LECTURE 5
Spike timing–dependent plasticity (STDP)

In the neuron models discussed so far each synapse is characterized by a single fixed real parameter w that determines the amplitude of the postsynaptic response to an incoming spike. However, a great deal of experimental evidence has accumulated that these synaptic weights can undergo persistent changes over time. Such changes are thought to be the neural correlates of learning and memory. Most experimental studies of synaptic plasticity have focused on refining the original theoretical postulate due to Hebb [**107**]: *When an axon of cell A is near enough to excite cell B or repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased*. A more modern interpretation of this postulate is that synaptic modification is driven by correlations in the firing activity of presynaptic and postsynaptic neurons.

In many regions of the brain, neurons are found to exhibit bidirectional plasticity in which the strength of a synapse can increase or decrease depending on the stimulus protocol [**109, 110, 111, 112, 113**]. Long term potentiation (LTP) is a persistent increase in synaptic efficacy produced by high-frequency stimulation of presynaptic afferents or by the pairing of low frequency presynaptic stimulation with robust postsynaptic depolarization. Long–term synaptic depression (LTD) is a long-lasting decrease in synaptic strength induced by low-frequency stimulation of presynaptic afferents. More recent experimental studies suggest that both the sign and degree of synaptic modification arising from repeated pairing of pre- and postsynaptic action potentials depend on their relative timing [**114, 115, 116**]. Long-term strengthening of synapses occurs if presynaptic action potentials precede postsynaptic firing by no more than about 50 ms. Presynaptic action potentials that follow postsynaptic spikes produce long-term weakening of synapses. The largest changes in synaptic efficacy occur when the time difference between pre- and postsynaptic action potentials is small, and there is a sharp transition from strengthening to weakening. This phenomenon of spike timing–dependent plasticity (STDP) is illustrated in figure 1.

One of the problems with standard formulations of Hebbian learning is that without additional constraints, the level of activity in a network can grow or shrink in an uncontrolled manner. The observation that LTP and LTD can occur at the same synapse depending on the pre-post correlations suggests that STDP can

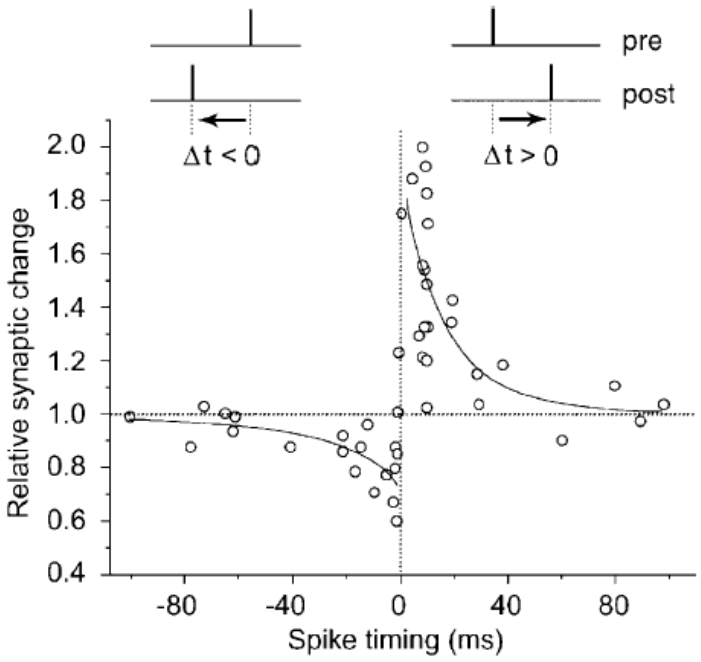

Figure 1. Spike-timing-dependent synaptic plasticity observed in hippocampal neurons. Each data point represents the relative change in the amplitude of evoked postsynaptic current after repetitive application of pre- and postsynaptic spiking pairs (1 Hz for 60 s) with xed spike timing Δt, which is dened as the time interval between pre- and postsynaptic spiking within each pair. Long-term potentiation (LTP) and depression (LTD) windows are each tted with an exponential function. Adapted from Bi and Poo [**108**]

provide a mechanism for self-stabilization of weights and rates in a network. A second important requirement of any learning scheme is that networks need to operate competitively. Thus, normally it does not make sense to drive all synapses into similar states through learning, because the network will then no longer contain inner structure. Instead, most real networks in the brain are highly structured, forming maps or other subdivisions. In this lecture we investigate the possible role of STDP in stability and competition by considering a simple phenomenological model of a stochastically driven neuron with adapting synapses. We then discuss a more detailed biophysical model of bidirectional synaptic plasticity that contains two essential components: a description of intracellular calcium dynamics with calcium influx mediated by NMDA receptors, and a mechanism for how calcium concentration can induce a change of synaptic efficacy.

5.1. Hebbian learning in a spiking neuron model

Consider a stochastically spiking neuron having N adapting synapses with weights w_j, $j = 1, \dots N$. Let $S_j(t) = \sum_n \delta(t - T_j^m)$ and $S(t) = \sum_m \delta(t - T^m)$ be the presynaptic and postsynaptic spike trains respectively. We assume that an instantaneous jump in the strength of a synaptic weight occurs whenever there is a presynaptic or

postsynaptic spike. The amplitude of each jump depends on the relative timing of previous spikes along the lines of figure 1 The rate of change of the synaptic weight w_j is then taken to be of the form [**117, 118, 2**]

$$\frac{dw_j}{dt} = c_0 + S_j(t)\int_0^\infty W(-s)S(t-s)ds + S(t)\int_0^\infty W(s)S_j(t-s)ds \tag{5.1}$$

where c_0 represents an activity–independent growth or decay and the kernel $W(s)$ for $s > 0$ ($s < 0$) gives the weight change if a postsynaptic spike is preceeded (followed) by a presynaptic spike with delay $|s|$. Experimental results such as those shown in figure 1 suggest that the kernel W defines a learning window for LTP or LTD. A simple choice for W is

$$W(s) = \begin{cases} A_+ e^{-s/\tau_2} & \text{for } s > 0 \\ -A_- e^{s/\tau_2} & \text{for } s < 0 \end{cases} \tag{5.2}$$

with $A_\pm > 0$. If the coefficients $A_\pm$ are constants then equation (5.1) is said to be an additive learning or update rule. Since excitatory synapses are positive and should not exceed some maximum value of, say, $w_j = 1$, it is necessary to impose additional constraints. One way to implement this is to consider a multiplicative rather than additive update rule with weight–dependent coefficients

$$A_- = w_j a_-, \quad A_+ = (1 - w_j)a_+ \tag{5.3}$$

Note that w_j is a step function of time with discontinuities whenever a spike occurs. In order to obtain a well-defined differential equation (5.1), we require that the amplitude of a step depends on the value of w_j immediately prior to a spike. The solution $w_j(t)$ is thus a left–continuous function of time t.

We assume that the input spike trains have stationary statistics with known mean and correlations. More specifically, input spikes at a given synapse j are generated by a doubly stochastic point process. First, spikes are generated by an inhomogeneous Poisson process with instantaneous rate $\nu_j(t)$. Thus,

$$\langle S_j(t)\rangle_E = \nu_j(t)$$

where $\langle\ldots\rangle_E$ denotes expectation with respect to the Poisson process. Second, the rate itself is drawn from a distribution with constant expectation and homogeneous correlations,

$$\langle \nu_j(t)\rangle_T = \nu^{pre}, \quad \langle \nu_j(t)\nu_k(t)\rangle_T = C_{jk}(t-t')$$

where $\langle\ldots\rangle_T$ denotes expectation with respect to the rate distribution. Finally, it is assumed that each presynaptic spike train has identical statistical properties in the sense that the mean input spike rate is j–independent and $\sum_k C_{jk}(t) = C(t)$ for all j. As we have seen in §3, determining the output response of a neuron to stochastic inputs is a non–trivial problem. Here we will assume that the output firing rate $\nu(t) = \langle S(t)\rangle_E$ is linearly related to the input rates according to [**119, 117**]

$$\nu(t) = \sum_j w_j \int_0^\infty \epsilon(s)\nu_j(t-s)ds \tag{5.4}$$

where $\epsilon(s)$ represents the time course of a postsynaptic potential. (Note the implicit assumption that the sum on the right–hand side remains positive). It then follows that

$$\nu^{\text{post}} \equiv \langle \nu(t)\rangle_T = \nu^{\text{pre}} \sum_j w_j.$$

