

A Neuromorphic System

David P. M. Northmore¹, John Moses², and John G. Elias³

¹Department of Psychology, University of Delaware; ²Biophysics Group, Los Alamos National Laboratory; ³Department of Computer and Electrical Engineering, University of Delaware.

Correspondence to:
John G. Elias
Department of Computer and Electrical Engineering
University of Delaware
Newark, DE 19716.

Introduction

Artificial Nervous Systems

In the quest for alternative forms of computing, especially computing that generates useful behavior in the real world, one naturally looks to biology for inspiration. While brains are economical in size and energy consumption, they depend for their computing power and speed on very large numbers of processing elements richly interconnected. At the present state of technological development, the best way of emulating brains, and the behavior they generate is by constructing “neuromorphs” (Mead, 1989), mimics of neurons fabricated in silicon with VLSI, and interconnecting them with a richness that approaches that of the central nervous system. Using current technology, neuromorphic systems could be the brains of smart structures - small devices that sense the real world and behave adaptively in it, or implanted in humans to repair or extend their capabilities.

Part of the impetus for building artificial nervous systems is to gather insights into how real nervous systems work. We argue that neuromorphic systems could foster understanding in ways that are not easily achieved, if at all, by conventional digital computer simulations. A common view is that in trying to understand intelligent systems a concern for the details of implementation (i.e. the hardware) is unnecessary. We do not subscribe to this view, but think that the nature of the neural machinery is closely bound up with the solutions that have evolved to perform perceptual, cognitive and motor tasks. The kinds of computing operations these most naturally support are therefore very different from those supported by a von Neumann architecture. Neuromorphs, such as ours, patterned after biological neurons, depend upon spike processing of information, giving them powerful signal processing capabilities, indeed more powerful than the typical sigmoidal units used in artificial neural networks (Maass 1997). Rather than review the field of neuromorphic engineering generally we refer the reader to some recent review articles and collected papers (Smith & Hamilton, 1998; Maass and Bishop, 1999; Douglas & Rasche, 2002), and instead focus on our neuromorphic system and the approaches taken in our laboratories to some promising applications.

At least initially, neuromorphic systems are likely to be used where compactness and low power consumption are at a premium, for example as the brains of autonomous vehicles and for neural prosthetics that interface between the nervous system and artificial effectors or a patient’s own musculature. Therefore in thinking about what kinds of networks to investigate, an underlying consideration is that small systems must make good use of their resources. Dynamic networks are of interest because modest numbers of interconnected neuromorphs should be able to store and process large amounts of information in the transition of states they undergo. In the control of autonomous vehicles, a fast, compact system is required that can efficiently utilize information about the environment. We are therefore exploring ways in which neuromorphic networks can learn to behave adaptively through sensorimotor experience.

The Neuron and the Neuromorph

The design of our artificial dendritic tree (ADT) neuromorph is based on what one might call the classical neuron as conceptualized in the mid 1950’s (Eccles, 1957). The input structures of this neuron are the branched processes forming the dendritic tree, and to a lesser extent the cell body or soma from which the dendrites extend. The neuron receives input in the form of impulses or spikes at numerous synaptic sites over these surfaces. The output structure is a spike-firing fiber or axon that may extend for considerable distances in the nervous system.

An input spike activating a synapse generates a brief postsynaptic potential within the dendrite by opening ion channels in the cell membrane. The potential may be of either polarity depending on which ion channels are involved. If it is positive going, the effect is excitatory because it will tend to increase the spike firing frequency of the axon; if negative it is inhibitory tending to reducing spike firing. The dendritic branches are considered to be passive cables over which the postsynaptic potentials mingle and diffuse. Their function, then, is to collect many synaptic inputs, delaying, attenuating and summing the synaptic potentials generated. The net

