which when multiplied by the synaptic weight, w_{XY} , gives the effect of a single spike. This is then further multiplied by the presynaptic population firing rate $\nu_{E,I}$ and the average number of synapses coming from presynaptic population, C_{XY} , giving the mean excitatory and inhibitory drive to the postsynaptic neuron. The recurrent noise equations break down into similar parameters.

In Eqs. (3.32,3.34) $\bar{\rho}$ is given by the mean synaptic value derived from our truncated Ornstein-Uhlenbeck process described by Eq. (3.15). This means that, for a given excitatory firing rate, ν_E , we can incorporate this value into the excitatory mean-field equation (Eqs. (3.29)) leading to a prediction of the network excitatory population firing rate under those conditions. We further couple the two network mean-field equations (Eqs. (3.29,3.30)), since the excitatory firing rate influences the inhibitory firing rate and vice-versa. Finally, we can optionally feed the excitatory firing rate back into the synaptic efficacy equation (Eq. (3.15)). This is numerically possible as the time scale of change of the synapse is orders of magnitude slower than the time scale of change of the network mean-field equations. In practice, we rarely needed to couple the synaptic efficacy equation to the dynamical firing rate but where we did we saw that the firing rate changed almost immediately in response to a change in input or mean synaptic efficacy, whereas the synaptic efficacy tracked the changes in firing rate more slowly. Note also that for the parameters studied in this paper the effect of heterogeneities in numbers of inputs (Amit and Brunel 1997a; Roxin 2011) have a negligible effect on the mean firing rates of the network.

3.3 Results

We now turn to using the analytical tools developed in Section 3.2 in order to analyse the behaviour of the Graupner and Brunel (2012) synaptic plasticity rule under sustained pre- and postsynaptic activity. In general the activity of neurons in the cortex is believed to be accurately modelled by a Poisson process. This allows us to use our Ornstein-Uhlenbeck derived theory for the activity of the synaptic plasticity model in order to predict the behaviour of the synapse under stationary Poisson activity.

We initially examine (Section 3.3.1) the application of the basic Ornstein-Uhlenbeck process to the model via comparison with the model incorporating a flat potential ($\partial U(\rho)/\partial \rho = 0$). Using simulations of the synaptic model with pre- and postsynaptic Poisson neurons as our reference, we demonstrate that $\bar{\rho}$ (Eq. (3.15)) is in fact an excellent predictor of the mean synaptic efficacy. We then move on to demonstrate that the decay of a ‘memory’ back to the stationary distribution is in fact an exponential decay with time scale τ_{eff} as described by Eq. (3.16). Finally, we demonstrate an analytically derived power law prediction for the memory time scale in the limit of low firing rates, which is matched by the behaviour we observe in the simulations.

In Section 3.3.2 we advance to examining the effects of the addition of a double-well potential (Eq. (3.2)) to the synaptic plasticity rule (Eq. (3.1)). As with the flat potential rule, for certain parameter sets and firing rates the basic Ornstein-Uhlenbeck predictions still hold. But important divergences from the flat potential theory appear for firing rates on the order of that observed in neocortical background activity when we use the *in vivo* parameter set. This leads us to extend the Fokker-Planck equation used in the flat potential theory (Eq. (3.14)) with an *effective* potential well (Eq. (3.20)) which combines the effects of the double-well with the ongoing potentiation/depression processes giving us a bistable synapse for a range of Poisson firing rates. Utilising Kramers theory for escape from a potential well, with this new Fokker-Planck formulation, we can predict the escape times from this effective potential well which dominate the time scales of convergence of synapses towards their lower potential steady-state.

Having demonstrated the power of our theoretically derived tools in the analysis of the behaviour of the synaptic plasticity rule when driven by pre- and postsynaptic Poisson processes, we move on to examining the behaviour of the rule when driven by pre- and postsynaptic leaky integrate-and-fire (LIF) neurons. The purpose of this is to move from independent processes to behaviour in a recurrent network. We begin (Section 3.3.3) by setting up a sparsely connected recurrent network of excitatory and inhibitory LIF neurons to run at a frequency of our choosing, whilst incorporating

synaptic plasticity on the excitatory-to-excitatory connections. We do this by making use of the mean-field equations presented in Section 3.2.8. In order to disentangle the recurrent effects of the network from the effects of the change from pure Poisson neuronal processes to LIF neurons we make comparisons with the behaviour of the synaptic plasticity rule when driven by independent pre- and postsynaptic LIF neurons, which have independent noisy external drives, which lead to mean firing rates of each neuron at our desired frequency. The results of this section indicate that while we have the ability to use our tools to setup a network into certain desirable states there is a non-Poisson aspect to the activity of LIF neurons which violates some of our theoretical precepts and leads to differences between predicted and observed behaviours. Having demonstrated both the capacities and limitations of our theoretical approach in setting up recurrent networks of LIF neurons we return, in Section 3.3.4, to the question of memory time scales. We examine two canonical cases, the decay of a synapse with a flat potential, using the *in vitro* parameter set, and the long term memory maintenance of a synapse with a double-well potential, using the *in vivo* parameter set. Finally in Section 3.3.5, closing the loop on memory plasticity in recurrent networks, we give a practical example of memory induction via high frequency neuronal bursting, which again is tightly tied to our theoretical predictions.

3.3.1 Memory behaviour for a synapse connecting two independent Poisson neurons

We begin by proceeding to study the time scales of synaptic decay to a stationary value. We start with the case of a synapse connecting two neurons firing according to uncorrelated Poisson processes, and compare the memory time constants in the flat and double-well potential cases. Simulations were performed using an event-based implementation of the synaptic plasticity model, which updates the synaptic efficacy only upon the occurrence of pre- and postsynaptic spikes (see Section 3.2 for details).

We initialise the synaptic efficacy to $\rho = 1$ and investigate the time constant of decay in the presence of an ongoing constant firing rate, initially for the flat potential synapse (Eq. 3.1, with $U(\rho) = 0$). Pre- and postsynaptic neurons emit uncorrelated spikes following Poisson statistics, both with a mean rate of 1/s. Under these conditions, a fully potentiated synapse progressively decays and eventually fluctuates around a value of 0.2. On average, this decay is well described by a single exponential function (Fig. 3.5,**A,B**). The time constant of this decay is much longer in the case of the *in vivo* parameter set (Fig. 3.5,**B**) than in the *in vitro* parameter set (Fig. 3.5,**A**). The decay time constant is 2.5 minutes for the *in vitro* case and

approximately 2 hours for *in vivo* in the presence of 1/s pre- and postsynaptic firing (Fig. 3.5,C).

The dynamics of the synaptic efficacy (Eq. 3.1) can be described by a truncated Ornstein-Uhlenbeck (OU) process if single calcium transients induce small changes in the synaptic efficacy and if the potential is flat (see Graupner and Brunel (2012) for the non-truncated case). Truncation of the process is induced by the bounds at $\rho = 0$ and $\rho = 1$. In such a process, the mean synaptic efficacy decays exponentially with a time constant, τ_{eff} (Eq. (3.16)), to an asymptotic average efficacy, $\bar{\rho}(\nu)$ (see Eq. (3.15) in Section 3.2). Both τ_{eff} and $\bar{\rho}(\nu)$ are dependent on $\Gamma_P(\nu) = \gamma_P \alpha_P(\nu)$ and $\Gamma_D(\nu) = \gamma_D \alpha_D(\nu)$, the net potentiation and depression rates, which depend in turn on the learning rates γ_p and γ_d as well as on the average fractions of time spent above the potentiation and depression thresholds, $\alpha_P(\nu)$ and $\alpha_D(\nu)$, respectively. The average fractions of time the calcium traces spend above the potentiation and depression thresholds α_P , α_D were derived in Eq. (3.5) based on $P(c)$, the probability density function of the calcium variable, which can be computed analytically in the case of independent pre- and postsynaptic Poisson firing (Gilbert and Pollak 1960) (see Section 3.2.2). The theory provides an excellent match for the dynamics of the mean synaptic efficacy – compare in Fig. 3.5,A,B the truncated OU theory (blue and red curves), with the simulation mean (green and light blue curves).

Synaptic efficacy decay becomes faster with increasing pre- and postsynaptic firing rates since the calcium trace spends more time above depression and potentiation thresholds (Fig. 3.5,C). At the same time, the asymptotic value of synaptic efficacy ($\bar{\rho}$) increases due to an increase in time spent above the potentiation threshold (Fig. 3.5,D). As a result of the smaller *in vivo* calcium amplitudes, the efficacy decay for the *in vivo* case is, at all firing rates, much slower than the decay *in vitro* (Fig. 3.5,C). The asymptotic efficacy value is lower, at small firing rates ($\nu < 1/\text{s}$), for the *in vitro* case since isolated postsynaptic spikes always cross the depression threshold ($C_{\text{post}} > \theta_D$) which results in a large net depression rate Γ_D , compared to *in vivo* (Fig. 3.5,D).

To get a deeper understanding of the dependence of the memory time scale on the firing rates of pre- and postsynaptic neurons, we set $C_{\text{pre}} = C_{\text{post}} = 1$. This simplification allows us to derive a power law relationship between the memory time scale and the firing rate $\tau_{\text{eff}} \sim 1/(\nu \tau_{\text{Ca}})^k$, where k is the number of (pre- and/or postsynaptic) spikes required to clear the depression/potentiation thresholds. To compute the memory time scale, we need to compute the fraction of times spent above the depression and potentiation thresholds, α_D and α_P . In the case $\theta_D < 1 < \theta_P$, one can show that at low rates $\alpha_P \ll \alpha_D$. Consequently it is only necessary to focus

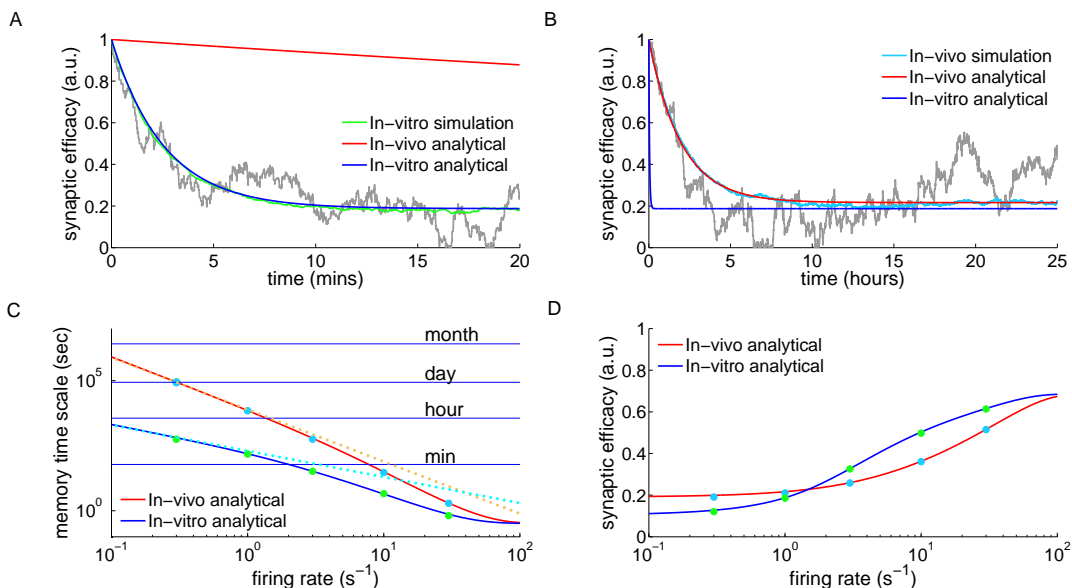


Figure 3.5. Memory decay for a single synapse with flat potential in the presence of uncorrelated pre- and postsynaptic Poisson firing. (A,B) Temporal evolution of the mean synaptic efficacy in the presence of pre- and postsynaptic Poisson firing at 1/s for the *in vitro* (green in *A*) and the *in vivo* (light blue in *B*) parameter sets (mean shown for $N = 1000$ synapses). Blue and red lines show the mean dynamics as predicted by the Ornstein-Uhlenbeck theory. Grey lines show example traces of synaptic efficacy evolution in time. **(C)** Decay time constant as a function of the firing rate for *in vitro* and *in vivo* parameter sets. The blue and red lines show the calculated decay time constant, τ_{eff} , from the OU theory. The points show exponential decay times obtained by fitting single exponential decay functions to the mean synaptic dynamics as shown in *A* and *B* illustrating that the OU theory correctly describes the full model behaviour. The cyan and orange dotted lines illustrate the derived power law behaviour, $\tau_{\text{eff}} \sim 1/\nu^k$, between memory time scales and low firing rates (see text). The power reflects the number of spikes required to cross the plasticity thresholds, that is, $k = 1$ for *in vitro* (cyan dotted line) and $k = 2$ (orange dotted line) for *in vivo* case. **(D)** Asymptotic synaptic efficacy as a function of the firing rate for *in vitro* and *in vivo* parameter sets. The lines show the calculated asymptotic value, $\bar{\rho}$, from the truncated OU theory ($\rho \in [0, 1]$) for *in vitro* (blue line) and *in vivo* (red line) cases. The points show steady-state values obtained by fitting single exponential decay functions to the mean synaptic dynamics as shown in *A* and *B* (green: *in vitro*; light blue: *in vivo*).

our analysis on α_D . When $\theta_D < 1$, the time spent above the depression threshold is

$$\alpha_D(\nu) = 1 - \frac{A\theta_D^{2\nu\tau_{Ca}}}{2\nu\tau_{Ca}}, \quad (3.36)$$

where ν is the firing rate of pre- and postsynaptic neurons ($\nu_{pre} = \nu_{post} = \nu$), τ_{Ca} is the decay constant for the calcium concentration and $A = \exp(-2\nu\tau_{Ca}\gamma)/\Gamma(2\nu\tau_{Ca})$ (see Eqs. (3.6)-(3.10)). This closed form solution allows us to perform an expansion for low firing rates ν

$$\alpha_D(\nu) \sim 2\nu\tau_{Ca} \log\left(\frac{1}{\theta_D}\right) - (\nu\tau_{Ca})^2 \left(2\log(\theta_D)^2 - \frac{\pi^2}{3}\right) + O((\nu\tau_{Ca})^3). \quad (3.37)$$

Similarly for $1 < \theta_D < 2 < \theta_P$ we have in the low rate limit,

$$\alpha_D(\nu) \sim (2\nu\tau_{Ca})^2 \left(\frac{\pi^2}{12} + \text{Li}_2(1 - \theta_D) + \log(\theta_D) \log(\theta_D - 1)\right) + O((\nu\tau_{Ca})^3) \quad (3.38)$$

where Li_2 is the dilogarithm, $\text{Li}_2(z) = \sum_{k=1}^{\infty} z^k/k^2$. Thus, in both cases we see that the leading term is $(\nu\tau_{Ca})^k$ where $k = \text{Ceil}(\theta_D)$ and as $\alpha_P \sim 0$ in this case we find that the memory time scale depends on the firing rate as

$$\tau_{\text{eff}} \sim \frac{\tau}{\gamma_D(\nu\tau_{Ca})^k}.$$

We expect this relationship to hold in general. Intuitively, this is due to the fact that we need k spikes arriving simultaneously on a time scale of order τ_{Ca} in order for the calcium concentration to cross the depression threshold θ_D , and that the probability of observing k spikes in a time interval τ_D is at low rates proportional to $(\nu\tau_{Ca})^k$. We also expect the result to hold in general for $C_{pre} \neq C_{post}$. In this case, we expect that $k = \text{Ceil}\left(\frac{\theta_D}{\max\{C_{pre}, C_{post}\}}\right)$.

The derived power law behaviour for τ_{eff} is plotted in Fig. 3.5, **C** together with the full analytical solution for τ_{eff} . We see that as expected, τ_{eff} scales as $1/\nu$ for the *in vitro* parameter set, where a single spike is enough to cross θ_D , while it scales as $1/\nu^2$ for the *in vivo* parameter set, where two spikes are needed to cross the depression threshold.

The implication of this theoretical result is that, at low firing rates, there is a direct relationship between the number of spikes required to clear the lower plasticity threshold and the memory time scale. Note that the full synaptic efficacy model with $C_{pre} \neq C_{post}$ (see Tab. 3.1) is considered in the following sections.

We have now shown that the truncated Ornstein-Uhlenbeck process (Eq. (3.14)) provides an excellent description of the synaptic efficacy model. For pre- and post-synaptic Poisson processes we can accurately predict the mean exponential decay of synapses from their initial values towards a stationary mean with a characteristic time scale. In addition, a power law relationship between the number of spikes required to cross the depression threshold and the firing rate has been shown to account for memory time scales at low rates. We will now move on to the introduction of a double-well potential to the synaptic plasticity rule.

3.3.2 Memory decay for a bistable synapse

We turn now to examine the effect of a bistability on memory time scales. The dynamics of the synapse is now described by Eq. (3.1), where the potential $U(\rho)$ is given by Eq. (3.2). This double well potential leads to a bistable synapse, that can take two possible efficacy states ($\rho = 0$ and $\rho = 1$) in the absence of activity. In the presence of background activity, transitions between these two states become possible. We investigate stability and transition times for *in vitro* and *in vivo* parameter sets as a function of pre- and postsynaptic firing rates.

The effect of background activity on the dynamics of ρ can be explained by the fact that it modifies the potential, $U(\rho)$, leading to an *effective potential*, U_{eff} , defined in Eq. (3.20). This modified potential combines the effects of the double-well potential with a depression term, driven by activity above the calcium-based depression threshold, which tends to stabilise the lower well (DOWN state) while causing the upper well (UP state) to become shallower, as $\alpha_D(\nu)$ increases. Finally, a potentiation term increases the minimum of the remaining lower well, towards larger values of ρ , as $\alpha_P(\nu)$ increases.

