IEEE World Congress on Computational Intelligence, 1994

Evolving Synaptic Connections for a Silicon Neuromorph

David P. M. Northmore and John G. Elias

Neuroscience Program &
Electrical Engineering
University of Delaware
Newark, DE. 19716

ABSTRACT

Our VLSI neuromorphs, possess extensive dendritic trees with hundreds of excitatory and inhibitory synaptic sites. Useful signal processing can be achieved by evolving the appropriate connections to the synapses using genetic algorithms. In order to reduce the very large solution space, schemes for specifying connections have been borrowed from neural development. Results show the evolution of connections to a single neuromorph for the discrimination of temporal and spatio-temporal patterns.

INTRODUCTION

Current techniques of VLSI fabrication put within reach the goal of making artificial nervous systems for “creatures” that are capable of behaving adaptively in complex environments. Although the elementary building blocks in our system, the neuromorphs, are electronic, they incorporate some of the essential functional properties of biological neurons. Inevitably, compromises had to be made in choosing their properties, but our neuromorphs have the advantages of simplicity of mass construction, compactness, low power consumption, and the facility to make a very large number of different connection patterns. It is this enormous range of connection possibilities that this work seeks to address by borrowing some principles of neural development in order to restrict the solution space, while achieving useful signal processing.

Each neuromorph is a VLSI circuit comprising a spatially extensive artificial dendritic tree [1] and a spike generating soma. Figure 1a is a simplified circuit diagram of a short, five-compartment section of silicon dendrite. Each compartment has a capacitor, C_m, representing a membrane capacitance, two programmable resistors, R_m and R_a, representing a membrane resistance and a cytoplasmic resistance, and several MOS field effect transistors that simulate synapses by enabling transient inward or outward transmembrane current. The resulting potential appearing at the soma, point S in Figure 1b, determines the rate of output spike firing. P-channel transistors (upper) produce excitatory effects on spike firing by increasing the membrane potential. Inhibition is mediated by two interleaved populations of n-channel transistors (lower). Half have their source terminals connected to ground and exert inhibitory effects by lowering membrane potential; the other half have their source terminals connected to a programmable voltage and exert inhibitory effects by pulling the membrane potential towards this voltage. When this voltage is set near the membrane resting voltage these transistors behave like shunting or silent inhibitory synapses [2].

The synapse transistors are turned on by an impulse signal applied to their gate terminals. Although synapse transistors are on for only 50 nsec, the resulting impulse response measured at the input to the soma (point S in Figure 1b) may last thousands of milliseconds. The peak amplitude of the response is largest for synaptic activation nearest the soma and diminishes rapidly for sites farther away, while the latency to peak increases with distance from the soma. Thus, as in passive dendrites of biological neurons [3], location on the dendrite determines the weighting of a synapse in its effects on amplitude and timing of the voltage response at the soma (see Figure 2).

The spike-generating soma is a resettable RC integrator. When the voltage across the capacitor, C, exceeds a programmable threshold voltage, V_{th}, the comparator generates a spike that discharges C and is captured in a register which is sampled by routing circuitry (“virtual wires,” [1]).

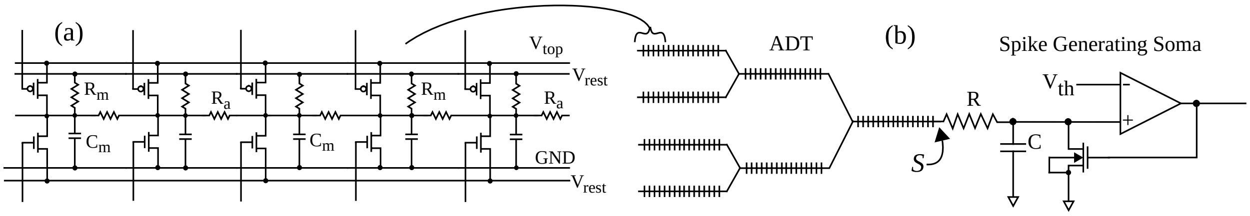

FIGURE 1. a) Five compartment segment of artificial dendrite. b) Diagram of VLSI neuromorph comprising a spatially extensive artificial dendritic tree (ADT) and a spike generating soma, The crosses on the ADT represent synapse locations. In most of our chips, dendritic trees are made of a number of branches each having 32 synapses.