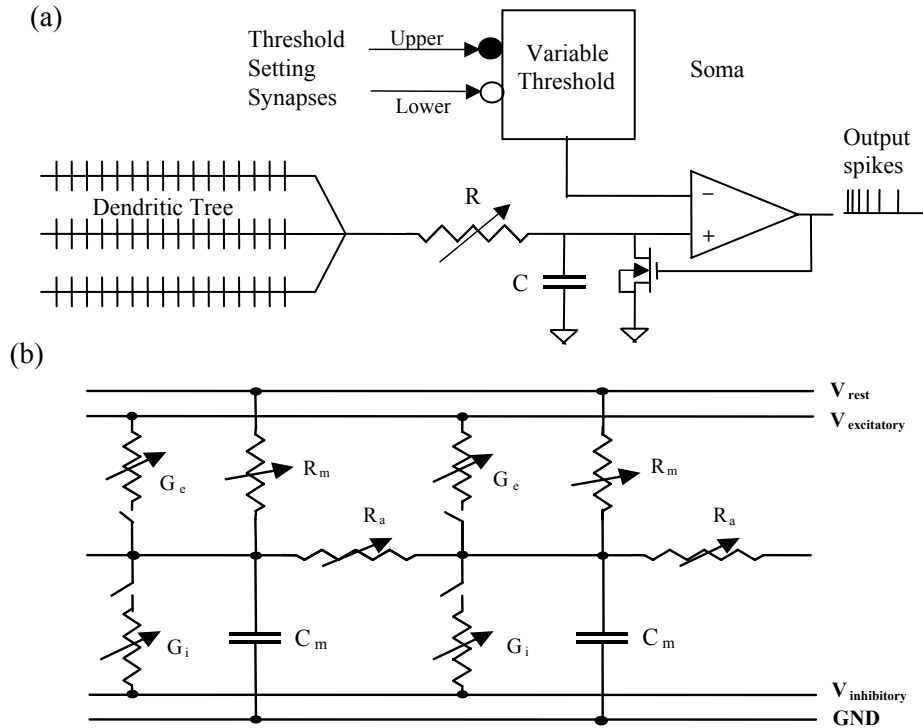


Figure 1. (a) Schematic of ADT neuromorph. Synapses are located at the cross points on the 3-branch dendritic tree and at the soma. Activating soma synapses sets the spike firing threshold for the integrate-and-fire spike generator whose integration time constant is determined by the programmable resistor, R , and fixed capacitor, C , which is discharged whenever a spike is generated. **(b)** A two-compartment section of dendrite. Each compartment contains a membrane capacitance (C_m), a membrane resistance (R_m) that connects to V_{rest} , and axial resistances (R_a) that connect to adjacent compartments. The excitatory (G_e) and inhibitory (G_i) synaptic conductances, which turn on momentarily when a synapse is activated, pull the compartment capacitor voltage towards $V_{excitatory}$ and $V_{inhibitory}$, respectively.

potential change that accrues at the junction of soma and axon determines the rate of firing of spikes emitted as output along the axon.

We now know that dendrites, which vary greatly across cell types in form, vary also in function, and are not usually passive, boosting the transmission of potentials along them with voltage-sensitive ion channels. Nevertheless, theoretical analyses show that even passive dendrites are able to perform useful spatiotemporal filtering allowing the discrimination of different input spike patterns (Rall, 1964). Experiments with our ADT neuromorphs, which are analogs of the classical neuron, showed that they could be connected so as to respond selectively to patterns of input spikes, for example to specific frequencies of an input spike train or to specific temporal orderings of spikes (Northmore & Elias, 1996). In designing the neuromorph,

we saw the role of the dendritic tree as very important to the computing power of the device, as we now know it to be in neurons (Mel, 1994). While the function of dendrites in the latter is complex, depending as it does on a multiplicity of molecular and ionic mechanisms, the simple passive dendrites that have been modeled form a starting point for exploiting the ingenuity of neuronal architecture. Fortunately dendrites in a variety of spatial configurations can be readily fabricated in VLSI. Neuromorphic modeling is, after all, very much the art of the possible,

The dendrites' filtering properties are strongly influenced by their dynamics – the resistances and capacitances that determine the time course of the postsynaptic potentials. Being able to control dynamics enables one to lengthen or shorten branches of the artificial dendrite (Elias & Northmore, 1995), providing a way of varying the functional properties of different pools of units. Another parameter of great importance to control is a unit's excitability or spike-firing threshold. It is desirable to be able to modulate the spike-firing threshold with the spiking activity of other units (Elias et al. 1997). A capability for learning by altering the efficacy of synaptic inputs is important, but at this stage of hardware development it is probably best done off-chip where different learning rules can be tested. In sum, our aim is to produce a neuron analog with sufficient flexibility that it could perform as a general-purpose unit, adaptable to many uses in a central neuromorphic system.

Hardware System

In our ADT neuromorphs fabricated by VLSI, the dendritic branches are modeled electrically by arrays of capacitors and resistors (Fig. 1b) (Elias, 1993; Elias & Northmore, 1995). The dendrites are composed of multiple compartments, each with a capacitor, C_m , representing a membrane capacitance, and two programmable resistors, R_m and R_a , representing a membrane resistance and an axial or cytoplasmic resistance. In most of our neuromorph chips (Figure 2), each dendritic branch has sixteen compartments, with 3-8 branches connected together to form a tree like that shown in figure 1a. Every compartment has an excitatory and an inhibitory synapse formed by MOS field-effect transistors that enable brief currents into or out of the compartment. Applying an impulse to the transistor gate turns on a synapse. The resultant "transmembrane" current depends upon the potential difference across the transistor, and upon the synaptic weight, which can be controlled (a) by the conductance of the transistor in the activated or on state, and (b) by the duration of the impulse supplied to the transistor's gate. Activation of excitatory and inhibitory synapses at different sites set up potentials that diffuse over the dendritic tree. The potential at the point where the branches converge (Fig. 1a) affects the rate of output spike firing. The dynamics, which determine the dendritic delays, are set by the programmable membrane and axial resistances (R_m) and (R_a), and by the fixed-value compartmental capacitance (Elias and Northmore, 1995). Whether the neuromorph fires or not and its rate of firing depends upon whether and how much the branch point potential exceeds a firing threshold implemented by a comparator in the "soma" (Fig. 1a). For this threshold potential to be influenced by the activity of other neuromorphs in a network, it is set by the ratio of spike frequency applied to two special synapses, the "upper" and "lower" in Fig. 1a (Elias et al., 1997).




Figure 2. Domain board. Dimensions: 14.5 cms square. The top half of the board is occupied by 8 chips (0.5 μ m CMOS), each containing 16 neuromorphs with 4 dendritic branches each. Most of the 2.2 x 2.2 mm chip area is filled by dendrites. The lower half of the board contains the spike routing system, connection memories and connectors for a host computer.

