if $[\mathrm{I_{Ca}}] \gg k_2/k_1$, transmitter release is independent of small changes in the calcium concentration and depends only on action potential duration.

In the other limiting case, where both reaction rates are slow, we can do a linear approximation to the exponential, and Equation 2.24 simplifies to:

$$\begin{aligned} [\mathrm{R}](t) &\approx \left(\frac{k_1[\mathrm{I_{Ca}}](t)}{k_1[\mathrm{I_{Ca}}](t)+k_2}\right)^4 \int_0^{\mathrm{D}}(t)((k_1[\mathrm{I_{Ca}}](t)+k_2)\tau)^4 d\tau \\ &= \frac{(k_1[\mathrm{I_{Ca}}](t))^4}{5}(\mathrm{D}(t))^5 \end{aligned} \tag{2.26}$$

In this case, we have a fifth order relationship between $[\mathrm{R}](t)$ and the action potential duration, as well as a fourth order relationship with the $\mathrm{Ca^{++}}$ concentration.

The experimental data in *Aplysia*, where a very small percent increase in action potential duration produces a large percent increase in the motor neuron EPSP (which presumably results from a large percent increase in transmitter released) provides evidence supporting the claim that vesicle bonds both form and decay slowly relative to the duration of an action potential. This also means that the proportion of vesicles released in a single action potential will be small relative to the total number of vesicles available.

One consequence of a fifth order dependence of transmitter release on action potential duration, assuming that the relationship between cyclic AMP and action potential duration is linear (which it may well not be), is that the learning rate[21] remains nearly constant even with large changes in synaptic strength.

[21] The learning rate is defined as the percent increase in synaptic strength resulting from a single paired stimulus presentation. In the model, this increase is proportional to the total $\mathrm{Ca^{++}}$ influx during a sensory neuron action potential. This influx determines the number of bonds between $\mathrm{Ca^{++}}$-calmodulin and adenylate cyclase, which determines the strength of the associative response.

Determining the ISI Effect

In studies of associative learning in *Aplysia*, the measured effect of the inter-stimulus interval on learning reaches its maximum at 0.5 seconds, and decays to about 25% of its maximum value after 2 seconds [1].

A possible functional reason for the delayed onset of the associative response could be to prevent the sensory neuron from learning that it is a predictor of itself. This could occur if the sensory neuron stimulates a facilitator interneuron that, in turn, makes contacts with the sensory neuron's presynaptic terminals. Facilitator interneurons exhibit recurrent inhibition [21], firing only at the onset of a sustained input. If interneuron firing ceases within half a second, the delay in the ISI effect will be sufficient to prevent auto-association.

Hawkins's model [20] uses a 10 second time step, which is sufficient to model the learning phenomena at a very rough scale, but prevents accurate modelling of ISI effects. In Byrne's model [8], where the associative response is modelled as directly proportional to the Ca^{++} concentration, the maximum response is reached almost immediately after an action potential and decays exponentially thereafter.

In the model presented here, the strength of the associative learning response is proportional to the increase in cyclic AMP caused by bonds between Ca^{++}-calmodulin and adenylate cyclase. Its time course is determined by the dynamics of the interactions between Ca^{++}, calmodulin, and adenylate cyclase. Bonding between Ca^{++} and calmodulin is assumed to be essentially instantaneous,[22] so that the dynamics of formation and dissolution of the Ca^{++}-calmodulin-adenylate cyclase bonds and the decay of the Ca^{++}-calmodulin concentration are the primary factors involved in determining the ISI.

The bonding reaction is

[22] This assumption is also made by both [8] and [20].

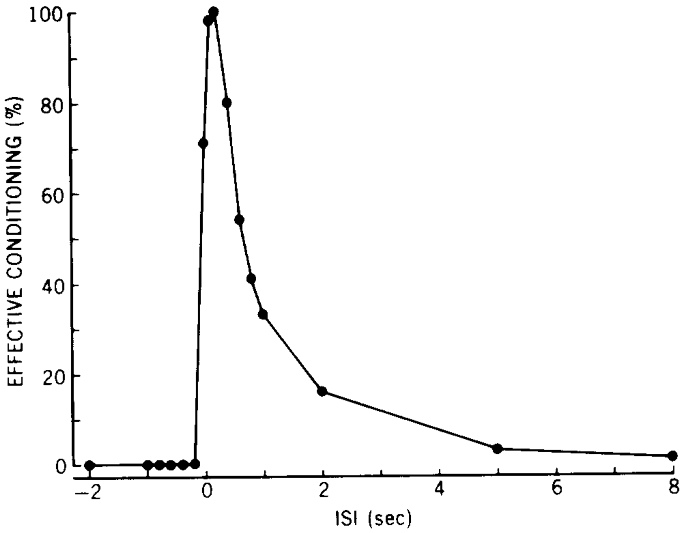

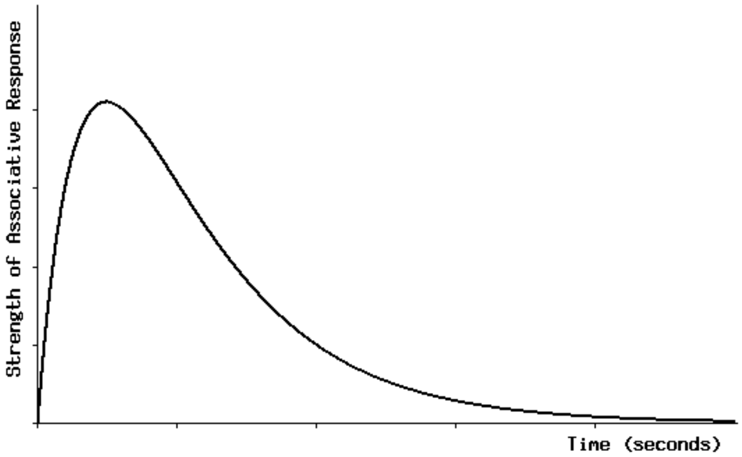

(b)

Figure 2-13: The ISI effect resulting from a simple Ca^{++} concentration dependency decays exponentially as the delay between the CS and the US increases. Two alternative models are shown here. a) ISI effect seen in Byrne's model, which has a simple dependency on the Ca^{++} concentration, but achieves a slight delay (170 msec.) in the peak response by modelling Ca^{++} diffusion within the synapse (From [8, page 69].) (b) ISI effect obtained with the two-stage reaction used in my model.

