

# Effects of LTP on Response Selectivity of Simulated Cortical Neurons

Karl Kilborn, Gary Lynch, and Richard Granger

University of California, Irvine

## Abstract

~ We report here on specific ways in which synaptic long-term potentiation (LTP) affects the response selectivity of primary sensory cortical cells. LTP increases synaptic efficacy by incremental "steps," up to a "ceiling" at which additional bursts of afferent stimulation cause no further potentiation. Endogenous and exogenous agents have been shown to modulate these two parameters of LTP, raising the question of the functional implications associated with the sizes of steps and ceiling. We provide an analytical treatment of the effects of these two physiological LTP parameters on the behavior of simulated olfactory (piriform) cortex target cells in response

to a range of inputs. A target cell's receptive field, i.e., the set of input patterns to which the cell responds, is broadened with potentiation of the cell's synapses, and is broadened more when the LTP step size is smaller, and when the LTP ceiling is higher. Moreover, the effects of step size and ceiling interact, and their joint relationship to receptive field breadth is non-linear. Values of step size and ceiling are identified that balance the tradeoff between learning rate and receptive field breadth for particular sensory recognition tasks, and these model values are compared to corresponding known and inferred physiological values. ~

## INTRODUCTION

Long-term potentiation (LTP) has been shown to occur in hippocampal archicortex (Bliss & Lomo, 1973), olfactory paleocortex (Jung, Larson, & Lynch, 1990) and neocortex (Kanter & Haberly, 1990; Iriki, Pavlides, Keller, & Asanuma, 1989; Kirkwood & Bear, 1994). Given that it appears to be a common characteristic of cortical architectures, the question arises as to the effects that LTP might have on cortical processing of sensory signals. Input to a network selectively activates some target cells more than others, on the basis of connectivity of the active inputs with the targets. Afferent stimulation patterns are "similar" to the extent that they share axons, which, when repetitively activated, can potentiate their target synapses. With repeated training, synapses shared among different inputs become more potentiated than unshared synapses. After training on a number of similar instances, this differential potentiation of shared synapses will cause the target cell that responds to one such instance to respond to others as well, since the contribution from the (familiar) shared, potentiated synapses will outweigh the contributions from the (novel) unshared, weaker synapses. With learning, then, target network responses to similar inputs will tend to become increasingly similar to each other. From the point of view of a given target cell participating in these responses, its receptive field is broadening, since it is responding after potentiation to inputs that it would not have responded to before.

It is worth noting that if these responses are the

representations of the inputs passed to other brain regions, the inputs will be treated as though they were more similar than they actually are; an effective learned broadening of the generalization gradient around inputs (reminiscent of the well-known psychological effects of input space distortion (Shepard, 1987) and categorical perception (Smith & Medin, 1981; Harnad, 1987)). However, the sensory and psychological interpretations of neural responses are remote, and it is unknown how generalization of cell population responses may be related to generalization of the organism's overall behavioral response to a stimulus. A neuron's response can be described more neutrally in terms of its receptive field, defined simply as the set of inputs to which it responds. For cell fields receiving topographic afferents, such as layer IV of primary visual, auditory and somatosensory cortices, the receptive field of a neuron has the additional property of preserving adjacency relations among the stimuli to which it responds. Nonetheless the notion of receptive field is equally well defined in non-topographic populations such as olfactory cortex (Jiang & Holley, 1992; Ezeh, Wellis, & Scott, 1993). Slight changes to an input will still tend to recruit responses from many of the same target cells, whereas very different inputs will not; the receptive field of those target cells thus includes the former inputs and excludes the latter.

In this paper we provide a theoretical treatment of the LTP parameters of step and ceiling in the context of a target layer of cortical pyramidal cells, analytically deriv-

ing relationships between these biological parameters and their receptive field characteristics.

## Input and Response

### *Physiological and Anatomical Characteristics Modeled*

Table 1 lists the physiological and anatomical features incorporated into the simulations used. Points 1-3 in the table constitute induction rules for LTP: the familiar "Hebbian" requirement for coactivity of afferent and target, together with the important extra constraint that most cell spiking will *not* induce potentiation: rather, the significantly higher (NMDA receptor) threshold must be exceeded for potentiation to be induced. The relatively stringent requirements for LTP induction ensure that most synapses do not become potentiated under most circumstances despite continuous cell spiking; rather, burst activity must occur in the presence of weakened

Table 1. Physiological and Anatomical Features Incorporated into LTP Parameter Simulations

LTP induction	
1.	Potentiation requires coactivity of afferent and target
2.	Potentiation requires exceeding a threshold above that of the spiking threshold
3.	Potentiation can be elicited by a single induction episode
LTP increment rule	
4.	Equally sized increments in efficacy occur in response to each induction episode
5.	Increment size can be manipulated pharmacologically
6.	Increment size can be manipulated endogenously (by NE, Ach)
LTP ceiling rule	
7.	After a fixed number of increments, further stimulation offers almost no increase in efficacy
8.	Ceiling level can be manipulated pharmacologically
9.	Ceiling level can be manipulated endogenously (by 5HT)
Local circuit anatomy	
10.	Sparse connectivity of afferents to targets
11.	No apparent topography (e.g., LOT to olfactory cortex)
12.	Short electrotonic length of excitatory cell dendrites
Time courses	
13.	Inputs arrive within relative synchrony with theta
14.	Once consolidated, potentiation is not reversible

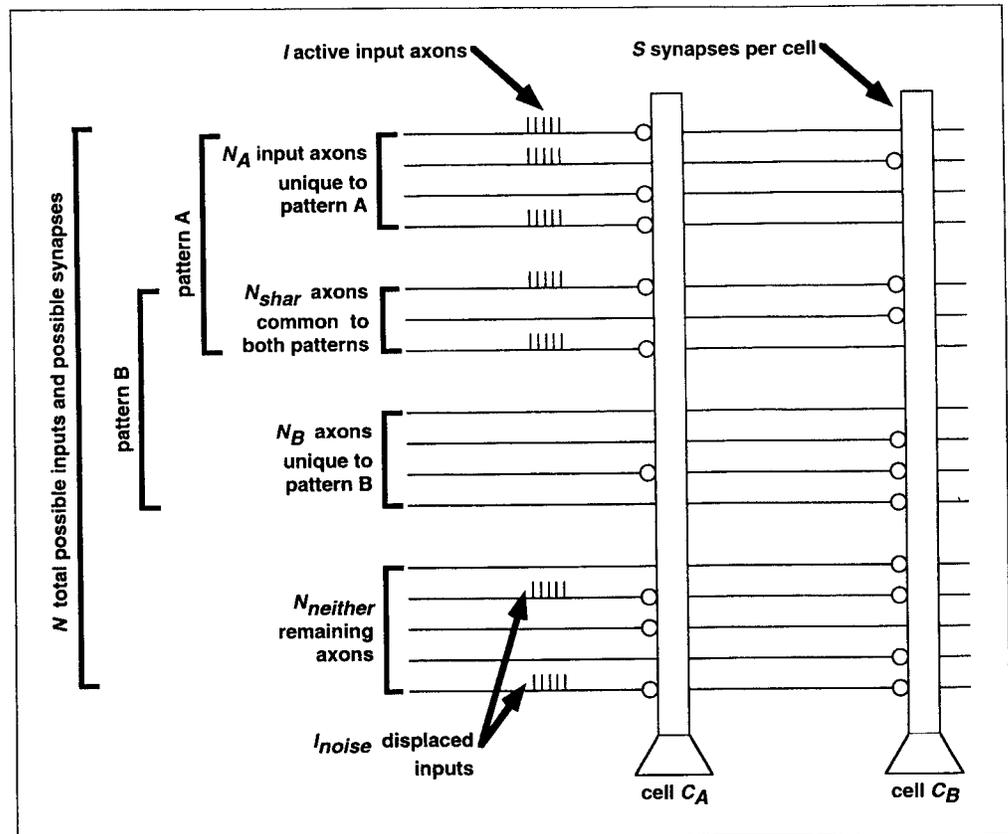
inhibition of the target neurons, as occurs with theta-burst stimulation. Numbers 4-6 and 7-9 specify the physiological details of LTP increments and the LTP ceiling, respectively, and their manipulation by endogenous and exogenous factors. These issues will be gone into in more detail. Items 10-12 describe the anatomies into which the synapses undergoing potentiation are embedded. Number 13 states the assumption, based on both *in vivo* recordings of freely moving, learning animals and *in vitro* studies of LTP induction, that during learning, inputs become relatively synchronized to the theta rhythm throughout the olfactory-hippocampal pathway.

Point 14 raises the issues of LTP reversal and LTD (Staubli & Lynch, 1987; Larson, Xiao, & Lynch, 1993; Bear & Malenka, 1994), neither of which is considered in the present treatment. Any form of synaptic reduction raises the more general question of capacity: i.e., can a system in which synapses only increase suffice as a memory? The obvious intuition, shared by many experimental and computational neuroscientists, is that what goes up must come down. It is clear that if all synapses are allowed to increase without some form of decrease or normalization, eventually all synapses will become saturated at their ceiling value, thus not only preventing any further learning, but also eliminating existing memories, since the differences among synaptic weights were the embodiment of those memories. A concomitant question rarely asked is *when* such saturation will occur—this is a question of capacity, which underlies that of synaptic saturation: how many "memories" can be stored via synaptic strengthening before saturation? In a particular formulation, an increase-only model has been shown to have very large capacity, which is a linear function of the size of the network. In that model, synapses increase without ever decreasing, and the system will thus inevitably saturate, but saturation does not occur until a very large number of memories are stored—such a large number that any reasonable lifetime of the organism would be long over before the limit was approached (Granger, Whitson, Larson, & Lynch, 1994). Thus the function of any LTD mechanism might not be directly related to capacity.