To build a network, one must provide a means for sending the output spikes of any neuromorph to the synapses of any other. In addition, the spikes must arrive at their destinations in a time short compared to network dynamics. To this end, we developed a multiplexing scheme that we call "Virtual Wires" (Elias, 1993). As with Mahowald's (1992) method of connecting neuron outputs to synapses, addresses are multiplexed in time over a few dedicated wires rather than through a much larger number of direct connections. In our scheme, a spike generated by a neuromorph activates destination synapses with delays that are programmable for each connection. The number of connections each neuromorph may have is limited only by the overall bandwidth (i.e., the multiplexing rate) of the system. Connections to synapses can be programmed to have one of 30 different weights, 15 inhibitory and 15 excitatory (Westerman et al., 1997). During activation of a synapse, its conductance changes from an essentially non-conducting state to one of the 30 conductance values for a duration of 50 nsec. Weights for each synapse are stored off-chip in a connection list along with the synapse address, thereby allowing multiple connections with independent weights to the same compartment.

The present system shown in Fig. 2 holds 128 neuromorphs. However, with higher density neuromorph chips it is capable of simultaneously interconnecting and running over 1000 neuromorphs. The entire network connectivity can be changed in a fraction of a second allowing rapid evaluation of a large number of different architectures. For this purpose, the spiking activity from a selected set of neuromorphs can be sampled at any time.

Neuromorphs in a Winnerless Competition Network

The coordinated firing of neurons in a network is likely to be a key feature of neural computation. Interconnected neurons with steady input behave as coupled oscillators, and such systems can display a range of coordinated behaviors. The phenomenon of synchronous firing of groups of neurons (Gerstner & Kistler, 2002) is probably the best known example of this. Other, more complex patterns of spikes can result from the interaction between neurons - the behavior depends on the details of the coupling. Such networks are strong candidates for computational mechanisms that biological systems may actually use for the representation and recognition of patterns (Natschläger et al., 2002). One that has attractive features for application to small networks of neuromorphs, is the winnerless competition network of Rabinovich et al. (2001). In the absence of a stimulus, these networks are quiescent or display unpatterned activity, but when stimulated, they generate cyclic patterns of spiking activity that are distinct and characteristic of the stimulus. These networks achieve spatiotemporal coding by executing heteroclinic orbits around saddles in a space that changes when the pattern of stimuli changes. The behavior, which captures some features observed of olfactory processing in the locust antennal lobe, was generated in a simulation with nine model neurons (FitzHugh-Nagumo) with strong asymmetric inhibitory interconnections.

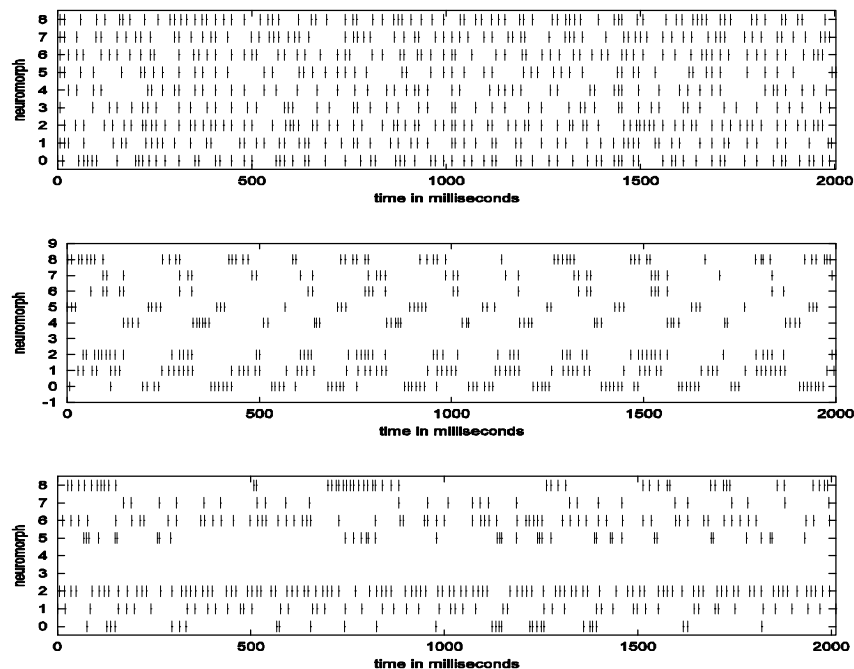


Figure 3. Spike rasters of nine neuromorphs in a winnerless competition network. **Top:** Spiking in the absence of stimulus and interconnections. **Middle:** Neuromorphs are connected via asymmetric inhibitory synapses and driven by a fixed pattern of stimuli. Periodic behavior is evident. **Bottom:** Same network, but with the stimuli rotated by one neuromorph. There is a different, weaker cyclic behavior.

