

## **Coupling an aVLSI Neuromorphic Vision Chip to a Neurotrophic Model of Synaptic Plasticity: The Development of Topography**

**Terry Elliott**

*te@ecs.soton.ac.uk*

*Department of Electronics and Computer Science, University of Southampton, Highfield, Southampton, SO17 1BJ, United Kingdom*

**Jörg Kramer**

*kramer@ini.phys.ethz.ch*

*Institute of Neuroinformatics, University of Zürich, and ETH Zürich, Winterthurerstrasse 190, CH-8057 Zürich, Switzerland*

**We couple a previously studied, biologically inspired neurotrophic model of activity-dependent competitive synaptic plasticity and neuronal development to a neuromorphic retina chip. Using this system, we examine the development and refinement of a topographic mapping between an array of afferent neurons (the retinal ganglion cells) and an array of target neurons. We find that the plasticity model can indeed drive topographic refinement in the presence of afferent activity patterns generated by a real-world device. We examine the resilience of the developing system to the presence of high levels of noise by adjusting the spontaneous firing rate of the silicon neurons.**

### **1 Introduction** ---

Exposing biological models to the full complexity of real environments provides far better and more rigorous tests of their underlying assumptions and mechanisms than simulating their environments on computers. Such simulations generally fail to model details of inherent noise and irregularities in the environment, some of which may have a significant influence on the development and functioning of the modeled biological system and the performance of the model. For the purposes of sensory-processing models in the neurosciences, it is therefore desirable to generate afferent electrical activity patterns for input into neuronal networks by using real-world sensory-transduction devices. Most work done according to this paradigm uses standard electronic sensing devices, such as video cameras and microphones, which do not bear much similarity to biological sensors beyond the fact that they transduce the sensory signal into an electrical one. A recent alternative to using such traditional sensors is provided by the relatively

new concept of neuromorphic engineering (Mead, 1989). Neuromorphic engineering typically uses analog VLSI (aVLSI) technology to implement, at the level of neuronal processing elements and the electrical signals they generate, key features of the low-level sensory processing found, for example, in the visual or auditory systems. Perhaps the best-known product of neuromorphic engineering is the so-called silicon retina, a device that generates retinal ganglion cell-like output in response to optical stimuli falling on its surface.

After an early attempt to build a retina model from discrete electronic components (Fukushima, Yamaguchi, Yasuda, & Nagata, 1970) that already showed the basic spatiotemporal filtering behavior observed in biological retinæ, a host of retina-like integrated circuits have been developed. These circuits differ in the types of retinal functions they implement and the detail with which they model these functions. Depending on the intended applications, some implementations focus on modeling spatial properties of the inner and outer plexiform layers (Boahen, 1999), while others, at the expense of detailed spatial modeling, put more emphasis on reproducing temporal dynamics, such as adaptation in the photoreceptors (Mahowald, 1991; Delbrück & Mead, 1994). For more extensive reviews of neuromorphic vision chips we refer to the literature (Etienne-Cummings & Van der Spiegel, 1996; Moini, 1999).

For the work examined here, a model has been used that restricts itself to temporal aspects of retinal processing in favor of high processing densities and low fixed-pattern noise (Kramer, 2002a). It implements the separation into ON and OFF channels found in the retina and the lateral geniculate nucleus, shows quickly decaying transient responses, as found in the magnocellular pathway projecting from the retina into the visual cortex, and employs the spike-based activity encoding of the retinal ganglion cells and later visual processing stages (Kramer, 2002b).

Silicon retinas have so far mainly been used to visualize the retinal output and to demonstrate its qualities. But they were also instrumental in driving the development of an event-driven, asynchronous communication infrastructure for neuromorphic chips and other systems (Lazzaro, Wawrzynek, Mahowald, Sivilotti, & Gillespie, 1993; Mortara & Vittoz, 1994; Deiss, Douglas, & Whatley, 1998; Boahen, 2000), which has been adopted for this work. A few attempts have been made to use this infrastructure to couple retina chips to other neuromorphic circuits implementing higher-level visual processing (Mahowald, 1994; Venier, Mortara, Arreguit, & Vittoz, 1997; Indiveri, Whatley, & Kramer, 1999; Higgins & Koch, 2000; Indiveri, Mürer, & Kramer, 2001; Liu et al., 2001). Recently, an autoconfigurable neuromorphic chip that models the activity-dependent release and uptake of a diffusible neurotropic factor has been developed. It uses a fixed set of gradient-sensing silicon growth cones, interspersed within a two-dimensional array of target neurons, to wire together neurons that fire together (Kwabena Boahen, personal communication, 2001). However, to our knowledge, no attempt has

been reported so far that uses a neuromorphic vision chip in conjunction with a detailed model of synaptic plasticity that permits synaptic sprouting and retraction.