### *Physiological Interpretation*

Consider a cortical network in which a set of input axons makes sparse nontopographic synaptic contact with the dendrites of a set of target cells, as in the olfactory (piriform) cortex (Price, 1973; Haberly, 1985; Lynch, 1986). Input activity is assumed to correspond to the stereotypical physiological activity found throughout the olfactory corticohippocampal pathway during learning and exploration (Macrides, 1975; Hill, 1978; Otto, Eichenbaum, Weiner, & Wible, 1991), consisting of brief relatively synchronized bursts of activity, on a sparse subset of the input lines. Inputs can thus be depicted as

**Figure 1.** Each cell has  $N$  potential synaptic locations and  $S$  actual synapses. Inputs consist of some active lines unique to one pattern or the other ( $N_A$  and  $N_B$ ), some lines that are shared between both patterns ( $N_{Shar}$ ), and some lines that do not signify either pattern ( $N_{neither}$ ). Noise ( $I_{noise}$ ) displaces some of the inputs in these shared and pattern-specific regions to either the other pattern's region or somewhere else along the cell. Note that these regions are depicted as contiguous for convenience only—inputs representing one pattern or the other may be distributed anywhere along the cell.



a set of active input lines arriving at target cells as in Figure 1. In general, inputs will be spatially distributed across the input lines. For simplicity, the illustrations here depict inputs consisting primarily of contiguous axons, as they would be if generated by stimulating electrodes rather than by environmental input; in the mathematics, no assumptions are made about axon distribution. Pictured are two distinct inputs, A and B, which can be presumed to be generated by two different stimulating electrodes. The spatial pattern of activation generated by a given electrode will differentially activate the target cells, as a function of the connectivity of the inputs to the cells. In particular, cells having more synapses with those axons active in input pattern A will be more depolarized in response to that input. The figure illustrates two target cells, one of which is better connected to the active axons in pattern A ( $C_A$ ), and one to the active inputs in B ( $C_B$ ). Thus cell  $C_A$  will be more depolarized in response to input A, via activation of stimulating electrode A, and vice versa for cell  $C_B$ .

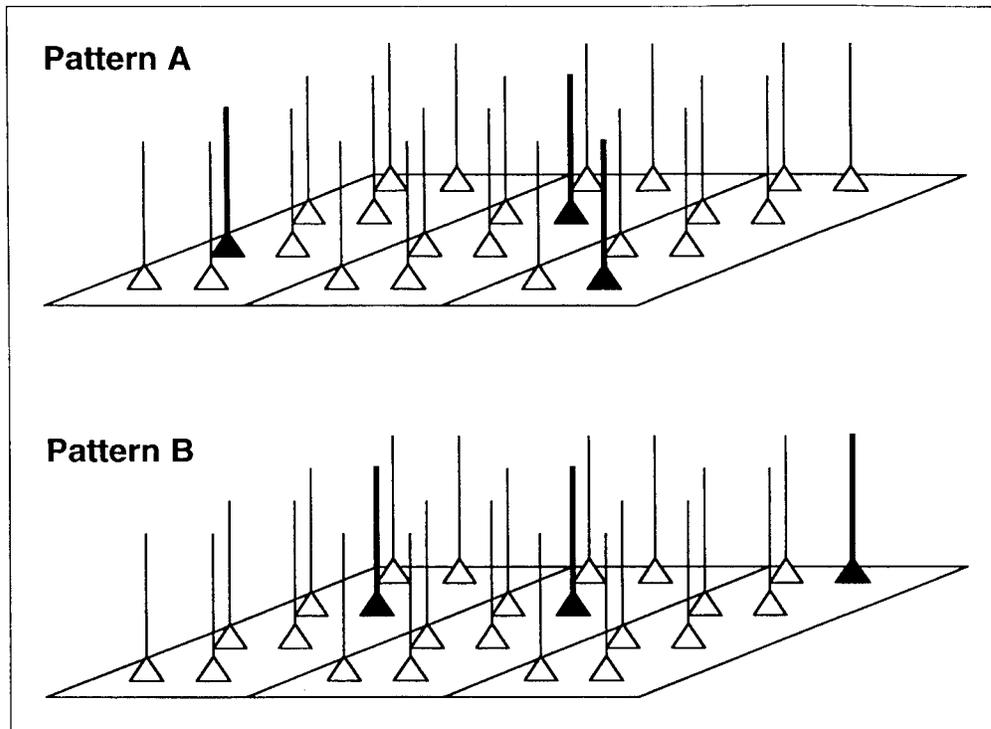
Note that although the responses of only two cells to two different input patterns are considered in this experiment, Figure 2 illustrates how this analysis treats one aspect of a larger distributed representation. In general, we assume that the relative activity of primary responding cells is significant, as in a patch of laterally inhibited cells where the cells with the highest activation will be the only ones to fire (Coultrip, Granger, & Lynch, 1992).

If these two inputs always occurred in the same way, then LTP simply causes the depolarization from the cells to become stronger, and the response to become more reliable. If, however, we make the more realistic assumption that the inputs are "noisy" in that some variation occurs from presentation to presentation, then LTP has a more complex effect, and its physiological induction parameters come into play in a nonintuitive fashion. This random variation within an input consists of probabilistic activation of afferents with components of omission and of commission: some synapses stimulated by electrode A will fail to become active on any given trial, and some synapses not stimulated by electrode A will spontaneously become active. Noise affects the reliability of the target cell response, and the step increment and ceiling level of potentiation during LTP both modulate the extent of this effect.

#### Geometric Interpretation

Inputs of this kind can be viewed as points or vectors in a space whose dimensionality is that of the total number of input lines. If every input made synaptic contact with every target cell, then the dendrites of target cells could be viewed as vectors in that same space. Actual targets correspond to vectors in a subspace of the input space, dictated by the sparseness of connectivity of the inputs with the target. Figure 3 illustrates the

Figure 2. The response of a layer of cells to two different input patterns. Within each patch (represented as bands within each layer), lateral inhibition ensures that only one cell (or a few cells) will fire. However, the response of the entire layer is distributed among many such patches. In some patches (such as the one in the center), the same cell will fire for both patterns. However, in other patches, different cells within each patch will fire for different patterns. The present study analyzes how parameters of LTP induction come to affect cell responses within patches that differentiate among patterns.



qualitative relationships among inputs and targets as mediated by synaptic connections. Depicted are the relative locations of two sets of inputs and two target cells, as projected onto the two-dimensional figure plane. Shown are two inputs, A and B, and their noisy variants ( $\alpha_1, \dots, \beta_1, \dots$ ) corresponding to the random variation of the inputs as described above.

Each input line corresponds to a dimension in the multidimensional input space; the large number of dimensions is necessarily reduced to just the two dimensions of the planar figure in this projection. Actual inputs are high-dimensional and thus the neighborhood of proximity is intuitively quite different from the case in two dimensions. Nonetheless, the figure captures some of the salient elements of the cases to be analyzed here.

A given cell (e.g.,  $C_A$ ) can "see" only that portion of active inputs to which the cell has synaptic connections. The portion of the input seen by the cell is (approximately) the projection ( $\Pi$ ) of the input vector into the lower-dimensional subspace of the cell vector. In the figure, the projection of an input  $x$  onto a cell  $C$  is the subset of the cell vector determined by the orthogonal "shadow" of that input onto the cell,  $\Pi(x, C)$ . The magnitude of response of a given cell to an input is the length of that projection vector. Intuitively, similar inputs in Figure 1 correspond to similar or "close" vectors in Figure 3. The closer an input vector is to a cell vector, the larger the projection of the former onto the latter, and the larger the magnitude of the cell response. The line  $\Pi(x, C_A) = \Pi(x, C_B)$  denotes the set of inputs whose

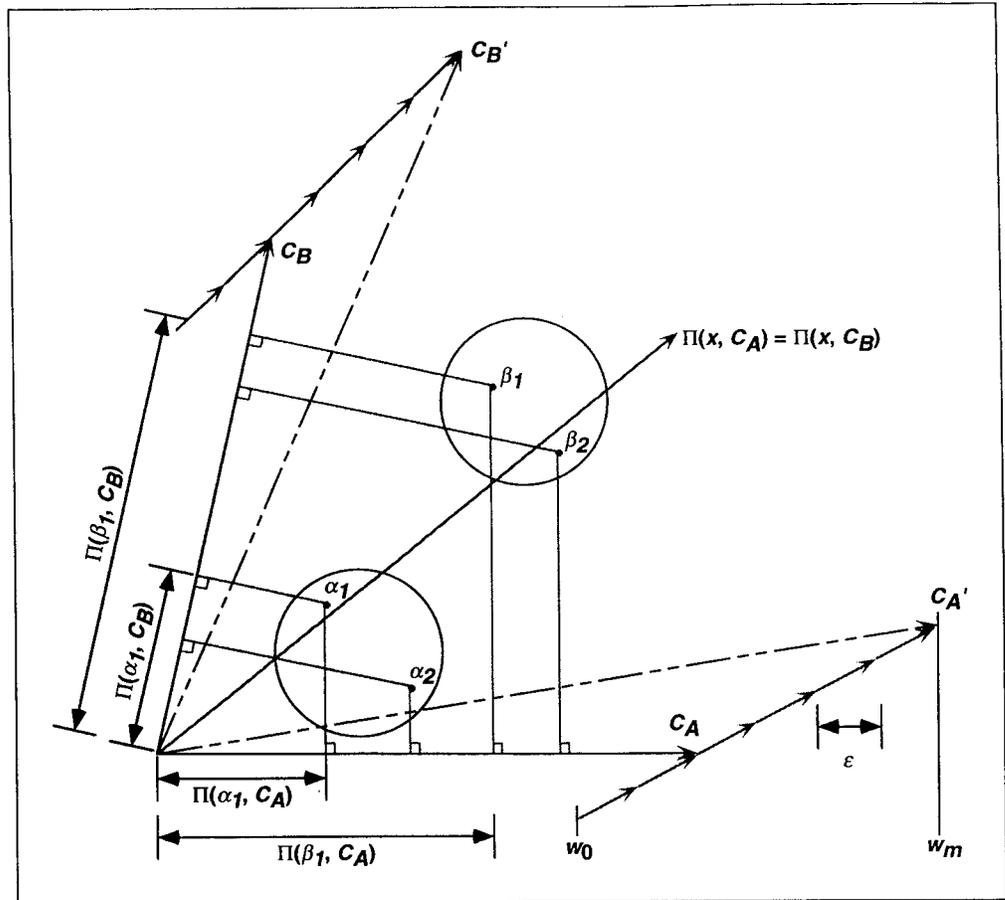
projection onto cell  $C_A$  is equal to that for cell  $C_B$ , i.e., inputs evoking equal-sized responses from both target cells. Inputs above the line will evoke a larger response from cell

below the line. The inputs in the figure have been placed such that most variants of A tend to generate larger responses (i.e., have longer orthogonal projections onto) in cell A than cell B, and vice versa for variants of B. In terms of physical characteristics, variants of A are much more similar to each other than any of them are to variants of B. Thus even a quite noisy or distorted version of A cannot be confused as being a B on the basis of physical characteristics. Nonetheless, the only internal "representation" of the inputs that is possible in this cell population is the activity trace of those cells. In the figure it can be seen that it is possible for the cells to "mistake" one for the other, i.e., for almost all of the variants of A, cell  $C_A$  responds more strongly than cell  $C_B$ , yet for a few variants of A, the reverse is true: cell  $C_B$  will respond more strongly than  $C_A$  since those variants are closer to the  $C_B$  vector. Misclassifications or errors of this kind occur as a function of the "fit" between the types of inputs that occur and the placement (connectivity) of cells in the input space.