We have produced similar behavior in a network of neuromorphs interconnected and stimulated similarly to the model neurons in the simulation of Rabinovich et al. The excitability of each neuromorph in the network is controlled by (a) supplying a 100 Hz spike train to the "lower" threshold synapse (Fig. 1A) and (b) by feeding back the neuromorph's output spikes to the "upper" threshold synapse. The ratio of "upper" to "lower" spikes determines the spike firing threshold; the negative feedback limits the firing rate, giving the desired background behavior (see Elias et al., 1997). The temporal patterns of spiking of the neuromorphic network is shown

in a series of raster plots in Fig. 3. The top plot illustrates quiescent behavior in the absence of any coupling or stimulus – the units are firing fairly steadily and independently of each other. When the neuromorphs are interconnected and driven with spike patterns corresponding to the inputs used by Rabinovich et al., a cyclic pattern of firing result, shown in the middle plot of Fig. 3. The neuromorphs fire in a pattern that repeats with a period of a little under 200 milliseconds. It is easily recognized by eye, and presumably could be easily recognized by a downstream neural network. Changing the stimuli by doubling or halving their spike frequencies compresses or expands the pattern, but does not change it qualitatively. Changing the stimuli by shifting the connections by one or two neuromorphs changes the pattern, as shown in the bottom plot of the Fig. 3. The network producing this pattern of spikes is interconnected as before, but the stimuli are displaced cyclically by one unit. The effect is dramatic. There is a weak appearance of periodic firing by several of the neuromorphs, but two or three seem to be firing at a roughly steady rate, and two are not firing at all. The network strongly selects between the two stimuli, and thus serves to identify the inputs, especially for the conditions of the middle plot.

The number of patterns that can be discriminated in this way is related to the number of combinations of neurons that can be formed by the network, and is far larger than the number of fixed points in a static network of the same size (Rabinovich et al., 2001). For computations with small networks - and with neuromorphs we are limited to networks that are small compared to biological networks - a dynamical network of this type is capable of much greater capacity as a memory or for pattern recognition. We may expect that networks exploiting this kind of dynamical behavior will play a significant role as applications are realized.

Sensorimotor development in a neuromorphic network

In this section we present an experimental system of simulated neuromorphs that develops sensori-motor capabilities. It was studied to find ways of using neuromorphs for controlling the behavior of an autonomous vehicle, but the same principles could be applied in a neuromorphic prosthetic where arbitrary patterns of spikes need to be interpreted and used to control artificial effectors or the patient's own muscles.

We have been investigating ways in which development could be performed in neuromorphic circuits using principles that are neurobiologically plausible. In animal nervous systems, the patterns of neural activity that emerge during development, influenced both endogenously and exogeneously, play a vital role in establishing normal processing. While there is no compelling reason to make the development of a neuromorphic "brain" or prosthetic neurobiologically plausible, using mechanisms evolved over eons would appear to be a good strategy.

If one can point to a single principle underlying both the development of neural connections and the subsequent adjustment of connections during learning, it would be the rule ascribed to the psychologist, D.O. Hebb. In essence, the rule requires the strengthening of the connection between a sending, or presynaptic neuron and a receiving, or postsynaptic neuron if the two fire together close together in time. Otherwise, weakening of the connection may result, depending upon which of several formulations of the rule is applied. Recent neurobiological work has shown that the relative timing of pre- and postsynaptic spiking is indeed responsible for potentiation and depression of synaptic efficacy or weight (Markram et al. 1997; Zhang et al., 1998).

To illustrate the neuromorphic approach, we present results on a simulation of a network of spiking neurons that approximate the behavior of our silicon neuromorphs. Previous theoretical work showed that the application of Hebb rules can lead to the development of adaptive connections and realistic sensory receptive fields (Linsker, 1990). Although the intent here is to develop a controller for an autonomous vehicle, the results demonstrate for a system of spiking neurons how inputs from sources with very different characteristics can combine automatically to generate order and useful functionality – in this case directional selectivity. In those sensory

systems that have topographic arrangements (i.e. visual, somatosensory, auditory), neurons are found that respond selectively to one direction of motion, whether it be in space, or in some other dimension like sound frequency. Such selectivity is important for higher neural processing of all kinds, for example our ability to use vision to guide ourselves through our environment, or to recognize speech sounds.

Simulated network

The simulated vehicle is equipped with two kinds of exteroceptors, whiskers and photoreceptors, and proprioceptors that signal the vehicle's own motion. Figure 4 shows the arrangement of sensors and the network of spiking neurons whose activity can be used to guide movements of the vehicle as it roams around an environment of obstacles.




Figure 4. Schematic of sensory system. Both whisker and photoreceptor arrays sample the frontal 120 degs of space. The spiking whisker units project topographically to both Hebb layers. The photoreceptors drive the transient and sustained spiking layers that are connected to the Hebb layers via modifiable, initially random connections. Each Hebb Layer is excited by proprioceptive spikes signaling left or right eye/head turns.

The six whiskers and the 10-20 photoreceptors are fixed to the vehicle and arrayed to sense the frontal 120° in the horizontal plane. Each whisker, when it contacts an obstacle generates a spike train at a frequency inversely proportional to the distance of intersection along the whisker length. The photoreceptors, which respond proportionately to light intensity integrated over a receptive field with Gaussian profile, drive two kinds of spiking units, the sustained and transient units, analogous to the main types of ganglion cells in the vertebrate retina. The receptive fields of these spiking units are about 10° wide.

