LETTER Communicated by Misha Tsodyks

A Statistical Theory of Long-Term Potentiation and Depression

John M. Beggs
National Institute of Mental Health, Lab of Neural Network Physiology, Bethesda, MD 20892-4075, U.S.A.

The synaptic phenomena of long-term potentiation (LTP) and long-term depression (LTD) have been intensively studied for over twenty-five years. Although many diverse aspects of these forms of plasticity have been observed, no single theory has offered a unifying explanation for them. Here, a statistical "bin" model is proposed to account for a variety of features observed in LTP and LTD experiments performed with field potentials in mammalian cortical slices. It is hypothesized that long-term synaptic changes will be induced when statistically unlikely conjunctions of pre- and postsynaptic activity occur. This hypothesis implies that finite changes in synaptic strength will be proportional to information transmitted by conjunctions and that excitatory synapses will obey a Hebbian rule (Hebb, 1949). Using only one set of constants, the bin model offers an explanation as to why synaptic strength decreases in a decelerating manner during LTD induction (Mulkey & Malenka, 1992); why the induction protocols for LTP and LTD are asymmetric (Dudek & Bear, 1992; Mulkey & Malenka, 1992); why stimulation over a range of frequencies produces a frequency-response curve similar to that proposed by the BCM theory (Bienenstock, Cooper, & Munro, 1982; Dudek & Bear, 1992); and why this curve would shift as postsynaptic activity is changed (Kirkwood, Rioult, & Bear, 1996). In addition, the bin model offers an alternative to the BCM theory by predicting that changes in postsynaptic activity will produce vertical shifts in the curve rather than merely horizontal shifts.

1 Introduction

Long-term potentiation (LTP) and long-term depression (LTD) are persistent changes in synaptic efficacy that can be induced by electrical stimulation. These phenomena have been intensively studied since 1973 (Bliss & Lomo, 1973), motivated by the hypothesis that synaptic changes affect information processing in the brain (reviewed in Bliss & Collingridge, 1993; Beggs et al., 1999). Recently, simple protocols have been developed so that either LTP or LTD can be reliably induced in field potentials in the mammalian hippocampus (Dudek & Bear, 1992; Mulkey & Malenka, 1992) and neocortex (Kirkwood, Dudek, Gold, Aizenman, & Bear, 1993).

Although these protocols are now widely accepted, an understanding of why they work has taken years to develop and is still incomplete. For example, a paradox surrounds the asymmetry of induction protocols. Induction of LTP requires only about 100 presynaptic stimulation pulses, while induction of LTD requires 900 pulses. This seems unbalanced, since these protocols typically produce amplitude or slope changes of nearly equal magnitude (Dudek & Bear, 1992; Mulkey & Malenka, 1992; Kirkwood, Rioult, & Bear, 1996). Existing models of synaptic plasticity would not predict this disparity (Hebb, 1949; Bienenstock, Cooper, & Munro, 1982; Bear, Cooper, & Ebner, 1987; Lisman, 1989). A seemingly unrelated phenomenon can be observed during LTD induction. As 1 Hz stimulation is delivered, the magnitude of the synaptic response declines in a decelerating manner (Mulkey & Malenka, 1992). No theory has been offered to explain this phenomenon or to relate it to other observations about LTD. Another question concerns the reduced magnitude of depression seen in slices of visual cortex taken from dark-reared rats. The well-known BCM theory (Bienenstock et al., 1982) has been used with some success to describe the results of LTP and LTD experiments in rat visual cortex (Kirkwood et al., 1996). The BCM theory does not predict that the magnitude of LTD observed in slices from dark-reared animals would be reduced, only that the frequency at which LTD is induced should be lower. Yet experimental results show that the magnitude of LTD is reduced in slices from dark-reared animals relative to controls (Kirkwood et al., 1996). Taken together, these issues suggest that there may be gaps in our knowledge about how synaptic plasticity is induced.

According to the bin model, all of these phenomena may be a consequence of the following hypothesis: *Long-term synaptic changes will be induced when statistically unlikely conjunctions of pre- and postsynaptic activity occur*. This implies that changes in synaptic strength will be proportional to information transmitted by conjunctions, subject to the constraint that synaptic strengths are bounded. This hypothesis is a novel fusion of two earlier lines of thought concerned with synaptic conjunctions and information theory.

The idea that conjunctions of pre- and postsynaptic activity could cause changes in synaptic strengths was first proposed by Hebb (1949) and has since been investigated extensively through both physiological (e.g., Kelso, Ganong, & Brown, 1986; Kirkwood & Bear, 1994) and modeling studies (e.g., Linsker, 1988; Zador, Koch, & Brown, 1990).

The hypothesis that synaptic strengths are related to transmitted information was originally proposed by Uttley (1966). Uttley argued that synaptic strength was proportional to the negative of the mutual information between the presynaptic input and the postsynaptic output. This negative relationship allowed for negative feedback and a highly stable system, with synaptic strengths often tending toward zero (Uttley, 1970, 1979). In contrast, Brindley (1969) and Marr (1970) proposed that synaptic strengths were related to the positive of the mutual information function, thus leading to

bistable synapses. Whenever the mutual information was above a threshold, the synapse would assume a single positive strength; whenever the mutual information was below this threshold, the synapse would be set to zero strength. In a slightly different vein, Linsker (1988) showed that if synapses obeyed a Hebbian rule, they would maximize information transfer.

The bin model described here proposes that synaptic strengths are proportional to the positive of the information transmitted by conjunctions (not mutual information). Further, the bin model shows that if synaptic strength is proportional to information transmitted by conjunctions, then the synapse will obey a Hebbian rule (different from Linsker's assertion). Here, these ideas are combined in a new way to account for the aggregate behavior of synapses seen in LTP and LTD experiments using field potentials in mammalian cortical slices.

2 The Bin Model

2.1 Probability and Information. Consider a model synapse with an identified presynaptic input and postsynaptic output, as shown in Figure 1A. Imagine that there is a time-stamped record of all the spikes that are fired on each of these axons over a given period of time. Imagine further that this record is divided into many small time bins of equal size. Let us define three temporal arrangements that could occur between a presynaptic and a postsynaptic spike:

- A **hit** will occur whenever a presynaptic and a postsynaptic spike occur in the same bin.
- A **near-miss** will occur whenever a presynaptic spike occurs in a vacant bin immediately after a postsynaptic spike.
- A **miss** will occur whenever a presynaptic spike occurs in a vacant bin and does not produce a near-miss.

We can make our definitions of these events arbitrarily precise by choosing the bin size to be suitably small (see the appendix for a more thorough treatment of binning).

Now let us calculate how often the presynaptic input and the postsynaptic output will produce hits. To determine chance performance, consider how often hits would occur if N_{pr} presynaptic and N_{po} postsynaptic spikes were randomly "dropped" into a row of N_b time bins, with the restriction that neither the presynaptic nor the postsynaptic cell may drop more than one of its own spikes into a bin (see Figure 1B). Given this model, the probability that n hits would occur by chance is described by the hypergeometric distribution (see Figure 1C), which is well known in the statistical literature

