

SPIKE TRAIN PROCESSING BY A SILICON NEUROMORPH: THE ROLE OF SUBLINEAR SUMMATION IN DENDRITES

David P. M. Northmore and John G. Elias

Departments of Psychology and Electrical Engineering
University of Delaware
Newark, DE 19716

A dendritic tree, as part of a silicon neuromorph, was modelled in VLSI as a multibranched, passive cable structure with multiple synaptic sites that either depolarize or hyperpolarize local "membran patches", thereby raising or lowering the probability of spike generation of an integrate-and-fire "soma". As expected from previous theoretical analyses, contemporaneous synaptic activation at widely separated sites on the artificial tree resulted in near-linear summation, as did neighboring excitatory and inhibitory activations. Activation of synapses of the same type close in time and space produced local saturation of potential, resulting in spike train processing capabilities not possible with linear summation alone. The resulting sublinear synaptic summation, as well as being physiologically plausible, is sufficient for a variety of spike train processing functions. With the appropriate arrangement of synaptic inputs on its dendritic tree, a neuromorph was shown to discriminate input pulse intervals and patterns, pulse train frequencies, and detect correlation between input trains.

1 Introduction

The goal of constructing artificial nervous systems in hardware for generating behavior in real environments can best be approached at the present time by using the techniques of VLSI. Mead (1989) has made the point that certain aspects of CMOS transistor physics naturally lend themselves to the emulation of neural structures and functions. Current technologies allow the construction of high densities of neuron-like elements, or neuromorphs to use Mead's term, for use in large-scale working systems for robots. Our aim was firstly to create VLSI neuromorphs that incorporate some of the basic operational principles of real neurons, and secondly to develop a system for interconnecting thousands of neuromorphs to form networks (Elias, 1993; Elias and Northmore, 1995). This approach is providing a useful modeling medium for testing ideas about brain function, particularly at the microcircuit and small system level (Elias and Northmore, 1996), and at the single neuron level, as in the work described here. Although computer modelling is supremely flexible, and can be made as realistic as needed (e.g., Koch and Segev, 1989), hardware models have the advantage of speed and compactness, making them better suited for robots that interact with the real world. While a number of VLSI implementations of neurons and neuronal processing systems

have been built (see Douglas et al. 1995 for review), most have focused on processing specific forms of sensory information, usually visual or auditory. Our approach is to develop "general purpose" neuromorphs that can be readily used in large-scale networks.

Signalling between our neuromorphs, as in much of biological nervous systems, is by means of impulses. Our electronic system takes full advantage of high-speed digital circuitry for multiplexing impulse signals on a limited number of wires, and importantly, makes it possible to route a neuromorph's output along pathways with programmable delays to an arbitrary set of synapses in a network. The utility of such a system, either as a modeling tool or for adapting a network to a practical task, requires that all these aspects of connectivity be easily and rapidly programmable.

Neuromorph design is, of necessity, an exercise in the art of the possible. With the objective of large-scale systems in mind, each neuromorph should be compact, dissipate little power, be insensitive to silicon process variations, yet perform with adequate reliability and precision. Its circuit components should be simple, yet allow as much programmability as possible. To meet these requirements, our present silicon neuromorphs (Elias, 1993; Elias and Northmore, 1995) were designed to embody the characteristics of the classical neuron (Eccles, 1957; Rall et al., 1992). Each comprises a spatially extensive, branched dendritic tree that receives and integrates pulsatile inputs at a number of synaptic sites, both excitatory and inhibitory. The resulting postsynaptic potentials diffuse passively to a "soma" that generates pulsatile outputs. Simplistic though it is, we show that the classical neuron, constructed with linear components for the most part, exhibits a number of useful signal processing capabilities.

Since the time that Rall (1964) showed that a passive dendritic branch could perform spatiotemporal filtering, a body of mostly theoretical work has led to a better understanding of how a neuron's response to afferent impulse trains are affected by the branching structure of dendritic trees and the spatial disposition of synaptic inputs (for reviews see Rall et. al., 1992; Mel, 1994). Because the same principles apply to our artificial dendritic tree neuromorphs, we are able to demonstrate that these devices can respond selectively to input spike frequencies, discriminate temporal patterns of spikes, and detect correlations between input spike trains, given the appropriate arrangement of synaptic connections upon the dendritic tree.

Because of recent interest in synchronous spiking activity in cortical neurons and the possibility that it may provide a solution to the binding problem (Singer and Gray, 1995), we also examined the response of our neuromorph to varying degrees of input synchrony. It might be assumed that a neuron could detect synchronous firing in a set of afferents simply by temporal summation: the more nearly coincident in time the afferent synaptic activations, the larger will be the summated EPSPs. However, a number of factors are likely to combine to make N highly synchronized input spikes less effective in exciting the recipient neuron than N spikes distributed over a longer time

period. Recently, simulations (Murthy and Fetz, 1994; Bernander et al. 1994) have examined the conditions under which a neuron's refractoriness following the firing of a spike could limit temporal summation of highly synchronized inputs, making them relatively less effective. In this study, we examine how sublinear synaptic summation in a neuromorph can influence its response to inputs with various degrees of synchrony, and compare these results with those produced by refractoriness alone. We also explore some synaptic arrangements that detect synchrony more efficiently than simple temporal summation.

2 Experimental Hardware System

Our silicon neuron, or neuromorph, comprises an artificial dendritic tree (ADT) (Elias, 1993), and a leaky integrate-and-fire "soma" (Fig. 1b). The dendritic branches are composed of a series of identical compartments each with a capacitor, C_m , representing a membrane capacitance, and two programmable resistors, R_a and R_m , representing the axial cytoplasmic and membrane resistances. Synapses at each compartment are emulated by a pair of MOS field effect transistors, one of which, the excitatory synapse, enables inward current, moving the potential of the membrane capacitance in a depolarizing direction. The other transistor, emulating an inhibitory synapse, enables outward current, hyperpolarizing the membrane capacitance. There is also a facility, for converting alternate inhibitory synapses into shunting or "silent inhibitory" synapses that pull the membrane potential towards a potential close to the resting value. The potential appearing at the soma (V_s , see Fig 1b) determines the output spike firing rate in conjunction with the integration time constant, RC , and the threshold, V_{th} . Synapses are activated by a brief impulse signal (50 nsec) applied to the gate terminals of the synapse transistors.