The whisker and the photoreceptor spike signals connect to two layers of simulated neuromorphs, Hebb Layer Left and Hebb Layer Right, which by a Hebbian rule are to acquire directionally selective properties to stimuli moving left or right respectively. These layers are composed of 8 units, each of which is based on the Spike-Response (S-R) model of Gerstner & Kistler (2002). This model provides an approximation to the functional properties of a silicon neuromorph limited to one operative synaptic site. A spike input (from unit i) to one of these S-R units (unit j) generates a postsynaptic potential that is positive-going (excitatory) or negative-going (inhibitory), the amplitude and sign depending on a numeric weight (w_{ij}) between -1.0 and $+1.0$. The rise and fall times of the postsynaptic potentials is fixed, unlike those in the ADT neuromorph in which it depends on the site of synaptic activation along the dendrite. Postsynaptic potentials are summed and fed to an integrate-and-fire spike generator, the firing threshold of which is controlled by a bias input. Refractoriness after spike firing is implemented by immediately raising the spike-firing threshold and allowing it to decay to the bias threshold over a fixed time course.

The connections from the whisker array to both Hebb layers are preset and not modifiable. The whiskers are mapped topographically across the Hebb layers so that stimulation of the leftmost whisker, for example, activates units at the leftmost end of the Hebb layers (see Fig 6A). Hebb layer units that are not related topographically to a whisker are inhibited by that whisker. The other preset and unmodifiable connections are the proprioceptive afferents to the Hebb layers. The left-turn proprioceptors fire when the vehicle rotates leftwards and tends to excite all the units in Hebb Layer Left. Similarly the right-turn proprioceptors excite Hebb Layer Right. The transient and sustained units of the “retina” are projected in a fully connected pattern with initially zero weights onto both Hebb layers. The latter connections are modifiable according to a Hebb rule akin to the mechanism that strengthens and weakens synapses in mammalian cortex (Markram et al. 1997) and in amphibian tectum (Zhang et al., 1998). Synaptic weight modification depends upon the relative timing of pre- and postsynaptic spikes: if the presynaptic spike occurs in a time window before the postsynaptic unit fires, the weight of the presynaptic synapse is increased; if it occurs after the postsynaptic spike, it is decreased; otherwise no change occurs (Fig. 5). The effect is that the synapses from any visual units that fire consistently just before the firing of a Hebb Layer units will have their connections strengthened or made more excitatory. If the visual units fire just after, their connections are weakened and eventually may become inhibitory. Visual unit firing that is unrelated to the activity of the Hebb layer units results in no net change of weight if the positive and negative phases of the window function (Fig. 5) integrate to zero.

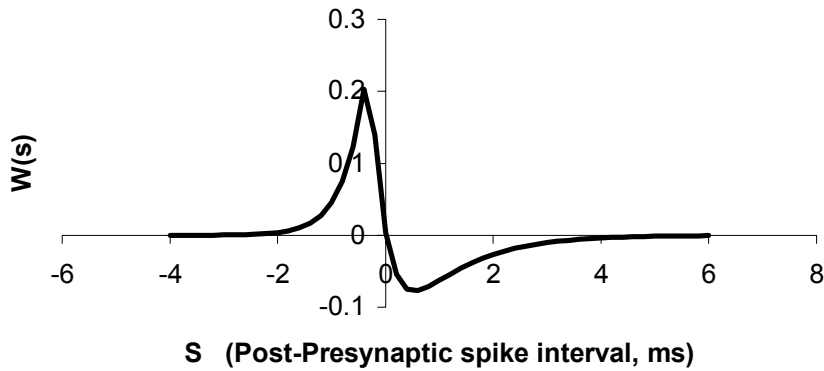


Figure 5. Hebbian window function. Synaptic weight is increased in proportion to $W(s)$. When a presynaptic spike precedes a postsynaptic spike by less than about 2 ms, $W(s)$ is positive; when a presynaptic spike follows a postsynaptic spike by up to about 4 ms, $W(s)$ is negative; otherwise vanishingly small. (See Gerstner & Kistler, 2002).

Weights onto unit i that has fired recently are updated by an amount Δw_{ij} according to the following rule:

$$\Delta w_{ij} = \epsilon(W(s_{ij}) - w_{ij})$$

Where ϵ is the learning rate, s_{ij} is the time interval between the firing of unit i and j , and $W(s)$ is the function shown in Fig. 5. Subtraction of the existing weight w_{ij} , ensures that weights are soft bounded i.e. they do not grow without limits (c.f. CPCA learning rule, O'Reilly & Munakata, 2000). The repeated application of this rule leads to stable weights but an excess of excitatory drive to the Hebb layer units. To maintain an equitable balance between excitation and inhibition, weight normalization is applied to make all the weights onto each unit sum to 1.

In the first phase of development, the vehicle traverses the field populated by obstacles of various sizes, advancing continuously in space by an equal distance every simulation time step. For the present demonstration, the vehicle's steering is disabled, so that it moves in a straight line, brushing against randomly placed obstacles with its whiskers while its "retina" generates visual impulses that have no effect on the Hebb layers. At this stage, we need only consider the sustained units and their connections. The connections from the whiskers, which are topographically mapped onto the Hebb layers (Fig. 6 A) are sufficiently strong to fire the Hebb layer units reliably. The application of the learning rule results in the strengthening of the connections from those visual units that were active shortly before the Hebb layer units spiked. Weights from the visual units are continuously updated so that the sustained visual units come to map topographically over both Hebb layers. This is shown in the weight matrices of Fig. 6B by the diagonal of excitatory weights. The effect of weight normalization is to form an inhibitory surround. If the spiking activity of the Hebb layers is now connected to the vehicle's steering, the vehicle avoids obstacles visually, without contact with the whiskers.