We now address the problem of making input connections to a single neuromorph of this design to perform elementary temporal and spatio-temporal discriminations. Because the neuromorph has a fixed number of synaptic sites, each with fixed properties, either excitatory or inhibitory, the problem becomes one of where to connect input signals to the dendritic tree [4]. Figure 2 shows impulse responses generated at the soma, illustrating the non-linear effect of shifting an excitatory synapse along a dendritic branch, and the linear effect of increasing the number of activated synapses at a given level on the dendritic tree. The activation of several synapses at different places and times will gives rise to complex waveforms representing the summation of many such impulse responses. Neural development suggests two approaches to “guiding” afferents to the appropriate synaptic sites. These, the Laminar and Topographic specifications, are encoded in chromosomes to be operated upon by a genetic algorithm (GA).

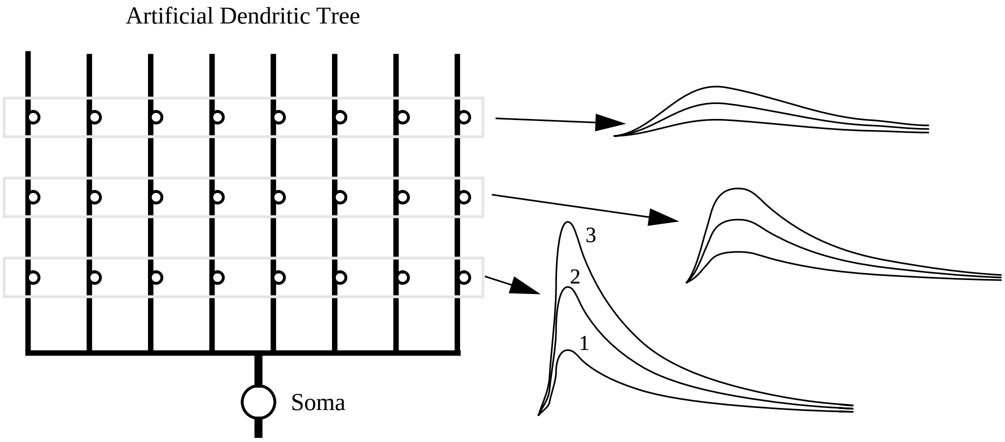

FIGURE 2. Impulse responses at the soma (voltage waveforms at right) due to activation of 1,2 or 3 excitatory synapses at three laminar levels on an eight-branch neuromorph.

METHODS

Synapse Specification - Laminar

Figure 3a shows a neuromorph with 8 primary dendritic branches contacted by 3 afferents. Each afferent is allowed to make synapses of only one functional type, excitatory, hyperpolarizing-inhibitory, or shunting-inhibitory, and these synapses are confined to one level or lamina of the dendritic tree. The effectiveness of an afferent is therefore proportional to the number of synapses it makes within a lamina. In one, straightforward approach that we are using, each afferent is encoded by (a) lamina, (b) number of synapses, and (c) synapse type in a chromosome by 4 + 3 + 2 bits. The arrangement is suited for a temporal discrimination application that involves no spatial information processing. In this case, a single input signal is distributed by several afferents, the minimum number required for adequate performance being determined by experiment.

Synapse Specification - Topographic

This scheme generates connections for topographically ordered arrays of afferents. Figure 3b shows two afferent arrays (e.g from sensor or other topographically organized structures) contacting an 8-branch neuromorph. Each afferent array activates only one type of synapse in a particular laminar pattern. Of the many ways in which the laminar pattern could be specified, the one illustrated has each member of an afferent array contact a number of neighboring synapses of a given type along one branch, starting at a specified laminar level - a columnar arrangement of connections. More varied patterns of synapses along a branch could be achieved by having the source of input (e.g. the sensors) drive several parallel afferent arrays, each of which contacts different levels along the dendritic branch. Connections are specified for an afferent array by (a) starting lamina, (b) number of synapses, and (c) synapse type (4 + 3 + 2 bits).