Thus, there are two distinct regions of firing rates in the bistable case with respect to the effective potential. For sufficiently low rates, the effective potential still has two minima (see Fig. 3.6, **A**, and the effective potentials for 0.1/s and 1/s, indicated by orange and magenta curves in the inset). There is a critical value of the rates at which the high efficacy minimum disappears through a saddle-node bifurcation. Beyond this rate, the synapse is no longer bistable, and synaptic efficacy has one stable state only (Fig. 3.6, **A**), equivalent to the asymptotic efficacy value for the flat potential (Fig. 3.5, **D**). Finally, at high firing rates, the ‘bare’ potential becomes negligible, and the effective potential approaches a quadratic potential with a single stable state whose location depends on the rate (green curve in the inset in Fig. 3.6, **A**). The transition from double-well to single well regimes occurs at different firing rates for the *in vitro* ($\sim 0.04/\text{s}$) and the *in vivo* ($\sim 1.3/\text{s}$) parameter sets due to the larger

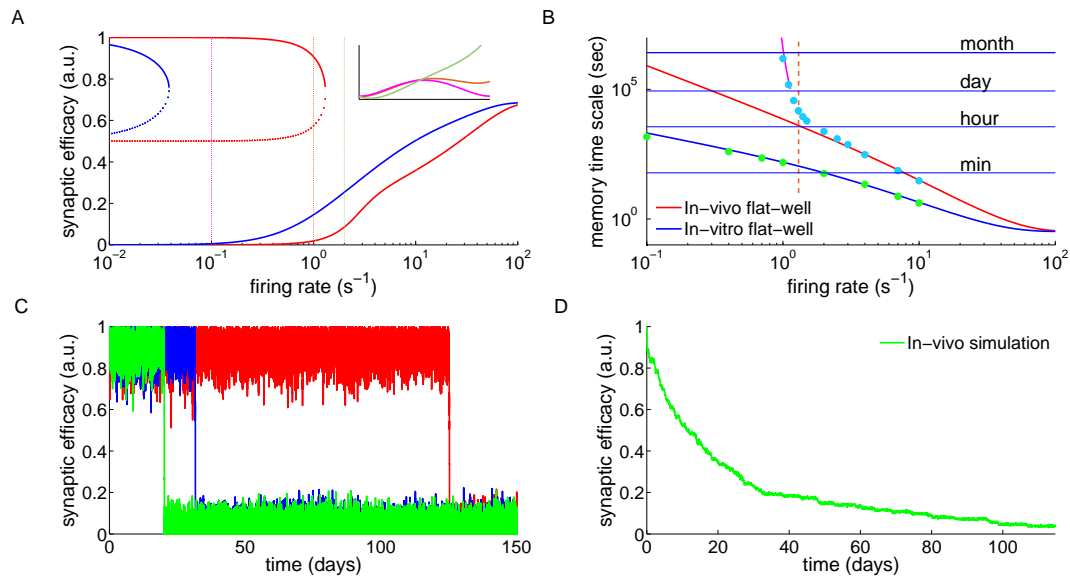


Figure 3.6. Memory decay for a bistable synapse in the presence of uncorrelated pre- and postsynaptic Poisson firing. (A) Steady-states of synaptic efficacy as a function of firing rate for the *in vitro* (blue) and the *in vivo* (red) parameter sets. Stable states are shown by solid lines and unstable states by dotted lines. Synaptic efficacy is bistable at low rates ($< 0.04/s$ for *in vitro* and $< 1.3/s$ for *in vivo*) and monostable at high firing rates. The effective potential of synaptic efficacy is shown for three firing rates (0.1/s - magenta line; 1/s - orange line; 2/s - green line) and the *in vivo* parameter set in the inset (firing rates indicated by vertical lines). (B) Decay time constant as a function of the firing rate for the *in vitro* and the *in vivo* parameter sets. For the *in vivo* parameter set below $\sim 1.3/s$, the bistability greatly extends memory time scale compared to a synapse with flat potential (red line) and can be predicted using Kramers escape rate (magenta line). The vertical dashed line illustrates the frequency at the *in vivo* bifurcation point. For the *in vitro* parameter set, the bistability has no influence on decay time constants for firing rates above 0.1/s. The points show exponential decay times obtained by fitting single exponential decay functions to the mean synaptic dynamics. (C) Individual synaptic efficacy traces for the *in vivo* parameter set at 1/s pre- and postsynaptic firing. The synapses remain in the upper potential well for a long time and stochastically cross the potential barrier to the low efficacy state. (D) Averaged synaptic efficacy trace of many synapses for the *in-vivo* parameter set at 1/s. The bistability extends the memory time scale from hours for a flat potential to days.

calcium amplitudes in the former.

For the *in vitro* parameter set, adding bistability to the synaptic efficacy has no influence on the decay time constant for firing rates larger than approximately 0.1/s (Fig. 3.6,B). In contrast, for the *in vivo* parameter set, bistability considerably prolongs memory decay times with respect to synapses with flat potential at firing rates below $< 1.4/s$. In the presence of two stable states, the decay of memory occurs only due to synaptic noise fluctuations that push the synaptic efficacy out of the upper well. The influence of the double well potential on the dynamics of the synaptic efficacy traps synapses in the UP state leading to long dwell times before crossing the potential barrier and converging to the low efficacy state (Fig. 3.6,C). The double-well has a prolongation effect on memory duration up to firing rates of about $3 - 4/s$ due to the transition between double-well and single-well regimes. At high firing rates, the potentiation and depression processes dominate and the effects of the double-well becomes negligible for both parameter sets, that is, the decay time constant is indistinguishable between flat and double-well potential synapses (see Fig. 3.6,B).

For low firing rates, we can accurately predict the increase in the decay time constant in the presence of bistability using Kramers escape rate for the mean first passage time across a potential barrier (Fig. 3.6,B; see Methods Eq. (3.21)). In this regime, we calculated an effective decay time constant using Kramer’s escape theory given by $\tau_{\text{eff}} \sim \exp\{2\Delta U/\sigma_{\text{eff}}^2\}$, where ΔU is the height of the effective potential barrier and the noise term, σ_{eff}^2 , drives the escape of the efficacy from the upper stable state (see magenta line in Fig. 3.6,B for the *in vivo* case). Both terms ΔU and σ_{eff}^2 are dependent on ν and are detailed, along with τ_{eff} , in Eqs. (3.17) and (3.21) (see Section 3.2 for more details). In the low rate limit, $\sigma_{\text{eff}}^2 \propto 1/\nu^k$ and therefore the memory time scale increases exponentially with the inverse of the rate to a power k , $\tau_{\text{eff}} \propto \exp(a/\nu^k)$, where k is again the number of simultaneous spikes needed to cross the depression threshold. This exponential dependence extends the time scale for synaptic decay at 1/s to the order of one month for a bistable synapse with the *in vivo* parameter set, up from hours for a synapse with flat potential.

3.3.3 Steady-state behaviour of networks of LIF neurons with plastic synapses

We next study the behaviour of the calcium-based synaptic plasticity model in a recurrent network of spiking neurons. We first examine the steady-state of synaptic efficacy and network activity. We again make use of the event-based implementation of the synaptic plasticity rule (described in Section 3.2.5) allowing us to simulate

much longer time scales than are normally attainable by a time stepping simulator.

The recurrent network consists of 8000 excitatory and 2000 inhibitory leaky integrate-and-fire (LIF) neurons. Each neuron receives an external input which consists of a constant (DC) term and a white noise term. External noise is independent from neuron to neuron. Each neuron also receives synaptic inputs from other neurons in the network. The connection probability between any two neurons is 0.05 and uniform in space and across neuron types. Synapses between excitatory neurons are plastic according to the calcium-based plasticity model (Eq. 3.1), while all synapses involving inhibitory neurons are fixed. Parameters of the network are chosen so that the network settles in a stable asynchronous irregular state (Brunel 2000). Hence, correlations between neurons are weak. See Section 3.2.6 for more details of the

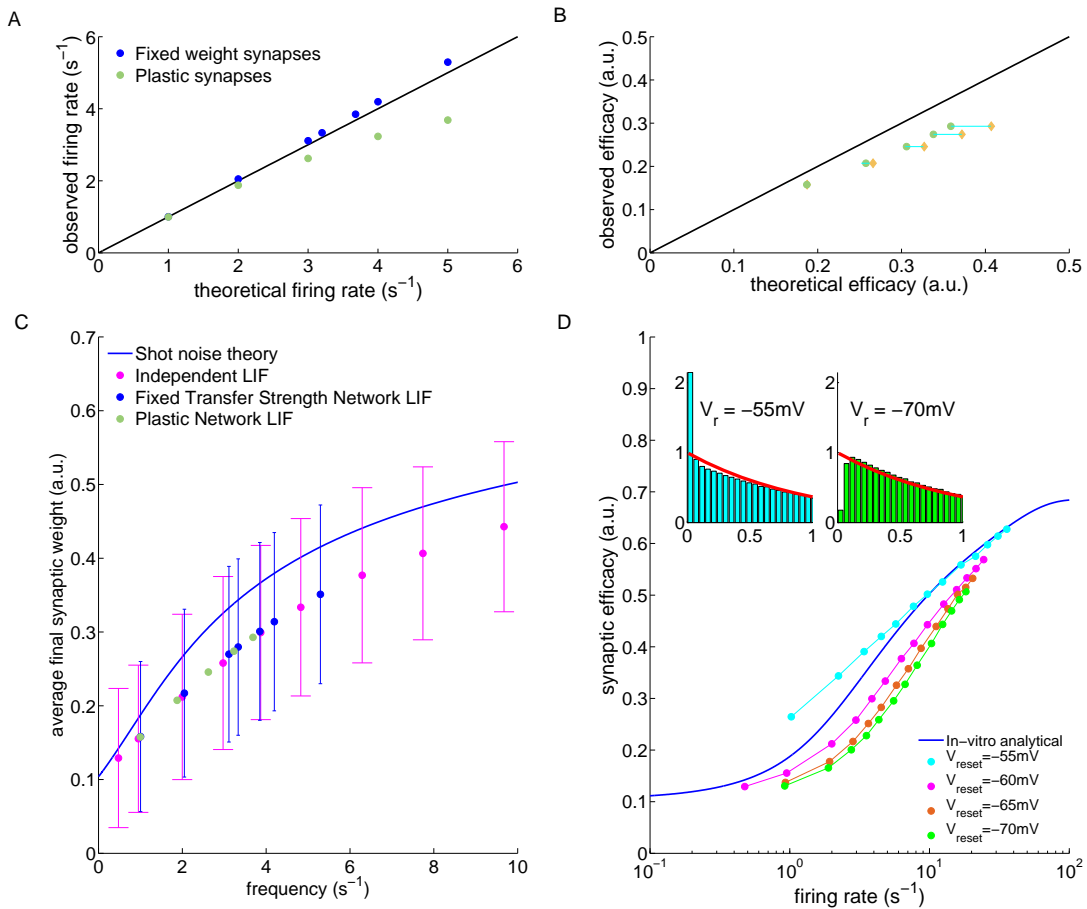


Figure 3.7. See caption overleaf.

Figure 3.7. Steady-state behaviour of a recurrent network with plastic synapses between excitatory neurons. (A) Firing rate mean-field predictions compared with network simulation results. The mean-field theory predicted firing rate is higher (black line) than the actual firing rate of the excitatory neurons (green dots) in the recurrent network of 8000 exc. and 2000 inh. LIF neurons. Network simulation with fixed synapses yield a good match with the theory (blue dots). (B) Average synaptic weight prediction compared with asymptotic average synaptic weights in the network simulation. The observed average synaptic efficacy of excitatory to excitatory connections is smaller (mustard dots) than the theoretical prediction (black line). Even when using the asymptotic firing rate of the network in the calculations (green dots), the average synaptic efficacy is overestimated by the theory. (C) Mean and standard deviation of synaptic weights vs. firing rate for independent LIF neurons (magenta), networked LIF neurons (green) and LIF neurons in a network in which actual weights are held constant but we examine how their efficacy would have evolved in the presence of observed firing (blue dots). Asymptotic synaptic weights for LIF neurons underestimate the efficacy predicted by the theory (blue line). (D) Average synaptic weight vs. firing rate for independent LIFs with different reset potentials. The analytical prediction of the asymptotic synaptic weight based on Poisson firing is shown by the blue line (same as in C). The reset potential in the LIF model, V_{reset} , has a marked influence on the observed average synaptic efficacy. Depolarised/hyperpolarised reset potentials (e.g. $-55/-70$ mV, cyan/green dots) lead to an over/under-representation of short ISIs (left/right inset) compared to Poisson neurons (red line in insets). ISI histograms in inset are shown for LIF neurons firing at $1/s$.

network model.

The fixed point of the network can be determined analytically by solving a set of three self-consistent equations for the excitatory and inhibitory mean rates as well as for the mean excitatory-to-excitatory (E→E) synaptic efficacy (see Section 3.2.8). Two of these equations give the stationary firing rates of excitatory and inhibitory populations (Eqs. (3.29)-(3.30)), as a function of the mean E→E synaptic efficacy (Amit and Brunel 1997b; Brunel 2000). The third equation gives the mean E→E synaptic efficacy as a function of the firing rate of the excitatory population, assuming Poisson firing statistics of neurons (Eq. (3.15)). Starting from the analytically determined initial conditions, the recurrent network converges to a steady-state of constant average firing rates of all neurons in the network, and constant average synaptic efficacy of the plastic connections. Figure 3.7,A shows how the firing rates

observed in the simulations compare with the analytically predicted firing rates. It shows that at sufficiently low rates, the analytical prediction gives a very good estimate of the observed rates; however, for rates above 3Hz the observed rates are significantly lower than the analytical prediction. Likewise, the analytical prediction for the mean E→E synaptic efficacy significantly overestimates the observed efficacies (green dots in 3.7,**B**).

What is the source of the difference between theory and simulations in predicting the steady network state? When synapses are fixed in the network at the efficacies predicted by the corresponding firing rate, the analytically predicted network firing rates provide a good approximation of the observed activity (blue dots in 3.7,**A**). This suggests that the underestimation of firing rates and synaptic efficacy emerges from the mapping of firing rates onto synaptic efficacy. Indeed, the average synaptic efficacies are smaller when spikes are generated from independent LIF neurons, compared to Poisson processes (3.7,**C**, compare magenta dots and blue line). Furthermore, independent LIF neurons firing with the same input statistics as neurons in the LIF network generated the same average final synaptic weight as observed in the LIF network (compare red and green dots in Fig. 3.7,**C**). Thus, differences between simulations and theory are due to the different spiking statistics of the LIF model compared to a Poisson process, and not due to correlations that are present in the network.

To investigate further how the spiking statistics of the LIF model and in particular the interspike-interval (ISI) distribution causes the differences seen in Fig. 3.7, we varied the ISI distribution of the LIF neuron by changing the reset potential (V_{reset} , see Section 3.2.6). This change had a strong effect on the average synaptic efficacy (3.7,**D**). A reset potential close to threshold ($V_{\text{reset}} = -55$ mV, $V_{\text{threshold}} = -50$ mV) yields an overrepresentation of short ISI compared to Poisson firing (3.7,**D**, inset) and in turn overestimates the average synaptic efficacy (3.7,**D**; cyan dots). Conversely, more depolarised reset potentials lead to an under-representation of short ISIs with regard to Poisson firing and consequently to an underestimation of the average synaptic efficacy (3.7,**D**; magenta, red and green dots). We use an intermediate value of $V_{\text{reset}} = -60$ mV in the following network investigations.

To conclude this section, the calcium-based synaptic plasticity rule does not affect the stability of the asynchronous irregular state in a large recurrent network of LIF neurons (see Fig. 3.4). Since LIF neurons in the network exhibit ISI distributions which deviate from those of Poisson neurons, the accuracy of our calculation of the average synaptic efficacy which is based on Poisson firing decreases with increasing firing rates up to a certain point. At high firing rate, calcium remains above the plasticity thresholds most of the time and the fraction of time spent above the thresholds

converges to one, irrespective of the underlying neuron model.

3.3.4 Memory decay in a recurrent network of LIF neurons

We now examine the decay of a memory trace in a network for the *in vitro* and the *in vivo* parameter set. We initialise all excitatory-to-excitatory synaptic weights at their theoretically predicted asymptotic weights, except for a randomly selected subset of 5% which are set to a weight of 1. With the *in vitro* parameter set, the potentiated synapses decay relatively quickly to their asymptotic value (Fig. 3.8,**B**). The time course of the average decay can be described by a single exponential function and the decay time constant is well approximated by the time constant, τ_{eff} , of synaptic decay from the truncated OU process (see Eq. (3.16); Fig. 3.8,**C**). This means that the average dynamics of synaptic decay in the network is equivalent to synapses driven by independent pre- and postsynaptic Poisson neurons firing at the same rate as the excitatory neurons in the network (compare to Fig. 3.5,**B**). The addition of the double-well potential does not change the decay time constant for the *in vitro* parameter set, as for a single synapse driven by independent pre- and postsynaptic Poisson firing (Fig. 3.8,**C** orange stars; compare with Fig. 3.6,**B**). The lack of short ISIs in LIFs compared to independent Poisson neurons leads to a small increase in observed decay times in the network as compared with the OU theory. This can be seen in Fig. 3.8,**C** where the simulation points all lie slightly to the right of the theory line, this is a stable reproducible effect for different random seeds and hence network conditions.

In contrast, when using the *in vivo* parameter set with the double-well potential, we observe that the potentiated synapses get locked in the UP state for the duration of the network simulation with an excitatory neuron firing rate of 1/s (Fig. 3.9,**B**). None of the synapses in the potentiated subset crosses the unstable fixed point and converges to the DOWN state during a network simulation of 120 min, neither does the reverse transition occur. We expect that the escape from the UP state will be predicted by Kramers escape rate (Eq. (3.21)) which correctly accounted for escape dynamics of an isolated synapses driven by independent pre- and postsynaptic Poisson processes (Fig. 3.6,**B**). There, the decay time constant for a firing rate of 1/s is on the order of a month, a time scale that cannot be reached by our network simulation.