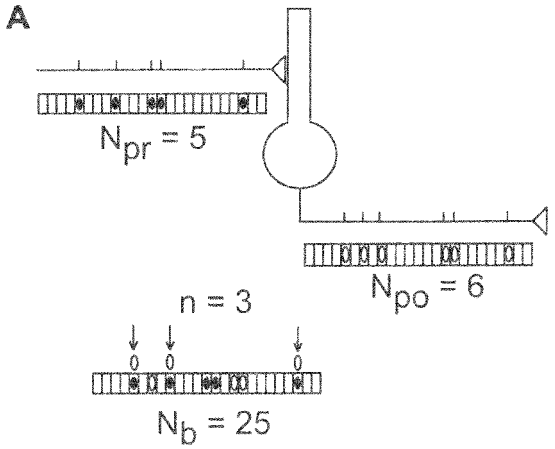
A
N_{pr} = 5
N_{po} = 6
n = 3
N_b = 25

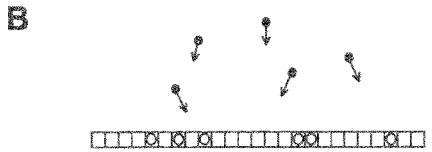
B

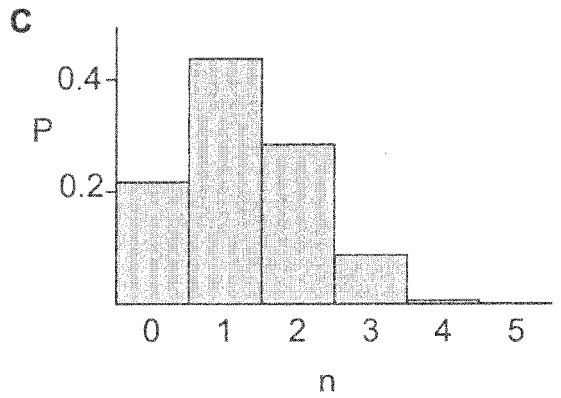
C
0.4
P
0.2
0
1
2
3
4
5
n

(e.g., Johnson & Kotz, 1969; Meyer, 1970):

$$P(n \mid N_{pr}, N_{po}, N_b) = \frac{\binom{N_{po}}{n}\binom{N_b - N_{po}}{N_{pr} - n}}{\binom{N_b}{N_{pr}}} \tag{2.1}$$

where:

$P(n \mid N_{pr}, N_{po}, N_b) =$ the probability that n hits will occur, given N_{pr}, N_{po}, N_b

$n =$ the number of times a pre- and postsynaptic spike will be in the same bin

$N_{pr} =$ the number of presynaptic spikes in the time period

$N_{po} =$ the number of postsynaptic spikes in the time period

$N_b =$ the number of bins in the time period

with the restriction that:

$$n \leq N_{pr} \leq N_{po} \leq N_b.^1$$

For large values of (N_{po}/N_b), the hypergeometric distribution approaches the binomial distribution (Meyer, 1970). The distribution will attain a maximum at n_{peak} (Johnson & Kotz, 1969), given by:

$$n_{peak} = \text{floor}\left[\frac{(N_{pr} + 1)(N_{po} + 1)}{(N_b + 2)}\right]. \tag{2.2}$$

Figure 1: *Facing page.* Bin model synapse. (A) N_{pr} represents the number of spikes fired by the presynaptic neuron (black ovals) in the given time period, N_{po} represents the number of spikes fired by the postsynaptic neuron (open ovals) in the same time period, and N_b is the number of bins in the time period. A hit occurs whenever the pre- and postsynaptic cells fire spikes in the same time bin. This is identified by an arrow. In this case, there are three hits ($n = 3$). (B) To calculate how many hits would occur by chance, the model assumes that spikes are dropped randomly into a row of time bins. In this example, five presynaptic and six postsynaptic spikes are dropped into 25 time bins. (C) The hypergeometric distribution describes the probability, P, of obtaining each number of hits, n, for the example chosen.

[1] For the mathematics to be correct, N_{pr} must always be less than or equal to N_{po}. Thus, N_{po} should always represent the cell that fired the most spikes. In this article, N_{po} will be assigned to the postsynaptic cell for the sake of clarity only. Of course, the model would work equally well if N_{po} were assigned to the presynaptic cell if it fired the most spikes, but that would make the nomenclature confusing.

This model implies that conditional probability is related to information transmission between the presynaptic input and the postsynaptic output. If the postsynaptic cell has access to the number of hits, n, along with knowledge of N_{pr}, N_{po}, and N_b, then it can determine the presynaptic spike train to some extent. If it does not have access to n, however, the postsynaptic cell cannot make any distinctions among the set of possible presynaptic spike trains given by N_{pr} and N_b. Therefore, n can be thought of as a message that conveys information about the presynaptic spike train to the postsynaptic cell. Using the formulation developed by Shannon (1948), the amount of information about the presynaptic spike train conveyed by n to the postsynaptic cell is given by:

$$I(n, N_{pr}, N_{po}, N_b) = -\ln(P(n \mid N_{pr}, N_{po}, N_b)).$$

To normalize this expression, we may add a constant so that no information will be conveyed when the system is in its most probable state (when $n = n_{peak}$):

$$I(n, N_{pr}, N_{po}, N_b) = -\ln(P(n \mid N_{pr}, N_{po}, N_b)) + C,$$

where $C = \ln(P_{peak})$.

This suggests the substitution:

$$W = \left(\frac{P(n \mid N_{pr}, N_{po}, N_b)}{P_{peak}} \right),$$

where W is the normalized probability of occurrence: $0 \leq W \leq 1$. Now the information conveyed by n may be rewritten as:

$$I(n, N_{pr}, N_{po}, N_b) = -\ln(W). \tag{2.3}$$

Note that $W = 1$ when $n = n_{peak}$, while W will approach zero when n is at either tail of the distribution. Intuitively, equation 2.3 means that when n is much larger or much smaller than the peak value n_{peak}, then W is small and n conveys much information. If n is near n_{peak}, however, then W is relatively large and the information conveyed by n is low. This is illustrated in Figure 2, where the information content is shown for each number of hits in the case where there are 5 presynaptic spikes, 25 bins, and several different numbers of postsynaptic spikes. As can be seen from the figure, the most extreme values of n convey the most information.

2.2 Probability Related to Changes in Synaptic Strength. Now plausible values for N_b, N_{pr}, n, and N_{po} will be chosen so that probability values produced by the bin model can be related to changes in synaptic strength reported in field potential experiments of LTP and LTD. The exact values of these parameters are relatively unimportant, since it will be shown later

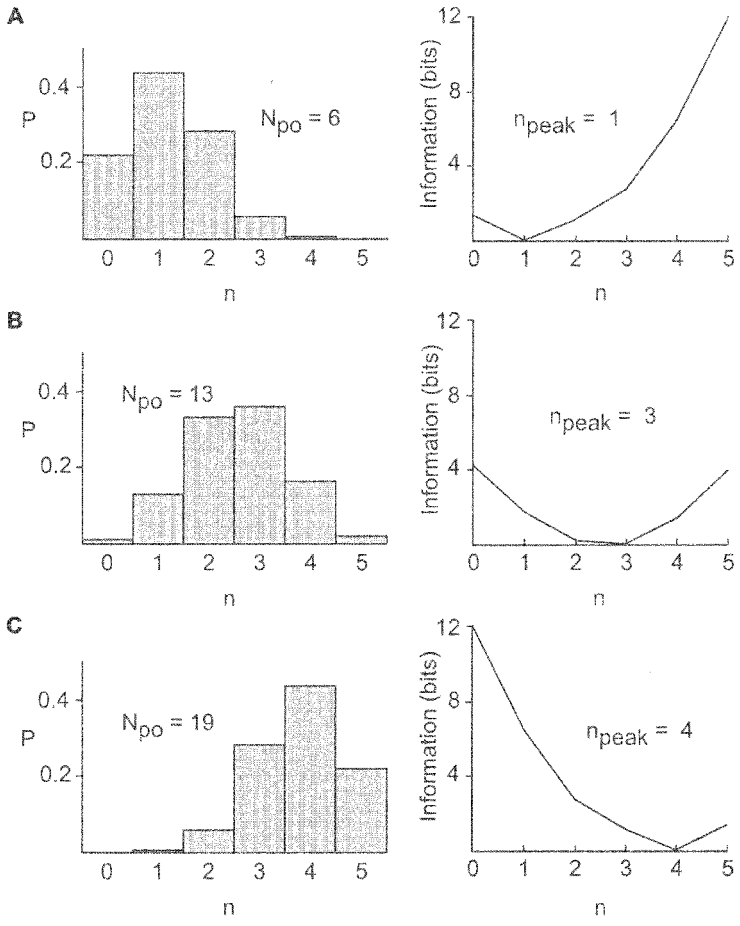

Figure 2: Unlikely numbers of hits convey the most information. (A, left) Probability of occurrence, P, for each number of hits, n, is plotted when $N_{pr} = 5$, $N_{po} = 6$, and $N_b = 25$. (A, right) The information (from equation 2.3, but given in bits) conveyed about the presynaptic spike train is plotted for each number of hits. Note that the information is lowest when n is near the peak value, $n_{peak} = 1$, and highest when n is much greater than n_{peak}. (B, left) When the postsynaptic firing N_{po} is increased to 13, but N_{pr} and N_b remain the same, the distribution is shifted to the right. (B, right) Information is again lowest when n is near the peak value $n_{peak} = 3$, but highest when n is either greater than or less than n_{peak}. (C, left) When the postsynaptic firing N_{po} is further increased to 19 but N_{pr} and N_b remain the same, the distribution is shifted even farther to the right. (C, right) Information is highest when n is much less than n_{peak}. This also shows that changes in postsynaptic firing rate can alter the amount of information conveyed by a given number of hits.