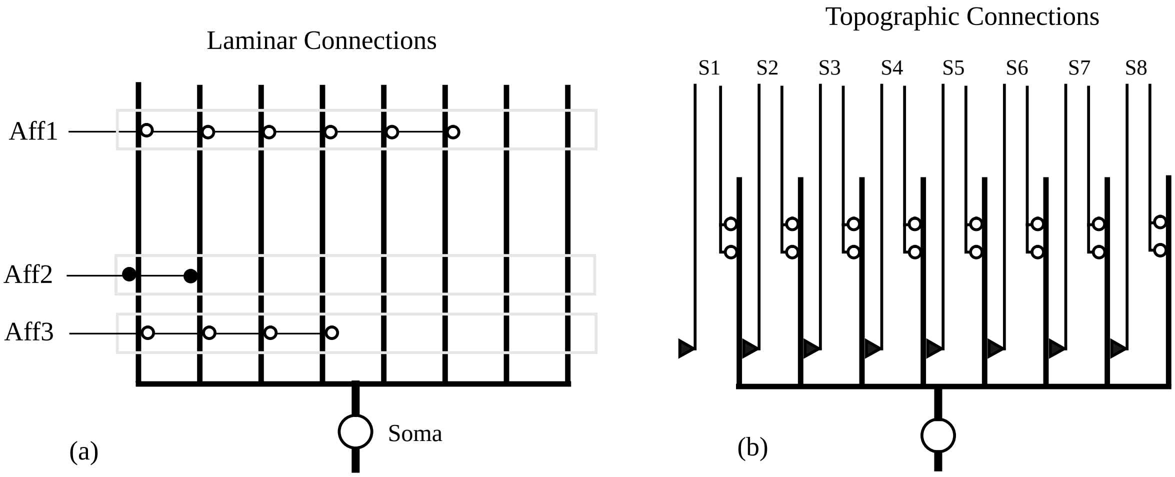

FIGURE 3. Synapse specification by (a) Laminar and (b) Topographic organization. Each neuromorph has eight equal-length dendritic branches. Open circles are excitatory synapses, filled circles are hyperpolarizing inhibitory synapses, and filled triangles are shunting inhibitory synapses.

The form of the laminar pattern specification for each afferent is likely to be critical to the task in hand. So, too, is the specification of the afferents' topographic projections, e.g. their orientation, magnification, boundaries etc. Thus in one of the problems investigated here, that of forming a connection pattern for generating directional selectivity, some spatial asymmetry between the excitatory and inhibitory connections is required. This could be achieved by introducing a lateral offset in, say, an inhibitory afferent array relative to another, excitatory afferent array, as shown in Figure 3b. In this case, pairs of afferents would be driven from the same sources (S1-S8 in Figure 3b).

Evolving with a Genetic Algorithm

Chromosomes were constructed of bit strings specifying afferent connections, by the Laminar scheme for pulse interval discrimination, or by the Topographic scheme for directional discrimination. Additional structural factors such as the number of afferents, afferent array offsets etc. were adjusted by hand. A component of the chromosome that was always adjusted genetically were 6-9 bits specifying the spike integrator threshold, V_{th}. An initial population of 200-400 chromosomes were generated randomly and decoded into, V_{th}, and the synaptic connections that were then written into the connection list memory of the virtual wire system [1]. The host computer delivered input spike patterns to the neuromorph in real time. Each chromosome was evaluated on the basis of the neuromorph's output spike train according to a fitness function. Two chromosomes were randomly selected for crossover, the offspring evaluated, and if fitter than the population mean, were used to replace less fit chromosomes in the population. Mutation rates of 0.5 - 5% were also employed.