Hence, as in case of independent Poisson neurons, the combination of a double-well potential with the *in vivo* parameter set leads to several orders of magnitude longer memory time constants, compared to the *in vitro* parameter set and a flat potential.

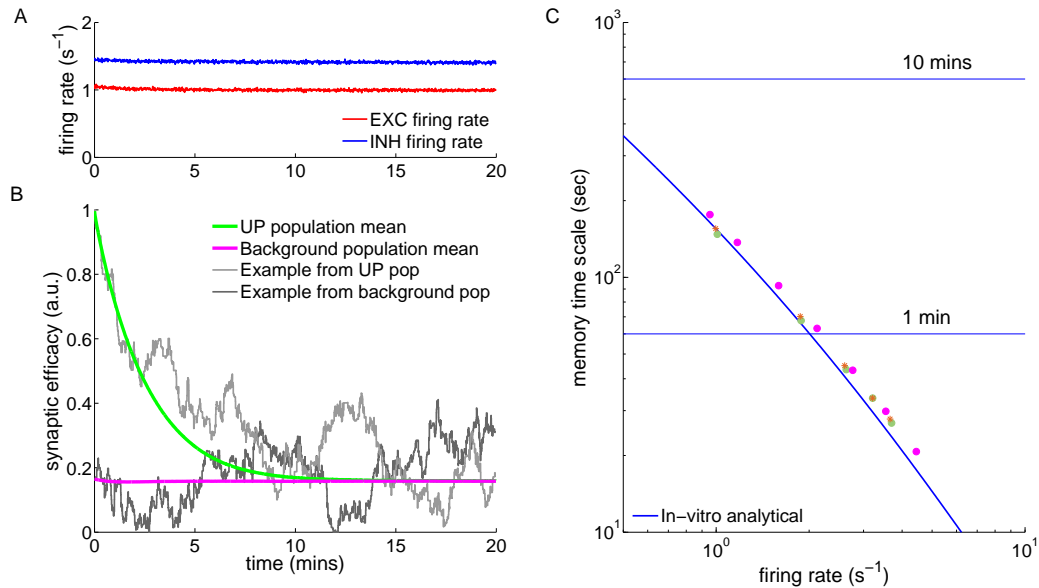


Figure 3.8. Memory decay for a subset of potentiated synapses in a recurrent network with the *in vitro* parameter set. (A) Temporal evolution of the average excitatory (red) and inhibitory (blue) firing rate. A network of 10,000 LIF neurons is initialised at the theoretically predicted steady-state and simulated for 20 min real time. (B) Temporal dynamics of synaptic efficacies in the network. The majority of synapses are initialised to the theoretically predicted asymptotic synaptic efficacy (mean: magenta; single synapse example: dark gray). A randomly selected subset of 5% are set to 1 at the beginning of the simulation (mean: green; single synapse example: light gray). (C) The exponential decay time constant of the potentiated synapses. The value obtained from fitting a single exponential to the mean decay (green dots) is well approximated by the analytically calculated decay time constant from the OU process (Eq. (3.16)). Introduction of a double-well potential does not modify the memory time constant for the *in vitro* parameter set (orange stars). The slight deviation of the decay time constants with respect to the OU theory, that is, the network decay time constants are slower, are due to the LIF firing statistics as can be seen from the comparison with independent LIF neurons (magenta dots).

3.3.5 Memory induction in a recurrent network of LIF neurons

Finally, we present an example of memory induction in a recurrent network. We initialise the network using the *in vitro* synaptic parameter set (Tab. 3.1) with a




Figure 3.9. Memory decay for a subset of potentiated synapses in a recurrent network with the *in vivo* parameter set and double-well potential. (A) Temporal evolution of the average excitatory (red) and inhibitory (blue) firing rate. A network of 10,000 LIF neurons is initialised at the theoretically predicted steady-state and simulated for 120 min real time. (B) Temporal dynamics of synaptic efficacies in the network. The average dynamics of the 95% initialised in the DOWN state (blue) and the 5% initialised in the UP state (red) is shown. The shaded grey region represents the range of values visited by synapses in the UP and in the DOWN state populations, indicating that no transition occurs.

mean excitatory firing rate of 1/sec (Fig. 3.10,B). The simulation is allowed to proceed for 200 secs, during which time the excitatory synapses are allowed to evolve dynamically following the synaptic efficacy rule (Eq. (3.1)) with a flat potential. A subset of 100 excitatory neurons are then selectively stimulated such that they fire at approx. 50/sec for 3 seconds (Fig. 3.10,A). This leads to a strong potentiation in the synapses joining the stimulated neurons (Fig. 3.10,C, magenta) but no change in the synapses not directly connected to stimulated neurons (Fig. 3.10,C, cyan). Following potentiation the directly stimulated synapses decay exponentially back to their 1/sec stationary values, following the time scale predicted for the *in vitro* parameter set

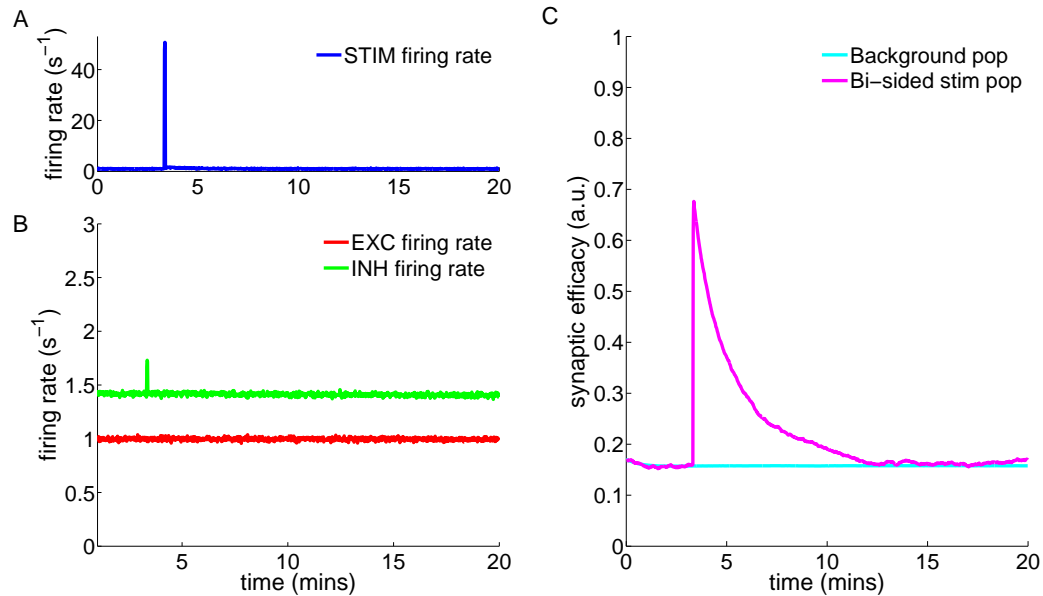


Figure 3.10. Memory induction for a subset of synapses in a recurrent network with the *in vitro* parameter set. A network of 10,000 LIF neurons is initialised at the theoretically predicted steady-state and simulated for 20 min real time. (A) Temporal evolution of the average firing rate for a randomly selected population of 100 excitatory neurons (STIM pop, blue). These neurons receive a 3-sec increase in external stimulation after 200 secs, leading to a burst of approx. 50/sec activity. (B) Temporal evolution of the average firing rate of excitatory neurons (excluding the STIM pop., red) and inhibitory neurons (green). (C) Temporal evolution of mean synaptic efficacies for synapses, incorporating the flat potential, which have no connection with the high frequency stimulation population (cyan) and synapses which receive both pre- and postsynaptic stimulation from the high frequency stimulation population (magenta). Stimulation at 50/sec increases mean synaptic efficacy to $\bar{\rho}(50)$ before a return to 1/sec background firing leads to an exponential decay to the 1/sec stationary efficacy with a time scale as predicted for the *in vitro* parameter set.

and previously demonstrated in Section 3.3.4.

In fact, the burst stimulation response can be predicted by the tendency of an Ornstein-Uhlenbeck process to respond with an exponential trend towards a new stationary value. In Fig. 3.11 (blue lines) we show the impulse response to a 3 second burst of the *in vitro* rule with initial values $\rho \in \{0, 0.1, \dots, 1\}$ for a range of

frequencies. At the 50/sec pre- and postsynaptic stimulation used in Fig. 3.10 we clearly expect the tendency to converge towards approx. $\bar{\rho} = 0.7$. The time scale of this response is predicted by the time scales in Figs. 3.5,C & 3.6,B. In future work it will be possible to use this impulse response analysis to predict the changes in synaptic weights to a series of repeated, shorter duration bursts at lower frequencies operating on multiple sub-populations of neurons.

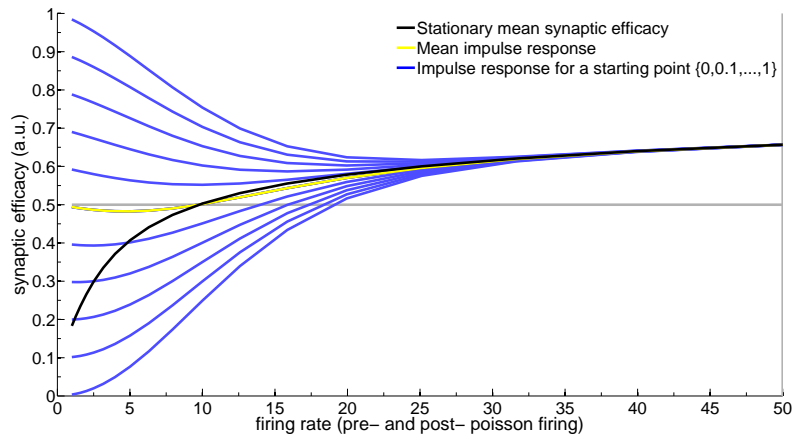


Figure 3.11. Impulse response to a 3 sec high frequency stimulus. Mean synaptic response for an initial value $\rho \in \{0, 0.1, \dots, 1\}$ to a burst of length 3 seconds for frequencies from 1 to 50/sec pre- and postsynaptic rates (blue lines). The average synaptic response assuming a uniform prior distribution of synaptic efficacies (yellow line). Truncated Ornstein-Uhlenbeck predicted stationary synaptic efficacy, $\bar{\rho}$, at each frequency for the *in vitro* parameter set (black line). At low frequencies the response curves show little deflection, whereas at high frequencies stimulation leads to a decrease in the effective time constant and consequent convergence to the long run synaptic efficacy.

3.4 Discussion

In this chapter, following development of our theoretical and numerical tools, we have focused our study on the stability of synaptic efficacy, in a plastic synapse subjected to background activity of pre- and postsynaptic neurons. We used a calcium-based plasticity model that has been shown to fit experimental data in hippocampal and

neocortical preparations (Graupner and Brunel 2012). The model was investigated numerically, using an event-based implementation of the plasticity rule, as well as analytically, using a diffusion approximation. Thanks to this formalism, we derived scaling laws that describe how memory time scale is related to the firing rates of pre- and postsynaptic neurons. At low firing rates, we find that, when synapses are monostable, synaptic efficacies decay to an equilibrium value with a time scale that depends on the firing rates as a power law, $\tau_{\text{eff}} \sim 1/\nu^k$, where k is the number of simultaneous spikes needed to cross the depression threshold. When synapses are bistable, memory decay is akin to diffusion of a particle out of a potential well, which leads to much stabler memories, with time scales that increase exponentially with the inverse of the firing rates, $\tau_{\text{eff}} \sim \exp(a/\nu)$, at low rates. We showed that these estimates accurately reproduce the results of simulations, both of a synapse connecting two isolated independent Poisson neurons, and of a large network of LIF neurons.

Previous studies have investigated memory maintenance in networks of neurons connected by synapses endowed with standard spike-timing dependent plasticity rules (Billings and Van Rossum 2009). Billings and Van Rossum (2009) demonstrated that the memory time scale depends dramatically on whether the rule is additive or multiplicative. In a multiplicative STDP rule, in which synaptic change depends on the current value of the weight such that synaptic changes decrease when the weights approach the bounds, distributions of weights are unimodal (Billings and Van Rossum 2009; Rubin et al. 2001; Van Rossum et al. 2000) and the memory of synaptic changes decay as $1/\nu^2$, since synaptic changes occur upon coincidence of pre- and postsynaptic spikes in the characteristic time window of the STDP rule. These behaviours are very similar to the behaviour of the calcium-based rule in the flat potential case, in the parameter region in which two spikes are needed to cross the depression threshold. This is due to the fact that the calcium-based rule defined by Eq. (3.1) is multiplicative. In the calcium-based rule however, the exponent describing the memory decay at low rates can be set to an arbitrary integer number, through an appropriate rescaling of the ratio between the amplitude of the calcium transients and the depression threshold. In additive STDP rules, the picture changes dramatically and the synaptic weight distributions become bimodal, with weights attracted either to the lower or upper bounds through a symmetry breaking mechanism (Billings and Van Rossum 2009; Song et al. 2000). In this situation, the memory time scales are much longer, and decay of synapses is similar to diffusion in a double well potential.

Several studies have shown that synaptic bi- or multi-stability can emerge from a number of mechanisms such as positive feedback loops in extensive protein signaling

cascades (Bhalla and Iyengar 1999), autophosphorylation of CaMKII (Castellani et al. 2009; Delord et al. 2007; Hayer and Bhalla 2005; Pi and Lisman 2008; Zhabotinsky 2000), self-sustained regulation of translation (Aslam et al. 2009), or modulation of receptor trafficking rates (Shouval 2005). Such mechanisms of bistability are effectively implemented here in the form of the double well potential. Miller et al. (2005) studied the stability of the up state in a model of the bistable calcium/calmodulin-dependent protein kinase II system with respect to stochastic fluctuations induced by protein turnover (Miller et al. 2005). They show that the CaMKII switch composed of a realistic number of CaMKII proteins is stable for years even in the presence of protein turnover, phosphatase as well as free calcium fluctuations. The transitions induced by background activity investigated here impose an upper limit on memory life-time which is typically lower, indicating that *in vivo* neuronal activity, not protein turnover, will be the limiting factor of memory life-times.

Distributions of synaptic weights have been examined in a number of studies (Barbour et al. 2007; Holmgren et al. 2003; Loewenstein et al. 2011; Markram et al. 1997; Sjöström et al. 2001; Song et al. 2005). In all of these studies, distributions of synaptic weights appear unimodal and skewed, and peak at a low weight. In some cases, the distribution has been shown to be well fitted by a lognormal distribution (Loewenstein et al. 2011; Song et al. 2005). This seems at first sight at odds with the distributions of weights shown in the present paper, which are either a truncated Gaussian in the flat potential case, lacking the fatter tail of the lognormal distribution, or bimodal in the double-well case. However, the model in the flat potential case can be made consistent with the data, by choosing synaptic efficacy variables which are an exponential of the ρ variable, rather than being linearly related to ρ . In this case, synaptic efficacies themselves become exponentiated Ornstein-Uhlenbeck processes, consistent with Loewenstein et al. (2011). The model with a double-well potential could also be made consistent with a unimodal distribution, provided the synaptic up and down states are highly heterogeneous from synapse to synapse. Finally, we should point out that the distributions we observe are asymptotic distributions under a statistically constant distribution of inputs. Synapses *in vivo* are typically subjected to highly non-stationary firing rates of pre and post synaptic neurons. These non-stationarities can also potentially strongly affect distributions of synaptic weights in our model.

A large number of distinct learning rules that capture quantitatively both spike-timing and firing rate effects have been proposed recently (Albers et al. 2013; Clopath et al. 2010; El Boustani et al. 2012; Graupner and Brunel 2012; Kumar and Mehta 2011; Pfister and Gerstner 2006; Yger and Harris 2013). Our rule can be distinguished from most of those rules by the fact that it includes calcium concentration as its

primary dynamic variable, which allows us to extrapolate parameters of the rule from *in vitro* to *in vivo* conditions, as we have explained here. Scaling laws derived here can be expected to hold also in those models: at low rates, the time scales of memory decay are expected to be inversely proportional to the rates to a power equal to the number of spikes needed to provoke plasticity. This power should be equal to 2 for standard STDP rules, triplet rules (Pfister and Gerstner 2006), and calcium-based rules in which 2 spikes are needed to cross the depression threshold (Kumar and Mehta 2011; Shouval et al. 2002); 1 for spike and voltage based rules (Clopath et al. 2010).

In this work, we have made the hypothesis that synaptic weights are altered during background activity, and that one can treat background activity as being essentially uncorrelated with the synaptic connectivity structure. Memory time scales could in principle be further extended by two factors. A first mechanism would be to gate plasticity by specific neuromodulator(s) that are present only during stimulus presentation. This idea is consistent with a growing body of experimental data showing how plasticity is modulated by dopamine (Zhang et al. 2009), acetylcholine (Couey et al. 2007; Seol et al. 2007), noradrenaline (Lin et al. 2003) (see also Pawlak et al. (2010) and references therein). However, we note that the model we have used here is built from *in vitro* plasticity data where these neuromodulators were present at very low concentrations, if at all. Hence, we believe that these neuromodulators are likely to *enhance* plasticity during behaviourally relevant epochs, but that the memory time scales discussed here are likely not to be affected if neuromodulators are not present at high levels during background activity.