It will be shown that

1. LTP-based learning broadens the radius of the effective receptive field around inputs, such that noisy versions of the input that fall within that radius will be correctly responded to after learning.

**Figure 3.** Points  $\alpha_1, \alpha_2, \beta_1,$  and  $\beta_2$  can be thought of as spatial representations of four different inputs to the network. If there are  $N$  afferents, then inputs are points in  $N$ -dimensional space where each point's value along any particular dimension depends on whether that afferent is active or not. The activation level of a cell is proportional to the projection of its input onto the vector representation of the cell itself. [For example,  $\Pi(\alpha_1, C_A)$  is the projection of input  $\alpha_1$  onto cell  $C_A$ .] Point  $\alpha_1$  has a larger projection onto cell  $C_B$  than cell  $C_A$ , so it is mistakenly associated with pattern B despite the fact that cell  $C_A$  was trained on inputs more similar to  $\alpha_1$ . **Network error**, therefore, is the chance of this situation happening. Point  $\beta_1$ , however, has a greater projection onto the cell on which B inputs were trained, so the network will perform a correct classification. As more inputs are presented to a cell, the cell will reduce its angle from the inputs by increasing its value along those dimensions common to both its synapse locations and the active afferents. Both the size of each increase,  $e$ , and the minimum and maximum value of a cell along any dimension where a synapse is present,  $w_0$  and  $w_m$ , are global parameters of the network (step and ceiling size). *Note: Two aspects of this diagram may be potentially misleading, but are necessary to visualize the spatial representation. First, all inputs from both patterns are actually the same distance from the origin. Here we see only a two-dimensional slice of the input space, so inputs appear to be of different lengths. Second, the activation of a cell is the inner product of the synaptic weights and the input, but here we instead show the inputs' projection onto each cell. Projection can also be used to compare cell activations as long as the cells' vector representations are the same length. Although this is only approximately the case in the model, projection is much easier to visualize in this diagram than inner product.*



2. The breadth of the receptive field is inversely related to the LTP increment (step size).

3. The breadth of the receptive field is proportional to the distance from naive to saturated (ceiling) synaptic weights.

### Theoretical Treatment

The target cortical network is assumed to have  $N$  afferent input lines making  $S \ll N$  (sparse) synaptic contacts with the dendrites of each of  $M$  target cells. The naive initial synaptic efficacy (weight) of each synapse is  $w_p$ . Input stimuli consist of  $I$  of the input lines being active at a given time. We consider the case of two distinct groups of inputs, corresponding to the variants of A and B as above, as depicted in Figure 1. Formally, these groups are defined such that each variant of input A

consists of the conjoint activation of three subsets of input lines: a set  $N_A$  unique to A, a set  $N_{\text{shar}}$  shared by all As and Bs, and random noisy inputs  $I_{\text{noise}}$ . Each input consists of a fixed number  $I$  of active afferents, such that  $N_A + N_{\text{shar}} = I$ , noise is added to each such input by random deletion of  $I_{\text{noise}}$  of these  $I$  active input lines, and activation of  $I_{\text{noise}}$  different randomly chosen lines. A "region" corresponds to the part of a cell that may contact a particular set of afferents. For instance, the  $N_{\text{shar}}$  set of input lines pass over and make sparse contact with the *shared* region of each cell, also depicted in Figure 1.

The sparse random connectivity of the input lines to the target cells causes specific inputs to differentially activate different targets, as a function of the fit between the  $I$  active input lines (out of the total  $N$  lines) and the  $S$  synapses present on a target dendrite. Before any learn-

ing takes place, the response of a cell to a particular input is  $a \cdot w_0$ , where  $a \leq I$  is the number of active inputs that contact the cell, i.e., the number of active synapses on this cell for this input. The probability that  $a$  equals a given value,  $i$ , follows the hypergeometric distribution:

$$\Pr[a = i] = \frac{\binom{S}{i} \binom{N-S}{I-i}}{\binom{N}{I}} \quad (1)$$

The expected value and variance of  $a$  are as follows:

$$E[a] = I \frac{S}{N} \quad (2)$$

$$\text{Var}[a] = I \left( \frac{S}{N} \right) \left( \frac{I-S}{N} \right) \left( \frac{N-I}{N-1} \right) \quad (3)$$

When the variance is sufficiently large, the hypergeometric distribution can be approximated by the normal distribution with the same mean and variance ( $\mu = E[a]$ ,  $\sigma^2 = \text{Var}[a]$ ). For a given input pattern, we consider just the cell that responds most strongly to it (neglecting ties) among the cells that we are considering. We wish to find the mean and variance of the activation of the cell that most strongly respond to this input. We therefore seek the point  $z$  at which the cumulative distribution function of the normal distribution reaches the value of  $M/(M+1)$ , where  $M$  is the total number of potential target cells:

$$\Phi(z) = \frac{M}{M+1}$$

[For instance, if there are 10 cells ( $M = 10$ ), we find the point where  $\Phi(z) = 0.91$ , i.e., where the area under the normal curve to the left of  $z$  represents 91% of the possible values of  $a$ . Then, the expected maximally responding  $a$  value will approximately be

$$E[a_{\text{win}}] \approx E[a] + zSD[a] \quad (5)$$

(Note that this is an approximation because we are using the normal approximation to the hypergeometric distribution.) This target cell activation is altered with potentiation. The mean and variance of the activation, after several potentiation episodes, depend on the LTP step size  $\epsilon$  and the distance,  $c$ , to the maximal or ceiling potentiation,  $w_m: c = w_m - w_0$ . The following derivation identifies the relations between these variables.

#### *Estimation of Mean and Standard Deviation of Cell Activation after Potentiation*

Every time an input is presented to a cell, the weights of activated synapses are increased by  $c$ . We will compare the differential effects of using different values of

this step size. To make the effects comparable, we calculate cell activation after training up to a given equal level regardless of step size, i.e., each cell receives  $(w_m - w_0)/\epsilon$  training trials. We now wish to find the mean and variance of the activation that will occur in each of the two target cells (A and B) in response to either one of the inputs (e.g., A). This will enable calculation of the probability of a "correct" response, that is, the probability that the A cell will respond to an input that is a variant of A more than the B cell will, and vice versa. For the purposes of this analysis, we make the simplifying constraint that the cell that responds most strongly to the first (noisy) version of A is trained on all subsequent noisy versions of A, and likewise a distinct target cell is trained on input B and its noisy variations. Figure 5 (whose details are discussed later) compares the empirical and theoretical performance of cells trained under our simplified paradigm versus cells trained under a strict winner-take-all unsupervised paradigm. Although in the latter case some pattern A inputs will potentiate the cell that first responds to a B input, and vice versa, under low to moderate noise conditions the difference between the two paradigms is not large, and for the qualitative conclusions reached in this paper, not important.

As before, each cell has  $N$  "positions," each of which may either receive a synaptic contact from an input line or not. (Since connectivity is sparse, many of the positions receive no contact.) We then divide the positions into three sets:  $N_{\text{shar}}$  represents the positions of the cell that coincide with the shared input lines,  $N_A$  coincides with the input lines of the pattern that will train this cell, and  $N_{\text{noise}}$  represents the remainder of the positions that will be trained only by the portion of the  $I_{\text{noise}}$  input lines whose activity is displaced into this region. (All equations use the size of these sets rather than their membership, so the name of the set and its cardinality may be used interchangeably.) From each of these sets, we wish to find the mean contribution to a cell's activation during training, and the mean and variance of the contribution given a test input once training has been completed. From these values, Appendix A derives the probability that a given cell will correctly respond to a version of an input similar to that on which it was trained:

$$p(Y - X < 0) = \Phi \left\{ \frac{E[X(\epsilon, c, \dots)] - E[Y(\epsilon, c, \dots)]}{\sqrt{\text{Var}[X(\epsilon, c, \dots)] + \text{Var}[Y(\epsilon, c, \dots)]}} \right\} \quad (6)$$

where  $X$  and  $Y$  are random variables for the activation of a cell in response to a version of the input on which it was trained versus a version of inputs on which it was not trained, respectively, and  $\Phi(\cdot)$  is the cumulative distribution function of the normal distribution (which approximates each cell's activation). The following equations show the dependency of this probability on the

values of the step size  $\epsilon$  and distance to maximum potentiation  $c = w_m - w_0$ , respectively.

Without loss of generality, we can estimate the mean and variance of a cell tested on (noisy versions of) the same pattern as it has been trained on, by estimating the mean and variance of a cell trained on noisy variants of A (a "pattern A" input) and tested on one such variant of A. Likewise, we can estimate the mean and variance of a cell's response to a different pattern than it was trained on by estimating the mean and variance of a cell trained on pattern B inputs but tested on a pattern A input. Let  $c = w_m - w_0$  be the distance to maximum potentiation,  $t_{\text{set}}$  be the number of active synapses in a set of synapses (e.g., a shared or pattern A-specific region) given a test input (e.g.,  $t_{\text{shar}}$  is the number of active synapses in the *shared* region),  $p(t_{\text{set}} = i)$  be the probability of a particular value of  $t_{\text{set}}$  (see Appendix A for the formulas for this and the following two terms),  $r_{\text{set}}$  be the percentage of synapses modified within a set on a single training input pattern presentation,  $E[I_{\text{set}}]$  be the expected number of active input lines to a region during training (once noise has deleted some of the active inputs in the  $N_{\text{set}}$  region),  $X_{\text{set}}$  be the activation level of a region trained and tested on pattern A inputs, and  $Y_{\text{set}}$  be the activation level of a region trained on pattern B inputs and tested on a pattern A input. Then,

$$E[X_{\text{set}}] = \sum_i p(t_{\text{set}} = i)(icr_{\text{set}} + iw_0) \quad (7)$$

for the shared, pattern B, and noise sets, and

$$E[X] = E[X_{\text{shar}}] + E[X_a] + E[X_{\text{noise}}] \quad (8)$$

Likewise, we can calculate the variance:

$$\begin{aligned} \text{Var}[X_{\text{set}}] = & \sum_i p(t_{\text{set}} = i) \left[ i^2 r_{\text{set}}^2 (1 - r_{\text{set}}) \left( \frac{N_{\text{set}} - E[I_{\text{set}}]}{N_{\text{set}} - 1} \right) + (icr_{\text{set}} + iw_0)^2 \right] \\ & - \left[ \sum_i p(t_{\text{set}} = i)(icr_{\text{set}} + iw_0) \right]^2 \end{aligned} \quad (9)$$

and

$$\text{Var}[X] = \text{Var}[X_{\text{shar}}] + \text{Var}[X_a] + \text{Var}[X_{\text{noise}}] \quad (10)$$