RESULTS

Temporal Discrimination

The aim was to evolve the connections on a neuromorph to respond when a pair of input spikes occurred with a particular time interval between them. During evolution, the fitness increased with the number of output spikes evoked by the target spike interval of 2 msec, and decreased when any spikes were evoked by shorter (1 msec) or longer (3 msec) intervals. The synaptic pattern was specified by Laminar coding with five afferents simultaneously activated by the input spikes. The soma integrator time constant was fixed (RC =7.5 msec), and V_{th} was encoded by 8 bits controlling a 1 volt range above resting potential. Figure 4a shows an evolved pattern of excitatory and hyperpolarizing synapses that in conjunction with a V_{th} (0.25 volts above rest) yielded on average two output spikes at each presentation of the input pair when separated by the target interval. Figure 4b shows the time-interval tuning achieved.

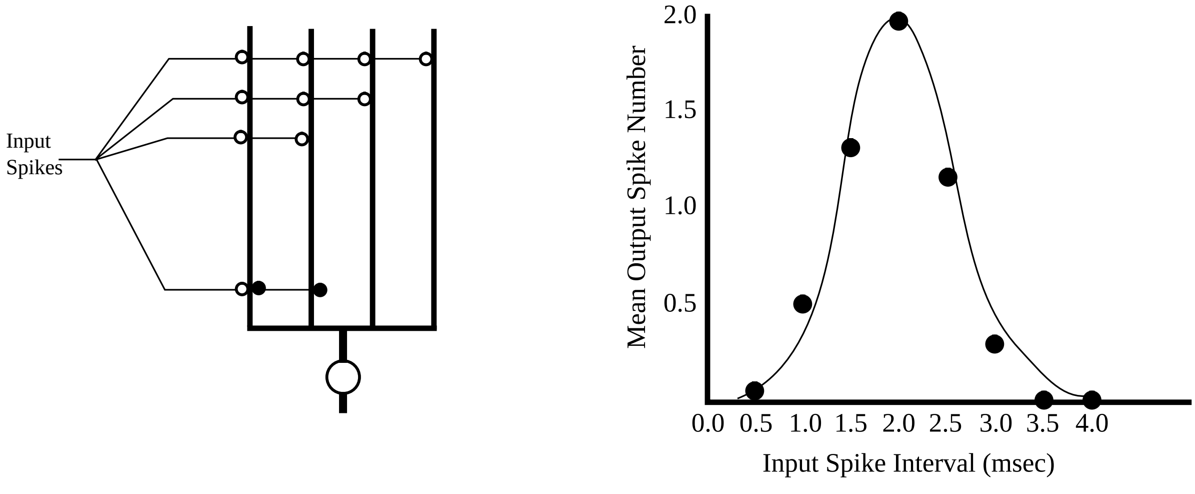

FIGURE 4. a) Evolved laminar connections for temporal interval discrimination. A 4-branch neuromorph is shown innervated by four excitatory afferents and one hyperpolarizing inhibitory afferent. b) Mean number of spikes evoked by different input spike intervals, averaged over 40 presentations. The target interval used for evolution was 2 msec, non-target intervals were 1 and 3 msec. The standard deviation associated with each point was 0.4 spikes.

Direction Discrimination

The controlling host computer simulated a row of 16 sensors, each of which fired an impulse when activated by a passing point object. The Topographic scheme was used to specify the connections between the receptor array and an 8-branch neuromorph. Connections were evolved to generate maximal spike output for one (preferred) direction of motion and no spikes for the opposite (null) direction. The types of afferent allowed were excitatory and either hyperpolarizing or shunting inhibitory. Evolving was done with a single moving point stimulus at one speed; testing was done at various speeds with single and multiple point stimuli, the latter corresponding to moving textures.

When the only form of inhibition allowed was hyperpolarizing, the GA was able eventually to find a connection pattern that yielded directional selectivity to the single moving spot. However, testing with texture stimuli led to a complete breakdown in discrimination. By contrast, connection patterns such as those illustrated in Figure 3b, in which shunting-inhibitory synapses very effectively gated excitatory input to the soma for one direction (right-to-left in Fig. 3b), yielded consistent directional selectivity for both punctate and textured stimuli (Figure 5).