$$\mathrm{C} + \mathrm{A} \underset{k_2}{\overset{k_1}{\rightleftharpoons}} \mathrm{B} \tag{2.27}$$

$$\mathrm{C} \xrightarrow{k_3} \textit{inaccessible forms} \tag{2.28}$$

where A is adenylate cyclase, C is Ca^{++}-calmodulin, B is the Ca^{++}-calmodulin-adenylate cyclase compound, and k_1, k_2 and k_3 are the rate constants for the various reactions. The second equation represents the removal of Ca^{++}-calmodulin from the synapse, through a combination of diffusion, uptake into intracellular structures, and expulsion of Ca^{++} from the cell through an ion pump.[23]

The resulting equations for the dynamics of the reactions are

$$\frac{d[\mathrm{B}](t)}{dt} = k_1[\mathrm{C}](t) \cdot [\mathrm{A}](t) - k_2[\mathrm{B}](t) \tag{2.29}$$

$$\frac{d[\mathrm{A}](t)}{dt} = -k_1[\mathrm{C}](t) \cdot [\mathrm{A}](t) + k_2[\mathrm{B}](t) \tag{2.30}$$

$$\frac{d[\mathrm{C}](t)}{dt} = -k_1[\mathrm{C}](t) \cdot [\mathrm{A}](t) + k_2[\mathrm{B}](t) - k_3[\mathrm{C}](t) \tag{2.31}$$

If we assume that there is an excess of adenylate cyclase in the system, so that only a small amount is bound by Ca^{++}-calmodulin at any given time, we can approximate this nonlinear system with the following linear one:

$$\frac{d[\mathrm{B}](t)}{dt} = k_1[\mathrm{A}] \cdot [\mathrm{C}](t) - k_2[\mathrm{B}](t) \tag{2.32}$$

$$\frac{d[\mathrm{C}](t)}{dt} = -(k_1[\mathrm{A}] + k_3)[\mathrm{C}](t) + k_2[\mathrm{B}](t) \tag{2.33}$$

[23] These mechanisms for Ca^{++} uptake may have different time constants, but it is assumed here that one rate dominates.

The general solution to these equations is given in Equations 2.3 and 2.4, with parameter values of:

$$c_1 = \frac{[\mathrm{B}](t_0) - a_2[\mathrm{C}](t_0)}{(a_1 - a_2)} \tag{2.34}$$

$$c_2 = \frac{[\mathrm{B}](t_0) - a_1[\mathrm{C}](t_0)}{(a_2 - a_1)} \tag{2.35}$$

and, defining $a = (k_1[\mathrm{A}] + k_2)/2$,

$$m_1 = -a + \sqrt{a^2 - k_2\, k_3} \tag{2.36}$$

$$m_2 = -a - \sqrt{a^2 - k_2\, k_3} \tag{2.37}$$

$$a_1 = \frac{m_1 + k_1[\mathrm{A}] + k_3}{k_2} \tag{2.38}$$

$$a_2 = \frac{m_2 + k_1[\mathrm{A}] + k_3}{k_2} \tag{2.39}$$

In the case encountered in the learning trials, $[\mathrm{B}](t_0) = 0$, so

$$[\mathrm{B}](t) = K\,(e^{m_1(t-t_0)} - e^{m_2(t-t_0)}) \tag{2.40}$$

$$K = [\mathrm{C}](t_0)\frac{a_1 a_2}{a_2 - a_1} \tag{2.41}$$

This function reaches its maximum at

$$\begin{aligned} \frac{d[\mathrm{B}](t_{max})}{dt} &= 0 \\ 0 &= K\,(m_1\, e^{m_1(t_{max}-t_0)} - m_2\, e^{m_2(t_{max}-t_0)}) \end{aligned}$$

$$t_{max} - t_0 = \frac{\ln|m_1| - \ln|m_2|}{m_2 - m_1} \tag{2.42}$$

If we want a maximum at 0.5 seconds, and a decay to 25% of its maximum at 2 seconds, we need to set m_1 and m_2 according to the following equations:

$$\begin{aligned} [\mathrm{B}](t_0 + 2) &= 0.25\,[\mathrm{B}](t_0 + 0.5) \\ e^{2\,m_1} - e^{2\,m_2} &= 0.25\,(e^{0.5\,m_1} - e^{0.5\,m_2}) \end{aligned} \tag{2.43}$$

and

$$\frac{\ln|m_1| - \ln|m_2|}{m_2 - m_1} = 0.5 \tag{2.44}$$

Solving these equations, we get the following parameter values, which are used in the simulation:

$$\begin{aligned} m_1 &= -2.93 \\ m_2 &= -1.28 \\ a_1 &= -0.34 \\ a_2 &= 0.52 \\ k_2\,k_3 &= 3.76 \\ k_1[\mathrm{A}] &= 4.22 - \frac{3.76 + k_2^2}{k_2} \end{aligned}$$

In order for the value of $k_1[\mathrm{A}]$ to be positive, $k_2\,k_3$ must fall within a fairly narrow range.

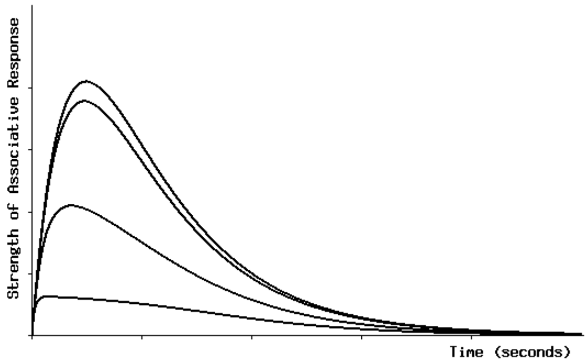

Figure 2-14: Nonlinear distortion of ISI when adenylate cyclase concentration is low. The value of $k_1[\mathrm{A}](t_0)$ is kept constant. The uppermost curve is without saturation; the saturated curves are (in descending order) for $[\mathrm{A}](t_0) = .5[\mathrm{C}](t_0)$, $[\mathrm{A}](t_0) = .05[\mathrm{C}](t_0)$, and $[\mathrm{A}](t_0) = .01[\mathrm{C}](t_0)$.

$$1.28 < k_2 < 2.93$$
$$1.28 < k_3 < 2.93$$

The maximum value for $k_1[\mathrm{A}]$ that matches the experimental results is 0.34, at $k_2 = k_3 = 1.94$. Using this value for $k_1[\mathrm{A}]$, and starting with $[\mathrm{B}](t_0) = 0$, the maximum concentration of bonds (at an inter-stimulus interval of 0.5) is $[\mathrm{B}]_{max} = 0.06[\mathrm{C}](t_0)$, which means we must have $[\mathrm{A}] \gg 0.06[\mathrm{C}](t_0)$, and $k_1 \ll 5.7[\mathrm{C}](t_0)^{-1}$, for our linear approximation to be valid.