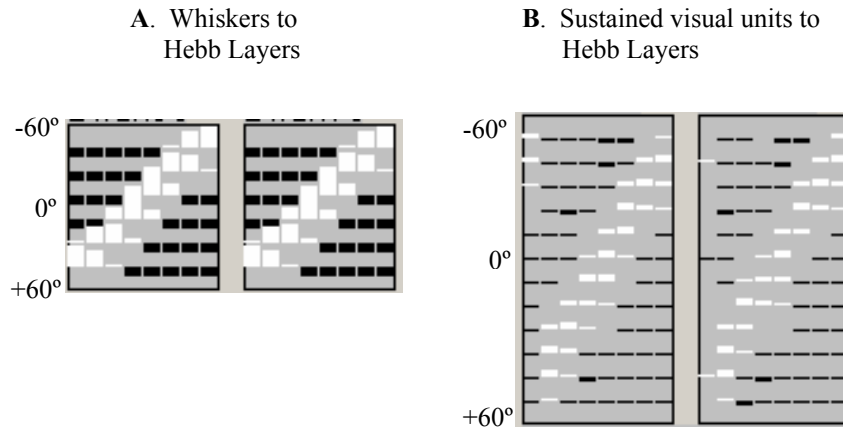


Figure 6. Weight matrices connecting (A) whiskers and (B) visual sustained units to the left and right Hebb Layers. Excitatory weights are shown in white, inhibitory in black. Height of bars represent weight values. Inputs in whisker and visual space (-60° to $+60^\circ$) are represented along rows; the different Hebb Layer units along columns.

Negotiating a complex environment requires that information other than just luminance is brought to bear; judging distances, reacting to impending collisions and guiding locomotion require the use of visual motion cues. The network of Fig. 4 and the application of Hebbian mechanisms during sensorimotor experience results in the development of motion selective units. For this, we added transient visual units that fire to changes in luminance. Because the extraction of velocity information depends on correlations between stimulus positions at different times, the spikes from the transient units are connected to the Hebb layers via different time delays. For the sake of illustration, the transient unit array is connected to the Hebb layers via four weight matrices, one each for delays of 0, 0.5, 1.0 and 1.5 ms (Fig. 7 A,B). In a neuromorphic system, these delays could be implemented by the facility for programmable axonal delays in the Virtual Wires system, or by dendritic delays that depend on the site of synaptic activation along a dendrite. Under appropriate conditions, the application of the Hebb rule in the presence of a moving stimulus image could selectively strengthen inputs that come from the four positions in space occupied by the image at times t , $t-0.5$, $t-1.0$ and $t-1.5$ ms. For Hebb Layer units to come to respond to the motion, they must be firing at time t . Therefore, they require an “instructive” input generated during image movement to bias them on so that they fire. In this example, the image motion is self produced, as by eye or head movements while looking at stimulus objects in space. The sources of the instructive inputs are proprioceptors that signal the vehicle’s own turning movements. The instructive inputs could also be derived from motor command signals. In the network of Fig. 4, proprioceptors for leftward or rightward movement fire a burst of spikes to all the units of their corresponding Hebb layer. Thus the convergence of transient and proprioceptive information onto Hebb layer units should result in their acquiring selectivity to one direction of motion, but without regard to position in the visual field. Extracting depth from parallax, however, likely requires that the spatial arrangement of different velocity vectors be preserved. As we have seen, the topography was established in the mapping of the sustained units over the Hebb layers, so if the sustained inputs also bias the Hebb units to fire, along with the proprioceptive inputs, directionally selective maps with topography should be formed.

Figure 7 shows examples of how weights from the transient units to the Hebb layer units developed. The four weight matrices at increasing delays show progressive shifts in their receptive fields, thereby generating a skewed spatiotemporal receptive field required for directional selectivity (Adelson & Bergen, 1985). To test directional selectivity, a single visible

object was rotated at a constant rate through the visual field. Fig. 7 C,D (lower half) shows that Hebb Layer Left responded with spikes to leftward motion but not to rightward motion, and vice-versa for Hebb Layer Right. Note also that the units have restricted receptive fields organized topographically, a first step towards analyzing motion flow fields. We can now envision some further steps that would be required to use this information for behavioral guidance: the development of selectivities for different velocities and different patterns of motion as exhibited by neurons in the visual areas of the cerebral cortex.