The next major step is to assume that weight changes are small during the time that it is necessary to approximately sample the input statistics. In this case we can separate the time-scale of learning from that of neuronal dynamics. The right–hand side of equation is then self–averaging [**119, 117**] so that it can be approximated by

$$\frac{dw_j}{dt} = c_0 + \int_{-\infty}^{\infty} W(s)\langle\langle S(t)S_j(t-s)\rangle\rangle ds \tag{5.5}$$

where double angular brackets mean averaging with respect to the double stochastic process. In order to combine both integrals in equation (5.1), we have assumed that $\langle\langle S(t)S_j(t-s)\rangle\rangle$ is approximately unchanged if $t-s \to t$ with s varying over the length of the learning window represented by $W(s)$. This is valid for the above linear model, for which it can be shown that [**119**]

$$\langle S(t)S_j(t-s)\rangle_E = \nu(t)\nu_j(t-s) + w_j\epsilon(s)\nu_j(t-s) \tag{5.6}$$

The first term on the right–hand side is the "chance level" of finding two spikes at $t, t-s$ respectively if the neurons were firing independently at rates $\nu(t)$ and $\nu_j(t-s)$. The second term describes the correlation arising from synaptic coupling; if the presynaptic neuron fires at time $t-s$ then the chance that the postsynaptic neuron fires at time t is increased by $w_j\epsilon(t-s)$ provided that $s>0$ (causality). It then follows that

$$\begin{aligned} \langle\langle S(t)S_j(t-s)\rangle\rangle &= \langle\nu(t)\nu_j(t-s)\rangle_T + w_{ij}\epsilon(s)\langle\nu_j(t-s)\rangle_T \\ &= \sum_k w_k(t)\int_0^{\infty}\epsilon(s')C_{jk}(s'-s) + w_j(t)\epsilon(s)\nu^{\rm pre} \end{aligned} \tag{5.7}$$

with $w_k(t)$ a slowly varying function of t. Substituting equation (5.7) into (5.5) then gives (after some rearrangement)

$$\frac{dw_j}{dt} = c_0 + W_0\nu^{\rm post}\nu^{\rm pre} + \sum_k w_k(t)\left[Q_{jk} + \delta_{jk}W_+\nu^{\rm pre}\right] \tag{5.8}$$

where

$$W_0 = \int_{-\infty}^{\infty} W(s)ds, \quad W_+ = \int_0^{\infty}\epsilon(s)W(s)ds \tag{5.9}$$

and

$$Q_{jk} = \int_{-\infty}^{\infty} dsW(s)\int_0^{\infty} ds'\epsilon(s')C^0_{jk}(s'-s) \tag{5.10}$$

with $C^0_{jk}(s)$ the input covariance function

$$C^0_{jk}(s) = \langle[\nu_j(t)-\nu^{\rm pre}][\nu_k(t-s)-\nu^{\rm pre}]\rangle_T \tag{5.11}$$

In order to gain further insight into the nature of the above Hebbian–like learning rule, it is useful to rewrite equation (5.8) in the equivalent form

$$\begin{aligned} \frac{dw_j}{dt} &= c_0 + \gamma\nu^{\rm post} \\ &\quad + \sum_k w_k(t)\left[Q_{jk} - \bar{Q}\right] + W_+\left[w_j(t)\nu^{\rm pre} - \frac{\nu^{\rm post}}{N}\right] \end{aligned} \tag{5.12}$$

where

$$\gamma = W_0\nu^{\rm pre} + N^{-1}W_+ + \bar{Q}/\nu^{\rm post} \tag{5.13}$$

and $\bar{Q} = N^{-1} \sum_k Q_{jk}$ with Q_{jk} given by equation (5.10). Suppose that the update rule is additive, that is, $W(s)$ is independent of the weights w_j. Multiplying both sides of equation (5.12) by $\nu^{\rm pre}$ and summing over j, we find that the terms on the second line of the right–hand side of equation (5.12) vanish and

$$\frac{d\nu^{\rm post}}{dt} = \gamma(\nu^{\rm post} - \nu^*) \tag{5.14}$$

with $\nu^* = -c_0/\gamma$. It follows that the mean output activity converges to the fixed point ν^* provided that $\gamma < 0$ and, hence, ilustrates how STDP can stabilize network activity. Since firing rates can only be positive, we also require $c_0 > 0$. Once it has converged to ν^*, the total sum of weights $\sum_j w_j(t)$ remains constant (a form of subtractive normalization [**120**]), and the terms on the first line of the right–hand side of equation (5.12) cancel. The remaining terms then determine the evolution of the weight vector $\mathbf{w}$ according to

$$\frac{d\mathbf{w}}{dt} = [\mathbf{Q} - \bar{\mathbf{Q}}]\mathbf{w} + W_+[\nu^{\rm pre}\mathbf{w} - N^{-1}\nu^*\mathbf{n}] \tag{5.15}$$

where $\bar{\mathbf{Q}}$ is the matrix with all its elements equal to $\bar{Q}$, and $\mathbf{n} = (1, 1, \ldots, 1)$. Dropping the term $N^{-1}\nu^*\mathbf{n}$ for large N, we see that the dynamics of $\mathbf{w}$ is dominated by the eigenvector of the matrix $\mathbf{Q} - \bar{Q}\mathbf{1}$ with the largest eigenvalue. In other words, the neuron is carrying out a form of statistical feature extraction along the lines of singular value decomposition. This form of correlation–based learning, when combined with spatially distributed synaptic interactions, has been used to model the development of feature maps in cortex.

Development of ocular dominance columns in V1

The primary visual cortex (V1) is characterized by a number of spatially distributed feature maps, in which local populations of neurons respond preferentially to stimuli with particular properties such as orientation and spatial frequency. Neurons also tend to respond more strongly to stimuli presented in one eye rather than the other, that is, they exhibit ocular dominance. Neurons sharing the same ocular dominance are grouped together into non-overlapping regions that form an alternating pattern of right and left eye preference across V1. Such regions have a characteristic periodicity and morphology that is species–dependent. For example, in the adult macaque monkey ocular dominance regions consist of branching stripes that have an approximately uniform width of $0.4mm$ whereas in cat they are more patchy. An example of the ocular dominance pattern in macaque is shown in figure 2.

Let $n_L(\mathbf{r})$ and $n_R(\mathbf{r})$ denote the synaptic densities of feedforward afferents from the left and right eyes to a point $\mathbf{r} \in \mathbf{R}^2$ in cortex. (For simplicity, we treat cortex as an unbounded two–dimensional sheet). Suppose that these densities evolve according to a spatially extended version of Hebbian learning with subtractive normalization:

$$\frac{\partial n_i(\mathbf{r}, t)}{\partial t} = \int_{\mathbf{R}^2} w(|\mathbf{r} - \mathbf{r}'|) \left[\sum_{j=L,R} Q_{ij} n_j(\mathbf{r}', t) - \mu \sum_j n_j(\mathbf{r}') \right] d\mathbf{r}' \tag{5.16}$$

Figure 2. Reconstruction of ocular dominance columns in primary visual cortex (V1) of macaque monkey shown in tangential section. Regions receiving input from one eye are shaded black and regions receiving input from the other eye are unshaded.

where $\mu = [Q_S + Q_D]/2$. Equation (5.16) is supplemented by the saturation constraints $0 \leq n_L(\mathbf{r}), n_R(\mathbf{r}) \leq N$. The term in square brackets represents the basic Hebbian learning rule (after averaging with respect to left and right eye inputs). The matrix $\mathbf{Q}$ specifies same–eye ($Q_{LL} = Q_{RR} = Q_S$) and opposite–eye ($Q_{LR} = Q_{RL} = Q_D$) input correlations with $Q_S > Q_D \geq 0$. The subtractive normalization term involving the constant factor μ ensures that there is competition between left and right eye afferents at each point in cortex. Finally, the function w represents the effects of synaptic interactions in cortex, which are assumed to be in the form of a mexican hat function (short–range excitation, long–range inhibition). These interactions generate a pattern forming instability that produces a spatial ordering of left/right eye competition, resulting in ocular dominance columns. In the following, we will explore the pattern formation of ocular dominace columns in more detail.

EXERCISE

Calculate the eigenmodes of the linear system of equations (5.16). Show that $n_L(\mathbf{r}) + n_R(\mathbf{r})$ is conserved pointwise in cortex. Determine the fastest growing eigenmodes and show that the resulting patterns are consistent with the formation of two–dimensional ocular dominance columns. What does linear theory predict for the width of an ocular dominance column?