Experiments were performed on a single neuromorph consisting of an 8-branch ADT fabricated with a 2 μ m CMOS double-poly n-well process on a 2 x 2 mm MOSIS Tiny Chip format. The dendrite dynamics was set by switched-capacitor circuits that emulated R_a and R_m , to give membrane time constants in the range of 5-20 msec. The dendritic trunk is terminated at a source follower to provide a low impedance input to the on-chip integrate-and-fire spike generator which has a programmable time constant (Fig. 1b). For most experiments, the spike generator had a relatively short RC time constant of 0.75 msec in order to produce reasonably high spike firing rates. The neuromorph under study was embedded in the "virtual wire system" previously described (Elias, 1993). This system allows spikes generated by neuromorphs to be distributed with programmable delays to arbitrary synapses throughout a network. A host computer set the spike firing threshold, V_{th} , (Fig. 1b), and generated the spatiotemporal patterns of spikes used as test stimuli for the neuromorph. It also recorded the membrane potential at the soma, V_s , and read the "spike history", which is part of the virtual wire system, to obtain the times of occurrence of the output spikes generated by the neuromorph.

Calculations show that the conductance of the synapse transistors, when activated for 50 nsec, is sufficient to

move the potential on the capacitor in the activated compartment almost instantaneously to a level close to the supply voltages, either 5 volts for excitatory activation, or 0 volts for inhibitory activation. The charge delivered during a synaptic activation depends upon the conductance of the transistor in the on state and the potential on the compartmental capacitor. The diffusion of charge through the dendritic tree leads to a response at the soma that can outlast the activating pulse by up to nine orders of magnitude at the longest membrane time constants. The shape and duration of the synaptic impulse response depends upon dendrite dynamics, which by means of the adjustable resistors R_a and R_m , is programmable over a range of six orders of magnitude (Elias and Northmore, 1995). Excitatory impulse responses typical of those obtained under the prevailing experimental conditions are shown in Figure 2b, f. Inhibitory impulse responses have similar shape, but are negative-going. The amplitude of an impulse response appearing at the soma falls off exponentially with the distance of the activated synapse from the soma while its latency to peak is approximately proportional to distance (Jack et al. 1975; Agmon-Snir and Segev, 1993; Elias and Northmore, 1995). Hence, the efficacy and delay accorded an input can be controlled to a certain extent by choosing the dendritic locus of synaptic activation. The overall effect at the soma, however, depends not only on its electrotonic distance from the soma but also on the local synapse voltage just prior to activation. This is the basis for the sublinear effect reported in this paper.

3 Results

3.1 *Two-pulse responses.* The classical modes of signal combination in neurons involve the spatial and temporal summation of postsynaptic potentials, and are often thought of as linear processes. Near-linear summation of two inputs can be obtained with the synaptic connections shown in Figure 2a, in which two input spikes, A and B activate excitatory synapses at two sites that are equidistant from the soma but on different branches of the dendritic tree. The smaller impulse responses in Figure 2b measured at the soma were evoked by spikes A or B delivered separately to their respective synapses. When both A and B were delivered with some intervening time interval (2 msec in Figure 2b), the resulting combined response (large response in Fig 2b) was very close to the sum of the individual impulse responses. Figure 2c shows that as the A-B interval lengthened, the integral of the combined response was always twice that of a single impulse response, while the peak amplitude of the combined response declined, an effect that was reflected in the mean rate of spiking of the soma (Fig. 2d). The precise form of the spike output function depends upon the setting of the spike generator threshold and its RC time constant. In the typical case shown in Fig. 2d, the neuromorph fired 0-3 spikes on each trial, the number depending upon the A-B interval. This accounts for the tendency to quantization shown by the plateaus on the mean spike response curves.

A departure from near-linear summation can be demonstrated by the synaptic connections shown in Figure 2e where spikes A and B depolarize the same dendritic "membrane patch" via excitatory synapses. The combined

response is significantly less than the sum of the individual responses for short A-B intervals, both in terms of the integral and peak amplitude (Fig. 2g). As a consequence, the neuromorph's soma potential and its spike output becomes tuned to a range of input spike intervals (Fig. 2h). Here the neuromorph fired 0-2 spikes depending on the A-B interval. Raising the spike generator threshold reduces the number of output spikes and increases the sharpness of tuning.

3.2 Dendritic saturating sublinearity. Sublinear summation of the type just described is due to a reduction in the local synaptic driving potential in dendrites (Rall, 1964; Shepherd and Koch, 1990; Koch and Poggio, 1992). In our present ADTs, each synaptic activation nearly saturates the potential of the activated compartment (V_{top} in the case of an excitatory synapse, GND for an inhibitory, see Fig. 1b). If, before this potential has fully decayed, a subsequent activation occurs at the same compartment it will deliver less charge than the first and add correspondingly smaller increments to the potential appearing at the soma. The resultant conditioning of an impulse response by prior activity can be influenced in various ways. The sublinear effects are promoted by any condition that slows the decay of potential from the activated dendritic compartment. For example, reducing the leakage of charge through the membrane resistance, by increasing R_m , lengthens the optimal interval of the two-spike response (e.g., Fig. 2g,h). Sublinear effects may also be prolonged by reducing the spread of charge along the dendrite by increasing the axial resistance, R_a . Thus, all spike-interval and frequency-selective effects due to sublinear summation can be readily tuned in our neuromorphs by varying the dendritic dynamics (Elias & Northmore, 1995). Diffusion of charge from a compartment can also be influenced by contemporaneous synaptic activations in neighboring compartments, so that the spatiotemporal arrangement of synaptic activations on dendrites can lead to a number of interesting and useful effects. Some of these will be illustrated by the neuromorphic responses to input spike trains of different frequencies.

3.3 Spike-Frequency selectivity. Figure 3 shows the results of experiments to demonstrate how an appropriate pattern of synaptic connections confers spike-frequency-selective properties on a single neuromorph, which for this simple illustrative case, requires two dendritic branches. Figure 3a shows the synaptic connections made by two afferent trains of input spikes, A and B, the spikes being Poisson distributed with mean frequencies varying from 0 to 500 spikes per second. Poisson distributed spike trains were used because they resemble natural spiking behavior more closely than constant frequency trains, although the latter gave qualitatively similar results. Each input spike simultaneously activated the specified synapses on its corresponding dendritic branch. Figure 3b shows the mean soma voltage, V_s , as a function of mean input spike frequency for different synaptic connection patterns. When a single excitatory synapse was driven at increasing frequency, V_s increased linearly at first and then with ever decreasing rate (curve A1), demonstrating sublinear summation at one dendritic compartment. As neighboring synapses were recruited for simultaneous activation by the train, V_s increased in amplitude, but with more pronounced sublinearity (curves A2, A3, A4). This effect is attributable to a decreased potential gradient in the

neighborhood of the activated region of dendrite. Similar effects on membrane hyperpolarization are obtained by activating a cluster of the hyperpolarizing inhibitory synapses.