(in Figure 6) that the bin model is quite robust and produces qualitatively similar results for a wide range of parameter values. In what follows, it will be assumed that hits will lead to potentiation, while near-misses and misses will lead to depression. It will be shown that highly unlikely events will lead to large changes in synaptic strength, while moderately unlikely events will lead to only moderate changes in synaptic strength. In the case of two relatively sparse spike trains (where bins are more often empty than occupied), this implies that if a given number of hits will lead to strong potentiation, then the same number of near-misses will lead to equally strong depression. In contrast, the same number of misses will lead to only mild depression, since misses are not as unlikely as hits or near-misses if the spike trains are sparse.

In choosing a bin size for the model, the work of Bi & Poo (1998) is informative. Using paired recordings from hippocampal neurons in culture, Bi and Poo have shown that potentiation is triggered when a presynaptic spike precedes a postsynaptic spike by 20 ms or less, an event that we may interpret as a hit (see the appendix for a more thorough treatment of binning and the calculation of probabilities associated with hits). In contrast, depression is triggered when a presynaptic spike follows a postsynaptic spike by 20 ms or less, an event that we may interpret as a near-miss. Although Bi and Poo did not closely examine misses, this type of event is prevalent in LTD experiments using field potentials (and will be discussed later). Motivated by the work of Bi and Poo, a bin size of 20 ms will be chosen for the model, and no more than one hit will be counted in a given bin.

How many bins (N_b) should be used? In experiments using hippocampal slices, Huang, Colino, Selig, and Malenka (1992) have shown that prior synaptic activity can influence the later induction of LTP. When a weak tetanus (30 Hz, 0.15 s) was delivered 20 minutes before 100 Hz stimulation, the induction of LTP was inhibited. In contrast, the same treatment given 80 minutes before 100 Hz stimulation was without effect. This suggests that if there is a finite time window during which spikes are integrated for long-term plasticity, that window is longer than 20 minutes but shorter than 80 minutes. In other experiments, Xu, Anwyl, and Rowan (1997) reported that transient exposure to stress could affect the induction of LTD in awake rats for 5 minutes after the stress was removed. When the stress had been removed for 20 minutes, however, no effects were seen. These data would suggest a time window at least 5 minutes long but shorter than 20 minutes. Taking into consideration data from both studies, a time period of 20 minutes will be chosen for the purposes of this model. When this 20-minute period is divided by 20 ms, it yields 60,000 bins ($N_b = 60{,}000$).

In the slice, the number of presynaptic spikes fired (N_{pr}) during the given time period will be determined by the number of stimulation pulses delivered. For the case of LTP induction, N_{pr} will be 120, since this is the number of presynaptic pulses delivered during "theta burst" stimulation, a widely accepted protocol that has proved effective in inducing LTP in mammalian

cortical structures (Kirkwood et al., 1993). In response to theta burst, it will be assumed that summation of postsynaptic potentials will cause the cell to fire 30 spikes—one spike for each burst of 4 pulses delivered at 100 Hz. This is motivated by the fact that postsynaptic potentials delivered at 100 Hz with near-threshold stimulation intensity frequently summate to produce spikes within 3 to 5 pulses (Beggs, 1998). Thus, n will be 30 for the case of LTP induction with 100 Hz. For the case of LTD, N_{pr} will be 900, since this is the number of pulses commonly used for induction of depression (Dudek & Bear, 1992; Mulkey & Malenka, 1992). In response to 900 pulses delivered at 1 Hz, it will be assumed that the postsynaptic cell fires no action potentials. This is reasonable, since field potentials recorded during LTD induction do not show population spikes, indicating that action potentials are not fired (Dudek & Bear, 1992; Mulkey & Malenka, 1992). Thus, n will be zero during LTD induction.

Next, the number of spikes for the postsynaptic cell (N_{po}) must be chosen. Although spontaneous activity in the slice is nearly zero, there is evidence to suggest that experimental conditions before slicing may nonetheless influence synaptic plasticity after slices are prepared. For example, days of light deprivation (Kirkwood et al., 1996) have been shown to influence synaptic plasticity dramatically in slices of visual cortex taken from experimental animals, relative to controls. From these data, it is plausible that some biophysical mechanism becomes tuned to the average firing rate and retains that tuning even after slices are prepared (see Quinlan, Philpot, Huganir, & Bear, 1999 for a candidate mechanism). Because light deprivation seems to have the effect of reducing the threshold for potentiation equally on all inputs to a given cell, we will assume that this effect is postsynaptically mediated and may be represented by N_{po}. In the unanesthetized mammal, the average firing rate of a cortical or hippocampal pyramidal neuron is typically 1 Hz or slightly higher (e.g., Thompson, Deyo, & Disterhoft, 1990; Czepita, Reid, & Daw, 1994; Collins & Pare, 1999). For the purposes of this model, a postsynaptic firing rate of 1.5 Hz will be chosen. In 20 minutes, this firing rate will produce 1800 spikes ($N_{po} = 1800$).

Now that we have candidate values for N_b, N_{pr}, n, and N_{po}, we may use them to calculate probabilities and relate them to observed changes in synaptic strength. We may develop a correspondence between the probabilities generated by the bin model and the amplitude changes observed during the induction of LTD. During LTD induction, 900 presynaptic pulses are delivered, and no postsynaptic spikes are produced. This is evidenced by the fact that the field potentials in studies of LTD do not show population spikes. As these 900 pulses are delivered, the magnitude of the synaptic response shows a progressive decline (Dudek & Bear, 1992; Mulkey & Malenka, 1992). This decline is seen in Figure 3A, which is an average of seven experiments and is typical of LTD induction observed with field potentials (Mulkey & Malenka, 1992). Note that 1 Hz stimulation initially causes enhancement, but this is a short-term effect and mostly rebounds at the end of stimula-