Distribution of synaptic weights

In the above analysis it was assumed that the weights evolve much more slowly than typical presynaptic and postsynaptic interspike intervals, so that statistical fluctuations in the weight dynamics can be ignored. However, if the synaptic weights can be changed significantly by only a few spikes then such fluctuations have to

be taken into account. One approach is to determine the steady state distribution of synaptic weights within the framework of a Fokker–Planck equation [**121, 118, 122**]. For the sake of simplicity, we consider a learning window with two rectangular regions so that the kernerl $W(s)$ is

$$W(s) = \begin{cases} A_+ & \text{for } 0 < s < d \\ -A_- & \text{for } -d < s < 0 \\ 0 & \text{otherwise} \end{cases} \tag{5.17}$$

The weights w_j are restricted to lie in the interval $[0, 1]$ by imposing either hard or soft bounds. Hard bounds means that the weights are simply nolonger increased (decreased) if the upper (lower) bound is reached. Soft bounds, on the other hand, gradually slow down the dynamics if a weight approaches one of its bounds. The latter can be implemented using the multiplicative rule (5.3).

Let $P(w,t)$ be the probability density for the population of synapses (assumed large) with normalization $\int_0^1 P(w,t)dw = 1$. Each presynaptic neuron fires independently at a constant rate ν^{pre} (we are now considering homogeneous Poisson processes). The evolution of $P(w,t)$ is then given by the master equation

$$\begin{aligned} \frac{\partial}{\partial t}P(w,t) &= -p_+(w,t)P(w,t) - p_-(w,t)P(w,t) \\ &\quad + \int_0^1 \delta(w - w' - A_+(w'))p_+(w',t)P(w',t)dw' \\ &\quad + \int_0^1 \delta(w - w' + A_-(w'))p_-(w',t)P(w',t)dw' \end{aligned} \tag{5.18}$$

Here $p_+(w,t)$ ($p_-(w,t)$) is the probability per unit time that the neuron fired at time t and a synapse with weight $w(t) = w$ received a presynaptic spike falling in the positive (negative) part of the learning window centered about t. Thus

$$\begin{aligned} p_+(w,t) &= \int_0^d \langle S(t)S_j(t-s)\rangle_E ds \\ p_+(w,t) &= \int_{-d}^0 \langle S(t)S_j(t-s)\rangle_E ds \end{aligned} \tag{5.19}$$

It can be shown that [**118, 123**]

$$\langle S(t)S_j(t-s)\rangle_E = \nu^{\text{post}}\nu^{\text{pre}}(1 + G(s)w(t)) \tag{5.20}$$

for some model–dependent $G(s)$ with $G(s) = 0$ for $s < 0$.

The master equation can be reduced to a Fokker–Planck equation by Taylor expanding to second order in the jump transitions $A_\pm$ along similar lines to §3 (assuming jumps are sufficiently small). This gives

$$\frac{\partial}{\partial t}P(w,t) = -\frac{\partial}{\partial w}[A(w)P(w,t)] + \frac{\partial^2}{\partial w^2}[B(w)P(w,t)] \tag{5.21}$$

with

$$\begin{aligned} A(w,t) &= p_+(w,t)A_+(w) - p_-(w,t)A_-(w) \\ B(w,t) &= p_+(w,t)A_+^2(w) - p_-(w,t)A_-^2(w) \end{aligned} \tag{5.22}$$

The Fokker–Planck equation can be solved numerically to find stationary solutions. One finds that the qualitative form of the distribution depends critically on how the bounds for the weights are implemented [**118, 122**]. With soft bounds the

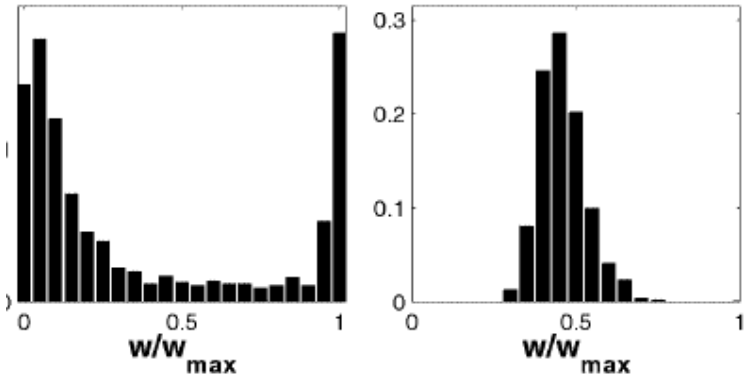

Figure 3. Stationary distribution of synaptic weights with (a) hard bounds and (b) soft bounds

distribution is unimodal whereas with hard bounds it is bimodal with peaks at the boundaries of the weight interval, see figure 3. We see that the use of hard bounds introduces competition between weights but is less stable, since weights tend to saturate at either boundary.

5.2. Biophysical model of calcium controlled bidirectional plasticity

Calcium has long been suggested as a major signaling agent for LTP and LTD [**124, 125, 110, 126**], and more recently for STDP [**116**]. As illustrated in figure 4, calcium can enter the cell through channels controlled by NMDA receptors. Two conditions are necessary for the opening of such channels: (i) arrival of an action potential at the presynaptic terminal releases the excitatory neurotransmitter glutamate which then binds to NMDA receptors and (ii) the postsynaptic cell is sufficiently depolarized so that magnesium ions are unblocked from the NMDA channels allowing the influx of calcium. One source of strong depolarization is the back propagation of an action potential into the dendritic tree, which is a signature that the postsynaptic neuron has fired a spike. Given that calcium ions can enter the cell only if glutamate has been released by presynaptic activity and if the postsynaptic membrane is sufficiently depolarized, it follows that the NMDA receptor is a natural candidate for the biophysical implementation of a Hebbian–like mechanism for detecting correlations between presynaptic and postsynaptic activity. In this section we describe a biophysical model that consists of two essential components: a model of calcium entry through NMDA synapses and a model of how intracellular calcium concentration modifies synaptic efficacy [**127**].

NMDA receptors and calcium influx

Suppose that a presynaptic spike arrives at an NMDA synapse at time T and the postsynaptic membrane potential is $u(t)$. A simple model for the calcium current through an NMDA receptor–controlled channel is

$$I_{\rm Ca}(t) = \bar{g}_{\rm Ca}\alpha(t-T)[u(t)-u_{\rm Ca}]B(u(t)) \tag{5.23}$$

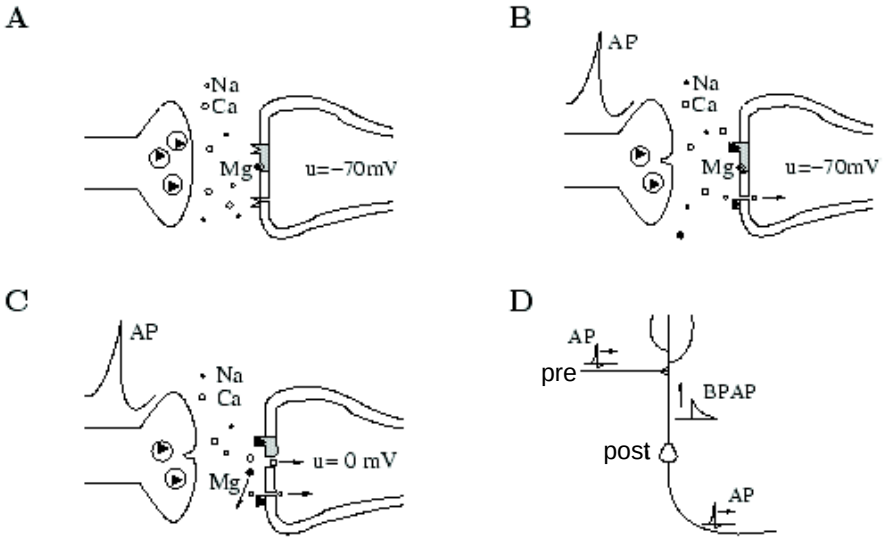

Figure 4. Different states of an NMDA synapse. A) Inactive state. B). Arrival of an action potential (AP) at the presynaptic terminal induces vesicles containing glutamate to merge with the cell membrane and releasethe neurotransmitter into the synaptic cleft. This binds to NMDA and AMPA receptors on the postsynaptic cell membrane. The AMPA associated channel opens but the NMDA channel is blocked by magnesium. C) Depolarization of the postsynaptic membrane unblocks magnesium allowing the influx of calcium. D) Depolarization may be caused by a back propagating action potential (BAP) that travels up the dendrites. [Taken from Gerstner and Kistler [**2**]]

where $\bar{g}_{\mathrm{Ca}}$ is the maximal conductance of the channel and u_{Ca} is the reversal potential of calcium. The function $\alpha(t-T)$ describes the time course of binding to NMDA receptors and is typically taken to be the sum of a fast ($\tau_f = 50ms$) and a slow ($\tau_s = 200ms$) exponential:

$$\alpha(t) = I_f \mathrm{e}^{-t/\tau_f} + I_s \mathrm{e}^{-t/\tau_s}, \quad t > 0 \tag{5.24}$$