A more linear response from a synaptic site can be obtained by increasing the local potential gradient. This is achieved by activating neighboring synapses of opposite polarity, as is illustrated on the lower branch of Figure 3a. The synaptic interaction in this case leads to an essentially linear relationship between V_s and input frequency (Fig. 3b, curve B). Note that because the inhibitory synapse is closer to the soma than is the excitatory synapse, the simultaneous activation of the pair results in net inhibition, so that curve B is actually negative but plotted positive for comparison. (The pairing of opposite type synapses could also be used to produce a near-linear increase in net excitation if the excitatory synapse were closer to the soma). The effect is significant in that a synapse cluster of this type can generate a response that overtakes the sublinear response of a cluster of same-sign synapses, depending upon the two clusters' relative proximity to the soma. Activating two such clusters on different branches of the dendritic tree with the same input train results in nearly linear summation between the two, giving the peaked frequency response of Figure 3c. With the threshold of the spike generator set appropriately, the neuromorph's spike output exhibits a pronounced spike-frequency tuning (Fig. 3d). The tuning characteristics can be altered by rearranging the synaptic connections. For example, by subtracting curve B from curve A2, the tuning peak can be shifted to lower frequencies. However, a wider range of tuning is available by adjustment of dendritic dynamics via R_a and R_m .

3.4 Interactions Between Multiple Sources. With an understanding of synaptic interactions, it is possible to design input connections to the dendritic tree to make a neuromorph respond to particular combinations of input patterns in time and space. The connections shown in Figure 4a deliver two different afferent spike trains to excitatory synapses on separate branches, and to inhibitory synapses on a third, common branch. The neuromorph's response to either train alone will be essentially zero because the excitatory and inhibitory effects sum nearly linearly and cancel. However, if spikes from the two trains arrive synchronously, or nearly so, the inhibitory responses summate sublinearly, with the result that excitation prevails over inhibition, and the neuromorph responds by firing. Figure 4b shows examples of pairs of random input trains of equal mean frequency. Over the first half of the record, the trains were uncorrelated; over the second half they became correlated and the neuromorph fired at an average rate of 42 spikes/sec. The membrane time constant, 20 msec in this case, determined the latency of build up in firing rate.

The contour plot of figure 5 shows another example of interactions between two sources. The synaptic connections of two Poisson-distributed spiking sources *A* and *B* shown in figure 5a are similar to those of Figure 4 in that some of the inhibitory synapses share dendritic sites, whereas none of the excitatories do. When one of the input sources is relatively inactive, activity of the other source generates approximately equal excitation and inhibition, and little output spike firing by the neuromorph. As both inputs simultaneously increase in frequency, the inhibitory

synapses that they share in common sum their effects sublinearly, failing to oppose the total excitation generated and increasing the rate of output firing. With the synaptic arrangements shown and the dendritic membrane constant of 7 msec, output firing reached a maximum when the mean frequencies of both inputs were about 170 spikes per second. As input frequencies increased still further, output firing declined. This effect was caused by excitation levelling off before inhibition. Because the excitatory synapses were located proximally (2nd compartment) on the dendritic tree, although on different branches, they were electrotonically closer to each other than were the inhibitory synapses. Consequently, the excitatory responses summed less than the inhibitory responses at high input frequencies.

The input spike frequency response shown in Fig. 5b can be shaped by an appropriate selection of synaptic connections, dendritic tree dynamics, and output spike-firing threshold voltage. As an example, if the excitatory and inhibitory synapses shown in Fig. 5a are switched, the highest response is evoked by A or B firing alone.

3.5 Pulse pattern discrimination. Similar principles of combining linear and sublinear summation can be applied to the discrimination of purely temporal patterns of input spikes. An illustration of how synaptic connections can be arranged to discriminate pairs of 3-pulse patterns is shown in figure 6. The problem posed was the discrimination of temporal patterns of pulses such as | | | versus | | |, where the number of pulses and their overall duration are equal. The neuromorph was required to fire spikes following only one of the two input trains, that designated the positive stimulus. In the solution shown in Figure 6a, the input pulses were delivered directly to an excitatory synapse on one branch, and a delayed version to an excitatory synapse at the corresponding site on a separate branch. The requisite delay was generated by the delayed synaptic connection facility provided by the "virtual wire" system.

Figure 6b shows results of experiments in which the interval between the first and third spikes was fixed at 15 msec and the middle spike was stepped by 2 msec increments between them. The neuromorph output spikes accumulated over 10 trials are shown on the same time scale as the input pulses. In the case shown in Figure 6b, spikes were generated only when the second input pulse occurred at or before the midpoint of the first and third pulses.

With the positive stimulus configuration (i.e., short intervals between first and second pulses), the delayed version of the second pulse coincided with the third input pulse on different branches, leading to near-linear summation. Coincidences that occurred between pulses on the same branches were subject to sublinear summation and were therefore less effective in firing the neuromorph. However, the coincidence between the delayed first pulse and the undelayed second pulse also takes place on different branches but occurred before the membrane potential had risen sufficiently to reach spike firing threshold. The spike generator threshold had to be adjusted appropriately to give good discrimination. Changing the delay to 5 msec reversed the discrimination, making | | | the positive stimulus.

3.6 Post-spiking Refractoriness. A possible role for the refractory period in a neuromorph's response to synchrony was investigated in the manner of Bernander et al. (1994) by delivering trains of a fixed number of spikes over a variable time period, T . In the present experiments, each spike of the input train was delivered to an excitatory synapse on a different dendritic branch of an 8-branch neuromorph, all synapses being located at the same distance from the soma (Fig. 7a). The purpose of separating successively activated synapses on different branches was to assess the effects of refractoriness in the absence of sublinear summation. Figure 7b shows that this was successful. With eight synaptic activations on separate branches, the number of spikes generated by the neuromorph was maximal at the lowest train duration, T ($f = 800$ spikes/sec), when there was maximal temporal summation of EPSPs. Spike output fell off monotonically as the train duration increased (i.e., as train frequency decreased), due to diminishing temporal summation. This is to be contrasted with activations on the same branch, when sublinear summation substantially reduced the response to the most synchronized trains, producing a tuned frequency response.

We used two different methods of making the neuromorph refractory following spike firing. In the first, neuromorph output spikes triggered a one-shot that held the capacitor, C , of the spike generator at ground potential for a time that determined the absolute refractory period. Because the spike generator was isolated from the dendritic tree by a source follower, this procedure had no effect upon the integrating function of the dendrites. Figure 7c shows that introducing a refractory period by this means limited the maximum number of output spikes generated, without achieving any spike-frequency tuning. The second method shunted the postsynaptic potentials within the dendrite by routing output spikes to proximal shunting inhibitory synapses on all eight branches. This had the effect of clamping the membrane potential to V_{rest} for a programmable time. Figure 7d shows that the maximum number of output spikes was again reduced, but the most synchronized trains (800 spikes/sec) were somewhat less effective than trains spread out over longer durations. This modest spike-frequency tuning effect may be contrasted with that produced by sublinear summation (Fig 7b).