A second mechanism that would extend memory time scales would be a scenario in which background activity is in fact strongly correlated with the connectivity structure, and wanders stochastically between network states that are strongly correlated with the states of the network during stimuli presentation. This idea is consistent with a growing experimental literature (Kenet et al. 2003; Luczak et al. 2009; Tsodyks et al. 1999) showing how spontaneous activity is transiently strongly correlated with sensory responses in visual and auditory cortices, and it is also consistent with the ubiquitous supra-Poissonian variability, potentially due to the doubly-stochastic process of combined rate stochasticity and individual neuronal Poisson spike processes, seen in background activity in cortex (Churchland et al. 2010; Litwin-Kumar and Doiron 2012). Recurrence of activity states resembling the network activity during stimulus presentation could refresh existing memory traces and therefore prolong their lifetimes.

We showed here that the low extracellular calcium concentrations *in vivo* could have a strong impact on plasticity. A first prediction of calcium-based rules is that

plasticity seen in standard protocols should be greatly reduced (and even possibly vanish altogether) at physiological calcium concentrations. A second prediction is that induced synaptic changes should be much more stable in the face of ongoing pre- and postsynaptic activity. These results emphasise the need for experimental studies at physiological calcium concentrations $\sim 1.5\text{mM}$ (Silver and Erecińska 1990), unlike most published studies that used concentrations in the range $2 - 3\text{mM}$. Our predictions could be easily tested in slice experiments, by providing background activity at a specified rate after the plasticity-inducing protocol. Our model would predict that in cortical slices, at 2.5mM calcium, induced synaptic changes should disappear on a time scale of minutes, while at 1.5mM calcium, they should be stable on a time scale of ~ 1 hour.

We provided here an event-based update scheme of plastic synapses which greatly accelerates simulations and should strongly facilitate future studies of the dynamics of recurrent networks with plastic calcium-based synapses. On the theoretical front, it would be interesting to extend the theory to non-Poissonian renewal processes (Takacs 1956) such as for leaky integrate-and-fire neurons used here, which would give a better approximation of average synaptic efficacies, especially at higher firing rates. It would also be of great interest to extend our initial work on the induction of memory states to examine how synaptic connectivity is modulated by non-stationary external inputs, particularly its response to patterned bursts of high-frequency activity, and how such changes in connectivity affect in turn the intrinsic dynamics of the network.

Our investigations show that realistic external calcium concentration and multistability of synapses might stabilise memory traces against the potentially deleterious effect of ongoing background activity. These results call for studies of synaptic plasticity induction and maintenance in more realistic conditions and ideally in the intact animal. They provide a glimpse of how plasticity results obtained *in vitro* might translate to the living organism.

Conclusion

In this thesis we have developed and analysed two models of synaptic plasticity which outwardly exhibit similarities in their formulation, but which will eventually be useful in very different contexts and to very different ends. The model of cerebellar plasticity is a huge step forward in potentially explaining the synaptic plasticity rule for this particular synapse. Previous explanations of cerebellar plasticity have focused on the *perceptron* learning rule (Albus 1971; Marr 1969) deriving predictions for the distribution of synaptic weights and memory capacity (Brunel et al. 2004; Clopath and Brunel 2013). This approach has worked well so far and has provided a potential explanation for the presence of a large number of silent synapses at this synaptic connection. Using our model, it is now possible to move beyond this formulation to a precise spike based approach. Indeed, recent work has shown that the complex spike in the Purkinje cell, rather than being all-or-none, as assumed in the perceptron model, has a graded duration and influence on synaptic plasticity (Yang and Lisberger 2014). Our model already implements a mechanism for variable duration calcium spikes via the depolarisation protocol part of the model allowing for easy extension to this situation once data becomes available. As our model relies on individual spikes to implement the plasticity rule it will allow us to explore the precise patterns of activity likely to be present in the cerebellar cortex and their implications for learning both at the single cell level and their influence on the output from the cerebellar cortex.

The neocortical model of synaptic plasticity (Graupner and Brunel 2012) for which we provide tools and analysis will serve a very different role. Many models already provide predictions of the associative learning rule typically found in neocortical synapses. However here we provide tools which allow us to link model features,

such as bistability and calcium influx level, back to practical questions of memory time scale and experimental conditions. This allows us, on the one hand, to explain how the experimentally demonstrated phenomenon of synaptic plasticity can lead to stable memories in the neocortex and, on the other, to extend the study of synaptic plasticity into large-scale networks with predictable responses to correlated external stimuli.

4.1 Cerebellar plasticity model

The cerebellar plasticity model developed in Chapter 2 was biologically inspired and phenomenological. It takes as inputs the pre- and postsynaptic spiking or depolarisation activity. This activity is translated into changes in internal variables based on calcium and NO pathway cascades. By combining the activation level of these cascades with a plasticity threshold we were able to reproduce both burst length and frequency dependent experimental results. By further introduction of a LTP-LTD threshold on the calcium pathway variable we reproduced the clear calcium dependence of the sign of plasticity. For low calcium concentration we saw LTP, whereas for high postsynaptic calcium concentration LTD occurred. Finally, a delay on the presynaptic inputs allows the model to easily fit an optimal delay for LTD and postsynaptic calcium influx seen in experiments.

We furthermore produced two alternative versions of the model, one involving the basic mechanics of purported parallel fibre production of NO via presynaptic NMDAR activation level, the other a more simplified linear model of NO production. Each model was initially fitted by hand, and later using numerical optimisation procedures, to a coherent set of experiments produced under comparable conditions by diverse groups (Bidoret et al. 2009; Safo and Regehr 2008, ‘*in house*’ collaborators). Both models fitted the data excellently. It should be stressed that the model parameters were actually highly constrained by the interactions of the parameters within the models with the experimental plasticity induction protocols. For example, an increase in one time constant would necessitate a decrease in the related influx parameter value (and vice-versa) but this would have significant effects on the tails of the frequency dependent protocols.

The question of which model is better is an interesting one. On all of our measures, χ^2 , AIC and BIC, we find that the NMDAR-based model better explains the data. Combined with experimental results from our colleagues (*under submission*) which demonstrate not only the presence of presynaptic NMDA receptors on the parallel fibre to Purkinje cell synapse but also the influx of calcium at these

receptors (suggesting that they are functional) the model of presynaptic NMDAR-based production of NO appears compelling. However, our models both look at the downstream effects of calcium and NO production jumping over the intermediate mechanisms. There may in fact be other sources of NO, which are equally influenced by presynaptic activity. A glutamate spillover model of NO production (Shin and Linden 2005) could potentially be modelled by our NMDAR-based model, as the method of production will be broadly similar. In such a case, however, the NO would be expected to have to diffuse further and would thus no longer be synapse specific. It would be particularly interesting to fit our model of NMDAR activation and NO production to imaging experiments which could determine the precise degree of NO production upon repeated spiking at a given frequency.

The utility of our cerebellar plasticity model is multi-fold. First and foremost we wanted to bring some order to an experimental field in which different groups disagreed fundamentally on even the most basic protocols for synaptic plasticity. We incorporated three aspects which are not under dispute, (i) high postsynaptic calcium leads to LTD, more moderate calcium probably leads to LTP (Coemans et al. 2004), (ii) NO is involved in both LTP and LTD (Lev-Ram et al. 1995, 2002), (iii) there is an optimal delay following parallel fibre activity before postsynaptic activity can drive LTD (Chen and Thompson 1995). These ideas are not in dispute but the actual stimulation protocols which lead to them are. We found that, at least on our relatively simple model dynamics, we could not reconcile protocols which claimed only a single parallel fibre stimulus was required for LTP versus protocols which required a high-frequency burst of 5 stimuli. The solution to this problem is most likely a combination of experimental factors such as, stimulation distance (Marcaggi and Attwell 2007) and strength (Hartell 1996) and indeed experimental operating temperature (32°C was typical for the burst results vs 22°C for the single pulse experiments). Our models can clearly also fit these protocols, which require only a single stimulus to achieve synaptic change, but would require refitting of the parameters. Also, our NMDAR-based model would be more difficult to reconcile with this context. Initially it might seem that no NO would be produced, as the NMDAR would be activated by the first spike but there would be no follow-up spike to induce presynaptic calcium influx. However, in the case of low experimental temperature this would completely change the dynamics of NMDAR deactivation, vastly prolonging the activation state. In the cases of strong or proximal stimulation we can imagine a greater release of glutamate from neighbouring synapses and a general activation of NMDA receptors leading to greater possibilities for NO production. This would explain the process of NO production but would add little to our understanding of synaptic plasticity for this synapse.

Being able to explain how different experimental contexts might change the model parameters is an important first step to reconciling the experimental results. In addition, for the experimental paradigm for which we have fitted the models we can now make predictions of the outcomes of specific experimental stimulation protocols. This is an extremely attractive result for experimentalists, although we would advise some caution until we know better how well each of the parameters are constrained by the existing fit. Due to the simple dynamics of both models we see a great utility in simulation and analytical studies of cerebellar learning on the scale of both the influence of multiple synaptic inputs onto a single Purkinje cell and also from a group of parallel fibres which contact multiple Purkinje cells. Both of these paradigms are experimentally tractable so predictions could be made either for the distribution of synaptic efficacies or the change in parallel fibre activity for a set of inputs, which could then be compared with experiments. This might allow for a huge leap forward in the theory of cerebellar learning, which would advance us beyond the *perceptron* model to a new theory of cerebellar function and dynamics. This approach would also be less susceptible to minor limitations in the current parameter fit than the predictions for the outcomes of precise stimulation protocols.

4.2 Hippocampal and neocortical plasticity model

The synaptic plasticity model of Graupner and Brunel (2012), analysed in Chapter 3, comprises two large advantages over alternative models, which make it particularly attractive for our use in developing a theory of recurrent cortical activity and synaptic plasticity. The first is that it is an analytically tractable model. Through our modelling of the average calcium activity of the model using shot noise processes we have been able to develop a theory of mean synaptic behaviour under any Poisson activity. The second advantage is that it is based on biological processes, thus we can directly relate our theoretical results to experimental conditions and we can make predictions of the impact of altering conditions such as extracellular calcium concentration (predictions of other manipulations such as blocking of NMDA receptors appear in Graupner and Brunel (2012)).

Much of our work in analysing the Graupner and Brunel (2012) model was in the development of analytical tools and the comparison of their predictions with large-scale simulations. We adapted a shot noise theory to predict the amount of time the calcium concentration would spend above the threshold of plasticity in the model for a given pair of pre- and postsynaptic Poisson processes. We then adapted the Ornstein-Uhlenbeck approximation of the model, first posited in (Graupner and Brunel 2012),

by incorporating a necessary truncation of the process at the bounds of synaptic efficacy. This proved to be a necessary adaptation without which our prediction of mean synaptic efficacy contained large errors. We then examined the tendency of the model to trend exponentially towards its mean value with a given time scale, under Poisson pre- and postsynaptic firing. This was a theoretical implication of the Ornstein-Uhlenbeck approximation which proved to be correct.

When we incorporated a double-well potential into the model we found that, at high frequency, it had little detectable effect on either the stationary synaptic efficacy or the time constant of decay towards the stationary efficacy. However, at low frequencies, the double-well creates a separation of synaptic efficacies between an UP state and a DOWN state. These two states create a massive separation in time scales of decay to the lower stationary value, from an approximately $1/\nu$ process in the absence of a double-well, where ν is the firing rate, to an approximately $\exp(1/\nu)$ process in the presence of the double-well. We developed a theory based on Kramers formula (Kramers 1940), which involved the incorporation of the double-well into the Fokker-Planck equation, to accurately predict this longer time scale.

After developing our analytical tools we applied them to the problem of memory stability under ongoing background cortical activity. Initially, using independent Poisson processes, we showed that memory time scales in a flat potential can be greatly enhanced by the lowering of extracellular calcium concentration from that seen *in vitro* to typical *in vivo* levels. The addition of a double-well potential to the *in vivo* conditions increased the memory time scale by several orders of magnitude. We then moved on to the, more typically used in simulation studies, LIF neuron and tested it in three separate situations. The simplest situation was independent LIF neurons, in which both pre- and postsynaptic processes are driven by implementations of LIF neurons which have independently noisy external drives which cause them to fire at a give frequency. In this case we did not allow the presynaptic spiking process to drive the postsynaptic neuron, hence the name *independent* LIF. This allowed us to examine the use of LIF neurons without the influence of correlations. LIF neurons differ from Poisson processes, if reset to a reset potential lower than the mean synaptic inputs they have an *effective* refractory period during which they are highly unlikely to spike as their membrane voltages are charging back up to their mean voltage value. If reset to a higher reset potential we rather increase the likelihood of a spike in the interval immediately following reset. This divergence from Poisson behaviour has an effect on the plasticity rule in our simulations, greatly reducing the likelihood of a short time scale double spike by any neuron, but happily this is more of an issue at higher frequencies, at frequencies around 1/sec the effects are relatively negligible.

Our next step was to establish stability of the synaptic plasticity rule in a recurrent network of LIF neurons. This proved relatively easy to achieve due to our analytical tools. Many groups have struggled to maintain stable network behaviour upon incorporation of a synaptic plasticity rule, the problem being that an increased firing rate typically leads to LTP which in turn leads to a further elevated firing rate in a run away process (LTD and decreased firing rate lead to a similar problem in the opposite direction). Strategies have been used which involve placing the network in a condition such that LTP induced by Hebbian ordered spikes is less than the LTD induced by anti-Hebbian ordered spikes (Song et al. 2000; Zenke et al. 2013). As seen in our examples, even increasing the synaptic efficacies of 5% of the synapses in our network to their maximal values does not lead to a run-away process. This is in part because of the nonlinearity imposed by the calcium thresholds on plasticity, which ensures stability at low firing rates, and also partly because of the inhibition dominating regime of our network, all of which we were able to analytically predict before running our simulations.

As a control, we implemented a fixed synaptic strengths network and simulated the effects of the activity seen in this network on a ‘shadow’ set of synapses. This allowed us to disentangle the feedback effects of the changes in synaptic efficacy on the network, which subsequently effect further changes in synaptic efficacy, from typical network spiking behaviour’s impact on synaptic plasticity. Perhaps surprisingly, we found that correlations in the network have no visible impact on mean synaptic efficacies. Instead the independent LIF simulations, the networked LIF simulations and the results from the shadow neurons based on a fixed network architecture were indistinguishable in their effects on the synaptic plasticity rule. This particular approach will be very useful in future studies when we will look at clustering in a network, allowing us to disentangle network spiking effects from actual plasticity effects.

Having established the stability of our recurrent network simulations we extended our analysis of memory decay time scales to this context. Similarly to the case of independent Poisson processes we found that we could use our theory to predict the exponential decay of mean synaptic efficacy for a population of synapses towards a stationary value. As with the independent case, we found that the decay time scale is directly dependent on the parameter set used, being much longer for the *in vivo* calcium concentrations than for the *in vitro* case. This was easily demonstrable for the *in vitro* case, where we were able to simulate the full cycle of memory decay to baseline with both a flat and double-well potential. In the *in vivo* case time scale predictions were much longer. At high frequencies the flat potential predictions were matched by the theory. We attempted *in vivo* flat potential simulations at low

frequencies, but while we did observe decay over time we were not able to simulate for long enough to accurately calculate a decay time constant. In any case, it was extremely long and clearly of the order predicted by the theory. Instead we chose to demonstrate memory maintenance in the *in vivo* double-well potential. There we saw that, to the limits of our ability to simulate, synapses which begin life in the UP state appear to behave as potential-locked noisy escape processes. That is they remain in the UP state for exceedingly long periods but will eventually get driven out of this state by noise. The time scale of this escape process while long is probably still the limiting factor of such memories, rather than the process of protein turnover modelled in (Miller et al. 2005).

Finally, we demonstrated a potential method of memory induction via high-frequency stimulation. At elevated firing rates the time scale of convergence of the Ornstein-Uhlenbeck process becomes much smaller, meaning that even a brief period of bursting may lead to potentiation of a synapse. If we combine this with a double-well potential and the *in vivo* parameter set we can quickly see how a memory can switch from a DOWN to an UP upon the occurrence of just one or two bursts of perhaps 500ms in length. In the case of the *in vitro* parameter set such a memory subsequently decays at the theoretically predicted rate of decay. It has not thus far been possible to perform sufficiently long simulations to thoroughly examine the *in vivo* case but our theoretical studies already give us strong predictions. First and foremost, we can expect long term maintenance of the memory. But there are also clustering effects which will lead to, at least temporarily, elevated firing rates in the cluster population (on the order of 2/sec in our tests), leading to subsequent decay of some of the potentiated synapses before the firing rate returns to a frequency in the bistable region of synaptic stability. This is a process which deserves much further study, in particular to determine whether it is an epiphenomenon or whether it in fact describes actual memory processes in the brain.

4.3 Outlook

The two models studied in this work occur in very different systems in the brain. Cerebellar plasticity is most closely associated with behavioural studies, including changes in VOR and skin based receptive fields. Neocortical plasticity is more typically associated with abstract processing, the most experimentally accessible of which being in the visual cortex and the decision making frontal cortex. Both models contain a lot of similarities. Indeed they both have a calcium dependence and both use thresholds on their internal variables in order to enforce a nonlinearity on synaptic

change. This can clearly be mapped to concentrations required in order to trigger postsynaptic cascades. In the case of the Graupner and Brunel (2012) model this leads to direct predictions of the results of manipulations of experimental conditions. For the cerebellar plasticity model the coupling of NO production to extracellular magnesium and calcium concentrations is less well understood so similar predictions are likely to contain a larger degree of error. A further commonality between the two models is their applicability to network predictions. In the case of the Graupner and Brunel (2012) model this means the neocortex while for our cerebellar model we can use it to predict input-output relationships in the cerebellar cortex. An important question in studying synaptic efficacy is how does learning occur? In the cerebellar model LTP occurs as a result of presynaptic bursts whereas LTD occurs when a coincident climbing fibre induced calcium spike occurs. Meanwhile in the Graupner and Brunel model, at low frequencies the coincidence of pre- post- spikes is important, whereas at higher frequencies plasticity outcomes are dominated by the frequency.