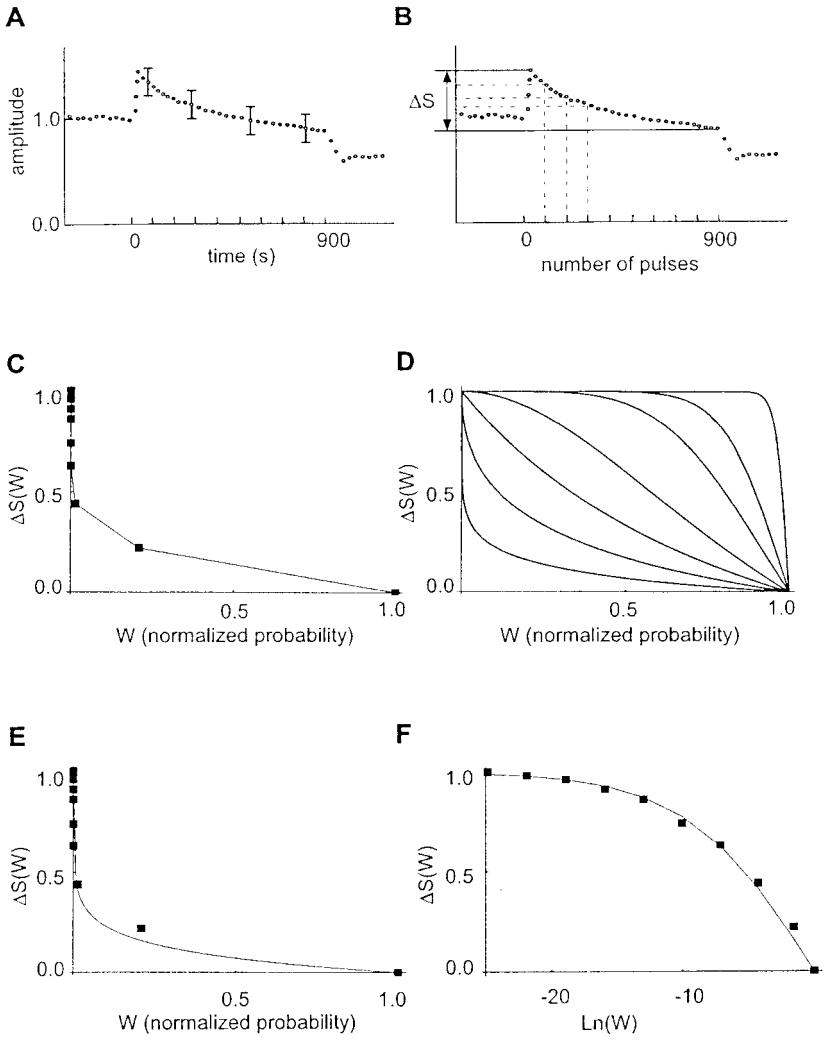

tion. Because this article describes a model of long-term potentiation and depression, we will not attempt to fit this short-term effect, but will focus only on effects that produce more enduring synaptic changes (see Varela et al., 1997, for a description of short-term plasticity in rat cortex).

The magnitude of depression for a given point along the declining curve can be matched to the number of pulses that have been delivered in the LTD induction protocol up to that point (see Figure 3B). Consider an excitatory synapse with initial strength S that is observed to show long-term depression of an amount ΔS by the end of an experiment with 900 pulses. Before

Table 1: Parameters for the LTD Induction Curve.

Pulses	W	ΔS
0	1.0	0.00
100	2.1×10^{-1}	0.22
200	1.4×10^{-2}	0.44
300	8.0×10^{-4}	0.63
400	4.4×10^{-5}	0.74
500	2.3×10^{-6}	0.86
600	1.2×10^{-7}	0.91
700	6.2×10^{-9}	0.96
800	3.2×10^{-10}	0.98
900	1.6×10^{-11}	1.00

the synapse begins its decline in efficacy, the observed depression is 0.00 ΔS. After 100 pulses, the observed depression is approximately 0.22 ΔS. After 300 pulses, the observed depression is approximately 0.63 ΔS. By the end of the experiment, the depression is 1.00 ΔS. These correspondences are taken from Figure 3B and listed in Table 1.

To relate these depressions to probabilities, the bin model can be used to calculate the probabilities of zero hits for each number of pulses, since no postsynaptic spikes are produced during the LTD induction protocol (because no population spikes are observed in the LTD field potentials). These

Figure 3: *Facing page.* The LTD induction curve reveals a relationship between probability and changes in synaptic strength. (A) Plot of field EPSP slope taken during the application of 1 Hz stimulation for 900 pulses (15 minutes) in hippocampus (data replotted from Mulkey & Malenka, 1992). The graph is a summary of seven experiments; some error bars are omitted for clarity. Zero marks the beginning of 1 Hz stimulation, and 900 marks the end. Note that 1 Hz stimulation initially causes enhancement, but this is a short-term effect and mostly rebounds at the end of stimulation. The bin model does not attempt to model these short-term effects. (B) The number of pulses at each point along the curve can be matched with the amount by which synaptic strength has decreased (ΔS). (C) A plot of changes in synaptic strength, ΔS, against normalized probability, W. Values for ΔS were taken from the data shown in A; values for W were calculated according to the bin model for each multiple of 100 pulses. Note that no change in strength is specified for $W = 1.0$, while the largest changes are associated with the lowest values of W. (D) A family of functions describing changes in synaptic strengths, $\Delta S(W)$, is plotted for different values of R ($R = 0.2$ lowest curve, $R = 50$, uppermost curve). (E) The function $\Delta S(W)$ (solid curve) provides the best fit to the data (black squares), when $R = 0.205$. This minimizes the squared error. (F) The function $\Delta S(W)$ (solid curve) is plotted against the data (black squares) on a log scale, showing the goodness of fit.

W values are listed in Table 1. From data in this table, it is possible to graph the magnitude of depression as a function of W. The resulting relationship, when plotted as magnitude of depression against W, is a sharply bending curve (see Figure 3C).

We seek a function that can fit this sharply bending curve. Let $\Delta S(W)$ be a function that would take as input the normalized probability of occurrence W ranging from 0 to 1, and produce as output the magnitude by which the synaptic strength is changed, $\Delta S(W)$, ranging from 1 to 0.

We can deduce some general qualities of the function $\Delta S(W)$ from looking at the numbers listed in Table 1. First, when W is very small, ΔS is large. This means that as W approaches 0, $\Delta S(W)$ should approach 1. Second, when W is relatively large, ΔS is very small. This means that as W approaches 1, $\Delta S(W)$ should approach 0. Third, as W increases from 0 to 1, the magnitude of $\Delta S(W)$ always decreases and never increases. Mathematically, this means that the magnitude of $\Delta S(W)$ should decrease in a strictly monotonic fashion as W ranges from 0 to 1. A function satisfying these general requirements is given by:

$$\Delta S(W) = \left(\frac{1 - W^R}{1 + W^R}\right), \tag{2.4}$$

where W is the normalized probability of occurrence of an event and $\Delta S(W)$ is the magnitude of change in strength caused by an event with probability W and

$$0 \leq W \leq 1$$
$$0 \leq \Delta S(W) \leq 1$$
$$0 \leq R \leq \infty.$$

Equation 2.4 has been plotted for several values of R in Figure 3D. The value of R can be any positive constant and will determine the exact shape of the curve. Note that with the appropriate choice of R, $\Delta S(W)$ can reach any point in the space defined by the requirements. For $R \ll 1$, the magnitude of $\Delta S(W)$ will be large only when W is near 0, while for $R \gg 1$, the magnitude of $\Delta S(W)$ will be small only when W is near 1. What value of R will best describe the data obtained in field potential experiments performed in slices of mammalian cortical structures?

Using numerical methods,[2] a value of R was chosen to minimize the squared error between the data points given in Table 1 and the function $\Delta S(W)$. It is found that $R \approx 0.205$ gives the best fit to the LTD induction curve (see Figures 3E and 3F). As the graph shows, this curve would not lead to much depression until $W < 0.2$, when the magnitude of ΔS values

[2] Numerical integration was performed using Mathematica 3.0.1.

would start to become greater than 0.2. This means that only unusually large numbers of misses would lead to substantial depression, while the vast majority of synaptic events would cause little change.

In this model we will assume that the relationship between probability and synaptic depression is one instance of a general mechanism that will be at work in synaptic potentiation as well. As will be shown later, this assumption provides a reasonable description of the data for LTP experiments. We may combine both relationships:

$$\begin{aligned} &\text{If } n \geq n_{peak}, \text{ then:} \quad \Delta S(W) = +\left(\frac{1 - W^{0.205}}{1 + W^{0.205}}\right) \quad \text{for hits.} \\ &\text{If } n < n_{peak}, \text{ then:} \quad \Delta S(W) = -\left(\frac{1 - W^{0.205}}{1 + W^{0.205}}\right) \quad \begin{array}{l}\text{for misses} \\ \text{and near-misses.}\end{array} \end{aligned} \tag{2.5}$$

Note the use of plus and minus signs to indicate the direction of synaptic change. These equations indicate synaptic potentiation for improbable numbers of hits and synaptic depression for improbable numbers of misses or near-misses, and no synaptic changes when $n = n_{peak}$.