4 Discussion

4.1 Dendrites, real and artificial. Dendrites of real neurons are much more complex than the electrically passive structure of the classical model. The present VLSI modelling of dendrites is further constrained by what can reasonably be done using integrated resistors, capacitors, and transistors. Despite the simplifications and expediences of their current design, our artificial dendrites have certain features, in addition to those responsible for the classical cable properties, that help them mimic physiological processes. The use of switched-capacitors allows a very compact VLSI implementation of resistors with extremely high values, so that in emulating the axial and membrane resistances of dendrites, it is possible to use membrane capacitances of small value, and consequently small area. An added bonus, is the capability of changing the axial and membrane resistances very easily with the clock frequencies

controlling the switched-capacitors. The ability to make rapid adjustments to the membrane resistance, and therefore time constant, provides a way of simulating the effects of background synaptic activity or modulatory influences that could affect excitability and temporal processing (Bernander et al. 1991; Agmon-Snir and Segev, 1993). A similar capability to adjust axial resistances, while not actively used by neurons as far as we know, provides a means of readily setting the electrotonic length of the dendritic branches.

While the phenomenon of sublinear synaptic summation has been fully described (e.g., Rall, 1964; Shepherd and Koch, 1990; Koch and Poggio, 1992) its role in dendritic information processing has not been demonstrated. Ferster and Jagadeesh (1992) have argued convincingly that sublinear summation occurs in pyramidal cells of visual cortex. Their measurements suggested that the dendritic membrane can be almost completely depolarized by normal visual stimulation, with a consequent reduction of individual EPSP amplitude. Simulations with anatomically characterized neurons also suggest that multiple excitatory activations result in a reduction of EPSP amplitude (Bush and Sejnowski, 1993; Mel 1993). Sublinear effects, including saturation, are likely to be important at points of high input resistance such as thin dendrites and dendritic spines, and for inhibitory synapses whose reversal potential generally lies close to the average membrane potential. The connections responsible for coincidence detection (Fig. 5), for example, depended upon sublinear summation of inhibitory inputs.

Our fixed-weight synapses, after one or two closely spaced activations, produce a membrane polarization somewhat equivalent to that in a real dendritic branch after an intense afferent bombardment of many weak synapses. The intensity of the bombardment required to produce this state of polarization, will depend upon the local input resistance. Given that such states obtain in real neurons, the rules governing linear and sublinear summation listed in the next section, would apply with the kind of consequences our results have already demonstrated.

4.2 Linearity and non-linearity of synaptic connection patterns. Here we summarize the possible interactions involving synchronous depolarizing and hyperpolarizing synaptic inputs on the ADT. Synchronous synaptic activation is taken to mean that the activations take place within approximately one membrane time constant.

(A) Near-linear summation is seen when i) synchronous synaptic activations occur at electrotonically distant sites, or when ii) synchronous depolarizing and hyperpolarizing activations occur at electrotonically neighboring sites.

(B) Sublinear summation occurs when a set of neighboring synapses of the same type, i.e. a cluster, is activated synchronously. This includes the case of a single synapse being activated at intervals less than about one membrane time constant.

These rules of thumb are useful in an intuitive approach to designing synaptic connectivities for various spike processing applications. Alternatively, systematic search procedures could be used (Elias, 1992; Northmore and Elias,

1994).

4.3 *Temporal processing of spike trains.* Frequency-selective neural circuits play an important part in neural information processing, particularly in sensory systems where the temporal structure of environmental signals must be analyzed. For example, neurons in the auditory system of frogs fire preferentially to the frequency of amplitude modulation present in the mating calls of conspecifics (Rose, 1986). The mechanism whereby these cells respond to specific spike train frequencies could involve the dynamics of certain membrane conductances (Zucker, 1989; McCormick, 1990; Midtgaard, 1994), although the present work suggests that the appropriate pattern of excitatory and inhibitory synapses on passive dendrites is sufficient.

The basic processes underlying temporal discriminations are best understood from the two-pulse experiments. At synaptic activation intervals much longer than the membrane time constant a failure of temporal summation to exceed the spike firing threshold leads to "high-pass" characteristics with spike trains. As activation intervals shorten, the growth of the soma potential is limited by the saturation of the membrane potential in the neighborhood of the activated compartments, with a corresponding sublinearity in spike firing output. Given an appropriate setting of spike firing threshold, these two processes generate selectivity for input spike intervals; given control over axonal delays, discrimination of more complex temporal patterns of input spikes is possible. "Band-pass" characteristics of the kind required of modulation detectors in frogs depend upon the interplay of near-linear and sublinear summation across different regions of the dendritic tree. Opportunities for tuning the peaks include changing dendrite dynamics (via R_m and R_d) (Elias and Northmore, 1995) and adjusting synaptic connectivity.

An alternate mechanism for generating spike-frequency selectivity, refractoriness following spike firing, can also limit a neuron's responsiveness to high input frequencies if synaptic activations are wasted during the refractory period. Using a single compartment model, Bernander et al. (1994) showed that refractoriness resulted in fewer output spikes being generated in response to synchronized inputs than to relatively desynchronized inputs. Our neuromorph exhibited no such effect when its "soma" was made refractory (Fig. 7c) because it is isolated from the artificial dendritic tree where most of the integration of postsynaptic potentials takes place. Refractoriness in this case then merely reduces the gain of the process that translates the potential at the trunk of the dendritic tree into spikes, and does it regardless of input activation frequency. When, on the other hand, refractoriness was produced by shunting the dendritic tree after each output spike, a modest fall off in response to high frequencies (i.e. synchronized inputs) was observed (Fig. 7d). Similar considerations are likely to apply to real neurons in that refractoriness of the soma will be much less discriminatory against activations occurring distally on dendrites rather than proximally. If, however, somatic spikes propagate into the dendrite tree and activate conductances that shunt synaptic potentials there (Stuart and Sakman, 1994), some wastage of highly synchronous synaptic inputs will occur. Sublinear synaptic

summation, on the other hand, operates locally, and therefore with greater specificity as to input source.

Neurons possess a variety of membrane channels and intracellular mechanisms that confer highly non-linear properties (for reviews see Mel, 1994; Midtgaard, 1994), some of which facilitate synaptic transmission after repetitive synaptic activation, and some like sublinear summation, depress it. From modelling facilitatory synaptic processes (e.g., Mel, 1993; Buonomano and Merzenich, 1995) it has been shown that neurons could perform temporal discriminations and correlations of the kind explored in this work. That artificial dendritic trees, composed of simple, linear components, can also perform such discriminations, given the appropriate pattern of input connections, seems to suggest some fundamental properties of dendritic structure to which the complexities of active membrane channels may be considered functional add-ons.

Acknowledgments

Research supported by National Science Foundation Grant (# BCS-9315879), and by the University of Delaware Research Foundation.