Our cerebellar plasticity model accurately resolves the rule for synaptic plasticity in the parallel fibre to Purkinje cell synapse for a homogenous set of experiments. It would be interesting now to extend the model fit to experiments performed under *in vivo* concentrations of extracellular calcium and magnesium. In the Purkinje cell there are multiple sources of postsynaptic calcium. These include both voltage-dependent calcium channels (VDCCs) and internal calcium stores, activated by IP3. These calcium sources also likely operate on different microdomains. It is presently unclear whether the joint contribution of these microdomains to the downstream C variable in our model is a linear combination from each domain or rather the result of saturation effects from particular plasticity protocols. Therefore, while it is clear that the reduction of extracellular calcium concentrations to *in vivo* levels will lead to reduced calcium influx via VDCCs, it is not clear whether this reduction will be linear nor what the effects on intracellular stores will be. Finally, the effects of this reduction in calcium concentration, and indeed of an altered magnesium environment, on the induction of presynaptic NO production via NMDAR based calcium influx is completely unknown. Limited experiments to date have shown that simply reducing the calcium influx value in the model allows us to determine the correct sign of plasticity in this case, but not the magnitude.

Even without discovering the model parameters for the *in vivo* experimental condition, it would be interesting to use the model to explore the effects of different patterned inputs on synaptic plasticity and hence on the postsynaptic cell. Indeed even more interesting would be to move straight to simulating the network level, where we could attempt to implement an adaptive controller using a spiking network model of the cerebellum, including spike-timing based synaptic plasticity. Previous

work on this front has used a simple perceptron based learning rule (Luque et al. 2011). It would be of particular interest now to attempt to implement a more biologically accurate scheme. This would likely give rise to a whole new set of questions about how other aspects of the cerebellum work.

Now that we have a set of analytical tools for the Graupner and Brunel (2012) plasticity rule, what can we do with them? We have developed a suite of analytical tools and clearly demonstrated their applicability to this synaptic plasticity rule, indeed we have used them to make predictions for how long memories might last in the presence of ongoing pre- and postsynaptic activity. These predictions are within the realm of experimental testability. But are there other tasks to which we can now turn our attention? In fact, this is just a beginning. We have already developed mean-field firing rate equations for more than one excitatory population allowing us to predict the impact on firing rate of target stimulation of a sub-population of neurons. We have further coupled these equations to the synaptic plasticity equations. This will allow us to predict the network response to a transiently increased selective input. Any such input will increase the firing rate for a subset of neurons, which should increase the internal network connection strengths within that population. Changed dynamics of clustered networks have previously been reported (Litwin-Kumar and Doiron 2012). To be discovered is whether this will lead to (i) enhancement of working memory states, (ii) learning, leading to faster recall when presented with a similar test stimulus, (iii) synfire chains, triggering overlapping populations into a spatial wave of enhanced firing rates. As the model is not only biologically inspired but actually fitted to electrophysiological experiments this makes it the ideal candidate for relating these macro phenomena back to underlying mechanisms.

In summary, we have contributed a new model of synaptic plasticity of the parallel fibre to Purkinje cell synapse to the literature. We have also provided a suite of analytical and numerical tools for an already existing model of neocortical and hippocampal synaptic plasticity. Both approaches have the benefits of being fundamentally derived from biophysical bases while scaling extremely well to allow for larger scale studies to be performed. We are now well placed to use both of these approaches to explore the dynamics of both the cerebellum and the neocortex under biologically realistic conditions of learning and behaviour.

Simulation parallelisation

The work presented in this thesis often involved the use of extremely long numerical simulations in order to validate the theoretical results. These simulations required the implementation of networks of typically 10,000 leaky integrate-and-fire neurons and 5 million connecting synapses. Such simulations would not have been possible without the ability to update the LIF neuron equations in a massively parallel fashion. In order to do this, two main technologies were available to us. The first, OpenMP (or similar) allows for the easy parallelisation of loops on a multi-core CPU architecture. This method would have allowed for relatively easy implementation, being broadly a serial implementation, in which only certain embarrassingly parallelisable sections are performed in parallel. These sections, surrounded by specific code tags for the compiler, perform normally subsequent loop iterations in parallel on different CPUs or CPU cores. However, even on a high-end workstation this could not have given more than a 7-8 fold speed increase of those sections over a purely serial implementation. The alternative technology was the use of a General Purpose GPU (GPGPU), capable of performing up to 256 operations completely in parallel, for the processing. This is the option we decided to pursue as this technique was likely not only to be faster than the OpenMP approach but also to scale well to larger networks.

Writing code for GPGPU gives two main choices of technology. Either one can write in a coding language specific to the manufacturer of the GPU (CUDA in the case of NVIDIA) or one can write in a device agnostic language (OpenCL). There are two further emerging technologies, OpenACC and OpenAMP, but their arrival was too recent to be considered for our work. Device specific language is typically faster than the device agnostic language, in fact the device agnostic language is usually

compiled into the device specific language upon program compilation or at run time. However, we chose to use OpenCL as it allowed us to use our code on devices from any manufacturer (NVIDIA, ATI, etc.). In addition, OpenCL is retargetable to CPU architecture allowing for full parallelisation of operation using CPU cores instead of GPU cores as the basic compute units. Finally, we are well placed for the emerging Intel technology, MIC, which co-locates a GPU chip on a CPU chip. We fully expect this to give enormous speed gains over our current simulations as access to memory will be much improved when compared with access from a dedicated GPU controller.

Speed was clearly our main aim in parallelising our simulations. Indeed, the programming effort involved in parallelisation, particularly when using the technology OpenCL, is not negligible. However, we could not have contemplated running the long network simulations in Chapter 3 without these speed advances. The primary speed barrier when running our simulations on GPU was the number of memory streams which we needed to write to and from the device on each time step. As we used a particularly small time-step (10^{-2} ms) in our simulations we required an enormous number of read/writes in order to simulate a few thousand seconds of real-time. We thus condensed memory writes wherever possible and used a random number generator which did not require any state variables (see Appendix B), in order to limit our streams. One technique which proved particularly useful was the use of ‘pinned’ memory on the GPU device, this is memory which cannot be swapped out to slower memory architectures. As the state space of a set of 10,000 LIF neurons is not particularly large we were able to store it all in pinned memory. Of course, with new technology nothing ever runs entirely smoothly, we were unfortunate to be the discoverers of a major bug in the NVIDIA implementation of OpenCL which leads to memory leaks; a work around was eventually found, but this architecture bug has not yet been repaired. In our work we concentrated on simulating LIF neurons. Ironically, due to our efforts in parallelisation, we could have simulated much more complex neuronal models with little or no slow-down. Due to the technology used the greatest delays in our simulation were the transmission of data to and from the GPU device, by comparison the LIF update equations were performed almost instantaneously.

For future work we have a number of strategies which we expect will greatly enhance the already extant speed gains in our simulations. One technique already used in many of the major commodity simulators is the iteration of multiple updates of the neuronal equations before a return to a spike detection and transfer loop. This technique relies on the interpolation of spike times and the ability to infer their result in a reasonable manner on the post-synaptic neurons. Depending on the question under examination this may be an appropriate technique. A more immediately useful

technique would be the parallelisation of the spike transfer process. For a randomly connected network, it appears that the optimum approach for such a method may be to examine each neuron in turn, using n parallel instances of our GPU code where n is the maximal out-degree of any of our neurons in the network. Then the i^{th} parallel instance, if that neuron has just spiked, will transfer a current to the neuron indexed by the i^{th} outgoing synapse from that neuron.

Generation of random numbers in parallel

Generation of random deviates (pseudo-random numbers) is a highly technical discipline which is often under-considered during the implementation of numerical simulations involving stochastic differential equations. In this work we have made considerable use of random numbers, the extremal case being our simulations of 10,000 LIF neurons in a recurrently connected random network for 2 hours using a time step of 10^{-2} ms. In this example, each LIF neuron required a gaussian distributed deviate on every single time step, in addition uniform deviates were required for the initial setup of network connectivity and further gaussian deviates were required upon each update of each synapse, a frequency dependent event. Counting only the deviates required for the LIF neuron update formula, this means that we needed greater than $7.2(10^{12})$ deviates for a single run of this simulation.

When considering a (pseudo-)random number generator two factors are considered particularly important. The first, the *randomness* of the deviates produced can be tested using a number of statistical libraries the most famous of which are the DIEHARD test suite by George Marsaglia (Marsaglia 1996) and a more recent test suite called TestU01 (L'Ecuyer and Simard 2007). Both of these test suites take a random number generator, or a list of numbers generated by one, and run a batch of statistical tests in order to attempt to detect patterns. In simple, non-recurrent, biological simulations it may not be so important for a generator to pass all of the tests, but as we expand our simulations to examine recurrent effects which may involve correlations as an actual feature of the phenomenon the randomness of the generator becomes much more important. The second major factor in a random

number generator is its average *period length*. This is the number of calls you can make to the generator before you see the same deviate value as one which you have previously seen and, at the same time, the generator has the same internal state variables; meaning that from this point onwards you will produce a deterministic looped series of random deviates. The gold-standard of random number generation is that a function be ‘cryptographically secure’, meaning that it is impossible to guess the internal state variables of a generator (and hence the next deviate) even from an excessively long series of random deviates produced by the generator. This standard is not applicable to computational biology however, it is not important to us whether one can predict the next deviate or not, only that there be no correlations or patterns in the output.

For non-scientific applications the POSIX programming standard (for C, etc.) provides for implementation of the function *rand()* which generates weakly random pseudo-random numbers. However, the definition of this function in the standard was never considered from the point-of-view of producing strongly random numbers and consequently should be avoided at all costs. It appears quite common amongst the scientific community to privilege the random number generators found in the book *Numerical Recipes in C* (Press et al. 1992), no doubt because of direct access to the source code. There are four random number generators published in this book, each of which has its own respective strengths and shortcomings. Within the cryptographic community criticism of these generators rests mainly on the view that the authors of the book knew little about random number generation and that they overestimate the period lengths of their generators. In any case, they claim a period length of 10^8 for their *ran1()* generator (the one most typically used) and this is clearly too short for our purposes. Their longest period length generator *ran2()* is claimed to have an average period length of $2.3(10^{18})$ (L’Ecuyer and Simard 2007) but it runs particularly slowly and may be susceptible to initial conditions. At present, the best option for serially coded simulations is the Mersenne Twister algorithm (Matsumoto and Nishimura 1998), which can be implemented in as little as 5 lines of C code and has a period length of order 10^{40} . Furthermore, it is implemented in libraries such as the GNU scientific library (GSL) (Galassi 2009) which means it can be used without risk of programming error.

Due to the highly parallelised nature of many of the problems solved in this work we decided to also try to generate our random deviates in parallel. None of the random generators discussed thus far were found suitable for parallelisation. Unfortunately they each either rely strongly on certain constants, which would be the same for each of the parallel generators and hence the streams of random deviates would likewise be identical, or in the case of Mersenne Twister, they have too many internal

state variables making the maintenance of upwards of 10,000 instances unwieldy and involving massive overheads when transferring these state variables to and from the GPU between time steps. We thus, initially, adopted the KISS generator by George Marsaglia (Marsaglia and Zaman 1993) and implemented a parallel version of it (our own invention). This generator has four state variables, which need to be maintained between calls to the generator. These are transferred to and from the GPU on each time step in our implementation. We used a different seed with each instance of the generator (a single seed and then an increment) and observed excellent results. We performed basic statistical tests on the output of the generator and could not detect correlations, despite exceedingly long sequences of random deviates (of order 10^{12}). We have not however, as yet, submitted the output to the full battery of tests in TestU01. No statistical holes have ever been found in the serial version of this generator, however it is sensitive to its internal constants (which are the same across all instances in our implementation) (Rose 2011); in the initial posting of the generator to an online discussion board George Marsaglia inadvertently swapped two of the constants, leading to significantly substandard results. In our final implementation we moved to the Random123 generator (Salmon et al. 2011) a recently discovered generator which belongs to a class of generators called "counter-based random number generators." The concept is similar to that of a cryptographic one-way hash function. In this case, the output of the function cannot be predicted by the result of the function on an input value with an arbitrary numerical relationship to the new input (i.e. even if you know the output of the function on a whole lot of other inputs, you cannot predict what it will do with the current, previously unseen, input.) The method is stateless and instead uses an index to generate a new random number, in our case we use the time step as one of our indices and the LIF ID as the second index (up to four independent indices are possible). This generator has passed all of the TestU01 tests and due to its statelessness is much faster where state must be transferred to and from GPU memory.

Once a random deviate from a uniform (0,1) distribution can be generated it is possible to use a number of techniques to transform this deviate into any desired distribution. Within our LIF update code we have tended to use a Box-Muller transform to generate gaussian deviates. This transform requires two uniform (0,1] deviates in order to generate two standard normal (0,1) deviates. The standard normal distribution can then be transformed to any arbitrary gaussian distribution by multiplication by the standard deviation. At other points, for more arbitrary distributions, we have used the acceptance-rejection and the inverse transform sampling methods.

Shot noise

A shot noise is a superposition of impulses occurring at times t_0, t_1, \dots , where the inter-impulse intervals, $t_i - t_{i-1}$, form an exponential distribution, and hence the shot noise events follow a Poisson event process. The impulses have a shape, $F(t)$, such that

$$I(t) = \sum_i F(t - t_i). \quad (\text{C.1})$$

For the case where the shapes are taken from a family of shapes, $F(a, t)$, parameterised by a , then

$$I(t) = \sum_i F(a, t - t_i). \quad (\text{C.2})$$

In keeping with the notation of Gilbert and Pollak (1960), and earlier Rice (1944), we define the amplitude (cumulative) distribution function

$$Q(I) = \text{Prob}[I(t) \leq I] \quad (\text{C.3})$$

and the probability density function

$$P(I) = Q'(I). \quad (\text{C.4})$$

Gilbert and Pollak (1960) obtain the following formula, from which all of their work is derived

$$IQ(I) = \int_{-\infty}^I Q(x) dx + n \int_{-\infty}^{\infty} Q[I - F(t)]F(t) dt \quad (\text{C.5})$$

Note: one alternative description of a shot noise process is of a poisson process passed through a linear filter (thus giving it its additive shape).

C.1 Mapping calcium dynamics to shot noise

In their work Gilbert and Pollak (1960) present a solution to the probability density function for the case of a shot noise event whose shape follows

$$F(t) = \begin{cases} 0, & t < 0 \\ \exp(-t), & t \geq 0 \end{cases} \quad (\text{C.6})$$

where t is time. In this case, for a Poisson process with parameter n Gilbert and Pollak (1960) obtain

$$IP'(I) - (n - 1)P(I) = -nP(I - 1). \quad (\text{C.7})$$

Eq. (C.6) can be mapped trivially to our calcium model, if the influx parameter $C_{\text{pre}} = C_{\text{post}} = 1$ and the time constant of decay is 1 then $F(t)$ follows the shape of a single spike event leading to a calcium influx followed by an exponential decay. Furthermore, for independent pre and postsynaptic Poisson processes, if the presynaptic process has parameter ν_{pre} and the postsynaptic process parameter ν_{post} , we can sum the pre- and postsynaptic processes to obtain a further Poisson process whose parameter is $n = \nu_{\text{pre}} + \nu_{\text{post}}$. This then allows us to solve Eq. (C.7) for $P(I)$ where I is the calcium concentration in our model, for a combined Poisson firing rate n .

In order to modify Eq. (C.7) in order to account for calcium amplitudes not equal to 1, that is, in order to obtain a solution for shot noise of the form

$$F(t) = \begin{cases} 0, & t < 0 \\ a \exp(-t), & t \geq 0 \end{cases} \quad (\text{C.8})$$

we resort to a master equation approach. If $I(t)$ is the concentration of calcium at time t then

$$I(t + dt) = \begin{cases} I(t) + a, & \text{with probability } ndt \\ I(t)(1 - dt), & \text{with probability } (1 - ndt) \end{cases} \quad (\text{C.9})$$

Assuming dt is small such that maximum one spike occurs within a single time interval, then the probability that we have calcium concentration I at time $t + dt$ is

$$P(I, t + dt) = ndtP(I - a, t) + \left(\frac{1 - ndt}{1 - dt}\right) P\left(\frac{I}{1 - dt}, t\right). \quad (\text{C.10})$$

As $dt \rightarrow \infty$ we can perform the expansion

$$P(I, t) + dt \frac{\partial P}{\partial t} = ndtP(I - a, t) + (1 - ndt)(1 + dt) \left(P(I, t) + Idt \frac{\partial P}{\partial I} \right). \quad (\text{C.11})$$

And by multiplication of the final term we get

$$P(I, t) + dt \frac{\partial P}{\partial t} = ndtP(I - a, t) + P(I, t) - ndtP(I, t) + dtP(I, t) + Idt \frac{\partial P}{\partial I}. \quad (\text{C.12})$$

The terms $P(I, t)$ cancel left and right, allowing us to divide across by dt obtaining

$$\frac{\partial P}{\partial t} = nP(I - a, t) - nP(I, t) + P(I, t) + I \frac{\partial P}{\partial I}. \quad (\text{C.13})$$

At stationary behaviour $\frac{\partial P}{\partial t} = 0$ giving us

$$I \frac{\partial P}{\partial I} = -nP(I - a, t) + nP(I, t) - P(I, t). \quad (\text{C.14})$$

We now trivially rewrite this in our previous notation as

$$IP'(I) - (n - 1)P(I) = -nP(I - a). \quad (\text{C.15})$$

We can reintroduce separate pre- and postsynaptic firing rates, letting $n = \nu_{\text{pre}} + \nu_{\text{post}}$ gives us

$$IP'(I) - (\nu_{\text{pre}} + \nu_{\text{post}} - 1)P(I) = -(\nu_{\text{pre}} + \nu_{\text{post}})P(I - a) \quad (\text{C.16})$$

where $a = C_{\text{pre}} = C_{\text{post}}$, the pre- and postsynaptic calcium influx parameters.