The equations for  $E[Y_{\text{set}}]$  and  $\text{Var}[Y_{\text{set}}]$  are similar to those for  $E[X_{\text{set}}]$  and  $\text{Var}[X_{\text{set}}]$ , respectively:

$$E[Y] = E[Y_{\text{shar}}] + E[Y_A] + E[Y_B] + E[Y_{\text{neither}}] \quad (11)$$

and

$$\begin{aligned} \text{Var}[Y] = & \text{Var}[Y_{\text{shar}}] \\ & + \text{Var}[Y_A] + \text{Var}[Y_B] + \text{Var}[Y_{\text{neither}}] \end{aligned} \quad (12)$$

except we further divide the cell into four regions by separating the *noise* region into regions *B* and *neither*. Since we are looking at a cell that is trained on one

pattern and tested on another, it helps to separate the cell into the region that gets the inputs shared by both patterns, the region that gets most of the pattern-specific training inputs (B), the region that gets most of the pattern-specific testing inputs (A), and the region that lies outside both pattern A and pattern B inputs except for noise. [These distinct "regions" are assumed to be distributed across the dendrite, since the inputs have no topographic organization. Moreover, as mentioned earlier, for the short electrotonic dendritic lengths identified in many modeling studies (Rall et al., 1992) the proximal vs distal location of specific inputs can be neglected.] When calculating the probability of active input synapses,  $p(t_{\text{set}} = i)$ , for cells trained on pattern B inputs and tested on A, the reversal of the number of input lines specific to each pattern must be taken into account (as discussed in Appendix A).

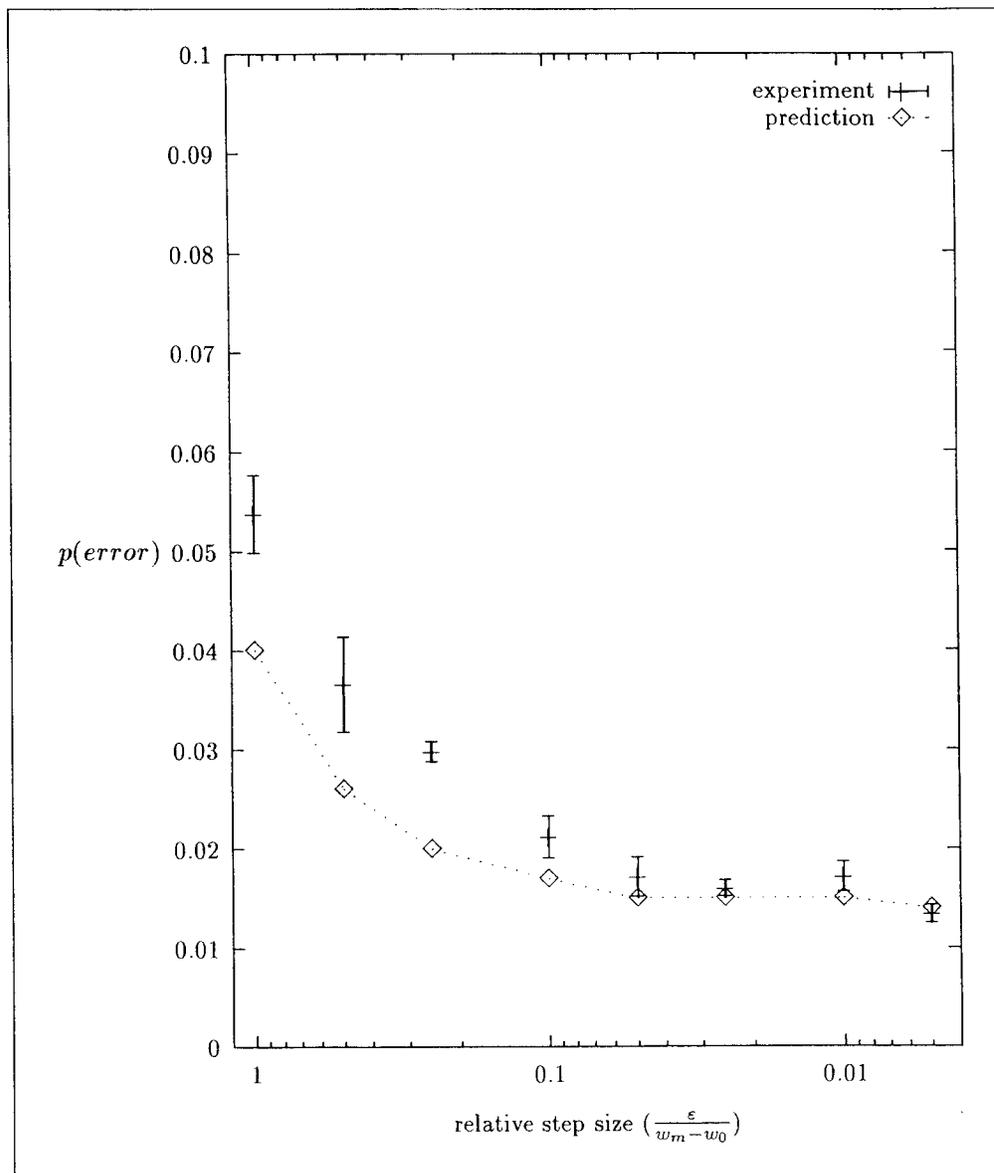
### Errors as a Function of Step and Ceiling

Figure 4 shows the effect of step size  $e$  on the probability of an error, i.e., the probability of the target network misrepresenting an A as a B or vice versa. Shown are both analytical results based on the above equations, and empirical results based on tests of a network implementing the described algorithm. It can be seen that the empirical results parallel the analytical results (although actual errors are slightly higher than analytical predictions for some large step sizes). Performance improves (error decreases) with decreasing step size. It is worth noting that the improvement in performance falls off once the step size reaches roughly 15 to 20% of the distance from the naive to the ceiling weight, i.e., the step size is such that there are roughly seven steps before the ceiling is reached.

A similar result is obtained for the recognition performance of the network as a function of the ceiling, i.e., the ratio of the maximal potentiated strength of a synapse to its naive strength (Fig. 5). Again, empirical and analytical results agree relatively closely; performance improves (errors decrease) as the ceiling increases. As before, the effect is highly nonlinear, and there is an obvious inflection point at which increases to the ceiling offer diminishing returns—there is little performance gained in raising the ceilings to a level more than three or four times larger than the naive synaptic strength.

Figure 6 shows the probability of error for a fixed step and ceiling [of  $0.1(w_m - w_0)$  and 2, respectively], as a function of the difficulty of the discrimination between the two inputs A and B. The more the two inputs share input lines, the more similar they are, and the more difficult it is to discriminate between them. (This corresponds to the points becoming closer and closer together in Figure 3.) As expected, more similar inputs yield higher probability of error. The effect interacts with the amount of random noise in the signal; if noise is restricted to only 10% of the input signal, error prob-

Figure 4. Predicted error and experimental results versus step size  $\epsilon$ . Both predictions and experiment observe the relative responses of two cells in a net of size twenty ( $M = 20$ ). Each cell has 300 synapses out of 1000 potential locations ( $N = 1000$ ,  $S = 300$ ) and is trained on inputs that have 100 active lines each ( $I = 100$ ). Input patterns are constructed from 30 active inputs lines common to both patterns and 70 unique to each ( $I_{\text{shar}} = 30$ ,  $I_A = I_B = 70$ ). Each training and testing input to the net is distorted by randomly displacing 40 input lines from meaningful positions ( $I_{\text{noise}} = 40$ ). The ratio of ceiling to naive synaptic weight ( $w_m/w_0$ ) is 4. Effects of step size are less noticeable as this ratio decreases. Each cell is trained as many times as the reciprocal of the step size times the difference between the ceiling and naive weights [ $(w_m - w_0)/\epsilon$ ]. Experimental results are obtained by determining the results of training and testing five sets of 20 nets at each step size. After training, each net is tested on 200 inputs from each pattern. The error represents the percentage of noisy inputs that produce a higher activation level in the cell not trained on that input's pattern than in the cell trained on that input's pattern. Error bars reflect the standard error among the average error probability within each of the five sets. Although dividing 100 test nets into five sets of 20 is arbitrary, we obtain similar error results from other divisions of the nets.



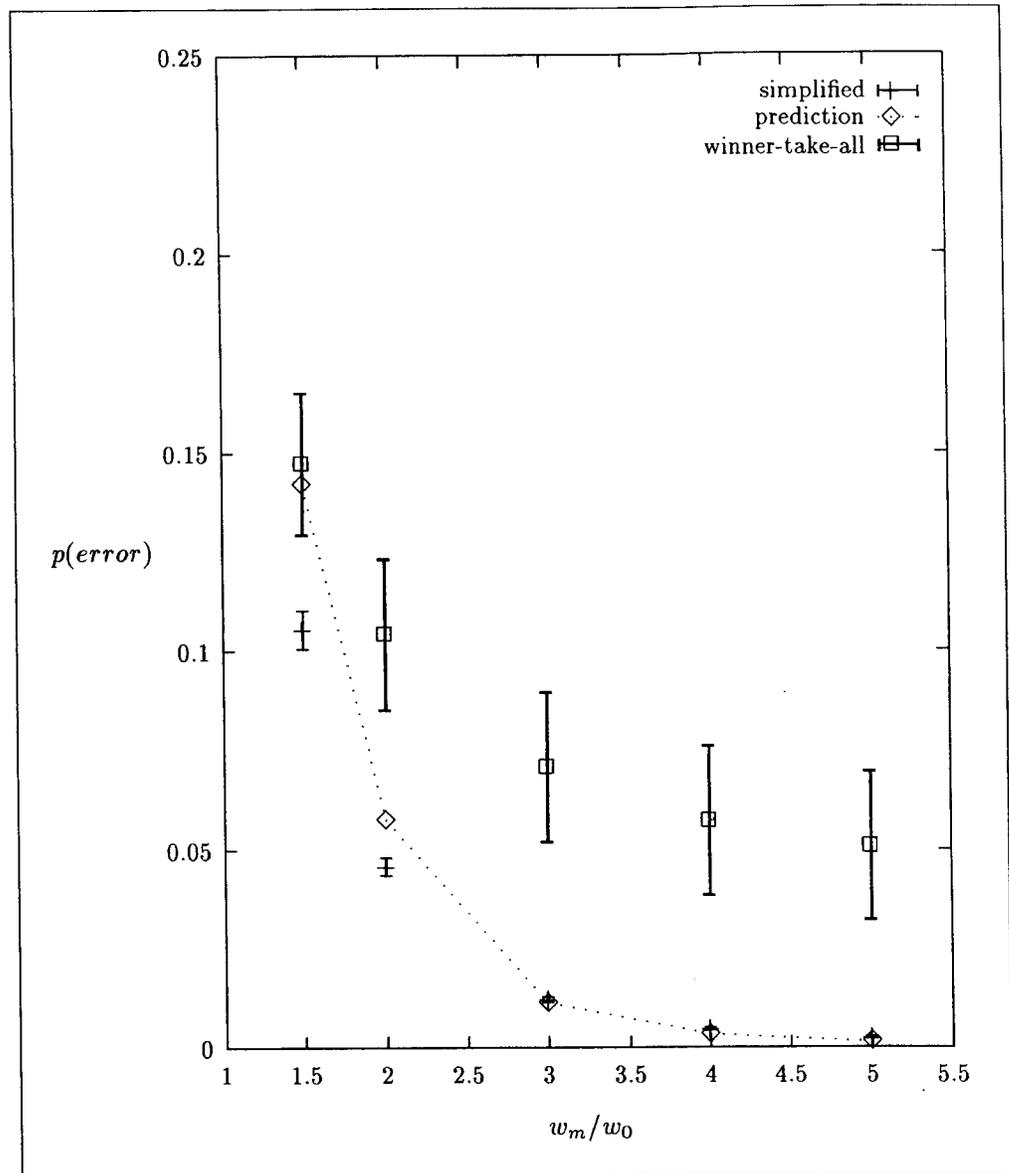
ability remains below 0.1 even when the two inputs overlap by almost 50%.