2.3 Information Related to Changes in Synaptic Strength. Let us explore the relationship between probability and changes in synaptic strength by introducing the relation:

$$\delta s(W) = -k \cdot \ln(W). \tag{2.6}$$

If the proportionality constant k is chosen correctly, $\delta s(W)$ can approximate $\Delta S(W)$ for most values of W (see Figure 4A). As W approaches zero, however, $\delta s(W)$ approaches infinity and departs from $\Delta S(W)$ (see Figure 4B). We will choose the value for k that minimizes the squared error, weighted by the probability of occurrence, over the interval from 0 to 1 (this minimizes the expected value of the squared error):

$$E = \int_0^1 W \cdot ((\delta s(W)) - (\Delta S(W)))^2 \cdot dW.$$

This may be rewritten as:

$$E = \int_0^1 W \cdot \left((-k \cdot \ln(W)) - \left(\frac{1 - W^{0.205}}{1 + W^{0.205}}\right)\right)^2 \cdot dW.$$

Using numerical methods,[3] E may be plotted as a function of k (see Figure 4C). As can be seen from the graph, E is minimized when $k \approx 0.101$. This value of k may now be substituted back into $\delta s(W)$, and $\delta s(W)$ may

[3] Numerical integration was performed using Mathematica 3.0.1.

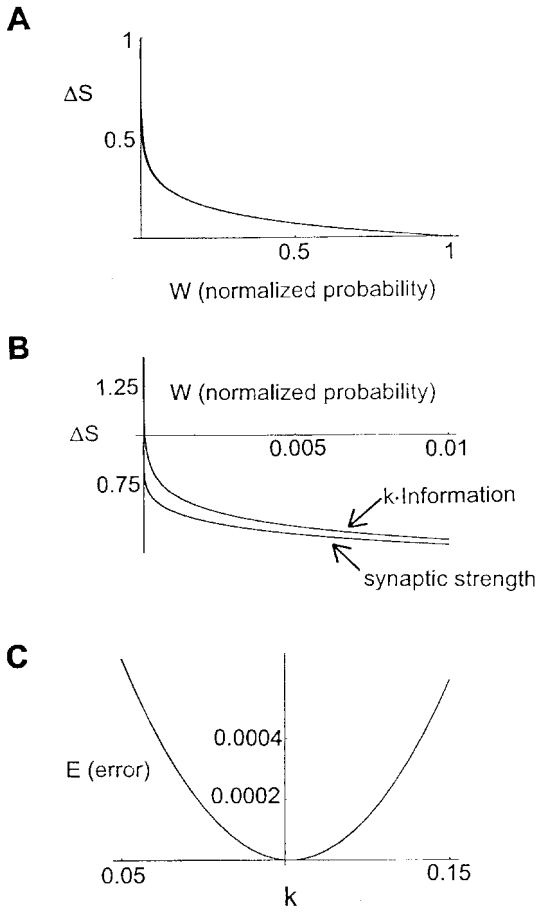

Figure 4: Actual changes in synaptic strength are nearly proportional to transmitted information. (A) The functions $\delta s(W)$ and $\Delta S(W)$ are plotted together. Note that they are indistinguishable at this scale. The function $\Delta S(W)$ fits the experimental data, describing changes in synaptic strength as a function of normalized probability. Real changes in synaptic strength must be bounded, so the function $\Delta S(W)$ has a maximum value of 1 as W approaches 0. The function $\delta s(W)$ is proportional to transmitted information. Information becomes infinite as W approaches 0, creating a difference between $\delta s(W)$ and $\Delta S(W)$. (B) A plot of $\delta s(W)$ and $\Delta S(W)$ at higher scale, showing that there are differences as W approaches zero. (C) When the proportionality constant in $\delta s(W)$ is chosen to minimize the error between $\delta s(W)$ and $\Delta S(W)$, the curves are in fairly close agreement. A plot of the expected value of the error (squared difference between $\delta s(W)$ and $\Delta S(W)$ weighted by the normalized probability) against the value of R. Note that E is minimized when $k \approx 0.101$.

be plotted together with $\Delta S(W)$. From the graphs, it is evident that $\delta s(W)$ and $\Delta S(W)$ are in fairly good agreement except at small values of W (see Figures 4A and 4B).

Let us now discuss the intuitive meaning of $\delta s(W)$ and $\Delta S(W)$. The information conveyed by an event with a normalized probability of occurrence W is given by $-\ln(W)$. This means that $\delta s(W)$ is proportional to the information conveyed by W. The function $\Delta S(W)$ was chosen to map the probabilities generated by the bin model onto changes in synaptic strength that were observed in experiments. In other words, $\Delta S(W)$ gives the magnitude of synaptic change as a function of probability. Because the magnitude of synaptic change must be finite, the output of $\Delta S(W)$ is mapped onto a bounded interval from 0 to 1. In contrast, the output of the function $\delta s(W)$ has no such restrictions and approaches infinity as W approaches zero. Thus, the function $\Delta S(W)$ closely approximates $\delta s(W)$ for events with high probabilities of occurrence and deviates from $\delta s(W)$ only for events with low probabilities of occurrence. From this it is clear that $\Delta S(W)$ accurately models changes in synaptic strength and closely approximates $\delta s(W)$ except at very small values of W. This leads to an interesting conclusion: Changes in synaptic strength are proportional to the information transmitted by conjunctions, subject to the constraint that the magnitude of synaptic change is bounded.

Excitatory synapses that obey this principle will also obey Hebb's rule. If synaptic strength is proportional to information, then increases in information transfer at excitatory synapses will lead to increases in synaptic strength, up to some maximum limit. The information conveyed by the number of hits, n, can be thought of as the amount by which knowledge of n will allow one to delimit the sample space of all possible presynaptic spike trains, given N_{pr}, N_{po}, and N_b. If an excitatory synapse were to change so as to increase the information given by n, then it would change so as to make the observed n larger than the peak value n_{peak} (recall Figure 2). This could be done if a synapse were to become strong enough to ensure that it fired the postsynaptic cell every time it received a presynaptic impulse. Under these conditions, n would be much larger than expected by chance, and the information conveyed by n would be high. Every postsynaptic period of silence could be used to predict when the presynaptic input did not fire. This is the case of a classic Hebb synapse (Hebb, 1949), where synaptic strength increases whenever the presynaptic and postsynaptic cell are active at the same time, thus driving $n > n_{peak}$ and increasing information transfer.

3 Application to Synaptic Plasticity

One of the most successful descriptions of long-term synaptic plasticity in mammalian cortex is the BCM theory (Bienenstock et al., 1982). In addition to its many predictions about the development of visual cortex, two well-known studies (Dudek & Bear, 1992; Kirkwood et al., 1996) have shown a

Table 2: Parameter Values for Normally Reared Rats.

Stimulus	n	N_{pr}	N_{po}	N_b	W	$\Delta S(W)$
0.067 Hz	0	80	1800	60,000	4.0×10^{-1}	−1.9
1 Hz	0	900	1800	60,000	1.6×10^{-11}	−19.8
10 Hz	6	120	1800	60,000	3.8×10^{-1}	+2.0
20 Hz	9	120	1800	60,000	3.3×10^{-2}	+6.7
100 Hz	30	120	1800	60,000	1.0×10^{-18}	+20.0

correspondence between the plasticity predicted by the BCM theory and the potentiation and depression observed in experiments with slices from mammalian cortical structures. We will examine whether the bin model can describe changes in amplitude observed by Kirkwood et al. (1996). In these experiments, field potentials were used to assess the synaptic plasticity of slices of visual cortex taken from normally reared and dark-reared rats. The normally reared rats will be considered first.