Using our understanding of the master equation it is similarly possible to provide a description of the case $C_{\text{pre}} \neq C_{\text{post}}$ by separating the pre- and postsynaptic ‘jump’ values in the final term $P(I - a)$ into two separate ‘jumps’. Clearly in this case, it is also necessary to associate the separate firing rates, ν_i , with the separate influx processes, C_i , giving us

$$IP'(I) - (\nu_{\text{pre}} + \nu_{\text{post}} - 1)P(I) = -\nu_{\text{pre}}P(I - C_{\text{pre}}) - \nu_{\text{post}}P(I - C_{\text{post}}). \quad (\text{C.17})$$

Now, the final development which will aid our analysis of the calcium concentration under pre- and postsynaptic Poisson firing in our model is the incorporation of the time constant of calcium decay into our solution. In Eq. (C.8) the decay constant is 1, however in our analysis the calcium decay constant is typically of order $\tau_{\text{Ca}} = 20\text{ms}$. We obtain this value by rescaling time via the Poisson frequency parameter. In fact, n defines n events per unit time, if we rescale time via $t' = t/\tau_{\text{Ca}}$ then we must rescale the number of events via $n' = n/t' = n\tau_{\text{Ca}}$. So for a shot noise process which decays with time constant τ_{Ca} we have

$$IP'(I) - (\nu_{\text{pre}}\tau_{\text{Ca}} + \nu_{\text{post}}\tau_{\text{Ca}} - 1)P(I) = -\nu_{\text{pre}}\tau_{\text{Ca}}P(I - C_{\text{pre}}) - \nu_{\text{post}}\tau_{\text{Ca}}P(I - C_{\text{post}}). \quad (\text{C.18})$$

C.2 Solving the shot noise equation

In the main text we present the solutions to Eq. (C.18) for certain special (Eqs. (3.6)-(3.9)) and more general (Eqs. (3.11)-(3.13)) cases. Solution follows the iterative scheme

$$\text{if } P'(I) = \phi(I) \quad (\text{C.19})$$

$$\text{then } P(I) = P(a) + \int_a^I \phi(x)dx, \text{ for } a < I. \quad (\text{C.20})$$

For simplicity let us assume that $C_{\text{pre}} < C_{\text{post}}$ and, as this is the case in the main text, that $\nu_{\text{pre}} = \nu_{\text{post}} = \nu$. We must now integrate the function

$$\phi(I) = \frac{(2\nu\tau_{Ca} - 1)P(I) - \nu\tau_{Ca}P(I - C_{\text{pre}}) - \nu\tau_{Ca}P(I - C_{\text{post}})}{I} \quad (\text{C.21})$$

Then in the first interval, $I \in [0, C_{\text{pre}}]$, we see that the terms $P(I - C_{\text{pre}})$ and $P(I - C_{\text{post}})$ both evaluate to 0. This leaves

$$P'(I) = \frac{(2\nu\tau_{Ca} - 1)P(I)}{I} \quad (\text{C.22})$$

which gives us

$$P(I) = AI^{2\nu\tau_{Ca}-1} \quad (\text{C.23})$$

where A is a constant of integration, which has a closed form solution only in special cases such as that presented in Eqs. (3.6)-(3.9). In all other cases we normalise the are under $P(I)$ to 1 numerically at the end of the calculation of the other intervals.

While it is possible to extend this analytical approach to further intervals, it is practical only for certain parameter values. It is much more fruitful to pursue a numerical integration scheme for Eq. (C.18). We do this using an exponential Euler method of integration. That is for

$$P'(I) = a(I)P(I) + b(I) \quad (\text{C.24})$$

we obtain a solution

$$P(I) = \exp\left\{\int_{I_0}^I a(y)dy\right\}\left(P(I_0) + \int_{I_0}^I \exp\left\{-\int_{I_0}^y a(z)dz\right\}b(y)dy\right) \quad (\text{C.25})$$

when $I - I_0$ is small. In the case of Eq. C.18 we define

$$a(I) = \frac{\nu_{\text{pre}} + \nu_{\text{post}} - 1}{I} \quad (\text{C.26})$$

$$b(I) = -\frac{\nu_{\text{pre}}P(I - C_{\text{pre}})}{I} - \frac{\nu_{\text{post}}P(I - C_{\text{post}})}{I} \quad (\text{C.27})$$

This means that

$$\exp\left\{\int_{I_0}^I a(y)dy\right\} = \left(\frac{I}{I_0}\right)^{(\nu_{\text{pre}}+\nu_{\text{post}}-1)} \quad (\text{C.28})$$

$$\exp\left\{-\int_{I_0}^y a(z)dz\right\} = \left(\frac{I_0}{y}\right)^{(\nu_{\text{pre}}+\nu_{\text{post}}-1)} \quad (\text{C.29})$$

giving us

$$P(I) = \left(\frac{I}{I_0}\right)^{(\nu_{\text{pre}}+\nu_{\text{post}}-1)} \left(P(I_0) + \int_{I_0}^I \frac{I_0^{(\nu_{\text{pre}}+\nu_{\text{post}}-1)}}{y^{(\nu_{\text{pre}}+\nu_{\text{post}})}} \left(-\nu_{\text{pre}}P(y - C_{\text{pre}}) - \nu_{\text{post}}P(y - C_{\text{post}}) \right) dy \right) \quad (\text{C.30})$$

We find that it is typically possible to numerically obtain the PDF, $P(I)$, by (i) setting $P(I < 0) = 0$, (ii) filling in the first interval, up to C_{pre} , using Eq. (C.23), (iii) by subsequent iteration of Eq. (C.30) up to a sufficiently high calcium concentration that the probability density is numerically indistinguishable from zero. We then numerically normalise the area under the curve to 1.

Bibliography

- Abramowitz, M. and I. A. Stegun (1970). *Tables of mathematical functions*. Dover Publications, NY (cit. on p. 75).
- Akaike, H. (Dec. 1974). “A new look at the statistical model identification”. In: *IEEE Transactions on Automatic Control* 19.6, pp. 716–723 (cit. on p. 49).
- Albers, C., J. T. Schmiedt, and K. R. Pawelzik (2013). “Theta-specific susceptibility in a model of adaptive synaptic plasticity”. In: *Front Comput Neurosci* 7, p. 170 (cit. on pp. 67, 103).
- Albus, James S. (Feb. 1971). “A theory of cerebellar function”. In: *Mathematical Biosciences* 10.1–2, pp. 25–61 (cit. on pp. 20, 21, 28, 64, 107).
- Amit, D. J and N. Brunel (1997a). “Dynamics of a recurrent network of spiking neurons before and following learning”. In: *Network: Computation in Neural Systems* 8.4, 373–404 (cit. on p. 85).
- (1997b). “Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex.” In: *Cerebral Cortex* 7.3, p. 237 (cit. on pp. 84, 95).
- Amit, D. J. and M. V. Tsodyks (1991). “Quantitative study of attractor neural network retrieving at low spike rates I: Substrate – spikes, rates and neuronal gain”. In: *Network* 2, pp. 259–274 (cit. on p. 84).
- Amit, D. J., N. Brunel, and M. V. Tsodyks (1994). “Correlations of cortical Hebbian reverberations: experiment vs theory”. In: *J. Neurosci.* 14, pp. 6435–6445 (cit. on p. 67).
- Amit, Daniel J. and Stefano Fusi (Sept. 1994). “Learning in Neural Networks with Material Synapses”. In: *Neural Computation* 6.5, pp. 957–982 (cit. on pp. 14, 67).

- Amit, Daniel J., Hanoch Gutfreund, and H. Sompolinsky (Sept. 1985). “Storing Infinite Numbers of Patterns in a Spin-Glass Model of Neural Networks”. In: *Physical Review Letters* 55.14, pp. 1530–1533 (cit. on p. 13).
- Aslam, Naveed, Yoshi Kubota, David Wells, and Harel Z. Shouval (2009). “Translational switch for long-term maintenance of synaptic plasticity.” In: *Mol Syst Biol* 5, p. 284 (cit. on p. 103).
- Barbour, Boris, Nicolas Brunel, Vincent Hakim, and Jean-Pierre Nadal (Dec. 2007). “What can we learn from synaptic weight distributions?” In: *Trends in Neurosciences* 30.12, pp. 622–629 (cit. on p. 103).
- Bell, C. and G. von der Emde (Oct. 1995). “Electric organ corollary discharge pathways in mormyrid fish”. In: *Journal of Comparative Physiology A* 177.4, pp. 463–479 (cit. on p. 31).
- Belmeguenai, Amor and Christian Hansel (Nov. 2005). “A role for protein phosphatases 1, 2A, and 2B in cerebellar long-term potentiation”. In: *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 25.46. PMID: 16291950, pp. 10768–10772 (cit. on p. 27).
- Bhalla, Upinder S. and Ravi Iyengar (Jan. 1999). “Emergent Properties of Networks of Biological Signaling Pathways”. In: *Science* 283.5400, pp. 381–387 (cit. on pp. 28, 67, 103).
- Bi, G Q and M M Poo (Dec. 1998). “Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type”. In: *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 18.24, pp. 10464–10472 (cit. on pp. 7, 8, 17, 65, 67, 72).
- Bidoret, Céline, Annick Ayon, Boris Barbour, and Mariano Casado (Aug. 2009). “Presynaptic NR2A-Containing NMDA Receptors Implement a High-Pass Filter Synaptic Plasticity Rule”. In: *Proceedings of the National Academy of Sciences* 106.33, pp. 14126–14131 (cit. on pp. 27, 39, 50, 51, 53, 108).
- Bienenstock, EL, LN Cooper, and PW Munro (Jan. 1982). “Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex”. In: *J. Neurosci.* 2.1, pp. 32–48 (cit. on pp. 12, 13, 16, 35, 60, 67).
- Billings, Guy and Mark C. W Van Rossum (June 2009). “Memory Retention and Spike-Timing-Dependent Plasticity”. In: *Journal of Neurophysiology* 101.6, pp. 2775–2788 (cit. on pp. 14, 15, 102).

- Bliss, T. V. P. and A. R. Gardner-Medwin (July 1973). “Long-lasting potentiation of synaptic transmission in the dentate area of the unanaesthetized rabbit following stimulation of the perforant path”. In: *The Journal of Physiology* 232.2, pp. 357–374 (cit. on p. 6).
- Bliss, T. V. P. and T. Lømo (July 1973). “Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path”. In: *The Journal of Physiology* 232.2, pp. 331–356 (cit. on pp. 5, 6, 10, 13, 17, 67).
- Bliss, Tim VP and Graham L Collingridge (1993). “A synaptic model of memory: long-term potentiation in the hippocampus”. In: *Nature* 361.6407, pp. 31–39 (cit. on p. 11).
- Brunel, Nicolas (2000). “Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons”. In: *Journal of computational neuroscience* 8.3, 183–208 (cit. on pp. 66, 82, 84, 94, 95).
- Brunel, Nicolas, Vincent Hakim, Philippe Isope, Jean-Pierre Nadal, and Boris Barbour (Sept. 2004). “Optimal Information Storage and the Distribution of Synaptic Weights: Perceptron versus Purkinje Cell”. In: *Neuron* 43.5, pp. 745–757 (cit. on pp. 64, 107).
- Burns, B. D. and A. C. Webb (1976). “The spontaneous activity of neurons in the cat’s cerebral cortex”. In: *Proc. R. Soc. Lond. B* 194, pp. 211–223 (cit. on p. 67).
- Cajal, S. Ramón (1911). *Histologie du Système Nerveux de l’Homme et des Vertébrés, Tome II*. Paris: Maloine (cit. on p. 21).
- Campanac, E. and D. Debanne (2008). “Spike timing-dependent plasticity: a learning rule for dendritic integration in rat CA1 pyramidal neurons”. In: *J. Physiol. (Lond.)* 586, 779–793 (cit. on p. 67).
- Casado, M, S Dieudonné, and P Ascher (Oct. 2000). “Presynaptic N-methyl-D-aspartate receptors at the parallel fiber-Purkinje cell synapse”. In: *Proceedings of the National Academy of Sciences* 97.21. PMID: 11016958, pp. 11593–11597 (cit. on p. 26).
- Casado, Mariano, Philippe Isope, and Philippe Ascher (Jan. 2002). “Involvement of Presynaptic N-Methyl-D-Aspartate Receptors in Cerebellar Long-Term Depression”. In: *Neuron* 33.1, pp. 123–130 (cit. on pp. 26, 37, 50).

- Castellani, Gastone C., Armando Bazzani, and Leon N. Cooper (Aug. 2009). “Toward a microscopic model of bidirectional synaptic plasticity.” In: *Proceedings of the National Academy of Sciences* 106.33, 14091–14095 (cit. on p. 103).
- Castellucci, Vincent, Harold Pinsky, Irving Kupfermann, and Eric R. Kandel (1970). “Neuronal mechanisms of habituation and dishabituation of the gill-withdrawal reflex in *Aplysia*”. In: *Science* 167.3926, 1745–1748 (cit. on p. 1).
- Chen, C. and R. F. Thompson (May 1995). “Temporal specificity of long-term depression in parallel fiber–Purkinje synapses in rat cerebellar slice.” In: *Learning & Memory* 2.3-4. PMID: 10467575, pp. 185–198 (cit. on pp. 23, 24, 109).
- Churchland, Mark M et al. (Mar. 2010). “Stimulus onset quenches neural variability: a widespread cortical phenomenon”. In: *Nat Neurosci* 13.3, pp. 369–378 (cit. on pp. 67, 104).
- Clopath, C., J. P. Nadal, and N. Brunel (2012). “Storage of correlated patterns in standard and bistable Purkinje cell models”. In: *PLoS Comput. Biol.* 8, e1002448 (cit. on p. 64).
- Clopath, Claudia and Nicolas Brunel (Feb. 2013). “Optimal Properties of Analog Perceptrons with Excitatory Weights”. In: *PLoS Comput Biol* 9.2, e1002919 (cit. on pp. 64, 107).
- Clopath, Claudia, Lorric Ziegler, Eleni Vasilaki, Lars Busing, and Wulfram Gerstner (Dec. 2008). “Tag-Trigger-Consolidation: A Model of Early and Late Long-Term Potentiation and Depression”. In: *PLoS Computational Biology* 4.12, pp. 1–14 (cit. on pp. 15, 16).
- Clopath, Claudia, Lars Busing, Eleni Vasilaki, and Wulfram Gerstner (Mar. 2010). “Connectivity reflects coding: a model of voltage-based STDP with homeostasis”. In: *Nat Neurosci* 13.3, pp. 344–352 (cit. on pp. 18, 67, 103, 104).
- Coesmans, Michiel, John T. Weber, Chris I. De Zeeuw, and Christian Hansel (Nov. 2004). “Bidirectional Parallel Fiber Plasticity in the Cerebellum under Climbing Fiber Control”. In: *Neuron* 44.4, pp. 691–700 (cit. on pp. 28, 33–35, 60, 109).
- Collingridge, G. L., S. J. Kehl, and H. McLennan (Jan. 1983). “Excitatory amino acids in synaptic transmission in the Schaffer collateral-commissural pathway of the rat hippocampus.” In: *The Journal of Physiology* 334.1. PMID: 6306230, pp. 33–46 (cit. on p. 11).

- Couey, J. J., R. M. Meredith, S. Spijker, R. B. Poorthuis, A. B. Smit, A. B. Brussaard, and H. D. Mansvelder (2007). “Distributed network actions by nicotine increase the threshold for spike-timing-dependent plasticity in prefrontal cortex”. In: *Neuron* 54, 73–87 (cit. on p. 104).
- Crepel, F. and D. Jaillard (Jan. 1991). “Pairing of pre- and postsynaptic activities in cerebellar Purkinje cells induces long-term changes in synaptic efficacy in vitro.” In: *The Journal of Physiology* 432.1. PMID: 1886056, pp. 123–141 (cit. on pp. 22, 23).
- Cummings, Jennifer A, Rosel M Mulkey, Roger A Nicoll, and Robert C Malenka (Apr. 1996). “Ca²⁺ Signaling Requirements for Long-Term Depression in the Hippocampus”. In: *Neuron* 16.4, pp. 825–833 (cit. on p. 17).
- Daoudal, Gaël and Dominique Debanne (Nov. 2003). “Long-Term Plasticity of Intrinsic Excitability: Learning Rules and Mechanisms”. In: *Learning & Memory* 10.6. PMID: 14657257, pp. 456–465 (cit. on p. 6).
- Debanne, D., B. H. Gähwiler, and S. M. Thompson (Feb. 1994). “Asynchronous pre- and postsynaptic activity induces associative long-term depression in area CA1 of the rat hippocampus in vitro.” In: *Proceedings of the National Academy of Sciences* 91.3, pp. 1148–1152 (cit. on pp. 7, 8, 10, 67).
- Debanne, D, B H Gähwiler, and S M Thompson (Oct. 1996). “Cooperative interactions in the induction of long-term potentiation and depression of synaptic excitation between hippocampal CA3-CA1 cell pairs in vitro.” In: *Proceedings of the National Academy of Sciences* 93.20. PMID: 8855337 PMCID: 38312, pp. 11225–11230 (cit. on p. 8).
- Delord, Bruno, Hugues Berry, Emmanuel Guigon, and Stéphane Genet (June 2007). “A new principle for information storage in an enzymatic pathway model.” In: *PLoS Comput Biol* 3.6, e124 (cit. on p. 103).
- Doi, Tomokazu, Shinya Kuroda, Takayuki Michikawa, and Mitsuo Kawato (Jan. 2005). “Inositol 1,4,5-Trisphosphate-Dependent Ca²⁺ Threshold Dynamics Detect Spike Timing in Cerebellar Purkinje Cells”. In: *The Journal of Neuroscience* 25.4, pp. 950–961 (cit. on pp. 28, 61).
- Dudek, S M and M F Bear (May 1992). “Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade”. In: *Proceedings of the National Academy of Sciences* 89.10, pp. 4363–4367 (cit. on pp. 6, 17).