If [A] is not sufficiently high, the system will saturate, and much of the available adenylate cyclase will be used up. The effect on the strength of the associative response is shown in Figure 2-14.

Controlling the ISI Effect

Although the numeric parameters used here are specific to the mechanism which determines the strength of the associative learning response, the general form of the equations may be shared with other processes involving chemical second messengers. The factors these reactions must have in common are that the second messenger only activates the protein complex for as long as it remains bound to it, and that its concentration decays over time at a regular rate.

The magnitude of the reaction's maximum effect is proportional to the initial concentration of the second messenger.[24] Its time course can be modulated by the cell either by modulating the temporal pattern of injection of the second messenger. It may also be able to modulate the time course if it has control over the reaction rates. For example, in this case the rate of decay of the Ca^{++}-calmodulin concentration is dependent on active Ca^{++} uptake. Increasing the concentration in the synapse of proteins responsible for the removal of Ca^{++} from the synapse will increase k_3, which changes the duration of the ISI effect.

2.3 Summary

This system demonstrates several specific examples of chemical processes which modulate synaptic responses within a cell. These processes implement short- and intermediate-term internal state changes in the synapse.

Traditional computational models of neurons consider only two time scales: the very fast time scale of the cell's electrical responses, and the "permanent" changes which result from long-term learning. A few exceptions can be found in real-time models of reinforcement learning (e.g. [2, 62, 63]), where an exponentially decaying internal memory trace is used to model the effect of the ISI on the learning response. The mechanisms found in *Aplysia*'s synapses, however, demonstrate the presence of an wide range of processes acting at different time scales which have a significant impact on the synaptic strength, and therefore the computational properties, of the reflex circuit.

[24] It also depends on the reaction rates, which are constant in the absence of other modulating substances.

These mechanisms can be broken down into distinct categories based on their time course. Electrical responses and the dynamics of transmitter release are very fast, taking place on a millisecond time scale. The processes which determine the ISI effect take place on a time scale of a few seconds, while short-term learning operates on a time scale of minutes to hours. Long-term learning in *Aplysia*, which is not modelled here, occurs over the course of several hours and lasts for at least several weeks. Processes which operate at similar time scales can have complex interactions, whereas a large difference in time scales leads to a decoupling of effects.

The model presented in this chapter reveals the importance of understanding the underlying chemical dynamics in *Aplysia*'s synapses in order to model the macroscopic behavior of the circuit for the gill withdrawal reflex. Two specific processes, which control the dynamics of transmitter release and of the ISI effect, are examined in detail. The results of testing the model against data from biological experiments reveal enough similarity to support the claim that these mechanisms play an important role in conditioning of the reflex, but they also differ in many details. The differences may be due to as yet undetected complexity in the underlying chemistry or to the contribution of parts of the reflex circuit that are not modelled.

Part II

Topographic Map Formation in the Visual System

Chapter 3

A Model of Topographic Map Formation

Topographically ordered projections from one region to another are a prominent feature in the brain. The presence of such maps has been known since the 1920's [31], when the spatial locations of electrical responses to sensory stimulation in the somatosensory cortex were discovered to lie in an orderly arrangement relative to the location of the stimulated body part. These maps have since been found throughout the brain.

The process of map formation in embryonic and early childhood development has been most extensively studied in the visual system, in projections from the retina to the thalamus and superior colliculus, and in projections from the thalamus to the primary visual cortex [66]. These connections form over long distances, and initially they are highly disordered. After the initial axonal projections form, they sort themselves over time to produce a detailed topographic correspondence between the retinal image and the image in the target structure. This map formation process makes use of a variety of chemical and electrical mechanisms.

A simulated model of the effects of chemical and activity-dependent mechanisms in topographic map formation is presented in this chapter. These mechanisms work

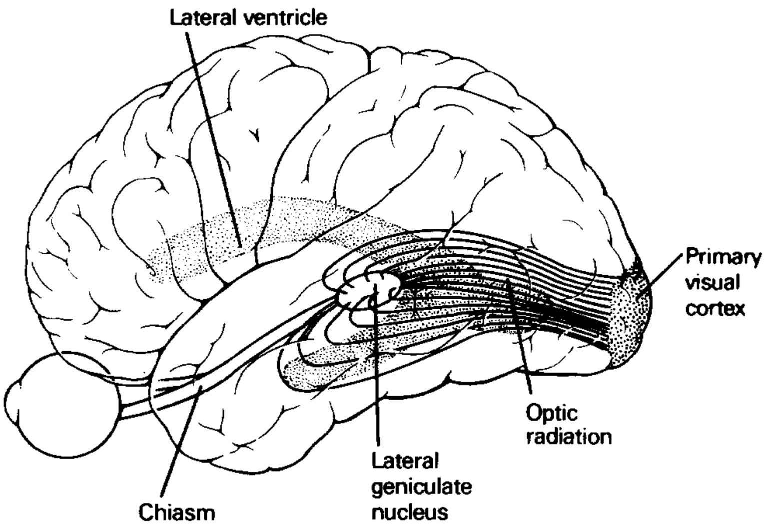

Figure 3-1: Axonal projections between brain structures in the early visual system. During development, these projections sort themselves to produce topographic mappings between the retina, the lateral geniculate nucleus of the thalamus, the superior colliculus (not shown), and the primary visual cortex (From [31, page 363].)

together to allow rapid and robust development of detailed maps. This model resolves a number of difficulties found in previous modelling attempts. These difficulties include controlling the final map orientation, attaining rapid convergence without loss of map continuity, and demonstrating a biologically plausible mechanism that produces the necessary coactivation pattern between source and target structures during activity-dependent map formation.

Experimental results on the ability to form rough, but not refined, maps in the absence of electrical activity, expansion of heavily stimulated areas, and map expansion and contraction following lesions are matched by this model. Although a highly simplified neural model is used, the goal is to explore organizational principles of biological development, so some care is taken to remain compatible with experimental results and with what is known about the underlying biological mechanisms.

3.1 Introduction

The motivation for this work comes from recent experimental findings about biological mechanisms that give rise to topographic map formation. A variety of such mechanisms have been studied in the early visual system, including the maintenance of axonal order within the optic tract, positional labelling via chemical markers, and activity-dependent synaptic plasticity [66].

The presence of chemical markers guiding map formation has been demonstrated in a variety of experimental preparations [66, 60]. Although some early researchers believed these markers might be sufficient to fully specify the final mapping, in practice this does not seem to be true. Evidence for the presence of other mechanisms comes partly from studies of expansion and compression of maps following changes in input activity or after lesions. Experiments described in [61] point to electrical activity in particular as an important contributor to map formation. In this study, blocking activity during development prevents the formation of a refined topographic map, although rough maps still appear.