The values of N_{pr} are determined by the induction protocols used. Test pulses were delivered once every 15 seconds (0.067 Hz), low-frequency stimuli were delivered with 900 pulses, and high-frequency stimuli were delivered with 120 pulses (see Table 2). The choice of n in several cases requires explanation. When 120 pulses of theta burst are applied at 100 Hz, it will be assumed that the neuron will spike 30 times, once for each burst of 4 pulses. When 120 pulses are applied at 20 Hz, it will be assumed that summation of postsynaptic potentials will be less, leading to only 9 spikes. When 120 pulses are applied at 10 Hz, it will be assumed that even less summation occurs, leading to only 6 spikes. The results of the model are fairly robust with respect to the exact choice of n (see Figure 6). It is important, however, that some gradual decline in the number of spikes occurs as stimulation frequency is reduced. It will be assumed that no postsynaptic spikes occur whenever stimulation is applied at frequencies of 2 Hz or lower, because postsynaptic summation is negligible in this range. Values for $\Delta S(W)$ were multiplied by 20 to scale to the maximum potentiation seen in the experiments (maximum LTP: 20%). As can be seen from the open circles in Figures 5A and 5B, the model predicts the general shape of the curve fairly well, with minor discrepancies for some of the points.

An intuitive explanation for the shape of this curve can be derived from the equations for the bin model. High-frequency stimulation produces several hits (about 30) that lead to large changes in synaptic strength because each hit is unlikely to have occurred by chance. Although the misses produced by low-frequency stimulation are not as unlikely, the fact that there are so many of them (about 900) is very unlikely, and this leads to large changes. Thus, the asymmetry of induction protocols for LTP and LTD arises because hits convey more information than misses, which arises because neurons are quiescent more often than they are firing spikes. The test

Table 3: Parameter Values for Dark-Reared Rats.

Stimulus	n	N_{pr}	N_{po}	N_b	W	$\Delta S(W)$
0.067 Hz	0	80	300	60,000	1.0×10^{0}	0.0
1 Hz	0	900	300	60,000	5.8×10^{-2}	−5.7
2 Hz	0	900	300	60,000	5.8×10^{-2}	−5.7
10 Hz	6	120	300	60,000	5.9×10^{-5}	+15.2
20 Hz	9	120	300	60,000	2.1×10^{-8}	+19.0
100 Hz	30	120	300	60,000	1.8×10^{-41}	+20.0

pulse will always produce a miss, but since it is delivered at such a low frequency, only 80 misses will occur in the 20-minute integration period. This small number of misses is not enough to convey as much information as 900 misses, so the depression induced by the test pulse is barely noticeable.

One of the predictions of the BCM theory is that the stimulation frequency at which depression crosses over to potentiation, called θ, should shift as a function of postsynaptic activity (see Figure 5C). Specifically, one prediction of the BCM theory is that as postsynaptic activity decreases, θ should slide to the left (Bienenstock et al., 1982). This means that a stimulation frequency that caused depression under normal conditions might cause slight potentiation under conditions of reduced postsynaptic activity. To test this prediction, Bear and colleagues (Kirkwood et al., 1996) prepared slices of visual cortex from rats reared in a completely dark environment. When these dark-reared slices were stimulated with a range of frequencies, the curve produced was indeed shifted with respect to the curve produced by control animals (see the filled circles in Figure 5A). Let us see if the bin model can describe the experimental data and whether the model can give some insight as to the nature of this shift.

For the dark-reared rats, it will be assumed that the average firing rate of neurons in visual cortex will be lower than the rate for neurons from rats reared in a fully lighted environment. Although the data are not available for rats, Czepita et al. (1994) found that the spontaneous firing rate for neurons in dark-reared cats could be less than half of the firing rate for neurons in normally reared cats activated by visual stimuli. For this model, a spontaneous firing rate of 0.25 Hz will be used for dark-reared rats, giving 300 spikes in the period of 20 minutes ($N_{po} = 300$). Otherwise, the values for n, N_{pr}, and N_b are the same as those previously given for the normally reared rats. The resulting probabilities and changes in synaptic strength are given in Table 3 and plotted as filled circles in Figure 5B. As can be seen from the figure, the model is qualitatively similar to the data.

In their original paper, Bear and colleagues made note of the fact that θ for the slices from the dark-reared rats is shifted to the left with respect to θ for slices from the normally reared rats. Interpreted this way, the data would agree with the BCM theory, which predicts a leftward shift. According to

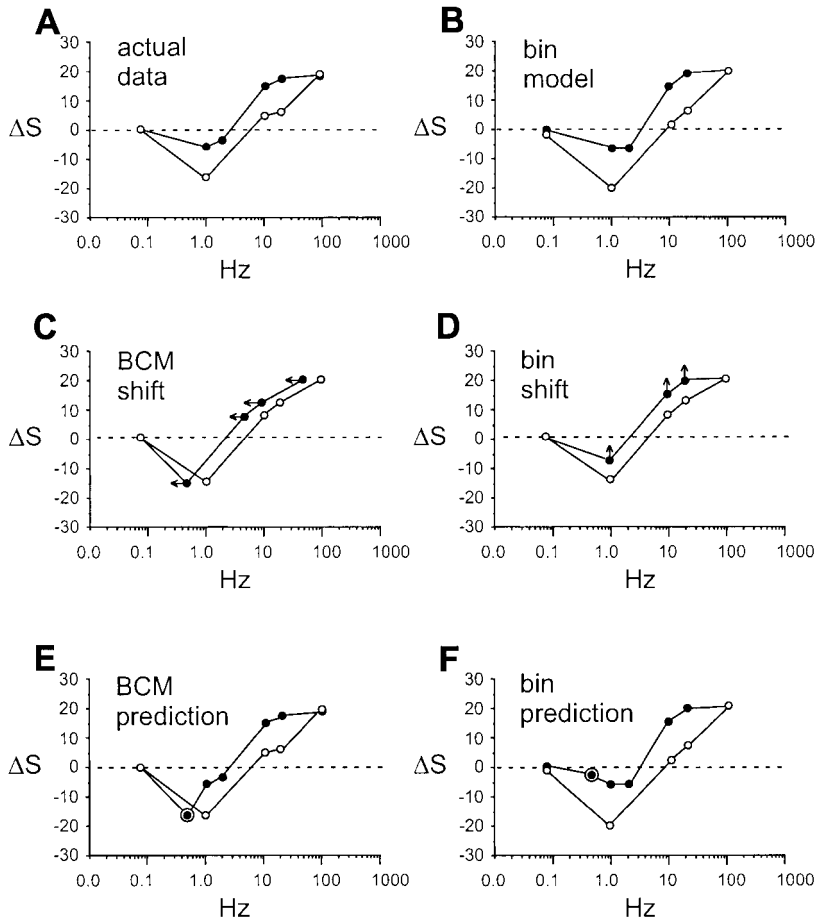

the BCM theory, the threshold frequency θ is a reflection of the average postsynaptic firing frequency. If a presynaptic input is activated at the same frequency as θ, then no synaptic change will result. If a presynaptic input is activated at a frequency below θ, then two things will happen. First, depression will result. Second, the average postsynaptic firing frequency will decline, causing θ to shift to the left, where it will reflect this lowered average firing rate. In a similar manner, a presynaptic input that is activated at a frequency above θ will cause potentiation and will cause θ to shift to the right. From this description, it should be clear that the BCM theory predicts that the frequency-response curve will shift left or right along the horizontal axis as a function of postsynaptic activity. No mention is made in the BCM theory about vertical shifts of the curve (Bienenstock et al., 1982; Clothiaux, Bear, & Cooper, 1991).

Now let us discuss what the bin model predicts. The shape of the frequency-response curve predicted by the model was described previously. How will this curve change as postsynaptic activity is reduced? If postsynaptic activity is reduced from 1.5 Hz (normally reared) to 0.25 Hz (dark reared), then a presynaptic spike that arrives at a synapse will be much less likely to encounter a postsynaptic spike. This fact has two consequences. First, since hits are much rarer, they will convey more information per spike than under normal conditions. Thus, 9 hits produced by theta burst stimulation at 20 Hz will convey more information in slices from dark-reared rats than in slices from normally reared rats. The potentiation seen in slices from dark-reared rats at 20 Hz will therefore be greater than that seen in slices from normally reared rats. Second, since misses are much more common, they will convey less information per spike than under normal conditions. Because of this, the 900 misses produced by LTD induction will convey less information in slices from dark reared rats than in slices from normally reared rats. Consequently, the LTD seen in slices from dark-reared rats will be less than that in slices from normally reared rats. The overall effect of a reduction in postsynaptic activity, then, will be to cause an upward shift in the frequency-response curve (see Figure 5D). This view is consistent with the data that show that LTP is greater and LTD is less in slices from dark-reared rats than in controls (for upward shifts under different conditions of reduced activity, see Wang & Wagner, 1999). Note that the BCM theory does not explicitly predict that the magnitude of LTD should be decreased in slices from dark reared rats, only that θ should shift to the left.