References

- Agmon-Snir, H., and Segev, I. 1993. Signal delay and input synchronization in passive dendritic structures. *J. Neurophysiol.* 70, 2006-2085.
- Buonomano, D.V., and Merzenich, M.M. 1995. Temporal information transformed into a spatial code by a neural network with realistic properties. *Science*, 267, 1028-1030.
- Bush, P.C., and Sejnowski, T.J. 1993. Simulations of synaptic integration in neocortical pyramidal cells. In *Computation and Neural Systems*. F.H. Eeckman and J.M. Bower, eds., pp. 97-101. Kluwer Academic Publishers, Boston.
- Bernander, O., Douglas, R., Martin, K., and Koch, C. 1991. Synaptic background activity influences spatiotemporal integration in single pyramidal cells. *Proc. Natl. Acad. Sci. USA* 88, 11569-11573.
- Bernander, O., Koch, C. and Usher, M. 1994. The effect of synchronized inputs at the single neuron level. *Neural Computation* 6, 622-641.
- Douglas, R., Mahowald, M., and Mead, C. 1995. Neuromorphic analogue VLSI. *Annu. Rev. Neurosci.* 18, 255-281.
- Eccles, J.C. 1957. *The Physiology of Nerve Cells*. Johns Hopkins Univ. Press, Baltimore, Md.
- Elias, J. G., 1992, "Genetic Generation of Connection Patterns for a Dynamic Artificial Neural Network," in *Combinations of Genetic Algorithms and Neural Networks*, eds. L. D. Whitley and J. D. Schaffer, IEEE Computer Society Press, Los Alamitos, CA, pp. 38-54.
- Elias, J.G. 1993. Artificial dendritic trees. *Neural Computation*, 5, 648-664.
- Elias, J.G., and Northmore, D.P.M. 1995. Switched-capacitor neuromorphs with wide-range variable dynamics. *IEEE Trans. Neural Networks*, vol. 6, No. 6, pp 1542-1548.
- Elias, J.G., and Northmore, D.P.M. 1996. Oscillatory networks of silicon neurons, in preparation.
- Ferster, D., and Jagadeesh, B. 1992. EPSP-IPSP interactions in cat visual cortex studied with in vivo whole-cell patch recording. *J. Neuroscience* 12, 1262-1274.

- Jack, J.J.B., Noble, D., and Tsien, R.W. 1975. Electric current flow in excitable cells. Clarendon Press, Oxford.
- Koch, C., and Poggio, T. 1992. Multiplying with synapses and neurons. In *Single Neuron Computation*. T. McKenna, J. Davis and S.F. Zornetzer, eds., pp 315-345. Academic Press, San Diego, California.
- Koch, C., and Segev, I. 1989. *Methods in Neuronal Modeling*. MIT Press, Cambridge, Massachusetts.
- Mead, C. 1989. *Analog VLSI and Neural Systems*. Addison-Wesley, Reading, Massachusetts.
- Mel, B.W. 1993. Synaptic integration in an excitable dendritic tree. *J. Neurophysiol.* 70, 1086-1101.
- Mel, B.W. 1994. Information processing in dendritic trees. *Neural Comp.* 6, 1031-1085.
- McCormick, D. A. 1990. Membrane properties and neurotransmitter actions. In *Synaptic Organization of the Brain*, G. M. Shepherd, ed. pp 32-66, Oxford Univ. Press, New York.
- Midtgaard, J. 1994. Processing of information from different sources: spatial synaptic integration in the dendrites of vertebrate CNS neurons. *Trends in Neuroscience* 17, 166-173.
- Murthy, V.N., and Fetz, E.E. 1994. Effects of input synchrony on the firing rate of a three-conductance cortical neuron model. *Neural Comp.* 6, 1111-1126.
- Northmore, D. P. M., and Elias, J. G. 1994, "Evolving synaptic connections for a silicon neuromorph," *Proceedings of the IEEE Conference on Evolutionary Computation, Orlando*, vol. 2, pp. 753-758.
- Rall, W. 1964. Theoretical significance of dendritic trees for neuronal input-output relations. In *Neural Theory and Modelling*. R.F. Reiss, ed., pp 73-97. Stanford Univ. Press, Stanford, California.
- Rall, W., Burke, R.E., Holmes, W.R., Jack, J.J.B., Redman, S.J., and Segev, I. 1992. Matching dendritic neuron models to experimental data. *Physiol. Revs.* 4 (Suppl): S159-S186.
- Rose, G. 1986. A temporal-processing mechanism for all species? *Brain Behav. Evol.* 28,134-144.
- Shepherd, G.M. and Koch, C. 1990. Dendritic electrotonus and synaptic integration. In *Synaptic Organization of the Brain*, G. M. Shepherd, ed. pp 439-473, Oxford Univ. Press, New York.
- Singer, W. and Gray, C. M. 1995. Visual feature integration and the temporal correlation hypothesis. *Annu. Rev. Neurosci.* 18, 555-586.
- Stuart, G.J., and Sakmann, B. 1994. Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature*, 367, 69-72.
- Zucker, R. S. 1989. Short-term synaptic plasticity. *Annu. Rev. Neurosci.* 12, 13-31.




Figure 1 (a) A 5-compartment segment of artificial dendritic tree (ADT). R_a , axial resistance; R_m , membrane resistance; C_m , membrane capacitance. Upper row of synapse transistors are excitatory, lower row inhibitory. Alternate inhibitory transistors shunt membrane voltage to V_{rest} for a programmable time period. (b) A dendritic tree with spike generating "soma" of the leaky integrate-and-fire type. V_s is "membrane voltage" at the soma; R and C form the integrator; V_{th} is spike firing threshold.