The function

$$B(u) = \frac{1}{1 + \eta[Mg^{2+}]\mathrm{e}^{-\gamma u}} \tag{5.25}$$

describes the unblocking of the channel at depolarized levels of membrane potential [**128**]. If there are several presynaptic spikes within 100ms then calcium accumulates inside the cell. The change in the calcium concentration $[\mathrm{Ca}^{2+}]$ can be described by the simple first order kinetic equation

$$\frac{d[\mathrm{Ca}^{2+}]}{dt} = I_{\mathrm{Ca}}(t) - \frac{[\mathrm{Ca}^{2+}]}{\tau_{\mathrm{Ca}}} \tag{5.26}$$

with $\tau_{\mathrm{Ca}} = 125ms$.

Calcium control hypothesis

Although the dynamics of NMDA synapses is reasonably well understood, much less is known about the biochemical signaling pathways that are triggered by calcium

influx and lead to persistent changes in synaptic efficacy. (Recent experimental evidence suggests that and important component of LTP and LTD is a change in the number of AMPA receptors in the postsynaptic membrane through receptor trafficking [**129**], see below). Instead of developing a detailed model of such processes, Shouval *et al* [**127**] adopted a phenomenological approach and assumed that the change in synaptic weight is fully determined by the intracellular calcium concentration. Under this so–called "calcium control hypothesis", the synaptic weight is taken to vary according to the simple first order scheme

$$\frac{dw}{dt} = \eta([\mathrm{Ca}])\,[\Omega([\mathrm{Ca}]) - w] \tag{5.27}$$

For constant calcium concentration, the weight w converges to the asymptotic value $\Omega([\mathrm{Ca}])$ at the rate $\eta([\mathrm{Ca}])$. Shouval *et al* [**127**] assume a particular form for the

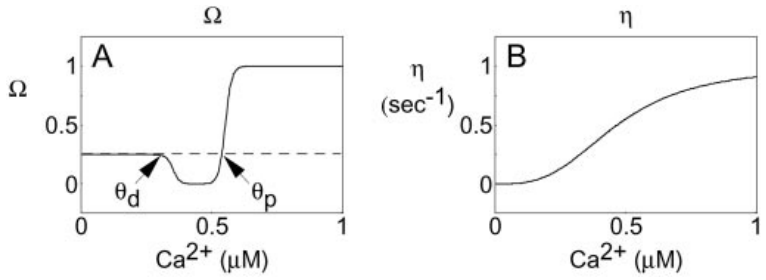

Figure 5. Calcium control hypothesis. The asymptotic weight value $\Omega([\mathrm{Ca}])$ and rate constant $\eta([\mathrm{Ca}])$ are shown as a function of calcium concentration [Ca]. Adapted from Shouval *et al* [**127**]

functions $\Omega([\mathrm{Ca}])$ and $\eta([\mathrm{Ca}])$, as shown in figure 5. For calcium concentrations below a lower threshold θ_d, the fixed point Ω assumes a resting value of $\Omega = 0.5$ (in appropriate units). For calcium concentrations in the range $\theta_d \leq [\mathrm{Ca}] \leq \theta_p$, the fixed point tends to decrease, whereas for calcium concentrations above the upper threshold θ_p, it increases. The rate constant η is taken to be a monotonically increasing function of calcium concentration.

Regulation of AMPA receptors and the calcium hypothesis

In figure 6(a) is shown an idealized model for the cycle of phosphorylation – dephosphorylation at two sites on the GluR1 subunit of an AMPA receptor. The model assumes two specific kinases (EK1, EK2) and two opposing specific phosphatases (EP1, EP2). It is assumed that high–frequency stimulation preferentially stimulates the activity of protein kinases, resulting in GluR1 phosphorylation, whereas low-frequency stimulation preferentially stimulates the activity of protein phosphatases, resulting in GluR1 dephosphorylation. The concentration of AMPA receptors containing GluR1 phosphorylated at sites S831 and S845 are denoted by A_{p1} and A^{p2} respectively. The corresponding concentration of receptors with phosphorylation at neither or both sites are denoted by A and A^{p2}_{p1}. $EP1$ and $EP2$ denote the rates of dephosphorylation at the two sites and $EK1, EK2$ denote the rates of phosphorylation. These rates are taken to be calcium–dependent. Finally, the effective weight of a synapse is assumed to be proportional to the linear combination

$W \sim A + 2[A_{p1} + A^{p2}] + 4A_{p1}^{p2}$. The synaptic plasticity equation can then be obtained from the kinetic rate equations associated with figure 6(a). However, There is growing evidence that AMPA receptor trafficking rather than phosphorylation underlies calcium mediated changes in synaptic strength. A simple model for the former is to assume two population of AMPA receptors, one (A_m) is inserted in the postsynaptic membrane and the other (A_I) is internal to the postsynaptic compartment, see figure 6(b). Denoting the calcium–dependent rate constants for insertion and removal by K_I and K_R respectively, we have the first order kinetics

$$\dot{A}_m = -K_R A_m + K_I A_I, \quad \dot{A}_I = K_R A_m - K_I A_I \tag{5.28}$$

Assuming conservation of the total number of receptors, $A_T = A_m(t) + A_I(t)$, and identifying synaptic strength as $W = \beta A_m$ where β is a proportionality constant, it is straightforward to derive the synaptic plasticity equation (5.27) assumed in the calcium hypothesis.

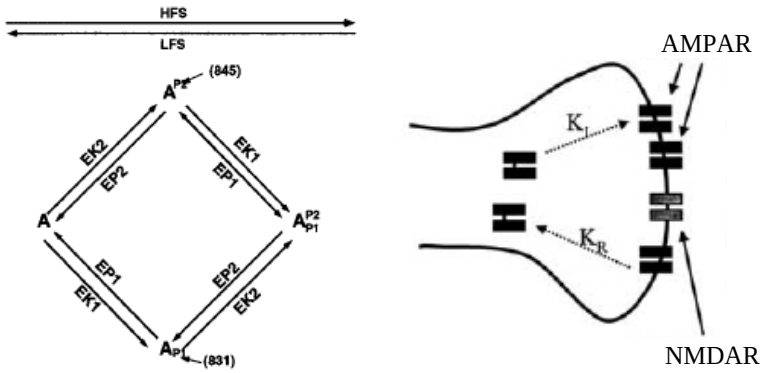

Figure 6. (a) Phosphorylation–dephosphorylation cycle for GluR1 subunit of an AMPA receptor with HFS (LFS) denoting bias under high (low) frequency stimulation (b) AMPA receptor trafficking

EXERCISE

Using the law of mass action, write down the set of first–order ODEs describing the evolution of the four receptor concentrations in figure 6(a). Assuming conservation of total amount of AMPA receptors, $A_T = A(t)+A_{p1}(t)+A^{p2}(t)+A_{p1}^{p2}(t)$, determine the equilibrium solutions for the concentrations as fractions of A_T. This then generates the fixed point Ω of the synaptic plasticity equation 5.27). How would one determine the rate of convergence η? Show how the pair of kinetic equations (5.28) in the recptor trafficking model generate equation (5.27).