The bulk of experimental evidence, therefore, points to cooperation among a variety of mechanisms, physical, chemical, and activity-dependent, for the formation and maintenance of topographic maps. Previous computational models, however, have attempted to find a single unified map formation process, and fail to exploit the advantages of combining multiple mechanisms.

3.2 Previous Models of Map Formation

There are four major problems to be overcome by any functional model of biological map formation. These are specifying the orientation of the map, maintaining global as well as local continuity, attaining a rapid convergence rate, and specifying a biologically plausible implementation.

Willshaw and von der Malsburg [68] simulated an early activity-based model. It solves the problems of map orientation and continuity, but requires the presence of a nucleation region in which the initial axonal connections are fully specified. Because of the need for a nucleation region, the process converges very slowly. There is no biological evidence for such a region. Their marker-based model [67] solves most of these problems, but does not match recent experiments establishing the dependence of the process on electrical activity.

Other activity-based models have been based on Kohonen's mapping algorithm [36, 37]. This model is mathematically successful, but is not well grounded in biology. The input to his algorithm is a vector specification of a single point in the input space, rather than a pattern of activity[1] over that space (which corresponds to the input cell layer), as is found in biology. The final map orientation is random, and the algorithm assumes a specific activity pattern in the target map without considering how that pattern could result from a biological process.

More recent modelling efforts have concentrated on the formation of ocular dominance and orientation columns, but some have included map formation as well. Obermayer [52] uses a variant of the Kohonen algorithm which demonstrates map formation using a biological input representation. Miller's model [48] displays activity-dependent map refinement, but the initial topography is fully specified, and axonal arbors narrow over time around a fixed location.

3.3 The Model

The model described in this paper is based on the following (biologically motivated) description of the developmental process: initial synaptic contacts from the source (presynaptic) cell layer are formed over a broad region in the target (postsynaptic) layer. Varying concentrations of chemical markers across the surface of the maps determine the affinity of synapses between particular source and target regions. This

[1] An activity pattern is the strength of activation of the neurons, as a function of their location in the map.

chemoaffinity provides a rough initial mapping.

Although some form of chemoaffinity has been found in all biological systems studied, the details vary considerably from one system to another. Instead of modelling the details of a particular chemical mechanism, therefore, an attempt is made to determine what constraints all chemical mechanisms must satisfy to ensure a correct and detailed final map when combined with an activity-based refinement process. The chemical mechanisms are assumed to produce a rough map topography, (possibly by the mechanism described in [67]), which is used to determine the initial connections between source and target layers. Chemical markers do not, in this model, play a role in the dynamics of the refinement process.

Activity-dependent mechanisms then act to refine the initial connections that result from chemoaffinity. Synapses are strengthened via a Hebbian mechanism that is patterned after long-term potentiation [34].[2] The coactivation pattern driving these changes results from correlations in the activity of near neighbors in the source map, combined with short-range lateral activation in the target to increase the activity of near neighbors that receive strong source inputs.

Correlated activity of neighboring neurons in the developing retina, prior to the maturation of photoreceptors, has been observed in rat [15, 41], and in cat and ferret [42]. The observed activity patterns consist of short correlated bursts of activity (lasting a few seconds) separated by relatively long periods of silence.

Synaptic growth is limited by competition between synapses projecting to the same target neuron. Competitive effects cause map expansion under heavy stimulation, and also cause the breaks in multiple uncorrelated source maps that result in ocular dominance column formation.

[2]See Appendix D for background on long-term potentiation.

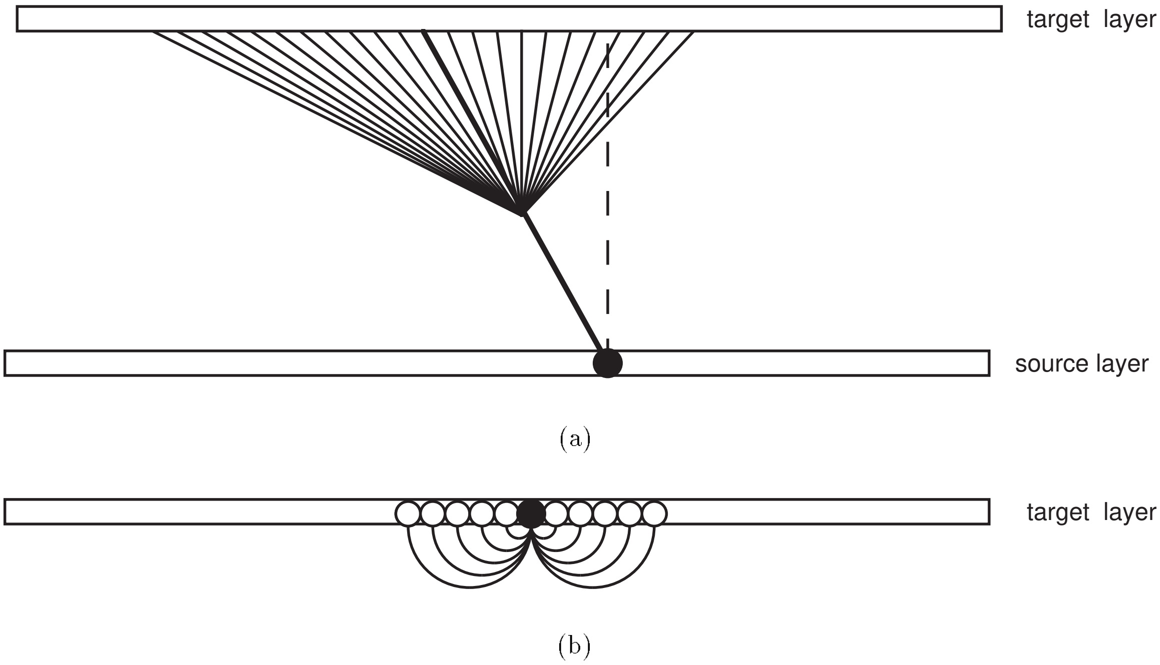

Figure 3-2: Connections used in the simulation. a) Modifiable projections from source (filled circles) to target (open circles) layers. The axonal arbor for a single source neuron is shown. b) Fixed-strength lateral connections within the target layer.

3.3.1 The Simulation

The simulation uses one-dimensional source and target layers to reduce computational complexity so that longer-range interactions can be studied. The algorithm does, however, generalize to two or more dimensions. Two types of synaptic connections are included, as shown in Figure 3-2: modifiable excitatory projections from the source layer to the target layer, and unmodifiable excitatory lateral connections within the target layer.