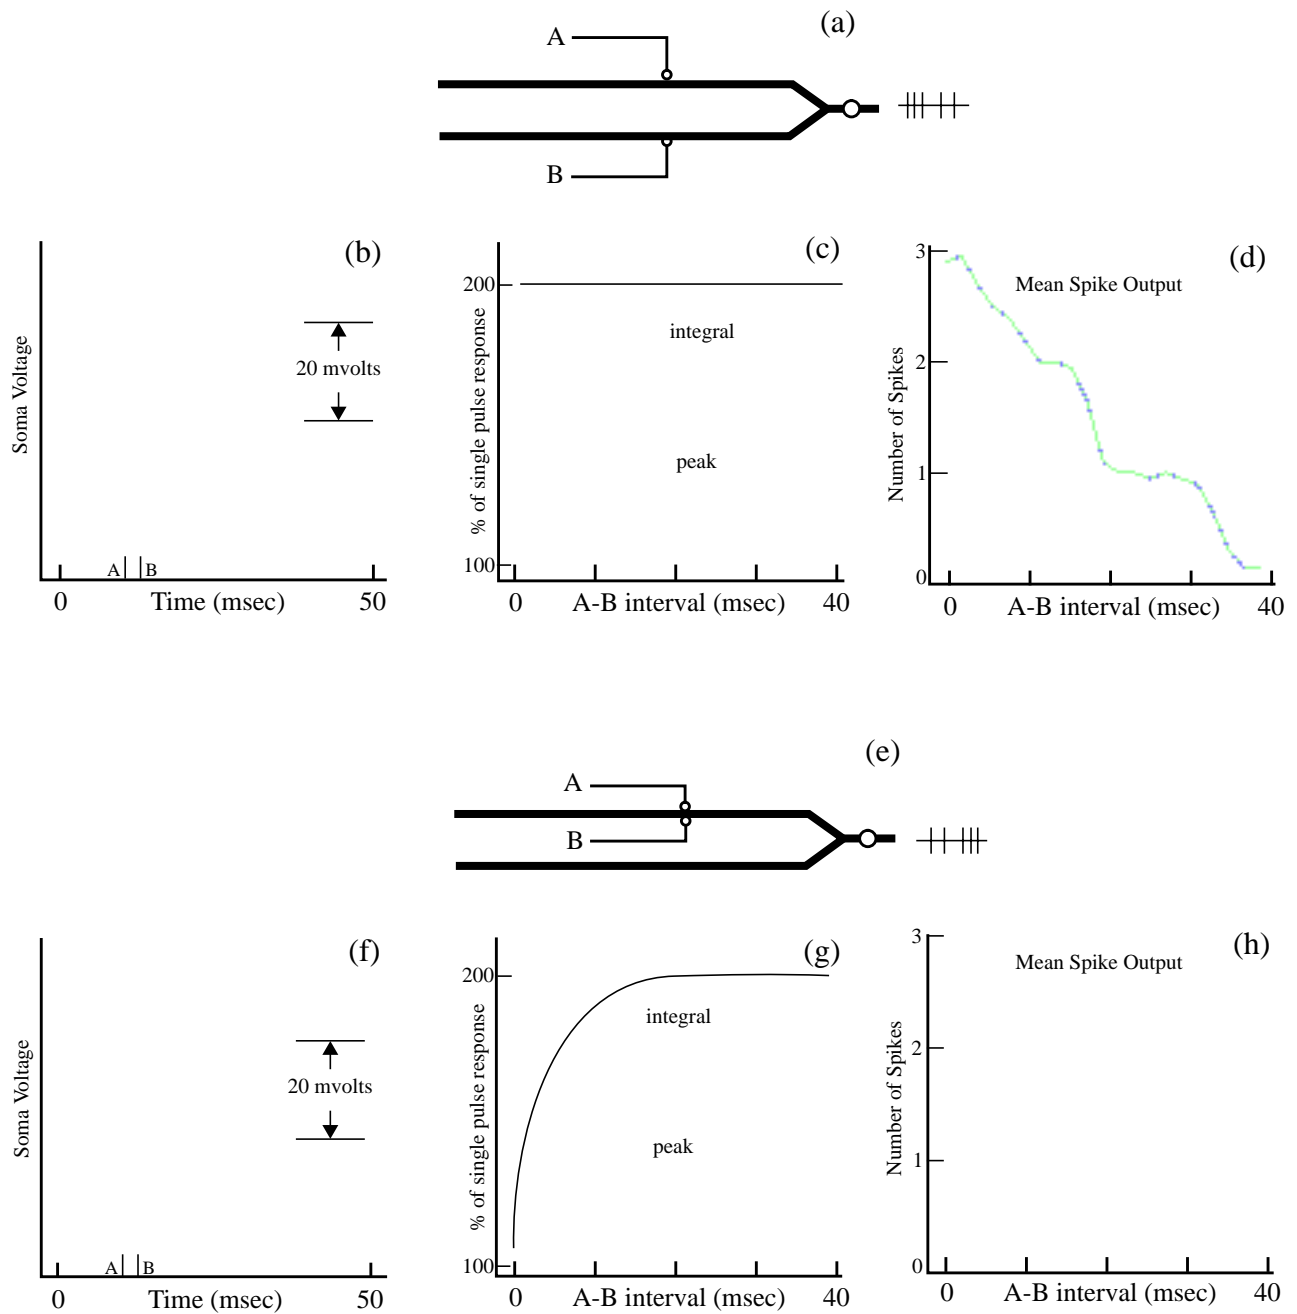


Figure 2 Measurements on a 2-branched ADT neuromorph demonstrating near-linear and sublinear summation. Input pulses, A & B, activated an excitatory synapse (#6) located 38% along the branch from soma on different branches, (Fig.2 a-d), or on the same branch (Fig. 2 e-h) of the ADT. Fig. 2 b, f: Impulse responses measured at the soma evoked by pulses A & B, individually (smaller curves), and together (larger curve). In this example, B followed A by 2 msec. Fig. 2 c, g: Peak impulse response and integral as a function of A-B time interval. 100% denotes response peak or integral to a single input pulse. Fig. 2 d, h: Mean number of spikes generated by the soma as a function of A-B interval.

