NMDA-Based Pattern Discrimination in a Modeled Cortical Neuron

Bartlett W. Mel
Computation and Neural Systems Program, Division of Biology, 216-76,
California Institute of Technology, Pasadena, CA 91125 USA

Compartmental simulations of an anatomically characterized cortical pyramidal cell were carried out to study the integrative behavior of a complex dendritic tree. Previous theoretical (Feldman and Ballard 1982; Durbin and Rumelhart 1989; Mel 1990; Mel and Koch 1990; Poggio and Girosi 1990) and compartmental modeling (Koch *et al.* 1983; Shepherd *et al.* 1985; Koch and Poggio 1987; Rall and Segev 1987; Shepherd and Brayton 1987; Shepherd *et al.* 1989; Brown *et al.* 1991) work had suggested that multiplicative interactions among groups of neighboring synapses could greatly enhance the processing power of a neuron relative to a unit with only a single global firing threshold. This issue was investigated here, with a particular focus on the role of voltage-dependent N-methyl-D-asparate (NMDA) channels in the generation of cell responses. First, it was found that when a large proportion of the excitatory synaptic input to dendritic spines is carried by NMDA channels, the pyramidal cell responds preferentially to spatially clustered, rather than random, distributions of activated synapses. Second, based on this mechanism, the NMDA-rich neuron is shown to be capable of solving a nonlinear pattern discrimination task. We propose that manipulation of the spatial ordering of afferent synaptic connections onto the dendritic arbor is a possible biological strategy for pattern information storage during learning.

1 Introduction

The cerebral neocortex and its connections account for almost 90% of the human brain by weight (Hofman 1989). Our understanding of its functions can be significantly advanced by an accurate input–output model for the individual pyramidal cell, the principal and most numerous neocortical cell type. How does an individual cortical pyramidal cell integrate its synaptic inputs over space and time? Which spatiotemporal patterns of synaptic activation in the dendritic tree lead to the firing of action potentials, and which do not? In short, what kind of "device" is the cortical pyramidal cell?

The response of a pyramidal neuron depends on a large number of variables, such as the detailed anatomy of the dendritic tree, numerous biophysical cell parameters, the number and distribution of synaptic inputs to the dendritic tree and their firing frequencies, and the strengths and time courses of the several excitatory and inhibitory synaptic conductances that are known to exist in these cells. The direct study of the input–output behavior of a dendritic tree under conditions of varying synaptic stimulation is not currently possible in the laboratory for technical reasons. The primary tool currently available for the study of complex dendritic trees is compartmental modeling, used to compute the time course of currents, voltages, and conductances in an electrical circuit model of an arbitrary neuron (Rall 1964; Perkel and Mulloney 1978; Traub 1982; Koch et al. 1983; Shepherd et al. 1985; Koch and Poggio 1987; Rall and Segev 1987; Shepherd and Brayton 1987; Shepherd et al. 1989; Brown et al. 1991).

Our initial hypothesis as to the integrative behavior of the neocortical pyramidal cell was derived from a simple abstract neuron type, called a sigma-pi unit (Rumelhart et al. 1986). The sigma-pi unit computes its response as a sum of contributions from a set of multiplicative clusters of inputs, and is in this sense a neural instantiation of a polynomial function. Three factors have recommended this single-unit abstraction as the seed for a biological model. First, direct monosynaptic connections between units with multiplicative input stages can encode a very general class of nonlinear associations (Feldman and Ballard 1982; Giles and Maxwell 1987; Durbin and Rumelhart 1989; Mel 1990; Mel and Koch 1990; Poggio and Girosi 1990). Second, the learning of such associations can be achieved with a simple Hebbian rule of a general type known to exist in the central nervous system (Bliss and Lømo 1973; Brown et al. 1988; Sejnowski and Tesauro 1989). Third, numerous suggestions have been made over the past decade that voltage-dependent membrane mechanisms in dendritic trees could underlie the multiplicative interactions among neighboring synapses needed for this type of model neuron (Feldman and Ballard 1982; Koch et al. 1983; Shepherd et al. 1985; Koch and Poggio 1987; Rall and Segev 1987; Shepherd and Brayton 1987; Brown et al. 1988; Shepherd et al. 1989; Durbin and Rumelhart 1989; Mel 1990; Mel and Koch 1990; Poggio and Girosi 1990; Brown et al. 1991).