Back propagating action potentials (BAPs) and STDP

In order to completely specify the model, it is necessary to determine the dynamics of the postsynaptic membrane potential $u(t)$. In the case of STDP, the main source of depolarization is thought to arise from a back propagating action potential (although this might not be the whole story [**130**]) . This can be modeled by

taking

$$u(t) = \eta(t - \widehat{T}), \quad \eta(t) = U\left(\widehat{I}_s e^{-t/\hat{\tau}_s} + \widehat{I}_f e^{-t/\hat{\tau}_f}\right) \tag{5.29}$$

where $\widehat{T}$ is the last firing time of the postsynaptic cell and $\eta(t)$ is the profile of the BAP, which is taken to be the sum of a fast and a slow exponential. The various components of the model given by equations (5.23), (5.26), (5.27) and (5.29) can now be combined to show how calcium influx leads to STDP [**127**]. The basic idea is that the relative timing of presynaptic and postsynaptic spikes alters the maximum level of calcium influx, which is proportional to the product of the NMDA glutamate binding factor α and the magnesium unblocking factor $B(u)$, see equation (5.23). The details are explained in figure 7. One potential problem with the model is that it predicts that LTD can also be induced by a sequence of pre–post spikes if the separation between spikes is larger than about $\Delta t = 40ms$. The reason is that in this case the removal of the magnesium block (induced by the BAP) occurs at a time when the probability of glutamate binding is reduced. As a consequence less calcium enters the cell so that only the LTD threshold is reached. There is not convincing experimental evidence for such a form of LTD. This has motivated a more detailed biophysical model in which the time course of the calcium current rather than just its level acts as signal for STDP [**131**].

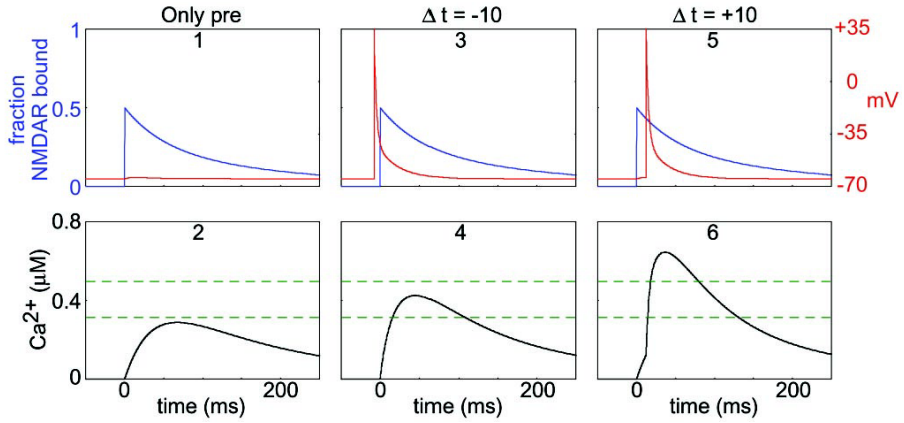

Figure 7. (1,2) Presynaptic stimulation alone results in the binding of glutamate to NMDA receptors but only a small depolarization so that there is only moderate calcium influx. (3,4) Post–pre stimulation ($\Delta t = -10ms$) results in a large brief depolarization due to a BAP, which partially overlaps with the onset of glutamate binding to NMDA receptors. The calcium influx is increased so that it crosses the threshold θ_d for LTD and there is a negative weight change. (5,6) Pre–post stimuluation ($\Delta t = +10ms$) results in a large brief depolarization that completely overlaps with the period of significant glutamate binding to NMDA receptors. The resulting calcium influx is well above the threshold θ_p for LTP and there is a positive weight change. Since the rate constant $\gamma([Ca^{2+}])$ is larger in the regime of LTP induction, the positive weight change is dominant even though the calcium concentration must pass through the LTD regime in order to reach the threshold θ_p. [Adapted from Shouval *et al* [**127**]]

5.3. Computational role of STDP

In this final section we briefly describe one potential computational application of STDP to prediction and reward learning [**132**]. Animals in the environment often have to react quickly to the earliest signs of harmful stimuli or potential prey. That is, they have to predict or anticipate future events. STDP provides a hint at how a simple form of predictive coding could be implemented at the cellular level. Consider as an example a single neuron receiving inputs from a set of N

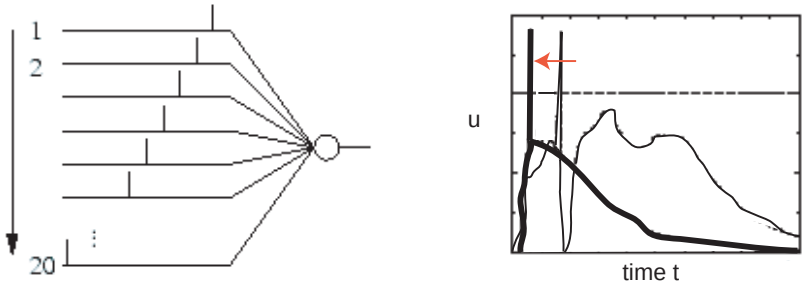

Figure 8. A) A sequence of presynaptic spikes is presented to a single neuron. Such a sequence could be generated by a moving stimulus. B) After many trials the synapses in the early part of the sequence have been strengthened so that the neuron fires earlier by an amount Δt. Initial response is shown by thin lines and adapted response by thick lines

presynaptic cells that fire sequentially (see figure 8). This could be the result of a moving stimulus. Initially all synapses have the same weight $w_j = w_0$ for all $j = 1, \ldots N$. Suppose that the postsynaptic neuron initially fires at time T during the presentation of the sequence. All synapses that have been activated prior to the postsynaptic spike are strengthened while synapses that have been activated immediately afterwards are depressed. After many trials, the firing time of the postsynaptic neuron within the sequence has shifted by an amount $-\Delta t$. Interestingly, this effective shift in the distribution of synaptic connections is consistent with experimental observations of a shift in the place fields of hippocampal neurons [**133**].

The shift of responses towards earlier predictors plays a central role in classical conditioning as exemplified by Pavlov's dog. An unconditioned simulus (US) – food – is preceded by a conditioned stimulus (CS) – a bell – at a fixed time interval ΔT. Before learning the US evokes an immediate response – salivation – wheres the CS evokes no response. After learning, the dog starts to salivate in response to the CS. Thus the reaction has moved from the US to the CS, which reliably predicts the US. STDP can replicate this result if the time difference ΔT between two stimuli is less than the width of the learning window. The mechanism is identical to the previous example illustrated in figure 8, except that the inputs are clustered into two groups corresponding to US and CS. A shift to early predictors is also observed experimentally in recordings from dopamine neurons in the basal ganglia concerned with reward processing [**135**], see figure 9. Before learning some dopamine neurons respond to an unpredicted reward. During learning this response decreases and

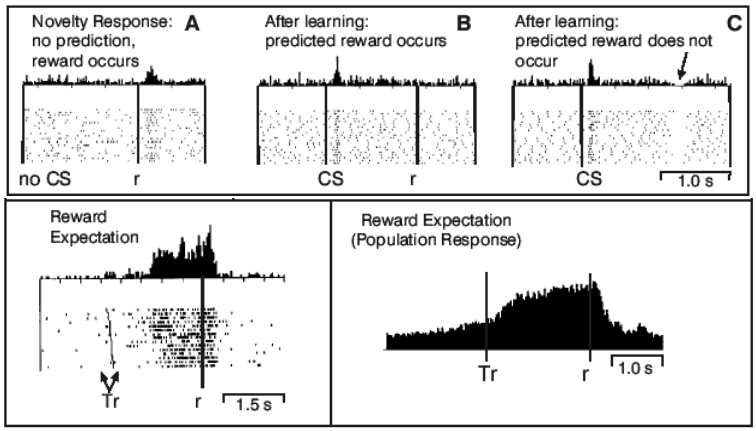

Figure 9. Neuronal response in the basal ganglia concerened with reward processing. A) Response of a dopamine neuron to the presentation of a reward (r) without a preceding CS. B) Response of same neuron after learning that the CS will predict the reward. C) If the expected reward fails to be delivered, the neuron is inhibited. D) Response of putamen neuron, which gradually increases and maintains its firing after a trigger (TR) until the reward is delivered. E) Population response of striatal neurons further illustrating expectation of reward response. [Adapted from [**134**]]

the neurons now increase their firing to a reward-predicting stimulus. The more predictable the reward becomes, the less strongly the neurons fire when it appears, and if it does not appear the neurons are inhibited. There are in fact a number of different types of reward–predicting dopamine neurons with distinct responses. For example, some neurons fire a prolonged discharge following a trigger stimulus until a reward is delivered. Dopamine neurons have a central role in guiding our behaviour and thoughts [**136**]. Their performance is impaired by every addictive drug and in mental illness. They are also lost in dramatically impairing illnesses such as Parkinsons disease. If dopamine systems are overstimulated, we may hear voices, experience elaborate bizarre cognitive distortions, or engage excessively in dangerous goal-directed behaviour. Dopamine function is also central to the way that we value our world, including the way that we value money and other human beings.