The effect of noise on the inputs, illustrated in Figure 7, exhibits a similar pattern. The graph shows the error probability for given amounts of noise, given two inputs that are difficult to discriminate (50% overlap). This effect interacts strongly with the ceiling size: for very low ceilings, noise dramatically impairs performance, but for higher ceilings, noise has less effect. It can be seen that increasing the ceiling past a certain point has diminishing returns. (To illustrate the asymptotic limit of the effect of the ceiling, the unrealistic ceiling value of 10,000 times the naive weight is included for comparison.)

Figure 8 summarizes some of these findings. Shown are percentage changes in receptive field radii corre-

sponding to various values of  $\epsilon$ , step size and ceiling. Changes to the ceiling have a much larger relative impact that changes to step size. Decreases in step size improve performance only by a few tenths of a percent; increases in ceiling size can improve performance by tens of percents. In the figure, receptive field radius is shown for three different step sizes and for four different ceiling heights, given a fixed error probability of 0.25 as discussed earlier. It can be seen that decreasing the step size barely expands the field radius, whereas increasing the ceiling expands the field by a much larger relative degree. The diminishing returns afforded by increased ceiling are also clear here: after increasing the LTP ceiling by a factor of 3, it does little good to increase it further at this error level.

**Figure 5.** Predicted error and experimental results versus ceiling to naive ratio ( $w_m/w_0$ ). Again,  $M = 20$ ,  $N = 1000$ ,  $S = 300$ ,  $I = 100$ ,  $I_{\text{shar}} = 30$ ,  $I_A = I_B = 70$ , and  $I_{\text{noise}} = 40$ . In the simplified experiment, each cell is trained on 10 inputs from the pattern it first prefers. In the strict winner-take-all experiment each cell is trained on inputs for which it has the highest activation in the patch. In both experiments the step size is one-tenth the difference between the ceiling and naive weights [ $\epsilon = 0.1(w_m - w_0)$ ]. Error is calculated in the same manner as described in the previous figure.



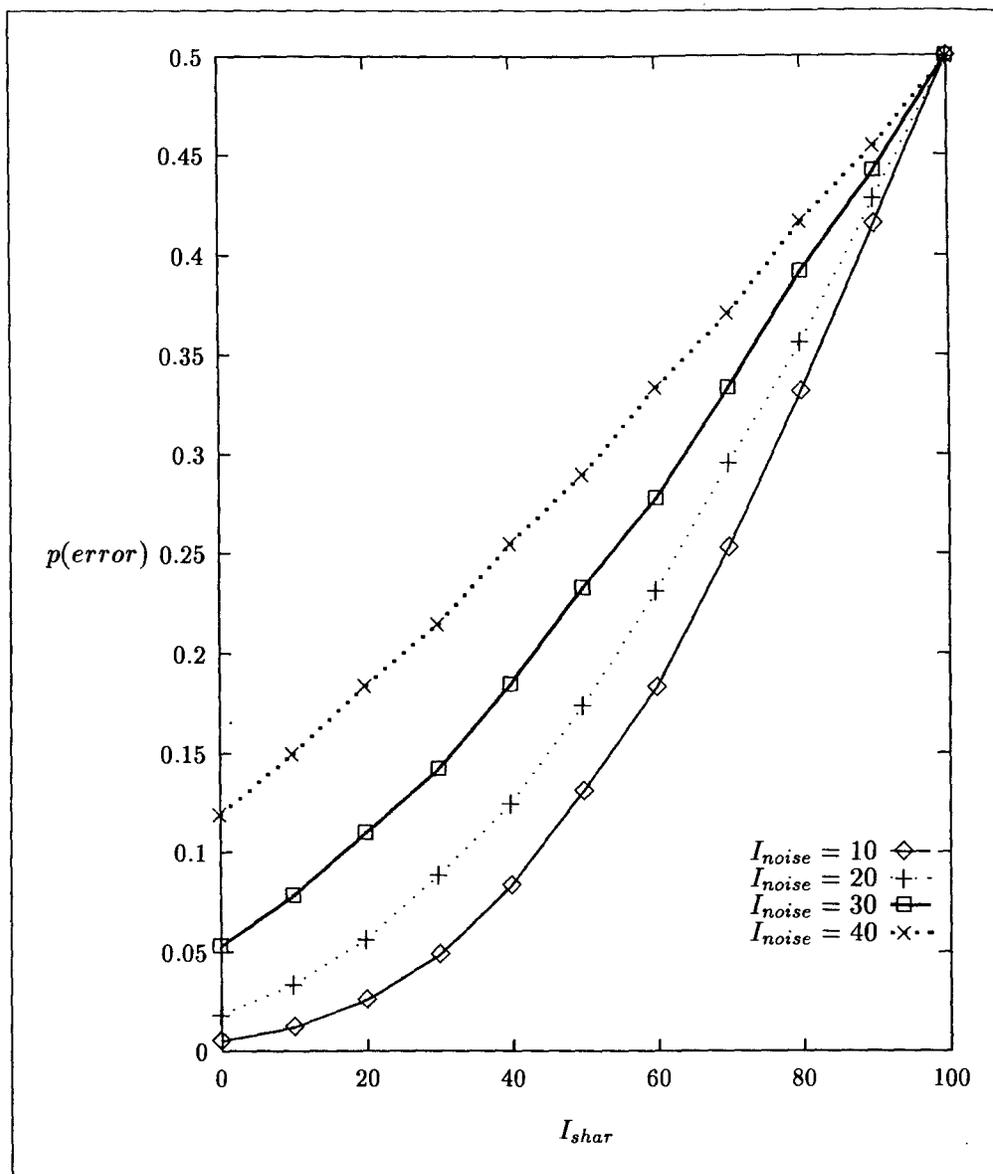
## DISCUSSION

Small LTP step size, then, yields two effects, one obvious and one less so: (1) many such steps must be taken to achieve a given level of potentiation, and (2) after learning, the target cell population will come to respond correctly to a somewhat wider range of inputs, effectively "recognizing" even noisy versions of inputs on which it has been trained. This latter effect corresponds to the receptive field of the target cells. The increment size and LTP ceiling each affect the receptive field breadth, and hence affect the restrictiveness of the response of the cell thus trained. If an input A generates a particular cortical response, how similar must input A' be to A to generate that same cortical response? If the LTP step is larger, or the ceiling lower, then the target receptive field is narrowed, and thus A must be very similar to A' to generate the same response. If the LTP

step is smaller or the ceiling higher, then the receptive field is broadened, and the required similarity between A and A' is reduced.

The effects of both step and ceiling are quite nonlinear, and a tradeoff between them can be identified. Increases to ceiling and decreases to step size both broaden receptive fields. Higher ceilings and smaller step sizes share the cost that they cause maximal potentiation to take longer to reach. Thus presumably training will be more protracted as ceiling is raised and step size decreased. The nonlinearities in the analysis provide inflection points that can be chosen to optimize this tradeoff. From the graphs it can be seen that for the parameters used in the simulation (such as connectivity and input size), the chosen ceiling size should be roughly 2 to 2.5 times larger than naive synaptic strengths, and the chosen step size should traverse roughly one-seventh of the distance from naive to maximal weight with each poten-

Figure 6. Predicted probability of error given various degrees of similarity between the two patterns. The predictions are for two cells which have 300 synapses out of 1000 potential locations each ( $S = 300$ ,  $N = 1000$ ). There are 100 active inputs ( $I = 100$ ),  $I_{\text{rar}}$  of which are common to both patterns and the rest of which are unique to each pattern ( $I_A = I_B = 100 - I_{\text{rar}}$ ). Each curve shows the predictions for a different level of noise,  $I_{\text{noise}}$ . The ceiling to naive ratio ( $w_m/w_0$ ) is 2, and the step size is one-tenth the difference between the ceiling and naive weights [ $\epsilon = 0.1(w_m - w_0)$ ]. Each cell is trained on 10 examples of one pattern.