2 The Biophysical Model

We sought evidence for sigma-pi–like behavior in modeling studies of a serially reconstructed layer 5 pyramidal cell from cat visual cortex (Douglas et al. 1991). The compartmental modeling program, called NEURON (Hines 1989), mapped lengths, diameters, and interconnectivity pattern of the dendritic branches into an equivalent electrical circuit that could be stimulated and probed at will. The dendritic tree shown in Figure 1 was
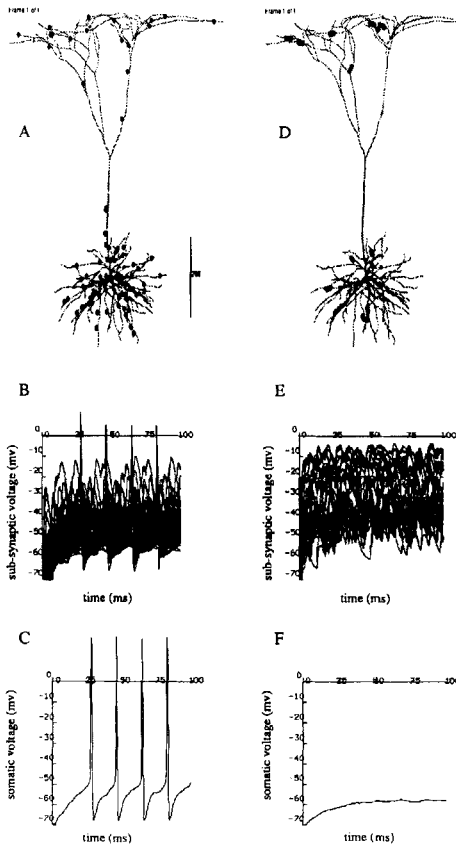
Figure 1: Testing for cluster sensitivity. Default membrane parameters were $R_m = 21,000$ Ω-cm^2, $C_m = 1.0$ $\mu F/cm^2$, $R_i = 306$ Ω-cm taken from a study of a layer 5 pyramidal cell in rat visual cortex (Stratford *et al.* 1990). (A–F) Zero-NMDA condition. In A, 80 spines were selected at random, designated by black circles. Each spine contained a synapse with a 0.8 nS AMPA conductance. In this and all other simulation runs, each synapse was stimulated asynchronously at 100 Hz for 100 msec, with a randomized interspike interval. Synaptic voltage traces are shown in B, and voltage at the cell body is shown in C where four spikes were counted. Experiment was repeated in (D–F) with 80 synapses in clusters of size 6 (plus one partial cluster of size 2). All synapses within a cluster normally fell within one 20 μm dendritic compartment. Subsynaptic voltage traces in E are larger on average than those in B, reflecting the fact that the postsynaptic depolarization is indeed boosted where synapses are closely grouped. However, each AMPA synapse injects less current postsynaptically in its strongly depolarized state. The response at the soma (F), a measure of total synaptic current, is thus attenuated. *Continued facing page.*
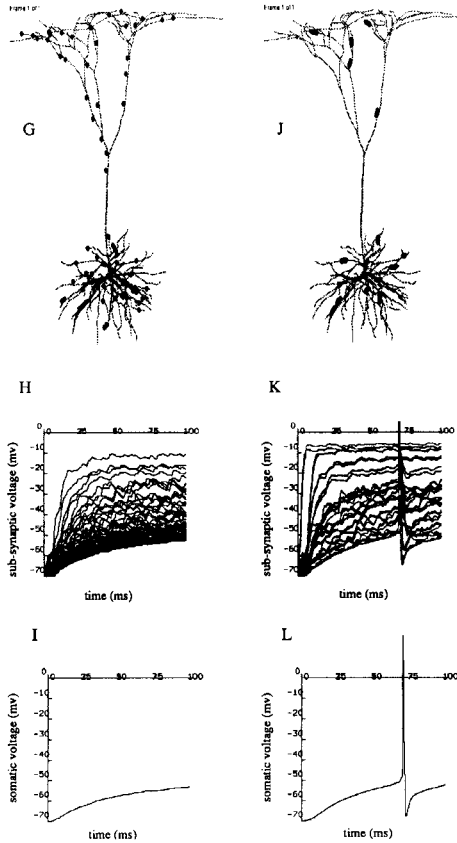
Figure 1: *Continued.* (G–L) High-NMDA condition. Experiments were repeated with NMDA peak conductance $\bar{g}_N = 0.9$ nS, and $\bar{g}_A = 0.1$ nS. A preference is now seen for clusters of size 6 (J–L) vs. clusters of size 1 (G–I) since the NMDA channels become more effective as they become more strongly depolarized.

thus converted into 903 coupled electrical compartments, one for each 20 μm length of dendrite. Membrane for unmodeled dendritic spines was incorporated into dendritic branches as in Stratford *et al.* (1990). All default simulation parameters are collected in Table 1.

In each simulation run, a set of locations (usually 80, but up to 200) was selected from across the dendritic tree, marking sites to be synaptically stimulated. At each location, an electrical compartment representing a cylindrical spine (1.0 x 0.15 μm) was attached to the associated dendritic compartment. The distal end of each spine contained an

excitatory synapse with both a fast non-NMDA, or AMPA, conductance, and a slow, voltage-dependent NMDA conductance (Mayer and Westbrook 1987). The voltage-independent AMPA conductance was modeled as an alpha-function (see Stratford *et al.* 1990):

$$G_A = \bar{g}_A \kappa t e^{-t/\tau}$$

$\kappa = e/\tau$, where peak conductance $\bar{g}_A$ varied between 0.1 to 1 nS, and time to peak was $\tau = 1$ msec. The NMDA conductance depended exponentially on membrane voltage V_m and time, modeled as follows (Jahr and Stevens 1990a,b; Zador *et al.* 1990):

$$G_N = \bar{g}_N \frac{e^{-t/\tau_1} - e^{-t/\tau_2}}{1 + \eta[\mathrm{Mg}^{2+}]e^{-(\gamma V_m)}}$$

with peak conductance $\bar{g}_N$ also in the range of 0.1 to 1.0 nS, $\tau_1 = 40$ msec, $\tau_2 = 0.33$ msec, $\eta = 0.33/$mM, $[\mathrm{Mg}^{2+}] = 1$ mM, and $\gamma = 0.06/$mV. The long NMDA time constant leads to the possibility of conductance summation across spikes in a high-frequency train; in most simulations, the NMDA