Figure 5: *Facing page.* The bin model compared to experimental data and the BCM theory. (A) Magnitude of synaptic change (ΔS) plotted against the log of stimulation frequency (Hz) for original data (modified from Kirkwood et al., 1996). Open circles represent slices of visual cortex prepared from rats reared in a lighted environment; filled circles represent slices of visual cortex prepared from dark-reared rats. (B) Points predicted by the bin model show fairly close agreement to the actual data. (C, D) Shifts in threshold compared for the BCM theory and the bin model. (C) The BCM theory models a decrease in postsynaptic activity by a threshold shift to the left, but does not specify that the amplitude of the curve should be changed. Thus, the peak magnitude of LTD for normal and deprived cortex is not predicted to differ. (D) In the bin model, a decrease in postsynaptic activity produces an upward shift in the curve. Note that the bin model predicts that the magnitude of LTD for deprived cortex will be less than that for normal cortex, in contrast to the BCM theory. (E, F) Experimental test to distinguish between the BCM theory and the bin model. (E) Data from Kirkwood et al. (1996) are plotted with the additional circled point that would be predicted by a leftward shift of the curve (as described by the BCM theory) if 900 pulses at 0.5 Hz were delivered to slices from dark-reared rats. (F) Values given by the bin model with the additional circled point that would be predicted if 900 pulses were delivered at 0.5 Hz to slices from dark-reared rats.

Which of these two theories better describes the data? A simple test could illuminate the issue: How much LTD will be produced in slices from dark-reared rats if they are stimulated at 0.5 Hz? The BCM theory, which describes only a leftward shift in the curve, would predict that the magnitude of LTD in slices from dark-reared rats at 0.5 Hz should be greater than that seen at 1 Hz and that this should be similar to the LTD seen in normal slices stimulated at 1 Hz (see Figure 5E). The bin model, however, describes a vertical shift in the curve and would predict that the LTD seen at 0.5 Hz would be less than that seen at 1 Hz in slices from dark-reared rats (see Figure 5F).

To explore the robustness of the bin model, three sets of parameters—integration period (N_b), average postsynaptic firing rate (N_{po}), and number of action potentials (n)—were either increased or decreased by 50% for slices from both normally reared and dark-reared rats. The results of these manipulations, shown in Figure 6, demonstrate that the model is extremely robust. Even under larger parameter variations, the qualitative predictions of the model remain unchanged.

4 Summary and Conclusions

Using only one set of constants, the bin model offers robust explanations for several disparate phenomena observed in LTP and LTD experiments. Many of these phenomena are not adequately explained by current theories. The asymmetry of induction protocols for LTP and LTD arises because hits convey more information than misses, which is a result of the fact that neurons are quiescent more often than they are firing spikes. The curving

Figure 6: *Facing page.* Robustness of the bin model. (A, B) Effects of a 50% increase or decrease in integration period. N_b (normally 60,000) was altered to either 90,000 or 30,000; all other parameters remained the same. (A) For slices from normally reared rats, the unaltered curve is plotted with open circles, the increase is plotted with squares, and the decrease is plotted with triangles. Note the minimal changes. (B) Slices from dark-reared rats. Note that no changes at all occurred in this case. (C, D) Effects of a 50% increase or decrease in postsynaptic firing rate. (C) For slices from normally reared rats, N_{po} (usually 1800) was altered to either 2700 or 900. Note the minor vertical shifts. (D) For slices from dark-reared rats, N_{po} (usually 300) was altered to either 450 or 150. Again, note minor vertical shifts. (E, F) Effects of a 50% increase or decrease in the number of action potentials fired in response to stimulation at 10, 20, and 100 Hz. For both types of slices, n (usually $n = 6, 9, 30$) was altered to either $(9, 14, 45)$ or $(3, 5, 15)$. (E) Minimal changes were brought at 10, 20, and 100 Hz, while no changes were brought at 1 Hz and at the test pulse, since no action potentials were fired there. (F) Again note the slight changes. Overall, the bin model retains the same qualitative predictions over large variations in parameter values.

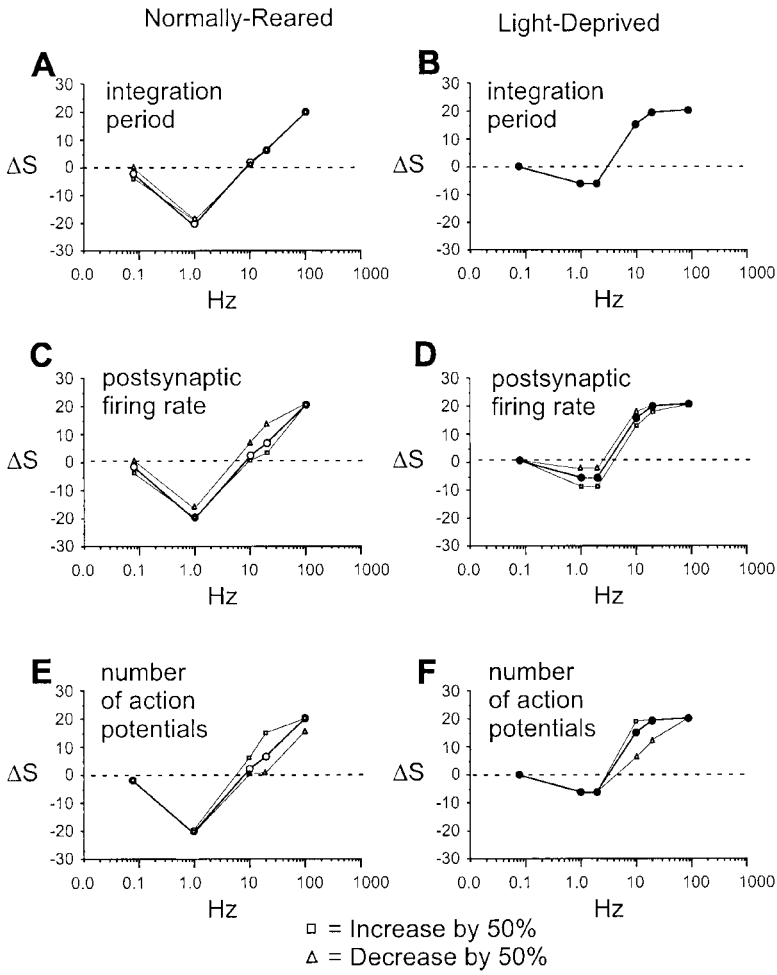

decline in synaptic strength seen during LTD induction is a consequence of the function $\Delta S(W)$, whose shape describes synaptic changes proportional to information transfer, subject to the constraint that synaptic changes must be bounded. The shape of the frequency-response curve predicted by the BCM theory and observed in experiments can also be explained in terms of the bin model. Unlikely numbers of hits cause potentiation, unlikely numbers of misses cause depression, and the misses caused by test pulses do not occur frequently enough in the integration time period to produce substantial depression. The bin model can also explain why reduced postsynaptic activity shifts this frequency-response curve. Reduced postsynaptic activity means that hits will convey more information per spike, so potentiation

will be greater at a given frequency. Reduced postsynaptic activity will also mean that misses convey less information, so the magnitude of depression will be less at a given frequency. The net effect of these changes will be an apparent upward shift in the frequency-response curve, an interpretation that is different from that offered by the BCM theory. Finally, the bin model offers a testable prediction so that these theories can be distinguished.

It should be noted that the bin model is not an attempt to explain directly the biophysical basis of LTP and LTD, nor is it an attempt to explain short-term synaptic plasticity. Rather, the bin model argues that a statistical view of synaptic interactions can offer a robust and unified description for previously unexplained phenomena seen in experiments of long-term synaptic plasticity. A major consequence of this view is that long-term synaptic changes encode transmitted information. We hope that such a description may inspire experimental tests of this idea and contribute to the formulation of more detailed biophysical models in the future.