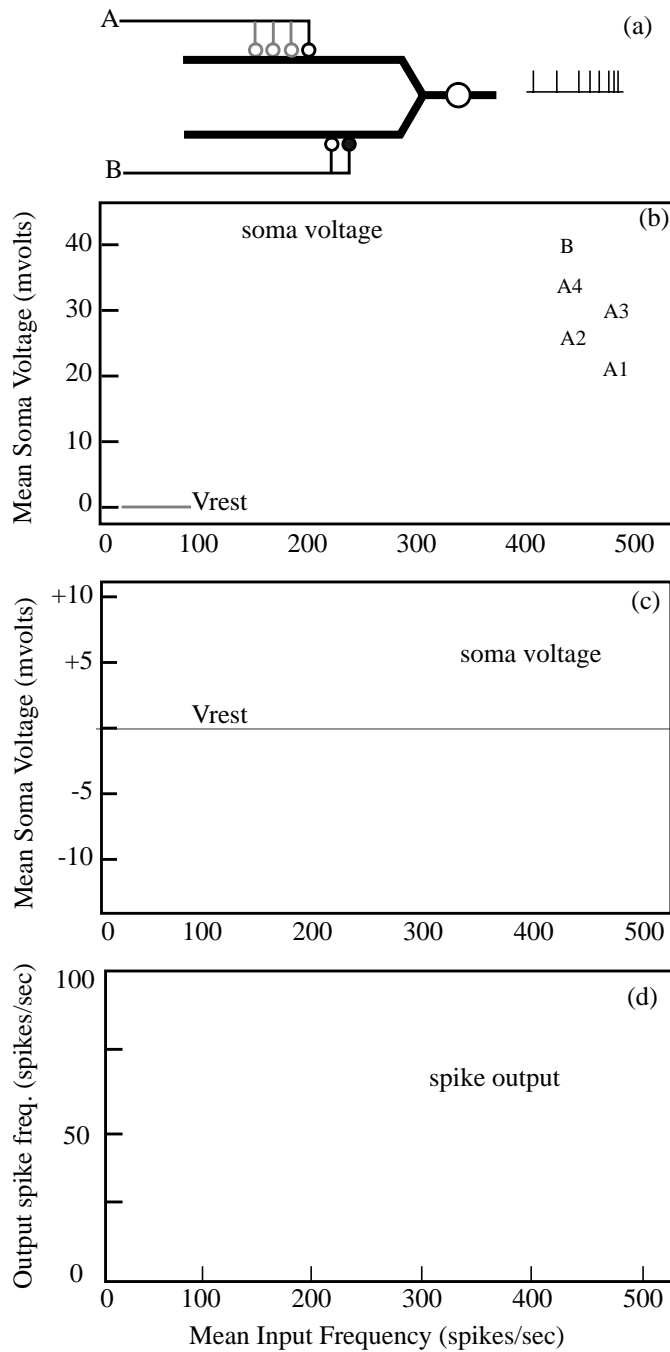


Figure 3 Measurements on a 2-branched ADT neuromorph showing impulse and spike responses evoked by Poisson-distributed spike trains of different frequencies. (a): Synaptic connections of input trains A and B. Open circles are excitatory synapses, filled inhibitory. (b): Mean soma voltage vs. mean frequency of input trains. Curves A1-A4: simultaneous activation of 1, 2, 3, or 4 excitatory synapses (synapse #s 14-20) on one branch. Curve B: simultaneous activation of one inhibitory and one neighboring excitatory synapse (#s 10, 9). Since the inhibitory synapse is the stronger, curve B is actually negative-going, but plotted positive-going for comparison. (c) and (d) show responses to simultaneous activation of all synapses shown in (a) by a single input train; (c): mean soma voltage relative to resting potential; (d): Mean frequency of soma spike output.




Figure 4 (a). Connections for coincidence detection. Two random Poisson trains (A & B) activate the excitatory (open circles) and inhibitory synapses (filled circles) on three dendritic branches as shown. (b) Four samples of trains A, B. For the first 60 msec (left of vertical dotted line), the two trains were uncorrelated; thereafter, train B was generated by randomly jittering the time of occurrence of each spike in train A by up to + 1.5 msec. (c) Histogram showing time of occurrence of resulting neuromorph output spikes accumulated over 100 trials. On the right side of (c) where the inputs are correlated, the average output firing rate is about 2.5 spikes in 60 msec, or about 42 spikes/sec.

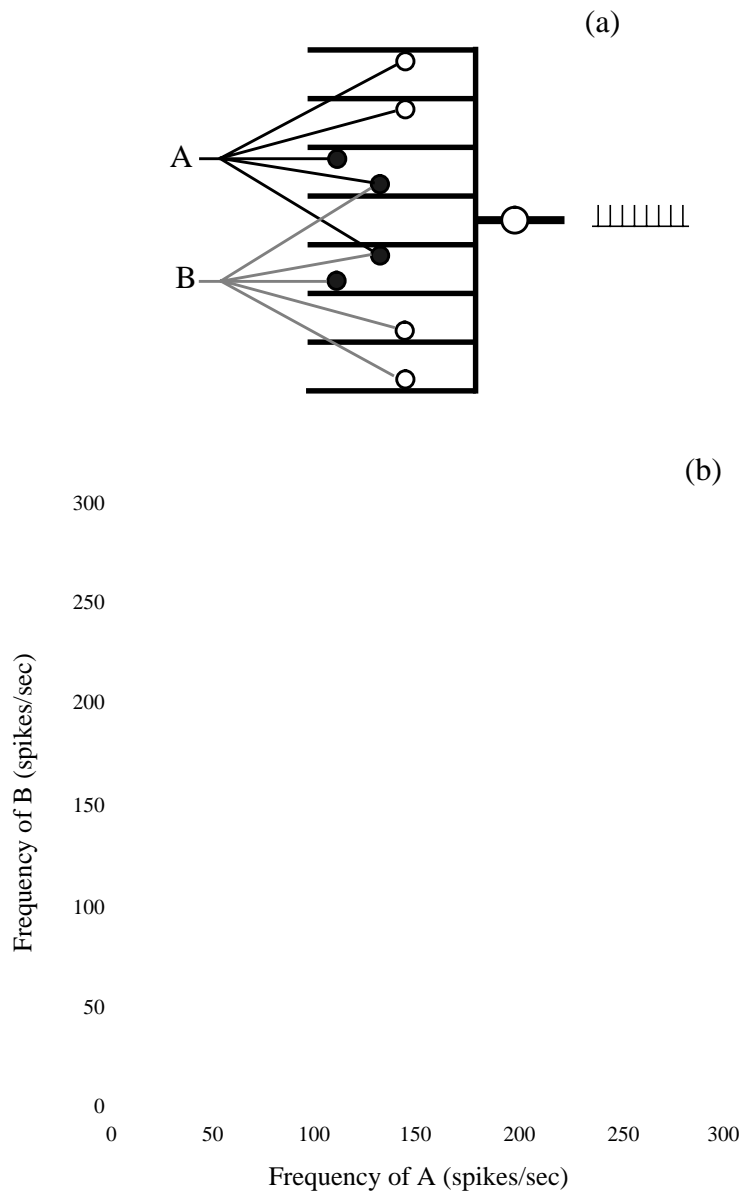


Figure 5 Neuromorphic response to combinations of two input train frequencies. (a) Synaptic connections of trains A and B. Open circles, excitatory synapses; filled circles, inhibitory synapses. (b) Contour plot showing neuromorph mean output spike firing frequency as a function of the mean frequency of the Poisson distributed trains, A and B. Contours are numbered in spikes/second.