Figure 2: *Facing page.* (A) Ten randomly generated synaptic layouts were run for each cluster size between 1 and 9 with only AMPA channels; summary plot shows predicted monotonic decline in cell response as a function of increasing cluster size. (B) In case of high dendritic NMDA, response of cell shows initially *increasing*, nonmonotonic dependence on cluster size. Extensive exploration of the biophysical parameter space indicated that the observed NMDA-based cluster sensitivity phenomenon is extremely robust. Parameters varied in plots C–F above are indicated in the upper left corner of each plot. Complete list of parameter variations (usually taken one or a few at a time) included membrane resistivity from 10,000 to 100,000 Ω-cm^2 (C), membrane capacitance from 0.9 to 2.0 μF/cm^2, cytoplasmic resistivity from 100 to 362 Ω cm (several configurations were taken from Stratford *et al.* 1990); synaptic conductance waveforms were changed from saturating to non-saturating curves (D), spine membrane correction was used or not, count of activated synapses was varied from 80 to 200 (E). Synaptic strengths were normally uniform across the entire tree; in one run (F), peak synaptic conductances were set proportional to local dendritic input resistance, leading to a factor of 30 variation in conductance from the finest to the thickest dendritic branches. Spine neck resistance was also increased in one run from 173 MΩ 6-fold to 1038 MΩ, which led to a significant reduction, but did not abolish, the cluster sensitivity (i.e. there persisted a 50% increase in average cell response as clusters grew to from size 1 to size 3). NMDA-based cluster sensitivity was also seen for both apical and basal dendrites individually, though absolute difference between clustered and unclustered activation patterns as measured at a passive soma was significantly larger for basal (40%) than apical (18%) dendrites. Finally, in a pilot study on a smaller layer 3 pyramidal cell (Douglas *et al.* 1991), cluster sensitivity was seen at spike train frequencies of 20 and 100 Hz, or when only a single spike was delivered to each synapse.

conductance was assumed to saturate to $\bar{g}_N$, corresponding to the case of many presynaptic transmitter molecules and few receptor sites. In one case (Fig. 2D), NMDA conductance was allowed to summate across presynaptic spikes, and could thus achieve values much larger than $\bar{g}_N$. In most experiments, all of the AMPA conductances were assigned the same peak value, as were all of the NMDA conductances, independent of location on the dendritic tree.

The NMDA channel is normally discussed for its role in synaptic plasticity (see Collingridge and Bliss 1987; Brown *et al.* 1988). For present
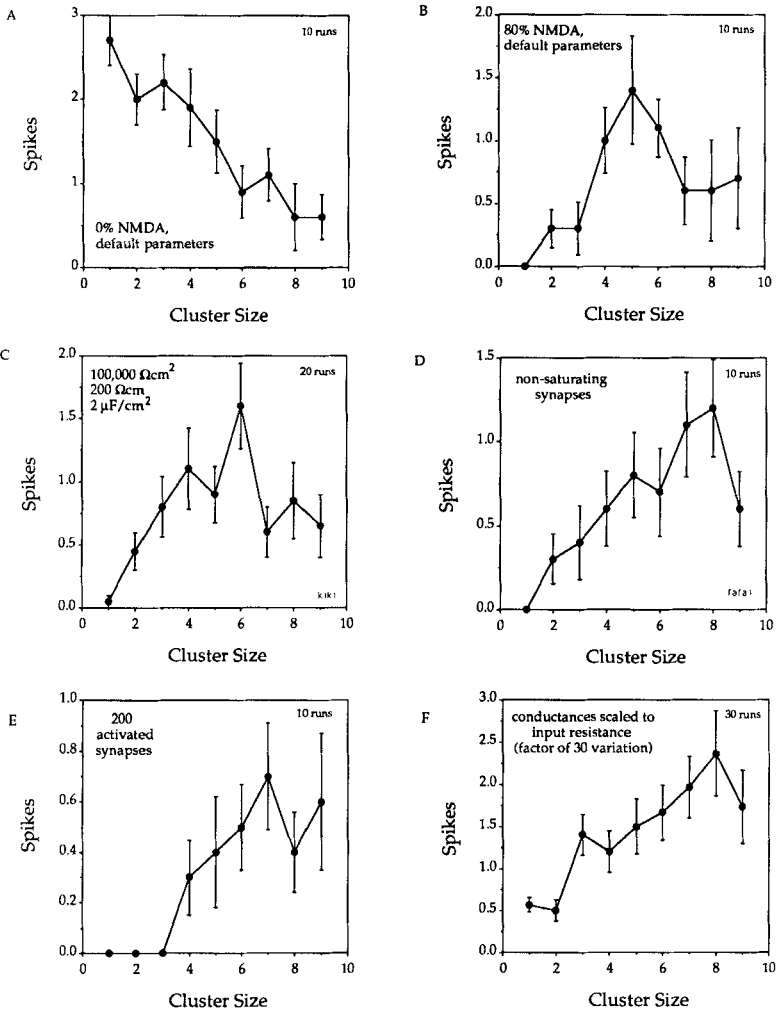
Table 1: Default Simulation Parameters.