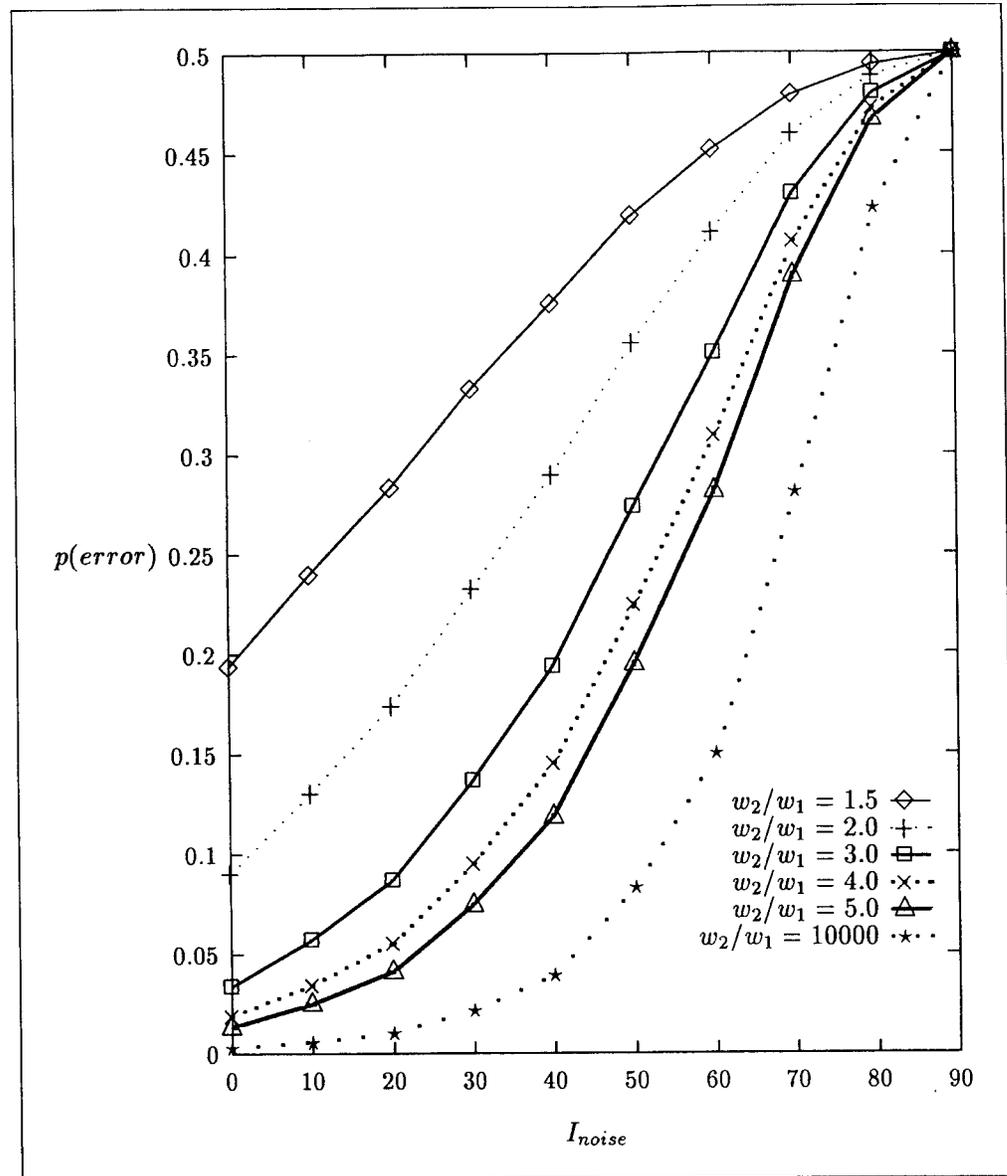
tiation episode. It is worth noting that these analytical results do not differ greatly from values experimentally identified in *in vitro* brain slices (Larson & Lynch, 1986; Larson, Wong, & Lynch, 1986).

Pharmacological agents that increase the step size and ceiling of LTP (Arai & Lynch, 1992) have been shown to affect the learning rate of animals to which these agents have been administered (Granger et al., 1993; Staubli, Rogers, & Lynch, 1994). From the findings presented here, it can be predicted that the resulting learned information will be formed with slightly more restrictive receptive fields than if learned via training with unenhanced step sizes. A larger effect should obtain for increase in ceiling size. That is, increase in LTP ceiling should broaden the receptive fields surrounding the learned items, and this effect will be large enough to overwhelm effects that may be due to step size. Thus if both step size and ceiling are increased by a drug, it is

predicted that the resulting target cell receptive fields will be broadened compared to learning that occurs without the drug. The behavioral consequences of this are unknown, although it might be loosely speculated that operational generalization gradients around learned stimuli may correspond to the receptive field sizes of the target cortical cells.

Convergent anatomical and physiological evidence suggests that step size and ceiling may be readily controlled endogenously via the effects of ascending modulatory systems on their GABA interneuron targets. Immunohistochemical work has indicated that ascending serotonergic (5HT) fibers preferentially terminate on GABA<sub>B</sub> interneurons (Seress et al., 1993) and pharmacological manipulation of the 5HT system suggests that its effect on GABA<sub>B</sub> cell activity is facilitatory. Blockade of the GABA<sub>B</sub>-mediated long hyperpolarization by forskolin has a dramatic effect on LTP: the magnitude of

Figure 7. Predicted probability of error versus noise. The predictions are once again for two cells, with  $N = 1000$ ,  $S = 300$ , and  $I = 100$ . There are 50 input lines similar to both patterns and 50 unique to each ( $I_A = I_B = I_{\text{shar}} = 50$ ). Each curve shows the predictions for a different ratio between the ceiling and the naive synaptic weight. The  $w_m/w_0 = 10000$  curve represents an "ideal" synapse, which has negligible initial weight. All predictions are for cells trained 10 times with  $\epsilon = 0.1(w_m - w_0)$ .

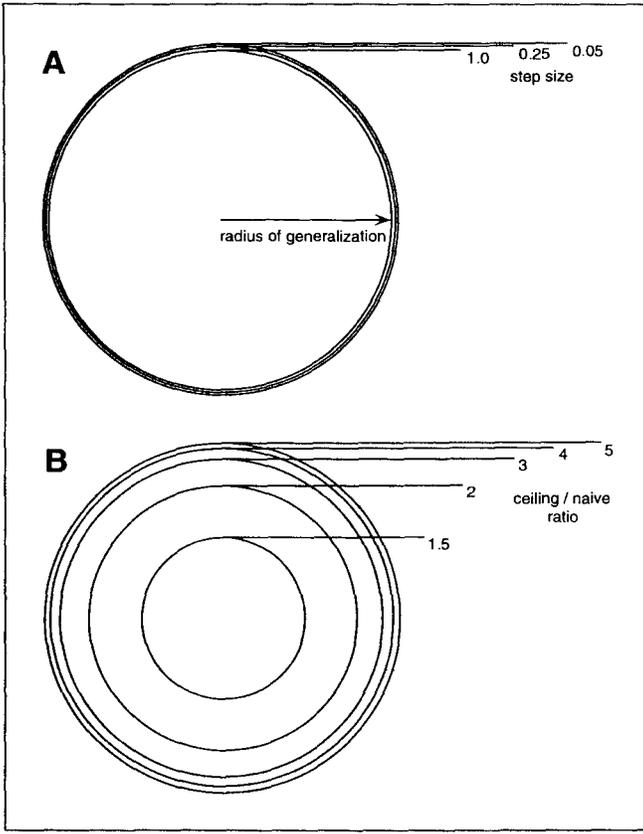


the maximal LTP nearly doubles (Arai & Lynch, 1992). The mechanism underlying this effect is not known, but it is reasonable to assume that the extra hyperpolarization due to 5HT-mediated GABA<sub>B</sub> enhancement will abbreviate any NMDA currents activated during a burst, since the duration of those currents are dependent on ongoing depolarization. Calcium is the trigger for LTP induction, and shortening of NMDA receptor opening would lessen the amount of the  $\text{Ca}^{2+}$  influx into the spine, possibly reducing the amount of LTP that can be induced. By this same reasoning, suppression of the 5HT system should allow more LTP to be induced, and this has recently been shown using the selective 5HT<sub>3</sub> inhibitor Ondansetron (Staubli, personal communication). If this modulator does raise the LTP ceiling, it is predicted that the resulting learning will exhibit broader receptive fields than comparable learning without the modulator.

Just as the 5HT system preferentially targets GABA<sub>B</sub>

interneurons, there is evidence that the ascending noradrenergic and cholinergic (NE and ACh) systems terminate on GABA<sub>A</sub> interneurons (Toth, Borhegyi, & Freund, 1993; Miettinen & Freund, 1992; Seress et al., 1993) enabling their modulation of GABA<sub>A</sub>-mediated fast IPSPs. This modulation should affect IPSP strength and may affect the strength or the latency of the refractory period for the GABA<sub>A</sub> synapse (Larson et al., 1986; Mott & Lewis, 1991). If the IPSP is strengthened, the amount of LTP that can be induced per step should be reduced; if the IPSP is weakened, more UP should be able to be induced per induction episode; i.e., the LTP step size would be increased. A modulator weakening the IPSP should thus enable faster learning, with the side effect of slightly narrowing the receptive fields resulting from the training.

Finally, it should be noted that the size of the UP ceiling varies in the same neuron across different den-



**Figure 8.** The radius of each concentric circle graphically represents the amount of generalization a network will tolerate for a given set of LTP parameters. Tolerance in (A) and (B) is arbitrarily set at 25%, i.e., we allow one-quarter of the examples that a cell prefers to belong to a pattern that the cell was not trained on. (A) Effect of step size  $\epsilon$  on generalization ability given  $N = 1000$ ,  $S = 300$ ,  $I_A = I_B = 50$ , and  $I_{\text{shar}} = 50$ . The ceiling to naive ratio ( $w_m/w_0$ ) is 5. (B) Effect of  $w_m/w_0$  on generalization ability given the same parameters as (A) and  $\epsilon = 0.1(w_m - w_0)$ .

driftic loci (Woodward, Chiaia, Teyler, Leong, & Coull, 1990; Arai, Silberg, & Lynch, 1995; Kolta, Larson, & Lynch, 1995) and it appears that this may hold true for the step size as well (Arai et al., 1995). This raises the possibility that networks form multiple representations of noisy inputs with some being restrictive and others broad in their receptive fields. Such a system, if coupled with variations in the stability of potentiation, could allow for the gradual sharpening of categories.