- Dudek, S. M. and M. F. Bear (July 1993). “Bidirectional long-term modification of synaptic effectiveness in the adult and immature hippocampus”. In: *The Journal of Neuroscience* 13.7, pp. 2910–2918 (cit. on pp. 7, 67).
- Eccles, John C, Masao Ito, and János Szentágothai (1967). *The Cerebellum as a Neuronal Machine*. Berlin, Heidelberg: Springer Berlin Heidelberg (cit. on pp. 20, 21).
- Eilers, J., G. Callewaert, C. Armstrong, and A. Konnerth (Oct. 1995). “Calcium signaling in a narrow somatic submembrane shell during synaptic activity in cerebellar Purkinje neurons.” In: *Proceedings of the National Academy of Sciences* 92.22, p. 10272 (cit. on p. 23).
- Ekerot, C.-F. and M. Kano (Feb. 1989). “Stimulation parameters influencing climbing fibre induced long-term depression of parallel fibre synapses”. In: *Neuroscience Research* 6.3, pp. 264–268 (cit. on pp. 21, 34).
- El Boustani, S., P. Yger, Y. Fregnac, and A. Destexhe (2012). “Stable learning in stochastic network states”. In: *J. Neurosci.* 32, 194–214 (cit. on pp. 67, 103).
- Finch, Elizabeth A. and George J. Augustine (Dec. 1998). “Local calcium signalling by inositol-1,4,5-trisphosphate in Purkinje cell dendrites”. en. In: *Nature* 396.6713, pp. 753–756 (cit. on p. 33).
- Froemke, Robert C. and Yang Dan (Mar. 2002). “Spike-timing-dependent synaptic modification induced by natural spike trains”. In: *Nature* 416.6879, pp. 433–438 (cit. on pp. 9, 17, 67).
- Froemke, Robert C., Ishan A. Tsay, Mohamad Raad, John D. Long, and Yang Dan (Mar. 2006). “Contribution of Individual Spikes in Burst-Induced Long-Term Synaptic Modification”. In: *Journal of Neurophysiology* 95.3. PMID: 16319206, pp. 1620–1629 (cit. on pp. 9, 17).
- Fusi, Stefano, Patrick J. Drew, and L. F. Abbott (Feb. 2005). “Cascade Models of Synaptically Stored Memories”. In: *Neuron* 45.4, pp. 599–611 (cit. on pp. 14, 15, 67).
- Galassi, Mark (2009). *GNU scientific library: reference manual*. Bristol: Network Theory (cit. on pp. 46, 122).
- Gardiner, C. W. (1986). *Handbook of stochastic methods for physics, chemistry and the natural sciences*. Springer-Verlag (cit. on p. 77).

- Gerstner, Wulfram, Raphael Ritz, and J. Leo van Hemmen (Sept. 1993). “Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns”. In: *Biological Cybernetics* 69.5-6, pp. 503–515 (cit. on p. 13).
- Gerstner, Wulfram, Richard Kempter, J. Leo van Hemmen, and Hermann Wagner (1996). “A neuronal learning rule for sub-millisecond temporal coding”. In: *Nature* 383.6595, pp. 76–78 (cit. on pp. 13, 67).
- Gilbert, E. N. and H. O. Pollak (1960). “Amplitude distribution of shot noise”. In: *Bell Syst. Tech. J* 39.2, 333–350 (cit. on pp. 73, 88, 125, 126).
- Gilson, Matthieu, Anthony Burkitt, and J. Leo van Hemmen (2010). “STDP in recurrent neuronal networks”. In: *Frontiers in Computational Neuroscience* 4, p. 23 (cit. on p. 18).
- Golding, Nace L., Nathan P. Staff, and Nelson Spruston (July 2002). “Dendritic spikes as a mechanism for cooperative long-term potentiation”. In: *Nature* 418.6895, pp. 326–331 (cit. on p. 9).
- Goodman, Dan F. M. and Romain Brette (Sept. 2009). “The Brian Simulator”. In: *Frontiers in Neuroscience* 3.2. PMID: 20011141 PMCID: PMC2751620, pp. 192–197 (cit. on p. 83).
- Graupner, Michael and Nicolas Brunel (Feb. 2012). “Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location”. In: *Proceedings of the National Academy of Sciences* (cit. on pp. 18, 19, 29, 65, 67–70, 72, 73, 76, 82, 86, 88, 102, 103, 107, 110, 114, 115).
- Hansel, C. (Dec. 2005). “When the B-team runs plasticity: GluR2 receptor trafficking in cerebellar long-term potentiation”. In: *Proceedings of the National Academy of Sciences* 102.51, pp. 18245–18246 (cit. on p. 27).
- Hartell, Nick A (Mar. 1996). “Strong Activation of Parallel Fibers Produces Localized Calcium Transients and a Form of LTD That Spreads to Distant Synapses”. In: *Neuron* 16.3, pp. 601–610 (cit. on pp. 28, 58, 109).
- Hayer, Arnold and Upinder S. Bhalla (July 2005). “Molecular switches at the synapse emerge from receptor and kinase traffic.” In: *PLoS Comput Biol* 1.2, 137–154 (cit. on p. 103).
- Hebb, Donald Olding (1949). *The organisation of behavior: a neuropsychological theory*. New York: Wiley and Sons (cit. on pp. 8, 11, 67).

- Ho, Victoria M., Ji-Ann Lee, and Kelsey C. Martin (Nov. 2011). “The Cell Biology of Synaptic Plasticity”. In: *Science* 334.6056, pp. 623–628 (cit. on p. 4).
- Holmgren, C., T. Harkany, B. Svennenfors, and Y. Zilberter (2003). “Pyramidal cell communication within local networks in layer 2/3 of rat neocortex.” In: *J. Physiol.* 551, pp. 139–153 (cit. on p. 103).
- Hopfield, J. J. (Apr. 1982). “Neural networks and physical systems with emergent collective computational abilities”. en. In: *Proceedings of the National Academy of Sciences* 79.8. PMID: 6953413, pp. 2554–2558 (cit. on p. 13).
- Ismailov, Iskander, Djanenkhodja Kalikulov, Takafumi Inoue, and Michael J. Friedlander (Nov. 2004). “The Kinetic Profile of Intracellular Calcium Predicts Long-Term Potentiation and Long-Term Depression”. In: *The Journal of Neuroscience* 24.44, pp. 9847–9861 (cit. on p. 68).
- Isope, Philippe, Romain Franconville, Boris Barbour, and Philippe Ascher (Feb. 2004). “Repetitive firing of rat cerebellar parallel fibres after a single stimulation”. en. In: *The Journal of Physiology* 554.3. PMID: 14634204, pp. 829–839 (cit. on p. 55).
- Ito, Masao and Laddawan Karachot (Aug. 1990). “Receptor subtypes involved in, and time course of, the long-term desensitization of glutamate receptors in cerebellar Purkinje cells”. en. In: *Neuroscience Research* 8.4, pp. 303–307 (cit. on p. 33).
- Ito, Masao, Masaki Sakurai, and Pavich Tongroach (Mar. 1982). “Climbing fibre induced depression of both mossy fibre responsiveness and glutamate sensitivity of cerebellar Purkinje cells”. In: *The Journal of Physiology* 324.1. PMID: 7097592, pp. 113–134 (cit. on pp. 21, 22).
- Jörntell, Henrik and Christian Hansel (Oct. 2006). “Synaptic Memories Upside Down: Bidirectional Plasticity at Cerebellar Parallel Fiber-Purkinje Cell Synapses”. In: *Neuron* 52.2, pp. 227–238 (cit. on p. 27).
- Takegawa, W. and M. Yuzaki (2005). “A mechanism underlying AMPA receptor trafficking during cerebellar long-term potentiation”. In: *Proceedings of the National Academy of Sciences of the United States of America* 102.49, p. 17846 (cit. on pp. 27, 34).
- Karachot, Laddawan, Raymond T. Kado, and Masao Ito (Dec. 1994). “Stimulus parameters for induction of long-term depression in in vitro rat Purkinje cells”. In: *Neuroscience Research* 21.2, pp. 161–168 (cit. on p. 22).

- Kauer, Julie A. and Robert C. Malenka (Nov. 2007). “Synaptic plasticity and addiction”. In: *Nature Reviews Neuroscience* 8.11, pp. 844–858 (cit. on p. 5).
- Kawaguchi, Shin-ya and Tomoo Hirano (Apr. 2013). “Gating of long-term depression by Ca²⁺/calmodulin-dependent protein kinase II through enhanced cGMP signalling in cerebellar Purkinje cells”. In: *The Journal of Physiology* 591.7, pp. 1707–1730 (cit. on p. 28, 61–63).
- Kempter, Richard, Wulfram Gerstner, and J. Leo van Hemmen (Apr. 1999). “Hebbian learning and spiking neurons”. In: *Physical Review E* 59.4, pp. 4498–4514 (cit. on p. 20).
- Kenet, T., D. Bibitchkov, M. Tsodyks, A. Grinvald, and A. Arieli (2003). “Spontaneously emerging cortical representations of visual attributes.” In: *Nature* 425, pp. 954–956 (cit. on p. 104).
- Khodakhah, Kamran and Clay M. Armstrong (Dec. 1997). “Induction of long-term depression and rebound potentiation by inositol trisphosphate in cerebellar Purkinje neurons”. In: *Proceedings of the National Academy of Sciences* 94.25. PMID: 9391143, pp. 14009–14014 (cit. on p. 27).
- Kohonen, T. (1982). “Self-Organized Formation of Topologically Correct Feature Maps”. In: *Biological Cybernetics* 43, pp. 59–69 (cit. on p. 11).
- Konnerth, A., J. Dreessen, and G. J. Augustine (1992). “Brief dendritic calcium signals initiate long-lasting synaptic depression in cerebellar Purkinje cells”. In: *Proceedings of the National Academy of Sciences* 89.15, p. 7051 (cit. on p. 22).
- Kramers, H. A. (1940). “Brownian motion in a field of force and the diffusion model of chemical reactions”. In: *Physica* 8, pp. 284–304 (cit. on pp. 77, 111).
- Kumar, A. and M. R. Mehta (2011). “Frequency-Dependent Changes in NMDAR-Dependent Synaptic Plasticity”. In: *Front Comput Neurosci* 5, p. 38 (cit. on pp. 67, 103, 104).
- Kuroda, Shinya, Nicolas Schweighofer, and Mitsuo Kawato (2001). “Exploration of Signal Transduction Pathways in Cerebellar Long-Term Depression by Kinetic Simulation”. In: *The Journal of Neuroscience* 21.15, pp. 5693–5702 (cit. on pp. 28, 36, 61).
- L’Ecuyer, Pierre and Richard Simard (2007). “TestU01: A C library for empirical testing of random number generators”. In: *ACM Transactions on Mathematical Software (TOMS)* 33.4, p. 22 (cit. on pp. 121, 122).

- Lev-Ram, Varda, Lewis R. Makings, Paul F. Keitz, Joseph P.Y. Kao, and Roger Y. Tsien (Aug. 1995). “Long-term depression in cerebellar Purkinje neurons results from coincidence of nitric oxide and depolarization-induced Ca²⁺ transients”. In: *Neuron* 15.2, pp. 407–415 (cit. on pp. 24, 25, 33, 35, 60, 109).
- Lev-Ram, Varda, Tao Jiang, Jason Wood, David S Lawrence, and Roger Y Tsien (June 1997). “Synergies and Coincidence Requirements between NO, cGMP, and Ca²⁺ in the Induction of Cerebellar Long-Term Depression”. In: *Neuron* 18.6, pp. 1025–1038 (cit. on pp. 35, 60).
- Lev-Ram, Varda, Scott T. Wong, Daniel R. Storm, and Roger Y. Tsien (June 2002). “A new form of cerebellar long-term potentiation is postsynaptic and depends on nitric oxide but not cAMP”. In: *Proceedings of the National Academy of Sciences* 99.12, pp. 8389–8393 (cit. on pp. 25, 26, 32–35, 60, 109).
- Lev-Ram, Varda, Samar B. Mehta, David Kleinfeld, and Roger Y. Tsien (Dec. 2003). “Reversing cerebellar long-term depression”. In: *Proceedings of the National Academy of Sciences* 100.26. PMID: 14671315, pp. 15989–15993 (cit. on p. 26).
- Lin, Y. W., M. Y. Min, T. H. Chiu, and H. W. Yang (2003). “Enhancement of associative long-term potentiation by activation of beta-adrenergic receptors at CA1 synapses in rat hippocampal slices”. In: *J. Neurosci.* 23, 4173–4181 (cit. on p. 104).
- Lisman, J. E. (May 1985). “A mechanism for memory storage insensitive to molecular turnover: a bistable autophosphorylating kinase”. In: *Proceedings of the National Academy of Sciences* 82.9, pp. 3055–3057 (cit. on pp. 14, 67).
- Lisman, John and Nelson Spruston (July 2005). “Postsynaptic depolarization requirements for LTP and LTD: a critique of spike timing-dependent plasticity”. In: *Nature Neuroscience* 8.7, pp. 839–841 (cit. on p. 8).
- Litwin-Kumar, Ashok and Brent Doiron (Nov. 2012). “Slow dynamics and high variability in balanced cortical networks with clustered connections”. In: *Nature Neuroscience* 15.11, pp. 1498–1505 (cit. on pp. 104, 115).
- Loewenstein, Yonatan, Annerose Kuras, and Simon Rumpel (June 2011). “Multiplicative Dynamics Underlie the Emergence of the Log-Normal Distribution of Spine Sizes in the Neocortex In Vivo”. In: *The Journal of Neuroscience* 31.26, pp. 9481–9488 (cit. on p. 103).

- Luczak, A., P. Bartho, and K. D. Harris (2009). “Spontaneous events outline the realm of possible sensory responses in neocortical populations”. In: *Neuron* 62, 413–425 (cit. on p. 104).
- Luque, N. R., J. A. Garrido, R. R. Carrillo, S. Tolu, and E. Ros (Oct. 2011). “Adaptive cerebellar spiking model embedded in the control loop: context switching and robustness against noise”. In: *International Journal of Neural Systems* 21.05, pp. 385–401 (cit. on p. 115).
- Ly, Romain, Guy Bouvier, Martijn Schonewille, Arnaud Arabo, Laure Rondi-Reig, Clément Léna, Mariano Casado, Chris I. De Zeeuw, and Anne Feltz (Dec. 2013). “T-type channel blockade impairs long-term potentiation at the parallel fiber–Purkinje cell synapse and cerebellar learning”. In: *Proceedings of the National Academy of Sciences* 110.50. PMID: 24277825, pp. 20302–20307 (cit. on p. 60).
- Lynch, G. S., T. Dunwiddie, and V. Gribkoff (Apr. 1977). “Heterosynaptic depression: a postsynaptic correlate of long-term potentiation”. In: *Nature* 266.5604, pp. 737–739 (cit. on p. 6).
- Lynch, Gary, John Larson, Stephen Kelso, German Barrionuevo, and Frank Schottler (1983). “Intracellular injections of EGTA block induction of hippocampal long-term potentiation”. In: *Nature* 305.5936, pp. 719–721 (cit. on p. 11).
- Malenka, Robert C, Julie A Kauer, Robert S Zucker, and Roger A Nicoll (1988). “Postsynaptic calcium is sufficient for potentiation of hippocampal synaptic transmission”. In: *Science* 242.4875, pp. 81–84 (cit. on pp. 11, 17, 68).
- Marcaggi, Païkan and David Attwell (Jan. 2007). “Short- and long-term depression of rat cerebellar parallel fibre synaptic transmission mediated by synaptic crosstalk”. In: *The Journal of Physiology* 578.2, pp. 545–550 (cit. on pp. 27, 31, 55, 109).
- Markram, Henry, Joachim Lübke, Michael Frotscher, and Bert Sakmann (Jan. 1997). “Regulation of Synaptic Efficacy by Coincidence of Postsynaptic APs and EP-SPs”. In: *Science* 275.5297, pp. 213–215 (cit. on pp. 8, 17, 67, 103).
- Marquardt, Donald W. (1963). “An algorithm for least-squares estimation of nonlinear parameters”. In: *Journal of the Society for Industrial & Applied Mathematics* 11.2, 431–441 (cit. on p. 41).
- Marr, David (June 1969). “A theory of cerebellar cortex”. In: *The Journal of Physiology* 202.2. PMID: 5784296, pp. 437–470 (cit. on pp. 20, 28, 64, 107).

- Marsaglia, George (1996). *DIEHARD: a battery of tests of randomness* (cit. on p. 121).
- Marsaglia, George and Arif Zaman (1993). *The KISS generator*. Tech. rep. Tech. rep., Department of Statistics, University of Florida (cit. on p. 123).
- Martens, James (2010). “Deep learning via Hessian-free optimization”. In: *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, 735–742 (cit. on p. 44).
- Matsuda, Shinji, Thomas Launey, Sumiko Mikawa, and Hirokazu Hirai (June 2000). “Disruption of AMPA receptor GluR2 clusters following long-term depression induction in cerebellar Purkinje neurons”. In: *The EMBO Journal* 19.12, pp. 2765–2774 (cit. on p. 27).
- Matsumoto, Makoto and Takuji Nishimura (1998). “Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator”. In: *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 8.1, 3–30 (cit. on p. 122).
- Miller, Paul, Anatol M. Zhabotinsky, John E. Lisman, and Xiao-Jing Wang (Apr. 2005). “The stability of a stochastic CaMKII switch: dependence on the number of enzyme molecules and protein turnover.” In: *PLoS Biol* 3.4, e107 (cit. on pp. 103, 113).
- Miyata, Mariko, Elizabeth A. Finch, Leonard Khiroug, Kouichi Hashimoto, Shizu Hayasaka, Sen-Ichi Oda, Minoru Inouye, Yoshiko Takagishi, George J. Augustine, and Masanobu Kano (Jan. 2000). “Local Calcium Release in Dendritic Spines Required for Long-Term Synaptic Depression”. English. In: *Neuron* 28.1. PMID: 11086997, pp. 233–244 (cit. on p. 33).
- Mizuno, Tomoyuki, Ichiro Kanazawa, and Masaki Sakurai (2001). “Differential induction of LTP and LTD is not determined solely by instantaneous calcium concentration: an essential involvement of a temporal factor”. In: *European Journal of Neuroscience* 14.4, 701–708 (cit. on p. 68).
- Mongillo, Gianluigi, Emanuele Curti, Sandro Romani, and Daniel J. Amit (2005). “Learning in realistic networks of spiking neurons and spike-driven plastic synapses”. In: *European Journal of Neuroscience* 21.11, 3143–3160 (cit. on p. 18).
- Mongillo, Gianluigi, Omri Barak, and Misha Tsodyks (Mar. 2008). “Synaptic Theory of Working Memory”. In: *Science* 319.5869. PMID: 18339943, pp. 1543–1546 (cit. on p. 14).