R_m	21,000 Ω-cm^2
C_m	1.0 μF/cm^2
R_i	306 Ω-cm
Input resistance (soma)	51 MΩ
Time constant (soma)	19 msec
Resting potential	-70 mV
Dendritic spines	L = 1.0 μm, D = 0.15 μm, neck resistance = 173 MΩ
Synaptic stimulation	100 Hz with randomized interspike interval
Active synapses	80, stimulated asynchronously
AMPA conductance	$\bar{g}_A$ = 0.1–1.0 nS
AMPA time constant	τ = 1 msec
NMDA conductance	$\bar{g}_N$ = 0.1 to 1.0 nS
NMDA time constants	τ_1 = 40 msec, τ_2 = 0.33 msec
NMDA Parameters	η = 0.33/mM, [Mg^{2+}] = 1 mM, and γ = 0.06/mV
Synaptic reversal potential	0 mV (both AMPA and NMDA)
Compartments	903 plus 1 per activated dendritic spine
Compartment length	20 μm
Integration time step	0.1 msec

purposes, however, the crucial difference between AMPA and NMDA channels relates to their dependence on voltage. Briefly, the synaptic current injected through an AMPA channel declines steadily as the membrane in which it sits is depolarized from rest. The depolarizing effect of one AMPA synapse thus reduces the effectiveness of neighboring AMPA synapses, and vice versa. In contrast, NMDA current *increases* steadily as the membrane is depolarized from -70 to -30 mV (Mayer and Westbrook 1987), such that NMDA synapses can be *more* effective when activated in groups. Except for NMDA channels, the dendritic membrane was electrically passive. The soma contained Hodgkin–Huxley-style sodium and potassium channels in addition to a passive leak conductance.

At the beginning of each run, the entire cell was set to a resting potential of -70 mV. A 100-Hz spike train was then delivered asynchronously to each of the preselected synapses, and the response at the soma was "recorded" over a 100-msec period. The primary measure of cell response was the number of spikes generated at the soma during the first 100 msec of synaptic stimulation, but passive-soma response measures were also used with similar results, including peak somatic potential, and time integral of somatic potential.

Our main hypothesis held that if sufficiently endowed with NMDA channels, the cell would respond selectively to patterns of stimulation in which the activated synapses were clustered, rather than uniformly distributed, across the dendritic tree. We first examined sensitivity to

cluster size in a dendritic tree with only AMPA channels (Fig. 1A–F). Sample runs are shown for clusters of size 1 (A,B,C) and 6 (D,E,F). As expected, clusters of 6 led to higher average subsynaptic depolarization (E vs. B), and hence less current injection per synapse. In consequence, the cell response was significantly reduced (F vs. C). This experiment was repeated with 90% of the total peak conductance assigned to NMDA channels.[1] In contrast to the pure AMPA case, a larger response was seen at the cell body when synapses were distributed in clusters of 6 (J, K, L), as compared to clusters of 1 (G, H, I). Note that "clustering" here refers to the spatial configuration of *activated* synapses, and not an underlying inhomogeneity in the anatomical distribution of dendritic spines.

To confirm these effects statistically, 10 random layouts were generated for each cluster size from 1 to 9, for both zero-NMDA and high-NMDA conditions. Figure 2A shows the predicted monotonic decline in cell response as a function of increasing cluster size when only AMPA channels were present. In contrast, Figure 2B shows for the high-NMDA case an initially increasing, nonmonotonic dependence on cluster size. The initial increase in cell response as clusters grow larger is explained by the average increase in membrane potential within each cluster, which leads to increased average effectiveness of the NMDA synapses. The subsequent fall-off of the response at larger cluster sizes is due to the diminishing driving potential as the membrane under each cluster saturates to the synaptic reversal potential. The specific cluster size at which the somatic response peaked, in this case 5, was not fixed, but typically varied over a range of 5 to 9 depending on the value of peak conductance assigned to the activated synapses. This basic effect of nonmonotonic cluster sensitivity was observed to hold under a very wide range of biophysical parameter manipulations (Fig. 2C–F).