Chemical markers are assumed to be present in source and target maps in a linear concentration gradient over the surface of the map. The chemical concentration is assumed to be detectable by the growing axons with some finite error. These assumptions are used to determine the initial connectivity of the map, before the activity-dependent process begins. However, the simulation results are valid not only for this particular hypothesized mechanism but also for any other mechanism which produces a similar pattern of connections. Initial connections from each source neuron cover a broad but finite region in the target, close to (and including) the desired final location.

Lateral connections are localized to a neighborhood around each cell, and their strength decreases as the distance between neurons increases.

Input activity is assumed to occur in spatially and temporally localized clusters of neurons. Each input activity cluster occurs at a random location within the source layer. For simplicity, all neurons within a cluster are given the same level of activity.

Each input activity cluster is simulated as a single "trial" of the learning algorithm. Within a trial, the source is activated, activation is spread through the target layer, and connection strengths are updated according to a learning rule which combines Hebbian learning (in the numerator) with input normalization as a competitive mechanism (in the denominator).

$$w_{ij}(T+1) = \frac{w_{ij}(T) + \alpha\, A_{ij}(T)\, o_i(T)\, a_j(T)}{\frac{1}{K}\sum_{j=0}^{m-1}(w_{ij}(T) + \alpha\, A_{ij}(T)\, o_i(T)\, a_j(T))} \tag{3.1}$$

$w_{ij}(T)$ is the connection strength between source neuron j and target neuron i, α is the learning rate, a_j and o_i are the source and target activity, $A_{ij} = 1$ if there is a synapse between source j and target i, and 0 otherwise, and K is the "competition limit", that is, the total synaptic strength over all inputs to a given target neuron.

The Hebbian component is assumed to arise from long term potentiation (LTP). The biological underpinnings of competition are not well understood, though some kind of competitive mechanism is necessary to explain experimental results such as expansion of a heavily stimulated map region [66].

High-level Modelling of LTP

When using such a highly simplified neural model, some care is necessary to ensure that the model maintains some correspondence with its known biological underpinnings.

Other computational models ignore the transient response to an input stimulus, and instead use a hypothetical steady-state frequency response as the target activity. However, LTP, the best candidate for a biological form of Hebbian learning, operates at a very fast time scale, on the order of 100 msec.[3] The dynamics of the transient response is therefore critical, as the system does not have time to settle into steady-state within that time window.[4]

In this model, the dynamics of the target response are simulated as follows:

1. Individual target neurons are modelled as simple thresholded elements. Their activation is a weighted sum of their inputs, and if it is above threshold they fire.

2. Each "trial" is divided into time steps. In each time step neurons whose input is above threshold fire a single action potential.

3. Inputs from the source layer arrive at the target neurons at the first time step of each trial, and initiate firing in neurons over some region.

4. Firing in the source layer ceases after the first time step, but activation continues to spread laterally in the target map. At each time step, inputs arrive through lateral connections from target neurons that fired in the previous time step.

[3]For example, in the model of LTP described in [7], the strength of the LTP response is assumed proportional to the peak Ca^{++} concentration; their simulation shows this concentration peaking and subsiding within 100-200 msec. after presentation of the input stimulus.

[4]The LTP response assumed by this model is proportional to $\int_0^\tau o_i(t)dt$, for a time window of length τ. Even if the response does reach steady state near the end of the time window, therefore, the transient response will still dominate.

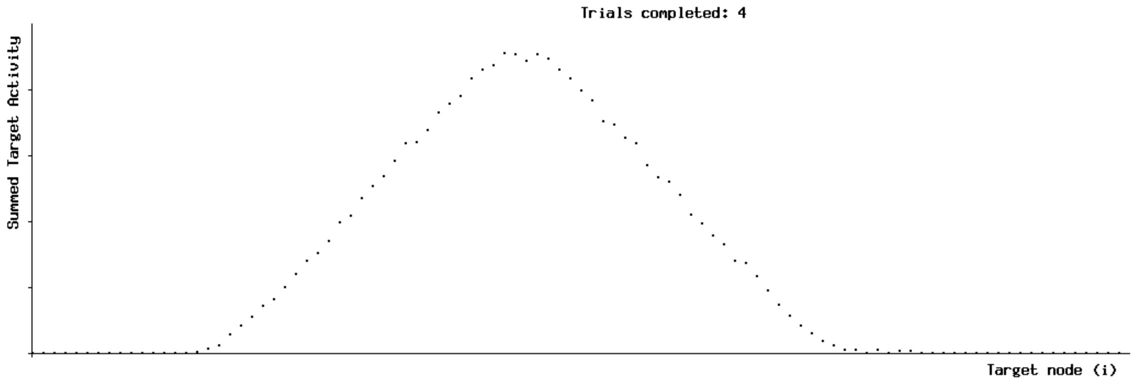

Figure 3-3: Typical target activity pattern in response to an input activity cluster.

Lateral synaptic strengths are set in such a way that the active region in the target shrinks over time.

5. Activation is summed over the duration of the target response to get the activity o_i for learning.

If most of the cells within some region in the target layer are activated, the activity will die out in a wave moving inward from the edges of that region. Neurons in its center therefore remain active longest. The reasons for assuming this pattern of target activity are discussed in Section 3.5.

Parameters in this simulation were set to give maximum burst lengths of about 10-20 spikes. At 100 Hz., this response would last no more than 200 msec., which is of the same order of magnitude as the LTP time window. This gives us a rough approximation of the actual target activity that would result from this pattern of lateral connections. Mathematical details and parameter constraints are given in Section 4.1.

3.4 Results

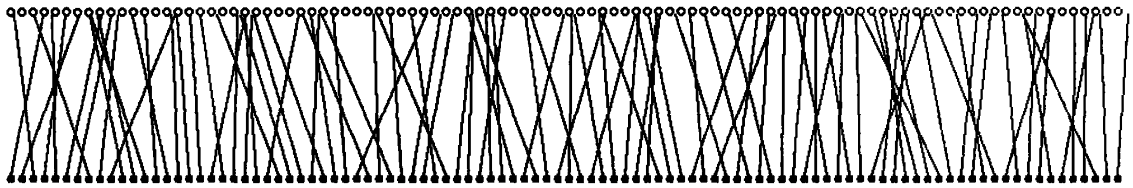

(a)