Temporal difference learning

It has been suggested that reward–predicting dopamine neurons contribute to some form of temporal difference (TD) learning [**136, 132**]. The simplest example of TD–learning arises in the problem of predicting some scalar quantity z using a single neuron with synaptic weights $\mathbf{w} = (w_1, \dots w_k)$. The presynaptic input is a sequence of vectors $\mathbf{x}_1, \dots \mathbf{x}_m$ and the output at discrete time step t is given by $P_t = \mathbf{w} \cdot \mathbf{x}_t$. The goal is to learn a set of weights such that the prediction P_t is as close as possible to the taget z. The weights are updated according to

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \lambda(P_{t+1} - P_t)\mathbf{x}_t \tag{5.30}$$

where λ is a learning rate and $P_{m+1} = z$. To motivate this TD–learning rule, suppose that the weights are initially zero so that $P_t = 0$ for all t. However, in the

final step $t = m$, there is a non–zero prediction error $P_{m+1} - P_m = z$, so that the weights are changed by an amount equal to $\lambda z \mathbf{x}_t$. Thus in the next trial P_m will be closer to z than before, and after several trials will tend to converge to z. A particularly interesting feature of the learning algorithm is that P_t acts as a training signal for P_{t-1} so that information about the target z is propagated backwards in time such that the predictions are corrected over many trials and will eventually converge to z. One way to interpret z is to view it as the reward delivered to an animal at the end of a trial. This idea can be generalized by assuming that a reward r_t can be delivered at each time–step. This leads to the standard model of TD–learning proposed by Sutton and Barto [**137**], in which P_t now predicts $\sum_{i>t} r_i$. Ideally

$$P_t = \sum_{i>t}^{m} r_i = r_{t+1} + \sum_{i>t+1}^{m} r_i = r_{t+1} + P_{t+1}$$

This suggests a prediction error of the form $\delta_t = r_{t+1} + P_{t+1} - P_t$ and a learning rule in which weights are updated to minimize this error

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \lambda(r_{t+1} + P_{t+1} - P_t)\mathbf{x}_t \tag{5.31}$$

Finally, note that the response illustrated in figure 9A-C suggests that the dopamine neuron is encoding a prediction error signal analogous to δ_t.

BIBLIOGRAPHY

1. P Dayan and L F Abbott. *Theoretical neuroscience*. MIT, Cambridge, MA, 2001.
2. W Gerstner and W Kistler. *Spiking neuron models*. Cambridge University Press, Cambridge, 2002.
3. E Izhekevich. *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*. MIT Press, Cambridge, MA, 2005.
4. C Koch. *Biophysics of Computation*. Oxford University Press, 1999.
5. J P Keener, F C Hoppensteadt, and J Rinzel. Integrate–and–fire models of nerve membrane response to oscillatory input. *SIAM J. Appl. Math.*, 41(3):503–517, 1981.
6. S Coombes and P C Bressloff. Mode-locking and arnold tongues in integrate-and-fire neural oscillators. *Phys. Rev. E*, 60(2):2086–2096, 1999; *ibid* 63: 059901, 2001.
7. J Guckenheimer and P J Holmes. *Nonlinear Oscillations, Dynamical Systems and Bifurcations of Vector Fields*. Springer-Verlag, New York, 1983.
8. W Gerstner and J L van Hemmen. Coding and information processing in neural networks. In E Domany, J L van Hemmen, and K Schulten, editors, *Models of Neural Networks II*, pages 1–93. Springer-Verlag, 1994.
9. W Gerstner and W Kistler. Population dynamics of spiking neurons: fast transients, asynchronous states and locking. *Neural Comput.*, 12:43–89, 2000.
10. W M Kistler, W Gerstner, and J L van Hemmen. Reduction of the Hodgkin-Huxley equations to a single-variable threshold. *Neural Comput.*, 9:1015–1045, 1997.
11. L F Abbott and C van Vresswijk. Asynchronous states in networks of pulse–coupled oscillators. *Phys. Rev. E*, 48(2):1483–1490, 1993.
12. G B Ermentrout and N Kopell. Parabolic bursting in an excitable system coupled with a slow oscillation. *SIAM J. of Appl. Math.*, 46:233–253, 1986.
13. G B Ermentrout. Type I membranes, phase resetting curves, and synchrony. *Neural Comput.*, 8:979–1001, 1996.
14. B Gutkin and G B Ermentrout. Dynamics of membrane excitability determine interspike interval variability; a link between spike generation mechanisms and cortical spike train statistics. *Neural Comput.*, 10:1047–1065, 1998.
15. X.J. Wang and J. Rinzel. Oscillatory and bursting properties of neurons. In M. Arbib, editor, *Handbook of Brain Theory and Neural Networks*. MIT Press,

1995.

16. F C Hoppensteadt and E Izhekevich. *Weakly Connected Neural Nets*. Springer-Verlag, New York, 1997.
17. Xiao-Jing Wang and John Rinzel. Alternating and synchronous rhythms in reciprocally inhibitory model neurons. *Neural Comput.*, 4:84–97, 1992.
18. G D Smith, C L Cox, S M Sherman, and J Rinzel. Fourier analysis of sinusoidally driven thalamocortical relay neurons and a minimal integrate-and-burst model. *J. Neurophysiol.*, 83:588–610, 2000.
19. A R R Casti et al. A population study of integrate–and–fire–or burst neurons. *Neural Comput.*, 14:957–986, 2002.
20. S Coombes, M R Owen, and G D Smith. Mode-locking in a periodically forced integrate-and-fire-or-burst neuron model. *Physical Review E*, 64:041914, 2001.
21. P C Bressloff and S Coombes. Physics of the extended neuron. *Int. J. Mod. Phys. B*, 11(20):2343–2392, 1997.
22. A Destexhe, Z F Mainen, and T J Sejnowski. Synthesis of models for excitable membranes synaptic transmission and neuromodulation using a common kinetic formalism. *J. Comput. Neurosci.*, 1:195–231, 1994.
23. A Fogelson and R S Zucker. Presynaptic calcium diffusion from various arrays of single channels. implications for transmitter release and synaptic facilitation. *Biophys. J.*, 48:1003–1017, 1985.
24. R Bertram, G D Smith, and A Sherman. Modeling study of the effects of overlapping calcium microdomains on neurotransmitter release. *Biophys. J.*, 76:735–750, 1999.
25. J H Byrne and J L Roberts. *From molecules to networks: an introduction to cellular and molecular neuroscience*. Elsevier, Amsterdam, 2004.
26. J Trommershauser, J Marienhagen, and A Zippelius. Stochastic model of central synapses: slow diffusion of transmitter interacting with spatially distributed receptors and transporters. *J. Th. Biolo.*, 198:101–120, 1999.
27. H Markram and M Tsodyks. Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature*, 382:807–810, 1996.
28. L F Abbott, J A Varela, K Sen, and S B Nelson. Synaptic depression and cortical gain control. *Science*, 275:220–224, 1997.
29. L F Abbott and E Marder. Modelling small networks. In C Koch and I Segev, editors, *Methods in Neuronal Modelling*, pages 361–410. MIT Press, 2nd edition, 1998.
30. R S Zucker. Short term synaptic plasticity. *Ann. Rev. Neurosci.*, 12:13–31, 1989.
31. G L Millhauser, E E Salpeter, and R E Oswald. Diffusion model of ion–channel gating and the origin of power–law distributions from single–channel recordings. *Proc. Natl. Acad. Sci. USA*, 85:1503–1507, 1988.
32. S Fusi, P J drew, and L F Abbott. cascade models of synaptically stored memories. *Neuron*, 45:599–611, 2005.
33. G Gilboa, R Chen, and N Brenner. History–dependent multiple-time-scale dynamics in a single-neuron model. *J. Neurosci.*, 25:6479–6489, 2005.
34. H C Tuckwell. *Introduction to Theoretical Neurobiology*, volume 1. (Linear cable theory and dendritic structure) of *Cambridge Studies in Mathematical Biology*. Cambridge University Press, 1988.