We conclude that a cortical pyramidal cell with a major proportion of its excitatory synaptic current carried by NMDA channels could respond preferentially to clustered distributions of activated synapses. In a pilot study on a smaller layer 2–3 pyramidal cell, a continuous increase in cluster sensitivity was seen as the proportion of peak NMDA conductance [i.e., $\bar{g}_N/(\bar{g}_N + \bar{g}_A)$] was increased from 0 to 100%. In the current study, it was observed that a small proportion of AMPA conductance was useful in rapidly boosting the local membrane potential, and hence neighboring NMDA channels, into the voltage range where the latter begin to open significantly. The actual proportion of NMDA conductance in cortical pyramidal cells has not yet been determined, though its contribution to

---

[1]Note that due to its dependence on voltage, the actual NMDA conductance achieved during synaptic stimulation will not normally approach the nominal maximum value $\bar{g}_N$. For example, with a 90% proportion of NMDA peak conductance, that is, $\bar{g}_N/(\bar{g}_N + \bar{g}_A) = 0.9$, the peak EPSC in response to a synchronous single-shock activation of 100 randomly placed excitatory synapses for the cell and parameters of Fig. 1 is only 30–40% due to current through NMDA channels. However, at least 85% of the total synaptic charge in this case is carried by the NMDA conductance, since its time constant is relatively much longer.

normal synaptic activation in these cells is known to be large (Miller *et al.* 1989; Fox *et al.* 1989; see also Keller *et al.* 1991 with respect to hippocampal granule cells).

## 3 A Basis for Nonlinear Pattern Discrimination

Regarding the biological utility of dendritic cluster sensitivity, we observe that any mechanism that allows a cell to distinguish certain patterns of stimulation from others can act as a basis for memory. For example, if the subset of afferents activated in stimulus condition A terminate in a clustered configuration on the dendrite of our pyramidal cell, then pattern A will fire the cell with high probability, and will thus be "preferred." Another pattern, B, for which no such clustering condition holds, will with high probability *fail* to fire the cell.

To establish a first empirical lower bound on the nonlinear pattern discrimination capacity of the single pyramidal cell, we sought to load a set of pattern "preferences" onto the dendritic tree in just this way. One hundred gray-level images of natural scenes were used as a convenient set of complex high-dimensional patterns. A simple visual preprocessing

---

Figure 3: *Facing page.* Gray-level images of natural scenes were used as a convenient set of complex high-dimensional training and test patterns. Each image was mapped onto a sparse binary pattern of activity over a set of orientation-selective visual filters. Four filters were centered at each of the 64 by 64 image pixel locations, sensitive to horizontal, vertical, and right- and left-oblique intensity edges, giving a total population of 16,384 orientation-selective visual "units." A global thresholding operation guaranteed that exactly 80 visual units were active in response to any presented image. Each visual unit gave rise to an afferent "axon" that could provide input to the dendritic tree. A training image was "loaded" by arranging that the 80 afferents activated by the image terminated in a clustered configuration on the dendritic tree that had been previously observed to fire the cell. After 50 images were loaded in this way, at most 4000 afferents made a synapse onto the dendrite—fewer in actuality due to feature overlap among training images. All remaining uncommitted visual afferents ($> 12,384$) were randomly mapped onto the dendrite as well, in order that any presented image should drive exactly 80 synapses in the dendritic tree. This resulted in a total of 16,384 synaptic "contacts" onto the dendritic surface, at a maximum density of one spine per micrometer of dendritic length. While a biologically unrealistic number of synapses for this single cell (5000 more probable), this control measure guaranteed that any observed preference for trained vs. control images in the experiments of Figure 4 was the result of the pattern of synaptic clustering, and not a result of total synaptic activation. Note that the response properties of the orientation-selective visual units and the pattern of afferentation described here are not intended as models of cortical visual processing; any other type of (nonvisual) patterns and preprocessing stage might have been used.

step was used to map each image onto a sparse binary pattern of "neural" activity, consisting of 80 active lines within a bundle of 16,384 orientation-selective visual "axons" (Fig. 3). Another source of patterns or another type of preprocessing could have been used for these experiments; the visual modality was chosen only for convenience, and was not intended as a model of any aspect of cortical visual processing.

Fifty images were chosen to act as a training set, 50 as a control set. For the first training image, a layout of 80 synapses in clusters of 8 was generated that produced at least one spike at the soma during 100 msec of 100 Hz stimulation (about 50% of clustered layouts of this kind generated at least one spike). The 80 orientation-selective visual afferents driven by the training image were then randomly mapped onto the 80 dendritic locations specified by the precomputed clustered layout (Fig. 3). The process was repeated for each training image in turn, except that a visual afferent, once mapped onto the dendritic tree, was never remapped in the