## APPENDIX A

### Activation of a Cell Trained and Tested on Pattern A Inputs

We divide all  $N$  potential locations of the cell (and input lines) into three sets:  $N_{\text{shar}}$  positions that are common to both patterns A and B inputs,  $N_A$  positions that are specific to pattern A inputs, and  $N_{\text{noise}}$  positions that correspond to the remainder:

$$N_{\text{noise}} = N - N_{\text{shar}} - N_A \quad (13)$$

Next, we calculate the expected number of active input lines to each set of potential synaptic locations. For the shared and pattern A sets,  $I_{\text{shar}}$  and  $I_A$ , respectively, this will equal the number of positions in that set (corresponding to a noiseless input pattern) less a proportional amount of positions that will be displaced by noise:

$$E[I_{\text{shar}}] = N_{\text{shar}} - I_{\text{noise}} \left( \frac{N_{\text{shar}}}{I} \right) \quad (1-1)$$

$$E[I_A] = N_A - I_{\text{noise}} \left( \frac{N_A}{I} \right) \quad (15)$$

If  $E[a_{\text{win}}]$  is the mean number of active synapses on the winning cell (the cell with the highest activation) for some noisy input, then we can assume that on average the shared and pattern A regions of that cell will have a slightly higher density of synapses than the remainder of the cell. Since a certain proportion of the entire input comes from the shared and pattern A input lines, an equal proportion of the active synapses must come from those sets. Where noise displaces the activity of some of the input lines to other regions, we can only assume that there is the same density of synapses as the remainder of the cell. In other words, for at least some of the positions in the shared and pattern A sets, we expect a higher probability of a synapse existing at any particular position than in an average cell, since this cell did produce a higher activation on a (noisy) version of the pattern than did other cells. The remainder of the cell will have a slightly lower probability of a synapse existing at any particular position to compensate. We can calculate the expected number of synapses ( $E[S]$ ) in each of these regions as follows:

$$E[S_{\text{shar}}] = E[a_{\text{win}}] \left( \frac{E[I_{\text{shar}}]}{I} \right) + I_{\text{noise}} \left( \frac{N_{\text{shar}}}{I} \right) \left\{ \frac{S - E[a_{\text{win}}] \left( \frac{E[I_{\text{shar}}] + E[I_A]}{I} \right)}{N - E[I_{\text{shar}}] - E[I_A]} \right\} \quad (16)$$

$$E[S_A] = E[a_{\text{win}}] \left( \frac{E[I_A]}{I} \right) + I_{\text{noise}} \left( \frac{N_A}{I} \right) \left\{ \frac{S - E[a_{\text{win}}] \left( \frac{E[I_{\text{shar}}] + E[I_A]}{I} \right)}{N - E[I_{\text{shar}}] - E[I_A]} \right\} \quad (17)$$

$$E[S_{\text{noise}}] = N_{\text{noise}} \left\{ \frac{S - E[a_{\text{win}}] \left( \frac{E[I_{\text{shar}}] + E[I_A]}{I} \right)}{N - E[I_{\text{shar}}] - E[I_A]} \right\} \quad (18)$$

To calculate the mean and variance of the activation of a trained cell on an input, we can calculate the means and variances of the contribution of each region. If  $X$  is the activation of a trained cell on an input, then

$$E[X] = E[X_{\text{shar}}] + E[X_A] + E[X_{\text{noise}}] \quad (19)$$

Since dependencies between the contributions of each region are not significant, we consider the contributions independent, and thus:

$$\text{Var}[X] = \text{Var}[X_{\text{shar}}] + \text{Var}[X_A] + \text{Var}[X_{\text{noise}}] \quad (20)$$

In presenting the following method of estimating variance, we will look at only one region at a time.  $N_{\text{set}}$  stands for  $N_{\text{shar}}$ ,  $N_A$ , or  $N_{\text{noise}}$ , depending on the region for which we are calculating variance. The same applies to  $E[S_{\text{set}}]$ ,  $E[a_{\text{set}}]$ , and  $E[I_{\text{set}}]$ . Every time a cell is trained on an input, the weight of a certain number of synapses is increased. For simplicity, we will assume that an average number,  $E[a_{\text{set}}]$ , synapses are increased during each training input pattern presentation (the mean number of active synapses in the set). We can think of weight increments as balls and synapses as bins: every time we train the cell, we randomly throw  $E[a_{\text{set}}]$  balls into  $E[S_{\text{set}}]$  bins. At first, some bins will be empty while others will contain a widely varying amount of balls. As we throw more balls in, however, the relative number of balls in each bin will tend to even out.

We calculate  $E[a_{\text{set}}]$  as the mean of the hypergeometric distribution of  $E[I_{\text{set}}]$  trials with  $E[S_{\text{set}}]$  good outcomes out of  $N_{\text{set}}$  total possible outcomes:

$$E[a_{\text{set}}] = E[I_{\text{set}}] \left( \frac{E[S_{\text{set}}]}{N_{\text{set}}} \right) \quad (21)$$

If the learning step size is large relative to the difference between the ceiling and naive weight ( $c = w_m - w_0$ ), then each of those balls represents a larger weight increment and we will throw in fewer of them. This will prevent us from getting to the point of having nearly identical weights at each synapse. However, if the step size is small, we will have more chance to even out the weights. Since the probability of an active input line in any position within a region is equal, the lower the variance among the synaptic weights, the lower the variance in that region's contribution to the activation level.

When we determine the activation of a cell on a test input pattern, we sum the weights of all active synapses. If there are  $t_{\text{set}}$  active synapses, and the cell has been trained on  $c/\epsilon$  inputs, then we are effectively making  $ct_{\text{set}}/\epsilon$  observations of a hypergeometrically distributed random variable, which represents whether or not a certain synapse was trained at a given training step.

The mean activation of a hypergeometric distribution is the product of the number of trials and the probability of success. In our case, we must scale the mean by dividing it by the number of steps, since each success represents a weight increment of only  $\epsilon$ . Furthermore, we must add any initial weight the synapses already had:

$$E[X_{\text{set}}] = \epsilon \left( \frac{ct_{\text{set}}}{\epsilon} \right) r_{\text{set}} + w_0 t_{\text{set}} = ct_{\text{set}} r_{\text{set}} + w_0 t_{\text{set}} \quad (22)$$

where  $r = E[a_{\text{set}}]/E[S_{\text{set}}]$ , the probability of a synapse being active within a region. Note that the mean activation does not depend on the step size according to this equation. (A large step size actually does have a very slight effect in terms of causing a few synapses to hit their ceiling weight early, but we can disregard this.)

The variance can be calculated as follows:

$$\begin{aligned} \text{Var}[X_{\text{set}}] &= \epsilon^2 \left( \frac{ct_{\text{set}}}{\epsilon} \right) r_{\text{set}} (1 - r_{\text{set}}) \left( \frac{N - E[I]}{N - 1} \right) \\ &= \epsilon t_{\text{set}} r_{\text{set}} (1 - r_{\text{set}}) \left( \frac{N - E[I]}{N - 1} \right) \end{aligned} \quad (23)$$

and once again we scale according to the actual size of each weight increment. Note that Eq. (23) depends only on  $\epsilon$  and not  $w_0$  directly, as the naive weight's value affects only the mean activation, not the variance.

We now have the mean and variance of a cell's activation given  $t_{\text{set}}$ , the number of active synapses from a test input. To estimate the mean and variance of the cell across all possible inputs within the category, we must find the probability of different values of  $t_{\text{set}}$  and combine variances according to those probabilities. This gives us a more accurate prediction than simply using  $E[t_{\text{set}}]$  in the above equation, as the number of active synapses in each region contributes largely to the variation of the overall activation of the cell.

Let  $X_{\text{set},i}$  be a random variable with hypergeometric distribution where  $E[X_{\text{set},i}]$  and  $\text{Var}[X_{\text{set},i}]$  are calculated according to the above formulas where  $i = t_{\text{set}}$ . Let  $p(t_{\text{set}} = i)$  be the probability that the number of active synapses is  $i$ , which follows the hypergeometric distribution:

$$p(t_{\text{set}} = i) = \frac{\binom{E[S_{\text{set}}]}{i} \binom{N_{\text{set}} - E[S_{\text{set}}]}{E[I_{\text{set}}] - i}}{\binom{N_{\text{set}}}{E[I_{\text{set}}]}} \quad (24)$$

We want to calculate the mean and standard deviation of the random variable  $X_{\text{set}}$ , where

$$X_{\text{set}} = \sum_i p(t_{\text{set}} = i) X_{\text{set},i} \quad (25)$$

Since

$$E[X_{\text{set}}] = \sum_i p(t_{\text{set}} = i) E[X_{\text{set},i}] \quad (26)$$

and

$$\begin{aligned} \text{Var}[X_{\text{set}}] &= E[X_{\text{set}}^2] - (E[X_{\text{set}}])^2 \\ &= \sum_i p(t_{\text{set}} = i) E[X_{\text{set},i}^2] - \left[ \sum_i p(t_{\text{set}} = i) E[X_{\text{set},i}] \right]^2 \end{aligned} \quad (27)$$

and we know the mean and variance of each  $X_{\text{set},i}$ , we can make the calculation. Note that although we do not

have an explicit formula for  $E[X_{\text{set},i}^2]$ , it falls out of the formula for variance:

$$E[X_{\text{set},i}^2] = \text{Var}[X_{\text{set},i}] + (E[X_{\text{set},i}])^2 \quad (28)$$

We have now incorporated two significant sources of variation into our calculations of a region's contribution to a cell's activation: varying synaptic weights and varying input line activation. These computations are identical to the ones for a cell trained and tested on pattern B inputs, except that the computations for the pattern A- and B-specific regions need to be switched accordingly.

Along with predictions for a cell trained on a different pattern than the test pattern (namely, trained on pattern B inputs and tested on a pattern A input), we can estimate the likelihood of an input generating a higher activation for a cell trained on inputs of the same pattern as opposed to a cell trained on inputs of a different pattern.

### Activation of a Cell Trained on Pattern B Inputs and Tested on a Pattern A Input

Suppose a cell has been trained on pattern B inputs and is tested on a pattern A input. Although we can expect the same behavior in the region of shared input lines, the only lines activating the pattern B region of the cell will be the result of noise in the pattern A test input. Likewise, the pattern A-specific input lines will synapse onto only a certain portion of what is the noise region of the pattern B-trained cell. Because of this, the noise region of the cell should be divided into two parts: the region that coincides with the pattern A-specific ( $N_A$ ) positions of the test input and the region that coincides with the positions of the test input that are part of neither pattern A nor pattern B inputs ( $N_{\text{neither}} = N_{\text{noise}} - N_A$ ). Note that the size of the two pattern-specific regions is equal ( $N_A = N_B$ ).