35. G B Ermentrout. The behaviour of rings of coupled oscillators. *Journal of Mathematical Biology*, 23:55–74, 1985.
36. C van Vreeswijk, G B Ermentrout, and L F Abbott. When inhibition not excitation synchronizes neural firing. *J. Comput. Neurosci.*, 1:313–321, 1994.
37. R E Mirollo and S H Strogatz. Synchronisation of pulse–coupled biological oscillators. *SIAM J. Appl. Math.*, 50(6):1645–1662, 1990.
38. N Kopell and G B Ermentrout. Symmetry and phase-locking in chains of weakly coupled oscillators. *Comm. Pure Appl. Math.*, 39:623–660, 1986.
39. G B Ermentrout and N Kopell. Multiple pulse interactions and averaging in systems of coupled neural oscillators. *J. Math Biol.*, 29:195–217, 1991.
40. D Hansel, G Mato, and C Meunier. Synchrony in excitatory neural networks. *Neural Comput.*, 7:2307–2337, 1995.
41. P C Bressloff aand S Coombes. Dynamics of strongly coupled spiking neurons. *Neural Comput.*, 12:91–129, 2000.
42. P C Bressloff aand S Coombes. Dynamical theory of spike train dynamics in networks of integrate-and-fire oscillators. *SIAM J. Appl. Math*, 60:828–841, 2000.
43. W Bair, C Koch, W T Newsome, and K H Britten. Power spectrum analysis of bursting cells in area mt in the behaving monkey. *J. Neurosci.*, 14:2870–2893, 1994.
44. W R Softky and C Koch. Cortical cell should spike regularly but do not. *Neural Comput*, 4:643–646, 1992.
45. F Rieke, D Warland, R van Steveninck, and W Bialek. *Spikes – exploring the neural code*. MIT Press, Cambridge, MA, 1997.
46. M N Shadlen and W T Newsome. Noise, neural codes and cortical organization. *Curr. Opin. Neurobiology*, 4:569–579, 1994.
47. S Thorpe, D Fize, and C Marlot. Speed of processing in the human visual system. *Nature*, 381:520–522, 1996.
48. M Rossumn, G G Turrigiano, and S B Nelson. Fast propagation of firing rates through layered networks of noisy neurons. *J. Neurosci.*, 22:1956–1966, 2002.
49. M Tsodyks and T J Sejnowski. Rapid switching in balanced cortical network models. *Network Comput. Neural Syst.*, 6:111–124, 1995.
50. G B Ermentrout. Neural networks as spatio-temporal pattern-forming systems. *Rep. Prog. Phy.*, 61:353–430, 1998.
51. P C Bressloff. Pattern formation in visual cortex. In C C Chow, B Gutkin, D Hansel, C Meunier, and J Dalibard, editors, *Les Houches 2003: methods and models in neurophysics*, pages 477–574. Elsevier, 2005.
52. J O'Keefe and M L Recce. Phase relationship between hippocampal place units and the eeg theta rhythm. *Hippocampus*, 3:317–330, 1993.
53. G Laurent. Odor encoding as an active, dynamical process: experiments, computation and theory. *Ann. Rev. Neurosci.*, 24:263–297, 2001.
54. K R Delaney, A Galperin, M S Fee, J A Flores, R Gervais, D W Tank, and D Kleinfeld. Waves and stimulus-modulated dynamics in an oscillating olfactory network. *Proc. Natl. Acad. Sci. USA*, 91:669–673, 1994.
55. C M Gray. Synchronous oscillations in neuronal systems: mechanisms and functions. *J. Comput. Neurosci.*, 1:11–38, 1994.
56. Y Dan, J M Alonso snd W M Usrey, and R C Reid. Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nature*

Neurosci., 1:501–507, 1998.
57. R C DeCharms and M Merzenich. Primary cortical representation of sounds by the coordination of action potential timing. *Nature*, 381:610–613, 1995.
58. W Singer and C M Gray. Visual feature integration and the temporal correlation hypothesis. *Ann. Rev. Neurosci.*, 18:555–586, 1995.
59. M N Shadlen and J A Movshon. Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron*, 24:67–77, 1999.
60. C M Gray. The temporal correlation hypothesis of visual feature integration: still alive and well. *Neuron*, 24:31–47, 1999.
61. E Salinas and T J Sejnowski. Correlated neuronal activity and the flow of neural information. *Nature Rev. Neurosci.*, 4:539–550, 2001.
62. H C Tuckwell. *Introduction to theoretical neurobiology*, volume 2. (Nonlinear and stochastic theories) of *Cambridge studies in Mathematical Biology*. Cambridge University Press, 1988.
63. van Kampen. *Stochastic processes in physics and chemistry*. North-Holland, Amsterdam, 1992.
64. C W Gardiner. *Handbook of stochastic methods, 2nd edition*. Springer, Berlin, 1997.
65. M J E Richardson. Effects of synaptic conductance on the voltage distribution and firing rate of spiking neurons. *Phys. Rev. E*, 69:051918, 2004.
66. A Omurtag, B W Knight, and L Sirovich. On the simulation of large populations of neurons. *J. Comput. Neurosci.*, 8:51–63, 2000.
67. L Sirovich. Dynamics of neuronal populations: eigenfunction theory; some solvable cases. *Network: Comput. Neural Syst.*, 14:249–272, 2003.
68. P Konig, A K Engel, and W Singer. Integrator or coincidence detector? the role of the cortical neuron revisited. *Trends Neurosci.*, 19:130–137, 1996.
69. N Brunel and V Hakim. Fast global oscillations in networks of integrate–and–fire neurons with low firing rates. *Neural Comput.*, 11:1621–1671, 1999.
70. W Gerstner and J L Van Hemmen. Coherence and incoherence in a globally coupled ensemble of pulse–emitting units. *Phys. Rev. Lett.*, 71(3):312–315, 1993.
71. C van Vreeswijk and L F Abbott. Self–sustained firing in populations of integrate–and–fire neurons. *SIAM J. Appl. Maths*, 53(1):253–264, 1993.
72. H R Wilson and J D Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys. J.*, 12:1–23, 1972.
73. H R Wilson and J D Cowan. A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13:55–80, 1973.
74. D H Hubel and T N Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Neurosci.*, 3:1116–1133, 1962.
75. N V Swindale. The development of topography in the visual–cortex: A review of models. *Network*, 7(2):161–274, 1996.
76. S Amari. Dynamics of pattern formation in lateral inhibition type neural fields. *Biol. Cybern.*, 27:77–87, 1977.
77. G B Ermentrout and J Cowan. A mathematical theory of visual hallucination patterns. *Bio. Cybern.*, 34:137–150, 1979.
78. P C Bressloff J D Cowan, M Golubitsky, P J Thomas, and M Wiener. Geometric Visual Hallucinations, Euclidean Symmetry and the Functional Architecture of Striate Cortex. *Phil. Trans. Roy. Soc. Lond. B*, 356:299–330,

2001.