80 features activated by each image

16,384 oriented feature detectors
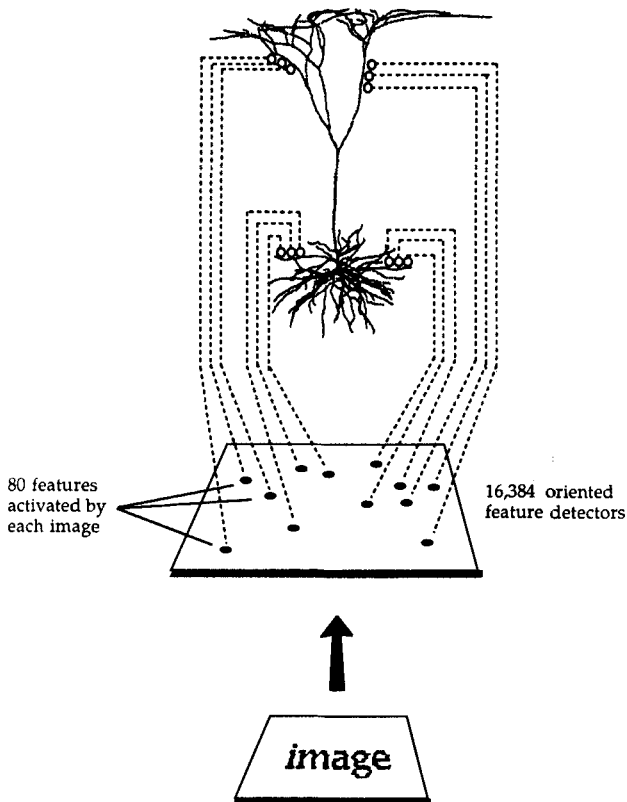
image

processing of a subsequent training image. When this loading procedure was complete, all remaining uncommitted visual afferents were randomly mapped to the dendrite as well, so that any image, trained or untrained, activated exactly 80 synapses of the same strength and with the same intensity.

Five different image types were then presented to the cell, including intact training images, three types of corrupted training images, and control images. Average cell responses to each during 100 msec are shown in Fig. 4. The most prominent feature of these results is the cell's overwhelming selectivity for the 50 trained vs. 50 control images. Only 1 in 100 control images presented to the cell in two complete runs generated a single spike. In contrast, 87% of the trained images generated at least one spike, with an average spike count of 1.25. Intermediate responses were seen for the three categories of corrupted training images. It is important to note that the procedure used for "storing" training patterns in the dendritic tree was designed to make the task of discriminating training vs. control patterns as difficult as possible for a linear neuron with only a global thresholding nonlinearity (e.g., a perceptron), since *every* presented pattern activated the same number of synaptic inputs, with the same synaptic strength. The much greater probability of firing an action potential in response to a trained (clustered) vs. control (diffuse) pattern is thus evidence that the dendritic tree is a more powerful pattern discriminator than a simple perceptron. Significantly reduced cell responses to linear superpositions of 2 training patterns is further evidence for this conclusion (Fig. 4).

We may relate the standard measure of pattern discrimination capacity used for the simple perceptron to the present situation. Cover (1965) proved that an N-input perceptron has a capacity of $2N$ patterns, counting both training *and* control patterns (see Hertz *et al.* 1991). This result means that for $2N$ randomly chosen patterns, some randomly designated as training patterns, others as controls, the probability that a perceptron can perfectly discriminate all training from control patterns is $1/2$. This simple result cannot be applied in the present situation, however, as we consider here instead the ability of the neuron to discriminate k training patterns (in this case 50) from *all other possible patterns* acting as controls. In this more demanding case, we do not require classification performance to be perfect, but rather that it fall within an acceptable error tolerance. For example, in the present nonoptimized experiment with 50 stored training images, the probability of misclassification was 14% (13% false negatives, 1% false positives).

4 Conclusions

The single statistical parameter, cluster size, has proved to be a remarkably good predictor of responses to complex patterns of synaptic in-
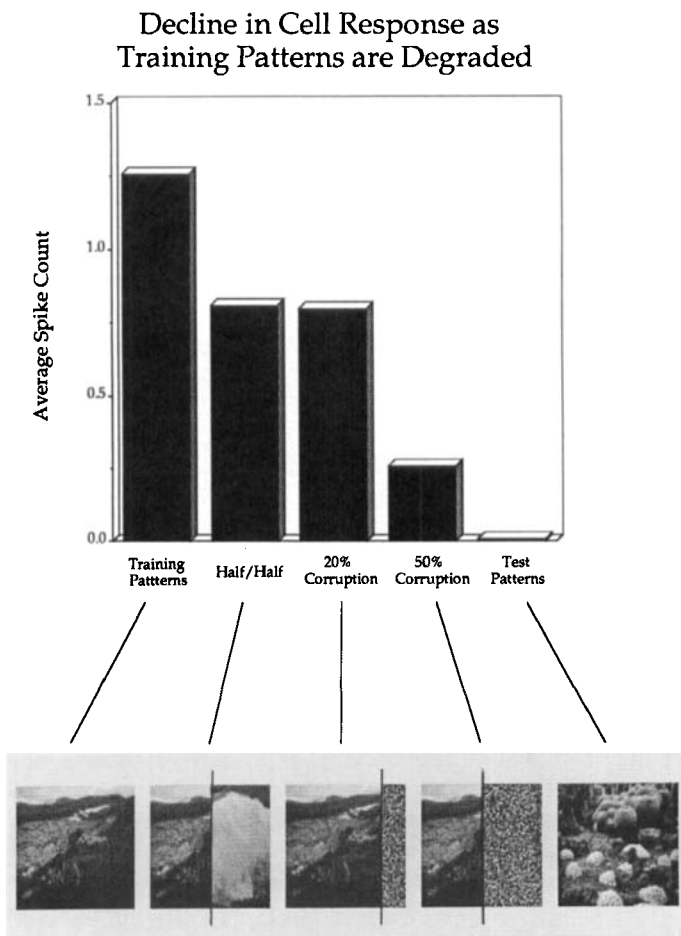
Figure 4: Average cell responses to five types of presented images, combined over two runs with 50 training and control images interchanged. Largest responses were to whole training images and smallest to control images. In three intermediate cases, progressively attenuated responses were seen to (1) random 50/50 composites of features taken from two training images, (2) training images with 20%, and (3) 50% of their visual features randomly replaced. Pictorial representation of each category is only suggestive of these manipulations; corruptions to training images were actually carried out within the internal orientation-selective visual representation. The 50/50 composite case provides direct evidence that the dendritic tree is capable of nonlinear discrimination, since a thresholded linear neuron would respond equivalently to any such superposition of training patterns.