During training, the pattern B-specific positions of the input will coincide with the pattern B region of the cell, so

$$E[I_B] = N_B - I_{\text{noise}} \left( \frac{N_B}{I} \right) \quad (29)$$

And the noisy lines of the training input will be divided between the pattern A region and the remaining region indicating neither pattern:

$$E[I_A] = I_{\text{noise}} \left( \frac{N_A}{N_A + N_{\text{neither}}} \right) \quad (30)$$

$$E[I_{\text{neither}}] = I_{\text{noise}} \left( \frac{N_{\text{neither}}}{N_A + N_{\text{neither}}} \right) \quad (31)$$

The expected number of synapses in the shared region ( $E[S_{\text{shar}}]$ ) will be the same as a cell trained and tested on

the same pattern. The mean number of synapses in the pattern B, pattern A, and neither pattern regions will be

$$E[S_B] = E[a_{\text{win}}] \left( \frac{E[I_B]}{I} \right) + I_{\text{noise}} \left( \frac{N_B}{I} \right) \left\{ \frac{S - E[a_{\text{win}}] \left( \frac{E[I_{\text{shar}}] + E[I_B]}{I} \right)}{N - E[I_{\text{shar}}] - E[I_B]} \right\} \quad (32)$$

$$E[S_A] = N_A \left\{ \frac{S - E[a_{\text{win}}] \left( \frac{E[I_{\text{shar}}] + E[I_B]}{I} \right)}{N - E[I_{\text{shar}}] - E[I_B]} \right\} \quad (33)$$

$$E[S_{\text{neither}}] = N_{\text{neither}} \left\{ \frac{S - E[a_{\text{win}}] \left( \frac{E[I_{\text{shar}}] + E[I_B]}{I} \right)}{N - E[I_{\text{shar}}] - E[I_B]} \right\} \quad (34)$$

During testing, the number of active input lines in the pattern A and pattern B regions reverses. Therefore, the calculation of  $p(t_{\text{set}})$  must use these new values:

$$E[I_A] = N_A - I_{\text{noise}} \left( \frac{N_A}{I} \right) \quad (35)$$

$$E[I_B] = I_{\text{noise}} \left( \frac{N_B}{N_B + N_{\text{neither}}} \right) \quad (36)$$

The estimation of the mean and standard deviation of the activation of a cell trained on pattern B inputs and tested on a pattern A input ( $Y$ ) can be made according to

$$E[Y] = E[Y_{\text{shar}}] + E[Y_A] + E[Y_B] + E[Y_{\text{neither}}] \quad (37)$$

$$\text{Var}[Y] = \text{Var}[Y_{\text{shar}}] + \text{Var}[Y_A] + \text{Var}[Y_B] + \text{Var}[Y_{\text{neither}}] \quad (38)$$

where  $E[Y_{\text{set}}]$  and  $\text{Var}[Y_{\text{set}}]$  use the same formulas as  $E[X_{\text{set}}]$  and  $\text{Var}[X_{\text{set}}]$  except for the above-mentioned switch of  $E[I_A]$  and  $E[I_B]$  in the calculation of  $p(t_{\text{set}})$ .

### Chance That an Input Will Generate a Higher Activation for the Cell Trained on the Same Pattern

To calculate the probability that a pattern A test input will generate a higher activation for the pattern A-trained cell (and similarly that a pattern B test input will generate a higher activation for the pattern B-trained cell), we make a normal approximation for the distributions of same pattern and different pattern cell activations. Let  $U$  be a random variable with normal distribution ( $\mu_U, \sigma_U^2$ ) where  $\mu_U = E[X]$  and  $\sigma_U^2 = \text{Var}[X]$  for a cell trained and tested on pattern A inputs. Let  $V$  be a random variable with normal distribution ( $\mu_V, \sigma_V^2$ ), where  $\mu_V = E[Y]$  and

$\sigma_V^2 = \text{Var}[Y]$  for a cell trained on pattern B inputs and tested on pattern A inputs. The former mean and variance generalizes to the activation of any cell tested with the same pattern as it was trained on, be that either pattern A or B. Likewise, the latter mean and variance generalize to the activation of any cell tested on a different pattern than it was trained on.

The probability of the cell firing that was trained on the same pattern as the test input is

$$p(V < U) = p(V - U < 0) \quad (39)$$

We can express this in terms of the cumulative normal distribution function:

$$p(V - U < 0) = \Phi \left( \frac{\mu_U - \mu_V}{\sqrt{\sigma_U^2 + \sigma_V^2}} \right) \quad (40)$$

This probability can be visualized as the lesser volume of the bivariate normal distribution of  $U$  and  $V$  after it has been bisected by the  $U = V$  plane. Note that due to the nature of the inputs (where noise displaces activity from one region to another), there is virtually no correlation between  $X$  and  $Y$ . In other input paradigms, such as purely additive noise, there may be a negative correlation ( $\rho$ ) that may be incorporated into the above equation:

$$p(V < U) = \Phi \left( \frac{\mu_U - \mu_V}{\sqrt{\sigma_U^2 + \sigma_V^2 - 2\rho\sigma_U\sigma_V}} \right) \quad (41)$$

## Acknowledgments

This work was supported in part by the Office of Naval Research under Grants N00014-94-C-0103 and N00014-92-1625. We would also like to thank the reviewers of this paper for their detailed and insightful comments.

Reprint requests should be sent to Richard Granger, Center for the Neurobiology of Learning and Memory, University of California, Irvine, Irvine, CA 92717-3800.

## REFERENCES

Arai, A., & Lynch, G. (1992). Factors regulating the magnitude of long-term potentiation induced by theta pattern stimulation. *Brain Research*, 598, 173-184.

Arai, A., Silberg, J., & Lynch, G. (1995). Differences in the refractory properties of two distinct inhibitory circuitries in field CA1 of the hippocampus. *Brain Research*, 704, 298-306.

Bear, M. E., & Malenka, R. C. (1994). Synaptic plasticity: LTP and LTD. *Current Opinion in Neurobiology*, 4(3), 389-399.

Bliss, T. V. P., & Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anesthetized rabbit following stimulation of the perforant path. *Journal of Physiology (London)*, 232, 334-356.

Coultrip, R., Granger, R., & Lynch, G. (1992). A cortical model of winner-take-all competition via lateral inhibition. *Neural Networks*, 5, 47-54.

Ezeh, P., Wellis, D., & Scott, J. (1993). Organization of inhibition

in the rat olfactory bulb external plexiform layer. *Journal of Neurophysiology*, 70, 263-274.

Granger, R., Staubli, U., Davis, M., Perez, Y., Nilsson, L., Rogers, G., & Lynch, G. (1993). A drug that facilitates glutamatergic transmission reduces exploratory activity and improves performance in a learning dependent task. *Synapse*, 15, 326-329.

Granger, R., Whitson, J., Larson, J., & Lynch, G. (1994). Non-Hebbian properties of LTP enable high-capacity encoding of temporal sequences. *Proceedings of the National Academy of Sciences USA*, 91, 10104-10108.

Haberly, L. B. (1985). Neuronal circuitry in the olfactory cortex: Anatomy and functional implications. *Chemical Senses*, 10(2), 219-238.

Harnad, S. (1987). *Categorical perception*. New York: Cambridge University Press.

Hill, A. (1978). First occurrence of hippocampal spatial firing in a new environment. *Experimental Neurology*, 62, 282-297.

Iriki, A., Pavlides, C., Keller, A., & Asanuma, H. (1989). Long-term potentiation in the motor cortex. *Science*, 245, 1385-1387.

Jiang, T., & Holley, A. (1992). Some properties of receptive fields of olfactory mitral/tufted cells in the frog. *Journal of Neurophysiology*, 68, 726-733.

Jung, M., Larson, J., & Lynch, G. (1990). Long-term potentiation of monosynaptic EPSPs in rat piriform cortex *in vitro*. *Synapse*, 6, 279-293.

Kanter, E. D., & Haberly, L. B. (1990). NMDA-dependent induction of long-term potentiation in afferent and association fiber systems of piriform cortex *in vitro*. *Brain Research*, 525, 175-179.

Kirkwood, A., & Bear, M. (1994). Hebbian synapses in visual cortex. *Journal of Neuroscience*, 14, 1634-1645.

Kolta, A., Larson, J., & Lynch, G. (1995). Comparison of long-term potentiation in the proximal versus distal stratum radiatum of hippocampal field CA1. *Neuroscience*, 66, 277-289.

Larson, J., & Lynch, G. (1986). Induction of synaptic potentiation in hippocampus by patterned stimulation involves two events. *Science*, 232, 985-988.

Larson, J., Wong, D., & Lynch, G. (1986). Patterned stimulation at the theta frequency is optimal for induction of long-term potentiation. *Brain Research*, 368, 347-350.

Larson, J., Xiao, P., & Lynch, G. (1993). Reversal of LIP by theta frequency stimulation. *Brain Research*, 600, 97-102.

Lynch, G. (1986). *Synapses, circuits, and the beginnings of memory*. Cambridge, MA: MIT Press.

Macrides, F. (1975). Temporal relationships between hippocampal slow waves and exploratory sniffing in hamsters. *Behavioral Biology*, 14, 295-308.

Miettinen, R., & Freund, T. (1992). Convergence and segregation of septal and median raphe inputs onto different subsets of hippocampal inhibitory interneurons. *Brain Research*, 594, 263-272.

Mott, D., & Lewis, D. (1991). Facilitation of the induction of long-term potentiation by GABA receptors. *Science*, 252, 1718-1720.

Otto, T., Eichenbaum, H., Weiner, S., & Wible, C. (1991). Learning-related patterns of CA1 spike trains parallel stimulation parameters optimal for inducing hippocampal long-term potentiation. *Hippocampus*, 1, 181-192.

Price, J. L. (1973). An autoradiographic study of complementary laminar patterns of termination of afferent fiber to the olfactory cortex. *Journal of Comparative Neurology*, 150, 87-108.

Rail, W., Burke, R., Holmes, W., Jack, J., Redman, S., & Segev, I. (1992). Matching dendritic neuron models to experimen-

- tal data. *Physiological Reviews (Supplement)*, *72*, S159-S186.
- Seress, L., Gulyas, A., Ferrer, I., Tunon, T., Soriano, E., & Freund, T. (1993). Distribution, morphological features, and synaptic connections of parvalbumin- and calbindin d28k-immunoreactive neurons in the human hippocampal formation. *Journal of Comparative Neurology*, *337*, 208-230.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317-1323.
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- Staubli, U., & Lynch, G. (1987). Stable hippocampal long-term potentiation elicited by 'theta' pattern stimulation. *Brain Research*, *435*, 227-234.
- Staubli, U., Rogers, G., & Lynch, G. (1994). Facilitation of glutamate receptors enhances memory. *Proceedings of the National Academy of Sciences USA*, *91*, 777-781.
- Toth, K., Borhegyi, Z., & Freund, T. (1993). Postsynaptic targets of GABAergic hippocampal neurons in the medial septum-diagonal band of Broca complex. *Journal of Neuroscience*, *13*, 3712-3724.
- Woodward, W., Chiaia, N., Teyler, T., Leong, L., & Coull, B. (1990). Organization of cortical afferent and efferent pathways in the white matter of the rat visual system. *Neuroscience*, *36*, 393-401.