Appendix: Binning and the Calculation of Probabilities

It may be argued that the definition of a hit presented in the bin model includes cases where LTD would be observed. For example, if a postsynaptic spike is assigned to a bin because it occurs at the beginning of a 20 ms period and a presynaptic spike occurs at any time less than 20 ms after this postsynaptic spike, it will be assigned to the same bin, resulting in a hit. However, this case should lead to LTD, as shown by Bi and Poo (1998), because the presynaptic spike occurs after the postsynaptic spike.

It can be shown that although binning the spike trains may destroy some information about precise temporal order, it does not alter the calculation of probabilities of hits or misses when compared to methods that do take precise temporal order into account. For example, consider a parallel recording where one presynaptic and one postsynaptic spike occur in a 1000 ms period. To keep track of temporal order, let us attach a "trailer" that is 10 ms long after the presynaptic spike. Let us also attach a "leader" that is 10 ms long before the postsynaptic spike. We will define a hit to occur whenever the presynaptic trailer and the postsynaptic leader overlap in time. Thus, a hit will occur only when the presynaptic spike is followed by the postsynaptic spike by less than 20 ms. There can be no hit if the postsynaptic spike occurs before the presynaptic spike. This interval of 20 ms constitutes a "hit zone" over which a hit will occur. The probability that the presynaptic spike would produce a hit with the postsynaptic spike can be calculated by taking the ratio of the length of the hit zone (20 ms) to the length of the recording period (1000 ms): $20/1000 = 0.02$. But this is exactly equivalent to calculating the probability with a bin method that does not account for precise temporal order. The probability that a randomly dropped spike would land in the same 20 ms bin already occupied by a previously dropped spike in a 1000 ms period is $1/50 = 0.02$. (The period is divided into 50 bins of 20 ms

length.) The case of multiple spikes can be handled in a similar manner, demonstrating that binning does not alter the calculation of probabilities, although it may destroy some temporal information.

Acknowledgments

This work was supported by a training grant from NIH awarded to Yale University and by a grant from NIH awarded to Thomas H. Brown. I thank Ted Carnevale, Dan Johnston, David Jaffe, Nigel Daw, Tom Carew, Josh Brumberg, John McGann, Jim Moyer, Patrick Tao, Tom Brown, and two anonymous reviewers for useful comments.

References

Bear, M. F., Cooper, L. N., & Ebner, F.F. (1987). A physiological basis for a theory of synapse modification. *Science, 237*, 42–48.

Beggs, J. M. (1998). *Electrophysiology of rat perirhinal cortex*. Unpublished doctoral dissertation, Yale University.

Beggs, J. M., Brown, T. H., Byrne, J. H., Crow, T., LeDoux, J. E., LeBar, K., & Thompson, R. F. (1999). Learning and memory: Basic mechanisms. In M. J. Zigmond, F. E. Bloom, S. C. Landis, J. L., Roberts, & L. R. Squire, (Eds.), *Fundamental neuroscience*. San Diego, CA: Academic Press.

Bi, G. Q., & Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience, 18*, 10464–10472.

Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience, 2*, 32–48.

Bliss, T. V. P., & Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology (Lond)., 232*, 331–356.

Bliss, T. V. P., & Collingridge, G. L. (1993). A synaptic model of memory: Long-term potentiation in the hippocampus. *Nature, 361*, 31–39.

Brindley, G. S. (1969). Nerve net models of plausible size that perform many simple learning tasks. *Proceedings of the Royal Society, 174*, 173–191.

Clothiaux, E. E., Bear, M. F., & Cooper, L. N. (1991). Synaptic plasticity in visual cortex: Comparison of theory with experiment. *Journal of Neurophysiology, 66*, 1785–1804.

Collins, D. R., & Pare, D. (1999). Reciprocal changes in firing probability of lateral and central medial amygdala neurons. *Journal of Neuroscience, 19*, 836–844.

Czepita, D., Reid, S. N. M., & Daw, N. W. (1994). Effect of longer periods of dark rearing on NMDA receptors in cat visual cortex. *Journal of Neurophysiology, 72*, 1220–1226.

Dudek, S. M., & Bear, M. F. (1992). Homosynaptic long-term depresssion in area CA1 of the hippocampus and effects of N-methyl-D-aspartate receptor blockade. *Proceedings of the National Academy of Sciences, USA, 89*, 4363–4367.

Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley.
Huang, Y. Y., Colino, A., Selig, D. K., & Malenka, R. C. (1992). The influence of prior synaptic activity on the induction of long-term potentiation. *Science, 255*, 730–733.
Johnson, N. L., & Kotz, S. (1969). *Discrete distributions*. Boston: Houghton Mifflin.
Kelso, S. R., Ganong, A. H., & Brown, T. H. (1986). Hebbian synapses in hippocampus. *Proceedings of the National Academy of Sciences, USA, 83*, 5326–5330.
Kirkwood, A., & Bear, M. F. (1994). Hebbian synapses in visual cortex. *Journal of Neuroscience, 14*, 1634–1645.
Kirkwood, A., Dudek, S. M., Gold, J. T., Aizenman, C. D., & Bear, M. F. (1993). Common forms of synaptic plasticity in the hippocampus and in neocortex in vitro. *Science, 260*, 1518–1521.
Kirkwood, A., Rioult, M. G., & Bear, M. F. (1996). Experience-dependent modification of synaptic plasticity in visual cortex. *Nature, 381*, 526–528.
Linsker, R. (1988). Self-organization in a perceptual network. *Computer Magazine, 21*, 105–117.
Lisman, J. (1989). A mechanism for the Hebb and the anti-Hebb processes underlying learning and memory. *Proceedings of the National Academy of Sciences, USA, 86*, 9574–9578.
Marr, D. (1970). A theory for cerebral neocortex. *Proceedings of the Royal Society, 176*, 161–234.
Meyer, P. L. (1970). *Introductory probability and statistical applications*. Reading, MA: Addison-Wesley.
Mulkey, R. M., & Malenka, R. C. (1992). Mechanisms underlying induction of homosynaptic long-term depression in area CA1 of the hippocampus. *Neuron, 9*, 967–975.
Quinlan, E. M., Philpot, B. D., Huganir, R. L., & Bear, M. F. (1999). Rapid, experience-dependent expression of synaptic NMDA receptors in visual cortex in vivo. *Nature Neuroscience, 4*, 352–357.
Shannon, C. E. (1948). A mathematical theory of communication. *Bell Systems Technical Journal, 27*, 623–656.
Thompson, L. T., Deyo, R. A., & Disterhoft, J. F. (1990). Nimodipine enhances spontaneous activity of hippocampal pyramidal neurons in aging rabbits at a dose that facilitates learning. *Brain Research, 535*, 119–130.
Uttley, A. M. (1966). The transmission of information and the effect of local feedback in theoretical and neural networks. *Brain Research, 102*, 23–35.
Uttley, A. M. (1970). The informon: A network for adaptive pattern recognition. *Journal of Theoretical Biology, 27*, 31–67.
Uttley, A. M. (1979). *Information transmission in the nervous system*. New York: Academic.
Varela, J. A., Sen, K., Gibson, J., Fost, J., Abbott, L. F., & Nelson, S. B. (1997). A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex. *Journal of Neuroscience, 17*, 7926–7940.
Wang, H., & Wagner, J. J. (1999). Priming-induced shift in synaptic plasticity in the rat hippocampus. *Journal of Neurophysiology, 82*, 2024–2028.
Xu, L., Anwyl, R., & Rowan, M. J. (1997). Behavioural stress facilitates the induction of long-term depression in the hippocampus. *Nature, 387*, 497–500.

Zador, A., Koch, C., & Brown, T. H. (1990). Biophysical model of a Hebbian synapse. *Proceedings of the National Academy of Sciences, USA, 87*, 6718–6722.

Received July 2, 1999; accepted April 20, 2000.