DISCUSSION

The solution to the temporal interval discrimination is one of many obtained by the GA operating upon the Laminar specification scheme. None of the other solutions achieved much higher peak firing rates to the target interval than the one shown in figure 4, although the addition of more afferents increased signal-to-noise ratio, a result expected of a linear summation of individual impulse responses. This is evidently a difficult discrimination, given only hyperpolarizing and depolarizing synaptic effects. We expect much more robust discrimination when the GA is allowed (a) to select variable delays in synaptic activation, an important feature of our virtual wire system [1], and (b) to exploit the non-linearities introduced by shunting inhibition.

In directional discrimination, the clear superiority of shunting synapses over hyperpolarizing is a vindication of prior theoretical work [2,5]. While connection patterns with hyperpolarizing synapses can be found that antagonize excitation more in the null direction, they are effective only under a restricted set of conditions. When challenged with a variety of moving stimulus patterns, the fine balance is disturbed because hyperpolarizing potentials make their effects felt throughout the dendritic tree. Shunting, on the other hand, exerts more localized effects.

The synaptic specification schemes used here are only a beginning, but they are shown to be effective operands for a GA. Not only can they greatly reduce the solution space but are biologically plausible. There are numerous instances of laminar termination in the brain. In structures such as olfactory bulb, cerebellum, cerebral cortex, the various classes of afferent fiber terminate preferentially at specific levels on the dendritic trees of neurons. Common too, particularly in cortical structures, are topographic, columnar patterns of innervation, where information is distributed vertically within the thickness of the cortex [6]. Thus, basing connections on the developmental principles that form similar patterns in the brain makes sense because guiding afferents into position is just as important for neuromorphs as it is for real neurons.