- Nabavi, S., H. W. Kessels, S. Alfonso, J. Aow, R. Fox, and R. Malinow (2013). “Metabotropic NMDA receptor function is required for NMDA receptor-dependent long-term depression”. In: *Proc. Natl. Acad. Sci. U.S.A.* 110, 4027–4032 (cit. on p. 68).
- Nadal, J. P., G. Toulouse, J. P. Changeux, and S. Dehaene (May 1986). “Networks of Formal Neurons and Memory Palimpsests”. In: *EPL (Europhysics Letters)* 1.10, p. 535 (cit. on p. 13).
- Nelder, John A. and Roger Mead (1965). “A simplex method for function minimization”. In: *Computer journal* 7.4, 308–313 (cit. on pp. 41, 45).
- Neveu, Dorine and Robert S. Zucker (Mar. 1996). “Postsynaptic Levels of $[Ca^{2+}]_i$ Needed to Trigger LTD and LTP”. In: *Neuron* 16.3, pp. 619–629 (cit. on pp. 11, 17, 68).
- Nevian, Thomas and Bert Sakmann (Oct. 2006). “Spine Ca^{2+} signaling in spike-timing-dependent plasticity”. In: *The Journal of Neuroscience* 26.43, pp. 11001–11013 (cit. on p. 68).
- O’Connor, Daniel H, Gayle M Wittenberg, and Samuel S-H Wang (2005a). “Dissection of Bidirectional Synaptic Plasticity into Saturable Unidirectional Processes”. In: *Proceedings of the National Academy of Sciences* (cit. on p. 10).
- O’Connor, Daniel H., Gayle M. Wittenberg, and Samuel S.-H. Wang (July 2005b). “Graded bidirectional synaptic plasticity is composed of switch-like unitary events”. In: *Proceedings of the National Academy of Sciences* 102.27. PMID: 15983385, pp. 9679–9684 (cit. on pp. 10, 68).
- Oja, Erkki (Nov. 1982). “Simplified neuron model as a principal component analyzer”. en. In: *Journal of Mathematical Biology* 15.3, pp. 267–273 (cit. on pp. 13, 67).
- Pawlak, V., J. R. Wickens, A. Kirkwood, and J. N. Kerr (2010). “Timing is not Everything: Neuromodulation Opens the STDP Gate”. In: *Front Synaptic Neurosci* 2, p. 146 (cit. on p. 104).
- Petersen, Carl C. H., Robert C. Malenka, Roger A. Nicoll, and John J. Hopfield (Apr. 1998). “All-or-none potentiation at CA3-CA1 synapses”. In: *Proceedings of the National Academy of Sciences* 95.8, pp. 4732–4737 (cit. on pp. 10, 68).

- Pfister, Jean-Pascal and Wulfram Gerstner (Sept. 2006). “Triplets of Spikes in a Model of Spike Timing-Dependent Plasticity”. In: *J. Neurosci.* 26.38, pp. 9673–9682 (cit. on pp. 18, 67, 103, 104).
- Pi, Hyun Jae and John E. Lisman (Dec. 2008). “Coupled phosphatase and kinase switches produce the tristability required for long-term potentiation and long-term depression.” In: *J Neurosci* 28.49, 13132–13138 (cit. on p. 103).
- Piochon, Claire, Carole Levenes, Gen Ohtsuki, and Christian Hansel (Nov. 2010). “Purkinje Cell NMDA Receptors Assume a Key Role in Synaptic Gain Control in the Mature Cerebellum”. In: *The Journal of Neuroscience* 30.45. PMID: 21068337, pp. 15330–15335 (cit. on p. 27).
- Polak, E. and G. Ribiere (1969). “Note sur la convergence de méthodes de directions conjuguées”. In: *ESAIM: Mathematical Modelling and Numerical Analysis - Modélisation Mathématique et Analyse Numérique* 3.R1, pp. 35–43 (cit. on p. 41).
- Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery (1992). *Numerical Recipes in C*. Cambridge University Press (cit. on p. 122).
- Rice, S. O. (1944). “Mathematical Analysis of Random Noise”. In: *The Bell System technical journal*, p. 282 (cit. on p. 125).
- Rose, Greg (2011). “KISS: A Bit Too Simple.” In: *IACR Cryptology ePrint Archive* 2011, p. 7 (cit. on p. 123).
- Roxin, A. (2011). “The role of degree distribution in shaping the dynamics in networks of sparsely connected spiking neurons”. In: *Front Comput Neurosci* 5, p. 8 (cit. on p. 85).
- Roxin, Alex and Stefano Fusi (July 2013). “Efficient Partitioning of Memory Systems and Its Importance for Memory Consolidation”. In: *PLoS Comput Biol* 9.7, e1003146 (cit. on p. 15).
- Rubin, J., D. D. Lee, and H. Sompolinsky (2001). “Equilibrium properties of temporally asymmetric Hebbian plasticity”. In: *Phys. Rev. Lett.* 86, pp. 364–367 (cit. on p. 102).
- Rubin, Jonathan E, Richard C Gerkin, Guo-Qiang Bi, and Carson C Chow (2005). “Calcium time course as a signal for spike-timing-dependent plasticity”. In: *Journal of neurophysiology* 93.5, pp. 2600–2613 (cit. on p. 17).

- Sabatini, B. L., T. G. Oertner, and K. Svoboda (2002). “The life cycle of Ca(2+) ions in dendritic spines”. In: *Neuron* 33, 439–452 (cit. on p. 70).
- Safo, Patrick and Wade G Regehr (Jan. 2008). “Timing dependence of the induction of cerebellar LTD”. In: *Neuropharmacology* 54.1. PMID: 17669443, pp. 213–218 (cit. on pp. 34, 36, 39, 40, 50, 55–57, 108).
- Sakurai, M. (Dec. 1987). “Synaptic modification of parallel fibre-Purkinje cell transmission in in vitro guinea-pig cerebellar slices.” In: *The Journal of Physiology* 394.1. PMID: 2832595, pp. 463–480 (cit. on p. 21).
- (May 1990). “Calcium is an intracellular mediator of the climbing fiber in induction of cerebellar long-term depression”. In: *Proceedings of the National Academy of Sciences* 87.9. PMID: 2159149, pp. 3383–3385 (cit. on p. 22).
- Salmon, J.K., M.A. Moraes, R.O. Dror, and D.E. Shaw (Nov. 2011). “Parallel random numbers: As easy as 1, 2, 3”. In: *High Performance Computing, Networking, Storage and Analysis (SC), 2011 International Conference for*, pp. 1–12 (cit. on pp. 83, 123).
- Sarkisov, D.V. and S.S. Wang (2008). “Order-dependent coincidence detection in cerebellar Purkinje neurons at the inositol trisphosphate receptor”. In: *J. Neurosci.* 28, 133–142 (cit. on pp. 27, 28).
- Schermelleh-Engel, Karin, Helfried Moosbrugger, and Hans Müller (2003). “Evaluating the fit of structural equation models: Tests of significance and descriptive goodness-of-fit measures”. In: *Methods of psychological research online* 8.2, 23–74 (cit. on pp. 47, 49).
- Schwarz, Gideon (Mar. 1978). “Estimating the Dimension of a Model”. In: *The Annals of Statistics* 6.2, pp. 461–464 (cit. on p. 49).
- Sejnowski, T. J (1977). “Storing covariance with nonlinearly interacting neurons”. In: *Journal of mathematical biology* 4.4, 303–321 (cit. on pp. 11, 67).
- Seol, G. H., J. Ziburkus, S. Huang, L. Song, I. T. Kim, K. Takamiya, R. L. Haganir, H. K. Lee, and A. Kirkwood (2007). “Neuromodulators control the polarity of spike-timing-dependent synaptic plasticity”. In: *Neuron* 55, 919–929 (cit. on p. 104).
- Shin, Jung Hoon and David J. Linden (2005). “An NMDA receptor/nitric oxide cascade is involved in cerebellar LTD but is not localized to the parallel fiber terminal”. In: *Journal of neurophysiology* 94.6, 4281–4289 (cit. on pp. 27, 32, 37, 109).

- Shouval, Harel Z. (Oct. 2005). “Clusters of interacting receptors can stabilize synaptic efficacies.” In: *Proceedings of the National Academy of Sciences* 102.40, 14440–14445 (cit. on p. 103).
- Shouval, Harel Z, Mark F Bear, and Leon N Cooper (Aug. 2002). “A unified model of NMDA receptor-dependent bidirectional synaptic plasticity”. In: *Proceedings of the National Academy of Sciences of the United States of America* 99.16, pp. 10831–10836 (cit. on pp. 16, 17, 67, 68, 104).
- Siebert, A. J. F. (1951). “On the first passage time probability problem”. In: *Phys. Rev.* 81, pp. 617–623 (cit. on p. 84).
- Silver, I A and M Erecińska (May 1990). “Intracellular and extracellular changes of [Ca²⁺] in hypoxia and ischemia in rat brain in vivo”. In: *The Journal of general physiology* 95.5, pp. 837–866 (cit. on pp. 63, 68, 70, 105).
- Sims, Robert E. and Nicholas A. Hartell (May 2006). “Differential Susceptibility to Synaptic Plasticity Reveals a Functional Specialization of Ascending Axon and Parallel Fiber Synapses to Cerebellar Purkinje Cells”. In: *The Journal of Neuroscience* 26.19, pp. 5153–5159 (cit. on p. 28).
- Sjöström, P. J., G. G. Turrigiano, and S. B. Nelson (2001). “Rate, timing, and cooperativity jointly determine cortical synaptic plasticity”. In: *Neuron* 32, pp. 1149–1164 (cit. on pp. 8–10, 17, 18, 29, 65, 67, 68, 70, 73, 103).
- Song, Sen, Kenneth D. Miller, and L. F. Abbott (2000). “Competitive Hebbian learning through spike-timing-dependent synaptic plasticity”. In: *Nat Neurosci* 3.9, pp. 919–926 (cit. on pp. 13, 18, 67, 102, 112).
- Song, Sen, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B Chklovskii (2005). “Highly nonrandom features of synaptic connectivity in local cortical circuits”. In: *PLoS Biol* 3, e68 (cit. on p. 103).
- Steiger, James H., Alexander Shapiro, and Michael W. Browne (Sept. 1985). “On the multivariate asymptotic distribution of sequential Chi-square statistics”. en. In: *Psychometrika* 50.3, pp. 253–263 (cit. on p. 49).
- Takacs, L. (1956). “On secondary stochastic processes generated by recurrent processes”. In: *Acta Mathematica Academiae Scientiarum Hungarica* 7, pp. 17–29 (cit. on p. 105).
- Tanaka, Keiko, Leonard Khiroug, Fidel Santamaria, Tomokazu Doi, Hideaki Ogasawara, Graham C.R. Ellis-Davies, Mitsuo Kawato, and George J. Augustine

- (June 2007). “Ca²⁺ Requirements for Cerebellar Long-Term Synaptic Depression: Role for a Postsynaptic Leaky Integrator”. In: *Neuron* 54.5, pp. 787–800 (cit. on p. 33).
- Toyoizumi, T, JP Pfister, K Aihara, and W Gerstner (Apr. 2005). “Generalized Bienenstock-Cooper-Munro rule for spiking neurons that maximizes information transmission.” In: *Proceedings of the National Academy of Sciences* 102.14, pp. 5239–44 (cit. on p. 16).
- Tsodyks, M., T. Kenet, A. Grinvald, and A. Arieli (1999). “Linking spontaneous activity of single cortical neurons and the underlying functional architecture.” In: *Science* 286, pp. 1943–1946 (cit. on p. 104).
- Tsodyks, M. V. and M. V. Feigel'man (May 1988). “The Enhanced Storage Capacity in Neural Networks with Low Activity Level”. en. In: *EPL (Europhysics Letters)* 6.2, p. 101 (cit. on p. 13).
- Tsodyks, Misha, Klaus Pawelzik, and Henry Markram (May 1998). “Neural Networks with Dynamic Synapses”. In: *Neural Computation* 10.4, pp. 821–835 (cit. on p. 14).
- Tsodyks, Misha V. and Henry Markram (Jan. 1997). “The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability”. In: *Proceedings of the National Academy of Sciences* 94.2, pp. 719–723 (cit. on p. 14).
- Turrigiano, Gina G and Sacha B Nelson (June 2000). “Hebb and homeostasis in neuronal plasticity”. In: *Current Opinion in Neurobiology* 10.3, pp. 358–364 (cit. on p. 16).
- Turrigiano, Gina G. and Sacha B. Nelson (Feb. 2004). “Homeostatic plasticity in the developing nervous system”. en. In: *Nature Reviews Neuroscience* 5.2, pp. 97–107 (cit. on p. 16).
- Van Rossum, M. C. W. and G. G. Turrigiano (June 2001). “Correlation based learning from spike timing dependent plasticity”. In: *Neurocomputing* 38-40, pp. 409–415 (cit. on p. 14).
- Van Rossum, M. C. W., G. Q. Bi, and G. G. Turrigiano (Dec. 2000). “Stable Hebbian Learning from Spike Timing-Dependent Plasticity”. In: *J. Neurosci.* 20.23, pp. 8812–8821 (cit. on pp. 14, 16, 102).

- Vogels, T. P., H. Sprekeler, F. Zenke, C. Clopath, and W. Gerstner (Dec. 2011). “Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks”. In: *Science* 334.6062, pp. 1569–1573 (cit. on p. 20).
- Wang, Huai-Xing, Richard C. Gerkin, David W. Nauen, and Guo-Qiang Bi (Feb. 2005). “Coactivation and timing-dependent integration of synaptic potentiation and depression”. In: *Nature Neuroscience* 8.2, pp. 187–193 (cit. on pp. 10, 18, 65, 67).
- Wang, Samuel S.-H., Winfried Denk, and Michael Hausser (Dec. 2000). “Coincidence detection in single dendritic spines mediated by calcium release”. In: *Nat Neurosci* 3.12, pp. 1266–1273 (cit. on pp. 24, 32, 34).
- Wang, Yu Tian and David J. Linden (Mar. 2000). “Expression of Cerebellar Long-Term Depression Requires Postsynaptic Clathrin-Mediated Endocytosis”. In: *Neuron* 25.3, pp. 635–647 (cit. on p. 27).
- Wittenberg, Gayle M. and Samuel S.-H. Wang (June 2006). “Malleability of Spike-Timing-Dependent Plasticity at the CA3?CA1 Synapse”. In: *The Journal of Neuroscience* 26.24, pp. 6610–6617 (cit. on pp. 10, 65, 67).
- Xia, Jun, Hee Jung Chung, Cornelia Wihler, Richard L Huganir, and David J Linden (Nov. 2000). “Cerebellar Long-Term Depression Requires PKC-Regulated Interactions between GluR2/3 and PDZ Domain-Containing Proteins”. In: *Neuron* 28.2, pp. 499–510 (cit. on p. 27).
- Yang, Shao-Nian, Yun-Gui Tang, and Robert S. Zucker (Feb. 1999). “Selective Induction of LTP and LTD by Postsynaptic $[Ca^{2+}]_i$ Elevation”. In: *Journal of Neurophysiology* 81.2, pp. 781–787 (cit. on pp. 11, 68).
- Yang, Yan and Stephen G. Lisberger (May 2014). “Purkinje-cell plasticity and cerebellar motor learning are graded by complex-spike duration”. en. In: *Nature* advance online publication (cit. on p. 107).
- Yger, Pierre and Kenneth D. Harris (Oct. 2013). “The Convallis Rule for Unsupervised Learning in Cortical Networks”. In: *PLoS Comput Biol* 9.10, e1003272 (cit. on pp. 67, 68, 103).
- Zenke, Friedemann, Guillaume Hennequin, and Wulfram Gerstner (Nov. 2013). “Synaptic Plasticity in Neural Networks Needs Homeostasis with a Fast Rate Detector”. In: *PLoS Comput Biol* 9.11, e1003330 (cit. on pp. 20, 112).

- Zhabotinsky, Anatol M. (2000). “Bistability in the $\text{Ca}(2+)/\text{Calmodulin-Dependent Protein Kinase-Phosphatase System}$ ”. In: *Biophys J* 79.5, pp. 2211–2221 (cit. on p. 103).
- Zhang, J. C., P. M. Lau, and G. Q. Bi (2009). “Gain in sensitivity and loss in temporal contrast of STDP by dopaminergic modulation at hippocampal synapses”. In: *Proceedings of the National Academy of Sciences* 106, 13028–13033 (cit. on p. 104).
- Zucker, Robert S. (1989). “Short-term synaptic plasticity”. In: *Annual review of neuroscience* 12.1, 13–31 (cit. on p. 14).