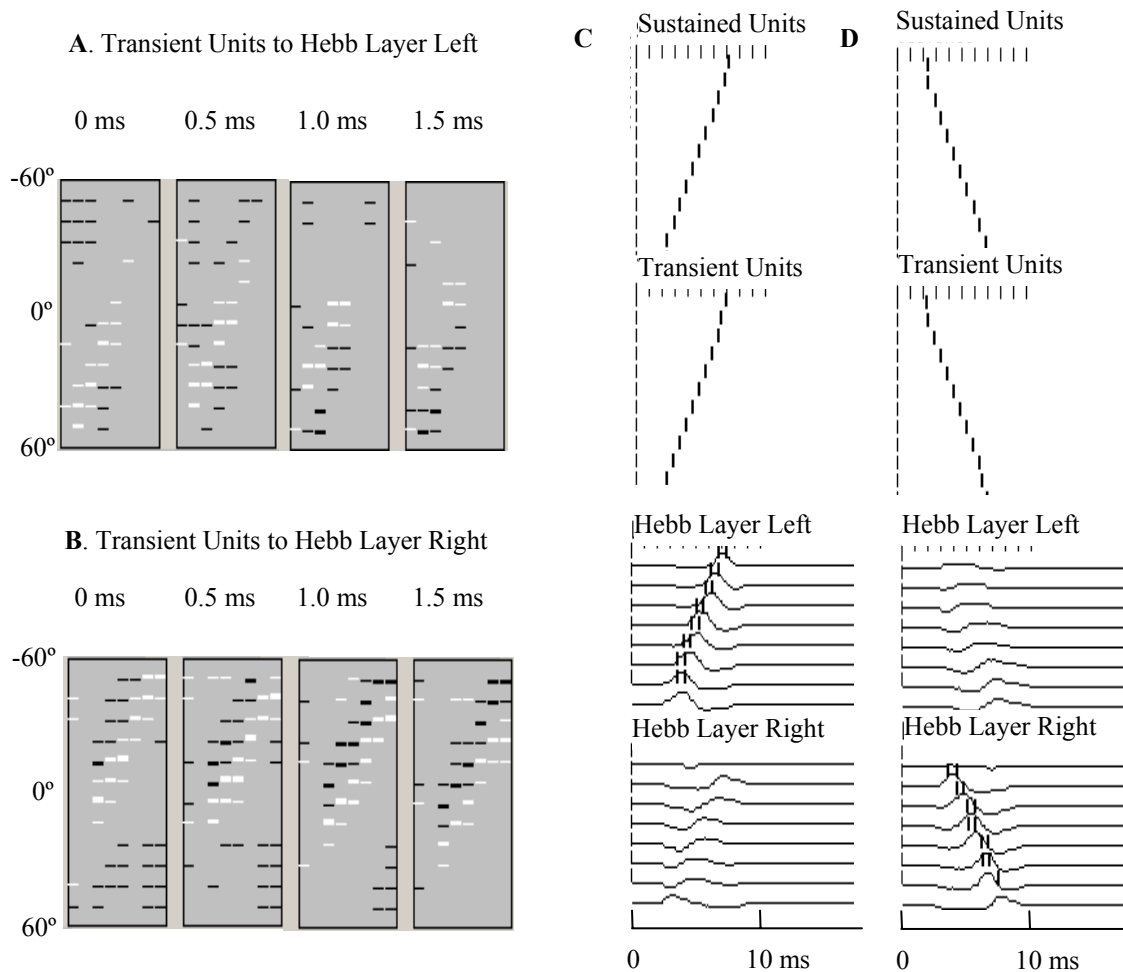


Figure 7. A & B. Weights of the transient units onto the Hebb Layers with delays of 0, 0.5, 1.0 and 1.5 ms. White indicates excitatory weights; black inhibitory. Inputs in visual space (-60° to 60°) are represented along rows; the different Hebb Layer units along columns.

C & D. Responses as a function of time. Spikes of sustained and transient visual units, and the two Hebb Layers in response to a stimulus rotating from right to left (**C**), and from left to right (**D**) through the entire visual field. Waves in the lower half show the units' internal potentials

Neuromorphs in neural prosthetics

Neuromorphic systems of the type we are developing may have a special advantage for neural prosthetics in that neuromorphs naturally deal with spike signals: they accept spike inputs and they generate spike outputs. A very small neuromorphic network implanted in a paralyzed patient could interpret spike signals from an array of electrodes embedded in the motor cortex (Serruya et al., 2002). The network would extract the relevant information contained in the frequency and timing of the cortical spikes and would generate output spike trains patterned so as to activate neurons, fiber tracts or muscles, restoring lost function to the patient. Neuromorphic prosthetics will certainly require the ability to adapt and learn because the input spikes are largely arbitrary in nature – they are whatever the implanted electrodes can pick up from the CNS – and they must be associated with patterns of spike outputs for controlling behavior (Taylor et al., 2002). While our neuromorphic development system (Domain board) can implement arbitrary learning schemes, it currently requires the assistance of an external computer.

Conclusions

The challenge to building neuromorphic systems is deciding what features to incorporate into the neuromorphic units and how to connect them to perform usefully. A process of design from biophysical principles is an estimable approach but difficult (Eliasmith & Anderson, 2002). Faced with the great variety of potential mechanisms that neurobiological research is revealing and the complex, nonlinear interactions between them, we prefer to seek neuromorphic systems that are relatively simple to make and able to self-organize and adapt dynamically. The results presented here show that valuable capabilities can emerge in networks through interaction with the environment, in effect, on-the-job training. To be sure, quite a bit of structure has to be imposed for this development to happen. How to arrive at appropriate structures is an interesting problem. One strategy is to borrow from what is known of structure in brains; an alternative is to evolve structure through evolutionary methods. Another problem is how to regulate excitabilities so that development can occur in an adaptive fashion. As we have seen, for anything sensible to be learned by Hebbian mechanisms, units must be biased so that they are driven to fire when learning is appropriate and inhibited when it is not. Brains seem to employ diffuse neuronal projections and humoral controls for this purpose and these could well be emulated in neuromorphic systems. Among the technical problems that need to be solved for neuromorphic systems to gain independence in small-scale applications is how to store and modify connections locally or “on chip” so that learning can occur without the intervention of external, digital computers. Notwithstanding these problems and unknowns, the hardware system described achieves, to greater or lesser extents, the following desirable features of a neuromorphic system:

- Parallelism of processing elements for fast processing of large amounts of information
- Distributed processing whereby many units simultaneously represent a sensory stimulus or execute a movement.
- Robust in the face of variation and faults in individual components, tolerant to damage and noise, performance degrading gradually, not catastrophically.
- Real-time processing with wide time-scale dynamics that can be matched to events in the real world, and to the dynamics of limbs and other body systems.
- Flexible, rapidly reprogrammable architecture and connections, capable of ontogenetic and evolutionary adaptation.
- Learning based on short-, medium- and long-term changes in synaptic efficacy, unit excitability, and persistent patterns of activity.
- Efficient, noise-immune, signaling by impulses with rich temporal coding possibilities.

- Powerful analog computation distributed throughout a dendritic tree and soma.
- Compact size and low power consumption, essential for mobility and portability.

References

- Adelson, E.H. & Bergen, J.R. (1985) Spatio-temporal energy models for the perception of motion. *J. Opt. Soc. Am. [A]*, 2: 284-299.
- Douglas, R. and Rasche, C. (2002) Silicon Neurons. In *The Handbook of Brain Theory and Neural Networks*, Michael Arbib (Eds.), MIT Press.
- Eccles, J.C. (1957) *The Physiology of Nerve Cells*. Johns Hopkins Univ. Press, Baltimore, MD.
- Elias, J.G. (1993) Artificial dendritic trees. *Neural Computation*, 5: 648-664.
- Elias J.G., Northmore, D.P.M. (1995) Switched-capacitor neuromorphs with wide-range variable dynamics. *IEEE Transactions on Neural Networks* 6:1542-1548.
- Elias J.G., Northmore, D.P.M. Westerman, W. (1997) An analog memory circuit for spiking silicon neurons. *Neural Computation*, 9:419-440.
- Eliasmith, C., Anderson C.H. (2002) *Neural Engineering*. The MIT Press, Cambridge, MA.
- Gerstner, W., Kistler, W. (2002) *Spiking Neuron Models*. Cambridge Univ. Press.
- Linsker, R. (1990) Perceptual neural organization: some approaches based on network models and information theory. *Annual Review of Neuroscience* 13: 257-281.
- Maass, W. (1997) Fast sigmoidal networks via spiking neurons. *Neural Computation*, 9: 279-304.
- Maass, W., Bishop, C.M. (1999) *Pulsed Neural Networks*. MIT Press, Cambridge, MA.
- Mahowald, M.A. (1992) Evolving analog VLSI neurons. In *Single Neuron Computation*, T. McKenna, J.Davis, S.F. Zornetzer, eds., Chap. 15. Academic Press, San Diego, CA.
- Markram, H. Lubke, M., Frotscher, B. & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic Aps and EPSPs. *Science* 275: 213-215.
- Mead, C. (1989) *Analog VLSI and Neural Systems*. Addison-Wesley, Reading, MA.
- Mel, B.W. (1994) Information processing in dendritic trees. *Neural Computation*, 6: 1031-1085.
- Natschläger, T., W. Maass, W. and Markram, H. (2002) The "liquid computer": A novel strategy for real-time computing on time series. *Special Issue on Foundations of Information Processing of TELEMATIK*, 8:39-43, 2002.
- Northmore, D.P.M., Elias, J.G. (1996) Spike train processing by a silicon neuromorph: The role of sublinear summation in dendrites. *Neural Computation* 8: 1245-1265.

- O'Reilly, R.C., Munakata, Y. (2000) *Computational Explorations in Cognitive Neuroscience*. MIT Press, Cambridge, MA.
- Rall, W. (1964) Theoretical significance of dendritic trees for neuronal input-output relations. In *Neural Theory and Modeling*. R.F. Reiss, ed., Stanford Univ. Press, 73-79.
- Serruya, M.D., Hatsopoulos, N.G., Paninski, L., Fellows, M.R. & Donoghue, J.P. (2002) Instant neural control of a movement signal. *Nature* 416: 141-142.
- Smith, L.S. and Hamilton, A. (1998) *Neuromorphic Systems: Engineering Silicon from Neurobiology*. World Scientific, Singapore.
- Taylor, D.M., Tillery, S.I. H., & Schwartz, A.B. (2002) Direct cortical control of 3D neuroprosthetic devices. *Science*, 296: 1829-1832.
- Rabinovich M, Volkovskii A, Lecanda P, Huerta R, Abarbanel HD, Laurent G. (2001) Dynamical encoding by networks of competing neuron groups: winnerless competition. *Phys Rev Lett* 87(6):068102-4
- Zhang LI, Tao HW, Holt, C.E., Harris, W.A. and Poo, M. (1998) A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, 395:37-44.