put in an NMDA-rich neuron. Further, a simple experiment has shown how the ability to respond selectively to statistically inhomogeneous distributions of activated synapses could act as a basis for nonlinear pattern discrimination within a dendritic tree. It must be noted, however, that other membrane mechanisms may exist in pyramidal dendrites in addition to NMDA channels that could mimic, enhance, alter, or abolish the cluster-based integrative behavior observed here. Nonetheless, the fact that NMDA channels can by themselves underlie a capacity for nonlinear pattern discrimination in dendritic trees is, in and of itself, useful knowledge.

Beyond the widely accepted proposition that modifiable synaptic weights are the likely sites for neural information storage, these results highlight the probable importance of synaptic *ordering* in the generation of cortical cell responses as well. Given the rich combinatorics that are possible in choosing spatial arrangements of 10^4 synapses in a dendritic tree, the true discrimination capacity of the single pyramidal cell is likely to be much greater than the level that has been empirically demonstrated here. In this vein, preliminary experiments have been conducted on an abstract cluster-sensitive neuron with 10,000 synapses in which the direct-loading procedure described above was replaced by a Hebb-type learning rule. Much larger discrimination capacities were easily achieved in this way for the abstract neuron (unpublished observations); application of a similar learning scheme within the present biophysical modeling context awaits further work (but see Brown *et al.* 1991). Elsewhere (Mel 1990; Mel and Koch 1990), it has been proposed that the three-dimensional geometry of intersecting axonal and dendritic arbors in the neocortical column are ideally suited to allow a given axon to establish synaptic contacts at many different dendritic loci as dictated by the constraints of a learning task. It must also be noted that the pyramidal cell is not alone in the cortex, but is a member of a group of cells within the cortical microcolumnar structure (Jones 1981; Peters and Kara 1987) within which a high degree of interaction is likely to occur. Within such a network of neurons, however, the issue of storage capacity is profound, and will need to be addressed in the future both analytically and through more sophisticated simulations.

Acknowledgments

References

Bliss, T. V. P., and Lømo, T. 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.*, **232**, 331–356.

Brown, T. H., Chang, V. C., Ganong, A. H., Keenan, C. L., and Kelso, S. R. 1988. Biophysical properties of dendrites and spines that may control the induction and expression of long-term synaptic potentiation. In *Long-Term Potentiation: From Biophysics to Behavior*, 201–264. Alan R. Liss, New York.

Brown, T. H., Mainen, Z. F., Zador, A. M., and Claiborne, B. J. 1991. Self-organization of hebbian synapses in hippocampal neurons. In *Advances in Neural Information Processing Systems*, Vol. 3, R. Lippmann, J. Moody, and D. Touretzky, eds., pp. 39–45. Morgan Kaufmann, Palo Alto.

Collingridge, G. L., and Bliss, T. V. P. 1987. NMDA receptors: Their role in long-term potentiation. *Trends Neurocsi.* **10**, 288–293.

Cover, T. M. 1965. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE Trans. Elect. Comput.* **14**, 326–334.

Douglas, R. J., Martin, K. A. C., and Whitteridge, D. 1991. An intracellular analysis of the visual responses of neurons in cat visual cort. *J. Physiol.* **440**, 659–696.

Durbin, R., and Rumelhart, D. E. 1989. Product units: A computationally powerful and biologically plausible extension to backpropagation networks. *Neural Comp.* **1**, 133.

Feldman, J. A., and Ballard, D. H. 1982. Connectionist models and their properties. *Cog. Sci.* **6**, 205–254.

Fox, K., Sato, H., and Daw, N. W. 1989. The location and function of NMDA receptors in cat and kitten visual cortex. *J. Neurosci.* **9**, 2443–2454.

Giles, C. L., and Maxwell, T. 1987. Learning, invariance, and generalization in high-order neural networks. *Appl. Opt.* **26**(23), 4972–4978.

Hertz, J., Krogh, A., and Palmer, R. G. 1991. *Introduction to the Theory of Neural Computation.* Addison-Wesley, Redwood City, CA.

Hines, M. 1989. A program for simulation of nerve equations with branching geometries. *Int. J. Biomed. Comput.* **24**, 55–68.