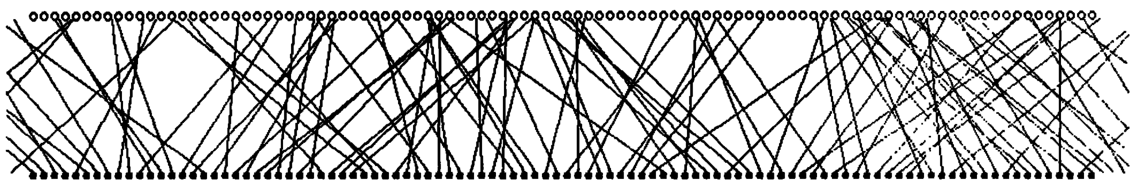

(b)

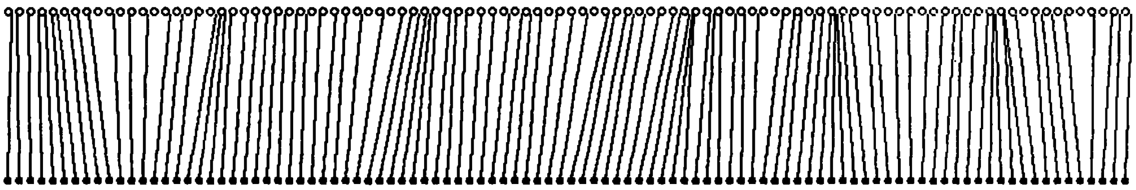

(c)

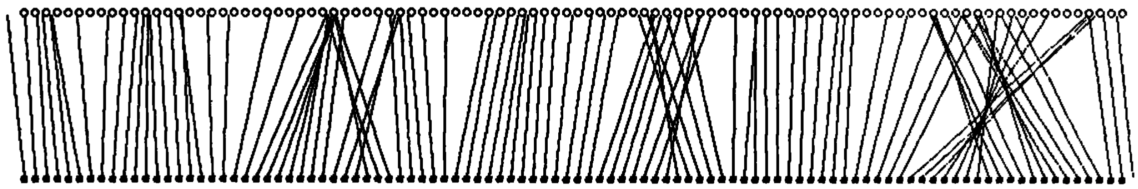

(d)

Figure 3-4: Examples of initial and final projections. Filled circles are source neurons, open circles are target neurons. Lines are drawn from each source neuron to the centroid of its axonal projection. (a–b) initial state, with position errors of ± 5 and ± 20. (c) typical result after 10,000 input stimulus presentations. Results from 15 of 16 tests were roughly identical. (d) worst case result (1 test out of 16), with an initial position error of ± 20, cluster size = 2 (discussed in text).

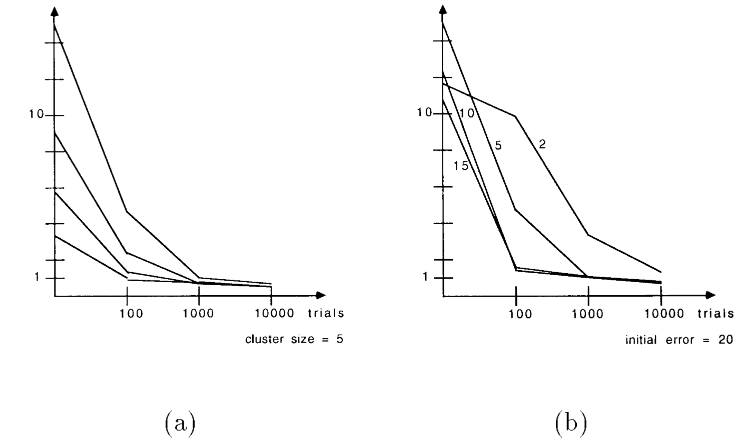

Figure 3-5: Spacing error. (a) fixed cluster size of 5, with several position errors. (b) fixed position error of ±20, with several cluster sizes.

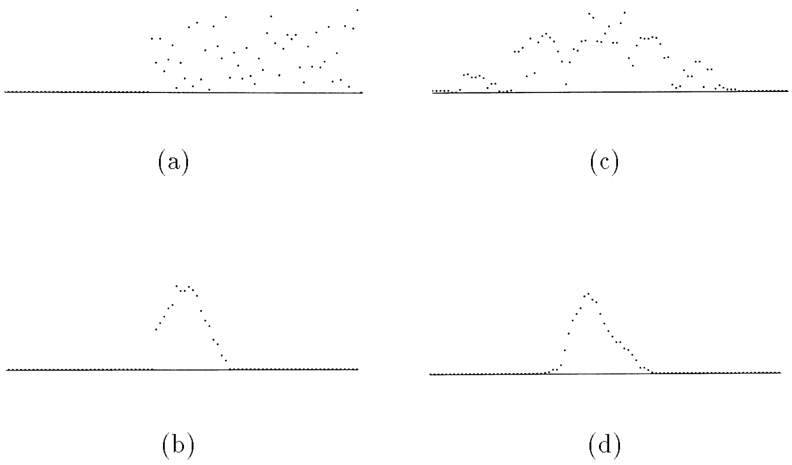

Figure 3-6: Typical arbor and receptive field strengths. (a–b) initial and final connection strengths across the target for a given source neuron. The "goal" location is indicated by an arrow. (c–d) initial and final receptive field strengths of a given target neuron, in response to activity in different locations across the source. A source activity cluster size of 5 was used, and activation was allowed to spread in the target as it does during the learning trials.

Tests were run on one-dimensional source and target layers, with 100 neurons in each layer. The ends were wrapped around to reduce edge effects. The location of the input activity cluster for each trial was chosen at random in a uniform distribution over the input map. Each test ran for 10,000 trials.

In all cases, the final map orientation is correct, and the map is continuous at a global level. Convergence to the correct location is extremely fast, being essentially complete within 1000 trials (an average of ten presentations of each input cluster).

The centroid of the axonal projection from each source neuron, weighted by connection strength, is used as a measure of its "location" in the target layer; the maximum strength connection is generally very close to the centroid. Examples of initial and final maps are shown in figure 3-4.

The average "spacing error", the variance in distance between centroid locations of neighboring source cells, is used to evaluate the accuracy of the local ordering in the map. If these locations are equally spaced in the target, the spacing error will be zero. At large variances, projections may be "crossed"; that is, their centroid locations may be ordered differently in the target than the neurons are in the source.

Figure 3-5 shows the results of two sets of tests. In one, the source activity cluster size is set at a constant, while the initial maximum position error of the projections is varied. The accuracy of the resulting map was found to be nearly independent of the position error, for values up to ± 20 neurons. Follow-on tests (results not shown) were done by doubling the size of the target layer (to 200 neurons) and using position errors of ± 30 and ± 40. Again, the final result is near perfect, though it takes slightly longer to converge.

The second set of tests uses a fixed position error, and looks at the effect of the source activity cluster size. Larger cluster sizes converged more rapidly, and with the smallest cluster size (number of simultaneously active neurons = 2), the final map was somewhat disordered. Reasons for this are discussed in Section 3.5.1.