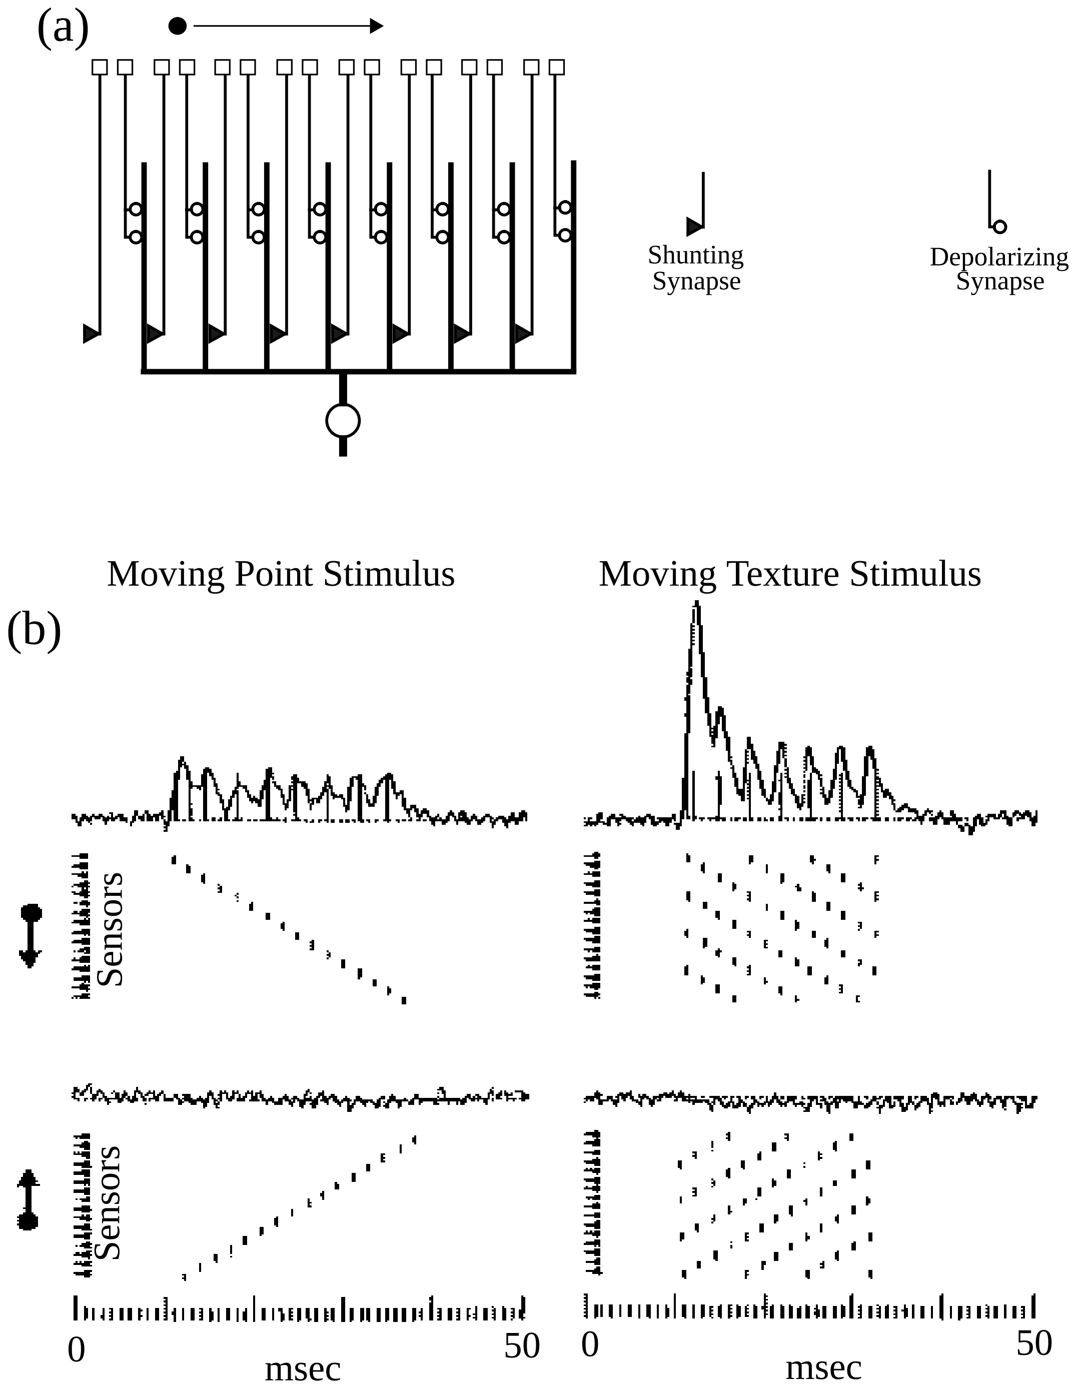

FIGURE 5. Directionally selective responses of a neuromorph driven by excitatory and shunting-inhibitory synapses in a topographically-specified pattern similar to that shown in figure 3b. Output spikes are superimposed upon the voltage waveforms at the soma. The 16 sensors generate the spatio-temporal sequence of input spikes shown below the soma waveforms. Left column: single spot movement up and down the sensor array. Right column: 4 spots moving together at the same speed.

ACKNOWLEDGMENTS

This work was supported by grants from the National Science Foundation (#BCS-9315879) and the University of Delaware Research Foundation.

REFERENCES

[1] J. G. Elias, “Artificial Dendritic Trees,” *Neural Computation*, vol. 5, pp. 648-664, 1993.

[2] V. Torre and T. Poggio, “A synaptic mechanism possibly underlying directional selectivity to motion,” *Proc. R. Soc. Lond.* B. vol 202, pp. 409-416, 1978.

[3] G. M. Shepherd and C. Koch, “Dendritic electrotonus and synaptic integration”, in *The Synaptic Organization of the Brain*, ed. G. M. Shepherd, Oxford University Press, appendix, 1990.

[4] J. G. Elias, “Genetic generation of connection patterns for a dynamic artificial neural network,” in *COGANN-92, 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, 1992

[5] C. Koch, T. Poggio, and V. Torre, “Nonlinear interactions in a dendritic tree: Localization, timing and role in information processing,” *Proc. Natl. Acad. Sci*. vol. 80 pp. 2799-2802.

[6] G.M. Shepherd. “The Synaptic Organization of the Brain: an Introduction”. Oxford University Press, New York, 1974.