Hofman, M. A. 1989. On the evolution and geometry of the brain in mammals. *Prog. Neurobiol.* **32**, 137–158.

Jahr, C. E., and Stevens, C. F. 1990a. A quantitative description of NMDA receptor-channel kinetic behavior. *J. Neuroscience* **10**, 1830–1837.

Jahr, C. E., and Stevens, C. F. 1990b. Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics. *J. Neurosci.* **10**, 3176–3182.

Jones, E. G. 1981. Anatomy of cerebral cortex: Columnar input-output relations.

In *The Organization of Cerebral Cortex*, F. O. Schmitt, F. G. Worden, G. Adelman, and S. G. Dennis, eds. MIT Press, Cambridge, MA.

Keller, B. U., Konnerth, A., and Yaari, Y. 1991. Patch clamp analysis of excitatory synaptic currents in granule cells of rat hippocampus. *J. Physiol.* **435**, 275–293.

Koch, C., and Poggio, T. 1987. Biophysics of computation: Neurons, synapses, and membranes. In *Synaptic Function*, G. E. Edelman, W. F. Gall, and W. M. Cowan, eds., pp. 637–698. Wiley, New York.

Koch, C., Poggio, T., and Torre, V. 1983. Nonlinear interactions in a dendritic tree: Localization, timing, and role in information processing. *Proc. Natl. Acad. Sci. U.S.A.* **80**, 2799–2802.

Mayer, M. L., and Westbrook, G. L. 1987. The physiology of excitatory amino acids in the vertebrate central nervous system. *Prog. Neurobiol.* **28**, 197–276.

Mel, B. W. 1990. The sigma-pi column: A model for associative learning in cerebral neocortex. CNS Memo #6, Computation and Neural Systems Program, California Institute of Technology.

Mel, B. W., and Koch, C. 1990. Sigma-pi learning: On radial basis functions and cortical associative learning. In *Advances in Neural Information Processing Systems*, Vol. 2, D. S. Touretzsky, ed. Morgan Kaufmann, San Mateo, CA.

Miller, K. D., Chapman, B., and Stryker, M. P. 1989a. Visual responses in adult cat visual cortex depend on N-methyl-D-aspartate receptors. *Proc. Natl. Acad. Sci. U.S.A.* **86**, 5183–5187.

Perkel, D. H., and Mulloney, B. 1978. Electrotonic properties of neurons: steady-state compartmental model. *J. Neurophysiol.* **41**, 621–639.

Peters, A., and Kara, D. A. 1987. The neuronal composition of area 17 of rat visual cortex. IV. The organization of pyramidal cells. *J. Comp. Neurol.* **260**, 573–590.

Poggio, T., and Girosi, F. 1990. Regularization algorithms for learning that are equivalent to multilayer networks. *Science* **247**, 978–982.

Rall, W. 1964. Theoretical significance of dendritic trees for neuronal input-ouput relations. In *Neural Theory and Modeling*, R. F. Reiss, ed., pp. 73–97. Stanford University Press, Stanford, CA.

Rall, W., and Segev, I. 1987. Functional possibilities for synapses on dendrites and on dendritic spines. In *Synaptic Function*, G. E. Edelman, W. F. Gall, and W. M. Cowan, eds., pp. 605–636. Wiley, New York.

Rumelhart, D. E., Hinton, G. E., and McClelland, J. L. 1986. A general framework for parallel distributed processing. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 1, D. E. Rumelhart, J. L. McClelland, eds., pp. 45–76. Bradford, Cambridge, MA.

Sejnowski, T. J., and Tesauro, G. 1989. The Hebb rule for synaptic plasticity: Algorithms and implementations. In *Neural Models of Plasticity*, J. H. Byrne and W. Berry, eds., pp. 94–103. Academic Press, Cambridge, MA.

Shepherd, G. M., and Brayton, R. K. 1987. Logic operations are properties of computer-simulated interactions between excitable dendritic spines. *Neurosci.* **21**, 151–166.

Shepherd, G. M., Brayton, R. K., Miller, J. P., Segev, I., Rinzel, J., and Rall, W. 1985. Signal enhancement in distal cortical dendrites by means of in-

teractions between active dendritic spines. *Proc. Natl. Acad. Sci. U.S.A.* **82**, 2192–2195.

Shepherd, G. M., Woolf, T. B., and Carnevale, N. T. 1989. Comparisons between active properties of distal dendritic branches and spines: Implications for neuronal computations. *Cog. Neurosci.* **1**, 273–286.

Stratford, K., Mason, A., Larkman, A., Major, G., and Jack, J. 1990. The modeling of pyramidal neurons in visual cortex. In *The Computing Neuron*, R. Durbin, C. Miall, and G. Mitchison, eds., pp. 296–321. Addison-Wesley, Wokingham, England.

Traub, R. D. 1982. Simulation of intrinsic bursting in CA3 hippocampal neurons. *Neurosci.* **7**, 1233–1242.

Zador, A., Koch, C., and Brown, T. H. 1990. Biophysical model of a hebbian synapse. *Proc. Natl. Acad. Sci. U.S.A.* **87**, 6718–6721.