Another important question in map formation is the shape of the axonal arbors (in terms of their connection strengths), and of the receptive fields in the target layer. These features are functionally important, and they are easily measured experimentally. Figure 3-6 shows the initial and final shape of a typical axonal arbor, and the strength of a typical receptive field of a target neuron to activity at different locations in the source. The receptive field is the combined response to the source and to secondary lateral stimulation, so it corresponds to what would be seen in a biological experiment.

As is seen in biology, the axonal arbors become more narrow over time, and the final receptive fields are strongest when input stimulation is at the center of the receptive field, and taper off in strength as input stimulation moves away from their center [31].

3.5 The Role of Target Activity

In constructing this simulation, we have considerable latitude in modelling the target activity pattern, $o_i(t)$. In the adult animal, patterns of activity tend to follow a characteristic "mexican hat" pattern, with short-range excitatory and mid-range inhibitory interactions between inputs. This has led many researchers to hypothesize a pattern of direct lateral connections which are excitatory at close range and inhibitory at slightly longer ranges. However, this pattern of activity may be due to factors other than lateral synaptic connections, and interactions between inputs during development may be quite different than in adults.

However, the nature of lateral interactions in the target is very important in controlling the outcome of map formation. Lateral connections determine the target response to a given input stimulus, and the relationship between patterns of activity in the source and target maps determines the strength of the synaptic connections when map formation reaches an equilibrium state. If input activity is uniformly distributed, the dynamics of the system have a fixed point at

$$w_{ij,eq} \propto \sum_{p \in P_j} o_{i,eq}(p) \tag{3.2}$$

where P_j is the set of all input activity patterns in which source node j is active, and $o_i(p)$ is the target response to input pattern p.[5] That is, the shape of the arbor for a given source node at equilibrium is proportional to the sum of the target responses to all input activity patterns in which that node is active. The form of the target activity patterns determines the final outcome.

This raises the question of what kind of activity is most useful for map formation. The next two sections look at this issue in more detail.

3.5.1 Noise Reduction

Map refinement can be thought of as a process of noise reduction. The location of the initial arbors provide a rough estimate of the desired target location (i.e. they are localized to a neighborhood around it, with some position error); near neighbors to a source neuron provide similarly rough estimates of nearby target locations.

When several source neurons fire in a cluster, their projections can be averaged to obtain a more accurate estimate of the correct target location than we could obtain from any single projection. When this process is repeated with different overlapping clusters, the estimate continues to improve. If we simply summed the target activity from the source layer, with no lateral connections, this is what we would get. However, although the centroid of the arbor converges to the correct location (in a map with full connectivity), the distribution of synaptic strength will spread out over time, until in the limit the strength is constant over the entire arbor.[6] To compensate for this effect, some kind of mechanism is needed to selectively enhance the target response in the center of the region activated by an input activity cluster.

[5] This result is derived in Section 4.3.3.
[6] See Section 4.3.4 for a derivation of this result.

A purely excitatory pattern of lateral connections was chosen because it is ideal for producing this type of activity. The only constraints on the source activity are that it be sufficient to initiate firing in the target over much of the extent of the axonal arbors, and that there be no large gaps in connectivity within an arbor. If these constraints are met, and the strength of lateral connections is such that neurons in the center of an active region continue to fire, while neurons bordering an active region cease firing, this will produce an activity pattern that is strongly enhanced in the center of the active region, with decreasing activity enhancement at greater distances from its center.

This kind of activity pattern is very difficult to obtain if a "mexican hat" pattern of lateral interactions is used. This pattern of connections tends to produce stable activity within a region whose breadth is determined by the extent of lateral excitation, and inhibit activity outside that region. Periodic activity can also be obtained with a mexican hat pattern. This kind of target activity pattern most naturally produces a discontinuous map; for example, the patterns of connections found in ocular dominance columns could be produced in this way.

Researchers using this pattern of lateral connections have utilized a variety of tricks to ensure a continuous final map. The neighborhood functions used with Kohonen maps [36] are initially very broad, and are manually reduced in size over the course of map formation. Synaptic arbors are in this way focussed very slowly into the correct region as map formation proceeds, which reduces the difficulty of maintaining map continuity.

Willshaw and von der Malsburg [68] also use a mexican hat pattern of lateral connections. In their model, continuity is ensured by the use of a nucleation region. A small part of the map is topographically organized as an initial condition, and synaptic strengths within this region are much stronger than in the surrounding tissue. This ensures that target activity is sufficient to induce learning only at the borders of the organized region, which prevents the formation of partial maps that might otherwise stabilize and form discontinuous sub-maps.

Because of the difficulty of maintaining map continuity in the presence of lateral inhibition, these results predict that map topography is largely determined before the development of an inhibitory surround. It is fairly well established that excitation predates inhibition in the developing nervous system [18, 40], so this seems a reasonable assumption.

In particular, experimental results on neurotransmitter metabolism in cat visual cortex during development [59, 14] were checked for compatibility with this hypothesis. The observed levels of glutamate metabolism during the stage of development in which map formation occurs are much higher than in adult animals, and GABA metabolism is still almost nonexistent. As glutamate is the most common excitatory transmitter in the cortex, and GABA is a well-known inhibitory transmitter, at least in the adult, this result supports the hypothesis that large amounts of excitatory activity are involved in map formation, and rules out one possible mechanism for lateral inhibition. However, as it is not actually known whether lateral inhibition in the adult is mediated by GABA, nor even whether GABA is inhibitory during development, this result is inconclusive.

Looking at the process as one of noise reduction also explains why the algorithm runs into problems with small cluster sizes and large position errors. The smaller the cluster size, the less information is available to each step of the noise reduction process, because there are fewer "samples" to average over. This means it will converge more slowly. However, as the arbors get narrower, due to the enhancement of activity in the center of the active region, their position becomes fixed. If this occurs before the noise reduction process produces an accurate estimate of the desired location, the map will remain somewhat disordered. This problem can be solved by increasing the size of the activity clusters. Another prediction of this model, then, is that activity clusters must be larger in early stages of map formation than they are later in development.

3.5.2 Convergence Properties

Another important effect of the target activity pattern is on the convergence properties of the learning algorithm. Erwin, Obermayer, and Schulten [13] have, in the case of Kohonen nets on a one-dimensional input space, determined the optimal target activity pattern[7] for ensuring a rapid convergence rate and for preventing the formation of map discontinuities. Although their results have not been demonstrated to generalize to this model, they nonetheless provide criteria that are useful for its analysis.