79. Switkes E Tootell R B, Silverman M S and DeValois R L. Deoxyglucose analysis of retinotopic organization in primate striate corte. *Science*, 218:902–904, 1982.
80. E Schwartz. Spatial mapping in the primate sensory projection: analytic structure and relevance to projection. *Biol. Cybern.*, 25:181–194, 1977.
81. J M Fuster and G Alexander. Neuron activity related to short-term memory. *Science*, 173:652, 1971.
82. Xiao-Jing Wang and John Rinzel. Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci.*, 24:455–463, 2001.
83. M Camperi and X-J Wang. A model of visuospatial short-term memory in prefrontal cortex: recurrent network and cellular bistability. *J. Comp. Neurosci.*, 5:383–405, 1998.
84. D Durstewitz, J K Seamans, and T J Sejnowski. Neurocomputational models of working memory. *Nature Neurosci.*, 3:1184–1191, 2000.
85. D Pinto and G B Ermentrout. Spatially structured activity in synaptically coupled neuronal networks: Ii. lateral inhibition and standing pulses. *SIAM J. Appl. Math*, 62:226–243, 2001.
86. S E Folias and P C Bressloff. Breathing pulses in an excitatory neural network. *SIAM J. Appl. Dyn. Syst.*, 3:378–407, 2004.
87. K Kishimoto and S Amari. Existence and stability of local excitations in homogeneous neural fields. *J. Math. Biol.*, 7:303–318, 1979.
88. C R Laing, W C Troy, B Gutkin, and G B Ermentrout. Multiple bumps in a neuronal model of working memory. *SIAM J. Appl. Math*, 63:62–97, 2002.
89. C R Laing and C C Chow. Stationary bumps in networks of spiking neurons. *Neural Comput.*, 13:14731494, 2001.
90. H Werner and T Richter. Circular stationary solutions in two-dimensional neural fields. *Biol. Cybern.*, 85:211–217, 2001.
91. E Aksay, R Baker, H S Seung, and D W Tank. Anatomy and discharge properties of pre-motor neurons in the goldsh medulla that have eye-position signals during xations. *J. Neurophysiol.*, 84:1035–1049, 2000.
92. H S Seung. How the brain keeps the eyes still. *Proc. Natl. Acad. Sci. USA*, 93:13339–13344, 1996.
93. H S Seung, D D Lee, B Y Reis, and D W Tank. Stability of the memory of eye position in a recurrent network of conductance-based model neurons. *Neuron*, 26:1259–271, 2000.
94. R D Chervin, P A Pierce, and B W Connors. Periodicity and directionality in the propagation of epileptiform discharges across neocortex. *J. Neurophysiol.*, 60:1695–1713, 1988.
95. D Golomb and Y Amitai. Propagating neuronal discharges in neocortical slices: Computational and experimental study. *J. Neurophysiol.*, 78:1199–1211, 1997.
96. G B Ermentrout and J B Mcleod. Existence and uniqueness of travelling waves for a neural network. *Proc. Roy. Soc. Edin. A*, 123:461–478, 1993.
97. L Zhang. On the stability of traveling wave solutions in synaptically coupled neuronal networks. *Diff. Int. Eq.*, 16:513–536, 2003.
98. S Coombes and M R Owen. Evans functions for integral neural field equations with heaviside firing rate function. *SIAM J. Appl. Dyn. Syst.*, 4:574–600,

2004.
99. S E Folias and P C Bressloff. Stimulus–locked traveling pulses and breathers in an excitatory neural network. *SIAM J. Appl. Math.*, In press, 2005.
100. D Pinto and G B Ermentrout. Spatially structured activity in synaptically coupled neuronal networks: I. traveling fronts and pulses. *SIAM J. Appl. Math*, 62:206–225, 2001.
101. G B Ermentrout. The analysis of synaptically generated travelling waves. *J. Comput. Neurosc.*, 5:191–208, 1998.
102. P C Bressloff. Synaptically generated wave propagation in excitable neural media. *Phys. Rev. Lett.*, 82:2979–2982, 1999.
103. D Golomb, X J Wang, and J Rinzel. Propagation of spindle waves in a thalamic slice model. *J. Neurophysiol.*, 75:750–769, 1996.
104. M Steriade. Coherent oscillations and short-term plasticity in corticothalamic networks. *Trends in Neurosci.*, 2:337–345, 1999.
105. D H Terman, G B Ermentrout, and A C Yew. Propgating activity patterns in thalamic neuronal networks. *SIAM J. Appl. Math*, 61:1578–1604, 2003.
106. S Coombes. Dynamics of synaptically coupled integrate–and–fire–or–burst neurons. *Phys. Rev. E*, 67:041910, 2003.
107. D O Hebb. *The organization of behavior*. Wiley, New York, 1949.
108. G Bi and M Poo. Synaptic modification in cultured hippocampal neurons: dependence on spike timing, synaptic strength and postsynaptic cell type. *J. Neurosci.*, 18:10464–10472, 1998.
109. T V P Bliss and T Lomo. Long–lasting potentiation of synaptic transmission in the dentate area of unanaeasthetized rabbbit following stimulation of the perforant path. *J. Physiol. (London)*, 232:331–356, 1973.
110. S M Dudek and M F Bear. Homosynaptic long–term depression in area ca1 of hippocampus and effects of nmda receptor blockade. *Proc. natl. Acad. Sci.*, 89:4363–4367, 1992.
111. T V P Bliss and G L Collingridge. A synaptic model of memory: long–term potentiation in the hippocampus. *Nature*, 361:31–39, 1993.
112. J E Lisman. Long-term potentiation: outstanding questions and attempted synthesis. *Phi. Trans. R. Soc. Lond. B*, 358:829–842, 2003.
113. R C Malenka and M F Bear. Ltp and ltd: an embarrassment of riches. *Neuron*, 44:5–21, 2004.
114. J C Magee and D Johnston. A synaptically contriolled associative signal forhebbian plasticity in hippocampal neurons. *Science*, 275:209–213, 1997.
115. H Markram, J Lubke, M Frotscher, and B Sakmann. Regulation of synaptic efficacy by conincidence of postsynaptic aps and epsps. *Science*, 275:213–215, 1997.
116. G Bi and M Poo. Synaptic modification of correlated activity: Hebb's postulate revisited. *Ann. Rev. Neurosci.*, 24:139–166, 2001.
117. W M Kistler and J L van Hemmen. Modeling synaptic plasticity in conjunction with the timing of pre- and post-synaptic action potentials. *Neural Comput.*, 12:385–405, 2000.
118. M C W van Rossum, G Q Bi, and G G Turrigiano. Stable hebbian learning from spike timing–dependent plasticity. *J. Neurosci.*, 20:8812–8821, 2000.
119. R Kempter, W Gerstner, and J L van Hemmen. Hebbian learning and spiking neurons. *Phys. Rev. E*, 59:4498–4514, 1999.

120. K D Miller and D J C Mackay. The role of constraints in hebbian learning. *Neural Comput.*, 6:100–126, 1994.
121. S Song, K Miller, and L F Abbott. Competitive hebbian learning through spike–time–dependent synaptic plasticity. *Nat. Neurosci.*, 3:919–926, 2000.
122. J Rubin, D D Lee, and H Sompolinsky. Equlibrium properties of temporally asymmetric hebbian plasticity. *Phys. Rev. Lett.*, 86:364–367, 2001.
123. H Cateau and T Fukai. A stochastic method to predict the consequence of arbitrary forms of spike–timing–dependent–plasticity. *Neural Comput.*, 15:598–620, 2003.
124. J E Lisman. Amechanism for the hebb and the anti–hebb processes underlying learning and memory. *Proc. Natl. Acad. Sci*, 86:9574–9578, 1989.
125. R C Malenka, J A Kauer, R S Zucker, and R A Nicoll. Postsynaptic calcium is sufficient for potentiation of hippocampal synaptic transmission. *Science*, 242:81–84, 1988.
126. S N Yang, Y G Tang, and R A Nicoll. Selective induction of ltp and ltd by postsynaptic calcium elevation. *J. Neurophysiol.*, 81:781–787, 1999.
127. H Shouval, M F Bear, and L N Cooper. A unified model of nmda receptor–dependent bidirectional synaptic plasticity. *Proc. Natl. Acad. Sci.*, 99:1069–1073, 2002.
128. C Jahr and C Stevens. A quantitative description of nmda receptor-channel kinetic behavior. *J. Neurosci.*, 10:1830–1837, 1990.
129. G L Collingridge, J T R Isaac, and Y T Wang. Receptor trafficking and synaptic plasticity. *Nature Rev. Neurosci.*, 5:952–962, 2004.
130. J E Lisman and N Spruston. Postsynaptic depolarization requirements for ltp and ltd: a critique of spike timing–dependent plasticity. *Nat. Neurosci.*, 8:839–841, 2005.
131. J Rubin, R C Gherkin, G Q Bi, and C C Chow. Calcium time course as a signal for spike–timing–synaptic–plasticity. *J. Neurophsyiol.*, 93:2600–2613, 2005.
132. P N R Rao and T J Sejnowski. Self-organizing neural systems based on predictive learning. *Phil. Trans. R. Soc. Lond.*, 361:1149–1175, 2003.
133. M R Mehta, M Quirk, and M Wilson. Experience–dependent asymmetric shape of hippocampal receptive fields. *Neuron*, 25:707–715, 2000.
134. F Worgotter and B Porr. Temporal sequence learning, prediction and control: a review of different models and their relation to biological mechanisms. *Neural Comput.*, 17:245–319, 2005.
135. W Schultz. Getting formal with dopamine and reward. *Neuron*, 36:241–263, 2002.
136. P R Montague, S E Hyman, and J D Cohen. Computational roles for dopamine in behavioural control. *Nature*, 431:760–767, 2004.
137. R S Sutton and A G Barto. *Reinforcement learning: an introduction*. MIT Press, Cambridge, MA, 1998.