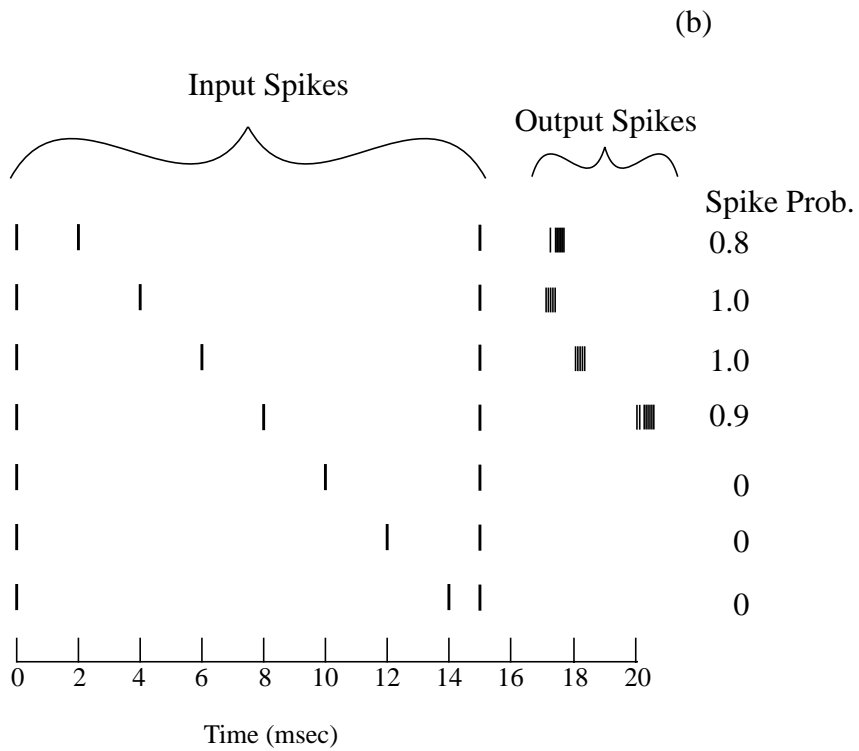



Figure 6 Discrimination of spike patterns: || | versus | | |. (a) The 3-spike input train was connected to a two-branch neuromorph as shown. Activation on one branch was delayed by 10 msec. (b) The timing of the second input spike was varied. Neuromorph output spikes are shown superimposed from 10 repetitions of each input pattern. Neuromorph fired a single spike with probabilities shown when the 2nd input spike occurred in the first half of the 15 msec period.

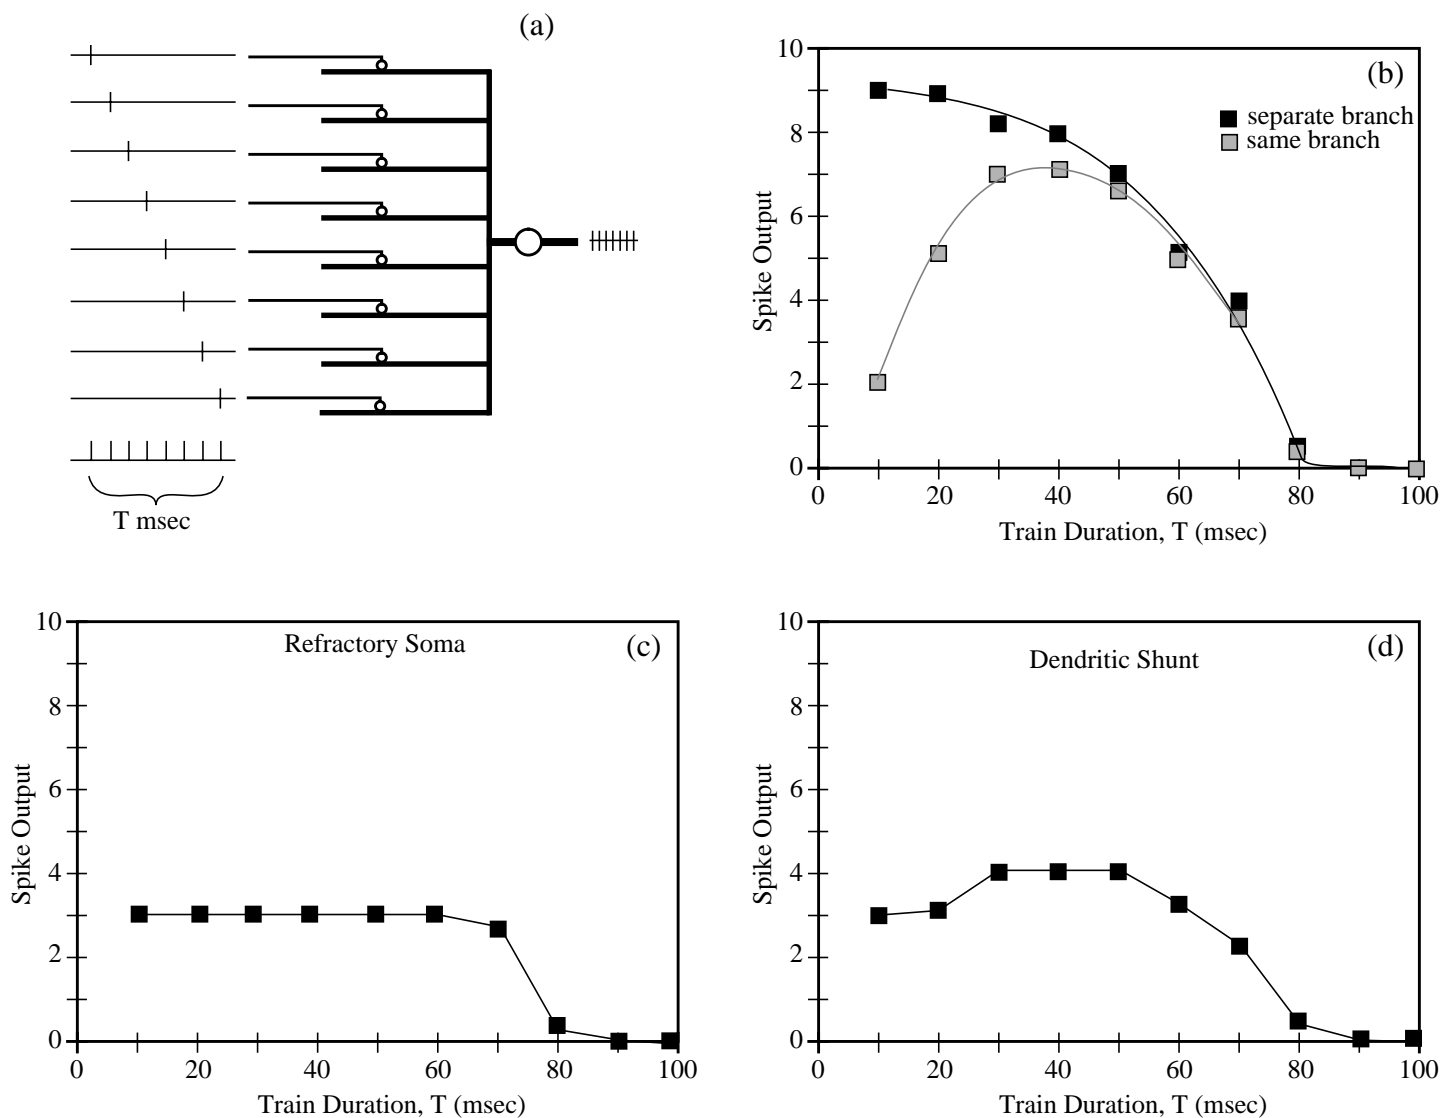


Figure 7 Comparison of sublinear synaptic summation with refractoriness. (a) To obviate the sublinear effect, each spike of an 8 input spike train was delivered to a different branch on an 8 branch ADT. To show sublinear summation effects, all 8 input spikes activated the same synapse on one branch (not illustrated). The 8 input spikes were distributed over a variable time period, T . (b) Mean number of spikes output by neuromorph as a function of T , for synaptic activation on separate branches and on same branch. (c) Mean spike output as a function of T when soma was made refractory for 2 msecs. (d) Mean spike output as a function of T when output spikes activated proximal shunting synapses for 2 msecs on all branches.