Two requirements are defined that are sufficient to ensure that the final map is continuous, and two more to ensure a rapid convergence rate. For a continuous final map, the first requirement is that the breadth of the activity pattern be greater than the uncertainty in position (in their model, this initially corresponds to the entire breadth of the map). The second is that it be "convex", a term they define to include a slightly larger class of functions than the usual meaning of the term. For rapid convergence, the requirements are that the breadth of the activity pattern be as narrow as possible without violating the continuity constraint, and that it be sloped toward the center, so that neighboring neurons have maximally differing activity levels.

Since the breadth of the target response in this model is determined directly by the breadth of the arbors, which is greater than their positional uncertainty, the first requirement is easily met. Also, as the arbors narrow down around their correct location, the breadth of the response also narrows, which not only meets the breadth constraint for rapid convergence, but also tailors itself over the course of learning to the current optimum, by narrowing as the position error narrows.

The shape of the response is a sum of direct activation from the source layer, which is (probabilistically) gaussian over the surface of the map, and the response

[7]The target activity pattern in this model corresponds to the "neighborhood function" used with Kohonen maps.

due to spreading activation through the lateral connections, which is triangular except for some rounding at its peak. The gaussian function is convex within a finite region about its center, and the rounded triangular function is convex over its entire nonzero region; the target activity is also (probabilistically) convex. The triangular component of this activity pattern also gives it a slope that meets the second criterion for rapid convergence.

The activity pattern which arises naturally from intercellular interactions in this model thus meets all the requirements given in [13] for obtaining a continuous final map with an optimal convergence rate.

3.6 Extensions to the Model

3.6.1 Biological Mechanisms for Competition

The competitive mechanism used in this simulation is a form of input normalization, where the total input to a given target neuron is held at a fixed value. Increases in synaptic strength due to the Hebbian learning mechanism are offset by decreases in the strength of other synapses. Some form of competition is necessary to limit the total arbor size, and to explain map expansion and compression in response to lesions or changes in the level of input stimulation.

However, although input normalization is commonly used as a competitive mechanism in topographic map models, it requires information which is not locally available at the synapse, and for this reason it would be very difficult to implement in biology. It is therefore unsatisfactory for a model in which biological plausibility is an important consideration.

Oja [53] showed that vector input normalization can be approximated (in the limit where α is small) by an activity-dependent decay, given the assumption that there are no lateral connections in the target map. This means that vector input normalization can be achieved, if there are no lateral connections, by decreasing the synaptic strengths at a rate which depends only on the target activity.

Activity-dependent decay suggests a biological mechanism whereby all synapses decrease in strength at a rate which is controlled by target depolarization. For example, a voltage-dependent process that affects the turnover rate of postsynaptic receptors might have this effect.

If we follow Oja's derivation, using the form of competition used in this model (which normalizes over the summed input strength instead of the vector magnitude), we obtain the following rule:

$$w_{ij}(T+1) \approx w_{ij}(T) + \alpha \, A_{ij}(T) \, o_i(T) \, a_j(T) - w_{ij}(T) \, o_i(T) \, \frac{\alpha N_i(T)}{K} \tag{3.3}$$

where

$$N_i(T) = \sum_{j=0}^{m-1} A_{ij}(T) \, a_j(T) \tag{3.4}$$

$N_i(T)$ is constant if we assume a constant input cluster size and full connectivity between source and target maps, and in this case the decay term becomes a pure activity-dependent decay. Its behavior in situations which violate this assumption is discussed in more detail in Section 4.2.

3.6.2 Map Expansion and Compression

A final set of tests was run to replicate biological experiments showing map expansion and compression. There are two ways to get these effects: one is to stimulate a localized input area very heavily (or, conversely, deprive an area of input), and the other is to introduce a lesion into the source (or target) map, and later observe the pattern of connections formed by the remaining tissue [26].

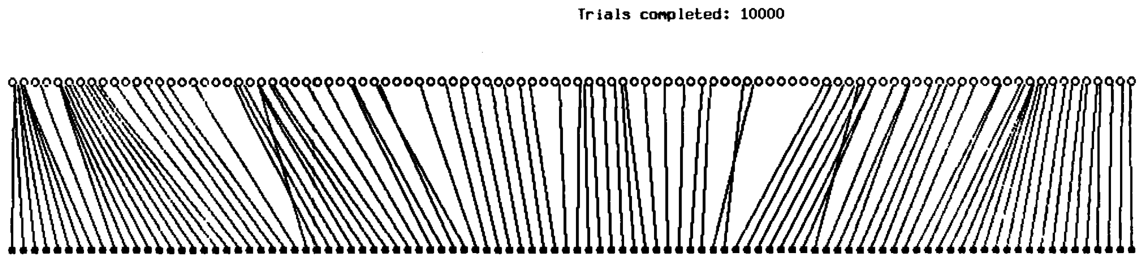

(a)

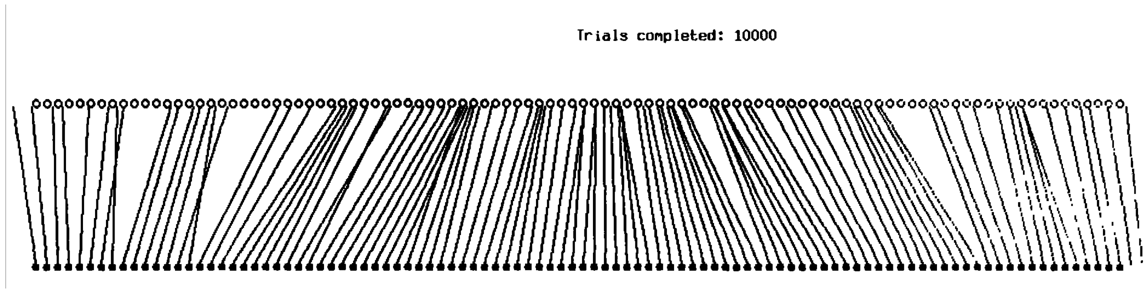

(b)

Figure 3-7: Final results for map expansion and compression after 10,000 stimulus presentations. (a) The central half of the input map was stimulated twice as often as the surround. (b) The center was stimulated half as often as the surround.

Usage-Dependent Changes

The stimulation-dependent experiments were run simply by changing the probability of occurrence of a source stimulus in different regions of the map. Results are shown in figure 3-7.

In normal cortical maps (where there have been no experimental manipulations), the size of the receptive field varies inversely with map magnification [31]. Properties of receptive fields within the expanded and compressed regions of the final maps were tested to determine how well they match this property of biological maps.

Receptive field properties were averaged over a set of 20 nodes in the center of the compressed region, and 40 nodes in the center of the expanded region. The receptive