Here, we explore the coupling of a silicon retina to a biologically inspired model of anatomical, activity-dependent competitive synaptic plasticity and neuronal development. This model is inspired by recent data implicating a class of molecular factors, called neurotrophic factors (NTFs), in activity-dependent synaptic competition in the developing vertebrate nervous system (for a review, see, for example, McAllister, Katz, & Lo, 1999). The model has been developed, extensively analyzed, and applied in simulation to the development of various systems, including the visual system and the neuromuscular junction (Elliott & Shadbolt, 1998a, 1998b, 1999; Elliott, Maddison, & Shadbolt, 2001). Coupling this model of synaptic plasticity to a neuromorphic chip not only permits us to explore further the properties of the chip but also to determine whether the model can successfully drive normal development in the presence of realistic, rather than simulated, evoked afferent activity patterns.

We consider in particular the competitive, activity-dependent development of an orderly topographic projection from an array of afferent neurons, representing the neuromorphic vision chip, to an array of target neurons. Such a system could represent, for example, the development of the retinotectal system of amphibians and fish, which is known to develop and regenerate in an activity-dependent, competitive manner (Meyer, 1983; Schmidt & Edwards, 1983; Schmidt & Eisele, 1985). We (Elliott & Shadbolt, 1999) and others (see, for example, Goodhill, 1993; Sirosh & Miikkulainen, 1997) have previously considered the development of retinotopic maps, although only in the presence of simulated afferent input.

## 2 The Neuromorphic Chip

---

The neuromorphic chip consists of a  $16 \times 16$  square array of pixels (picture elements) consisting of adaptive photoreceptors with rectifying and threshold differentiating elements. The chip responds to ON and OFF transient optical signals at different output terminals. In response to a sufficiently large ON or OFF transient, which in an imaging setup typically corresponds to a moving edge, a pixel generates a single binary spike or a train of binary spikes, depending on the shape of the transient and the pixel's refractory period, which can be globally set with a bias control. For short refractory periods, a pixel emits a burst of spikes, even for sharp transients, and is thus said to be in the bursting mode, while for long refractory periods, even a prolonged transient can evoke only a single spike, and the chip is then said to operate in the nonbursting or single-spike mode. The response threshold and the spontaneous firing rate of the pixel can be varied separately for the ON and OFF pathways with bias controls. For example, the thresholds can be set such that only one of the pathways shows any response, and sponta-

neous firing can be completely suppressed if desired. The spikes from the two outputs of the different pixels in the array are filtered through an arbiter system, determining the order in which they are multiplexed to be sent off chip. Communication with other devices obeys a standard protocol, according to which the spikes are placed on an asynchronous communication bus as binary addresses that code for the identity of the sending neurons. In our case, the address codes for the location of the pixel in the array and its output terminal (ON or OFF). Data traffic on the communication bus is regulated using a four-phase handshaking protocol. Bus addresses rapidly decay to zero in the absence of further spiking activity.

The chip is coupled to a standard PC running under Linux for data acquisition purposes. Due to unresolved difficulties in the rate of processing interrupts by the Linux kernel, it has been found simplest to set the chip in free running mode by shorting its handshaking pins, and reading the address bus, connected to the computer's parallel port, at a rate set by the speed of the computer simulation. Such an approach, however, entails two problems. First, without reading the status of the handshaking pins, it is impossible to determine whether a zero address ( $X = 0, Y = 0$ , these being pixel coordinates on the chip) on the bus corresponds to a decayed address in the absence of recent spiking activity or to a recent spike generated by the (0,0) pixel. All zero address reads are therefore discarded, so that the (0,0) pixel is ignored. Second, in free running mode, the computer can read multiple times the same address, generated from a single-spike event, before it decays to zero. In bursting mode, when each pixel generates several spike requests in response to one ON or OFF edge, it is therefore impossible, without examining the handshaking pins, to distinguish multiple reads of the same spike event from multiple reads of the same address generated by several spike events from the same pixel. The second problem can be eliminated by setting the chip to nonbursting mode, so that precisely one spike is generated in response to each ON or OFF edge, and by discarding subsequent consecutive reads of any new address on the bus. Using the nonbursting mode also practically eliminates a third possible problem: that spikes may be ignored by the PC due to interspike intervals being shorter than the PC's read period. This would happen in the presence of strong activity on the bus, given that the maximum communication rate of the chip is significantly larger than the maximum read rate of the PC. We find that our approach generates data that reflect very accurately the optical stimuli presented to the chip.

Optical stimuli for the chip are generated on the LCD monitor of the same PC as is used for data acquisition. The stimuli consist of eight separate sequences of images, each sequence corresponding to a white bar sweeping across a black background in a window on the monitor. The sequences are distinguished only by the orientation and direction of motion of the bar. A bar can be horizontal, vertical, or diagonal (on either diagonal) with respect to the orientation of the rows of the array and move in either direction

perpendicular to the bar's orientation. All bars are sufficiently long to fill the field of view of the chip (i.e., the bars are effectively infinitely long, although of finite width). These sequences are repeatedly played without interruption, each new sequence being selected randomly after the completion of the previous one. The speed of the bar is largely unimportant. Provided that the speed is neither so low nor so high that the chip does not respond faithfully to an edge, any variation in intermediate speeds merely tests the computer's capacity to process the captured data rapidly enough before another set of data is captured. The optical stimuli are focused onto the chip's light-sensitive surface by a variable-aperture, variable-focus lens. These stimuli consist of both a leading ON edge and a trailing OFF edge. However, because the chip's response to ON and OFF edges is not quite symmetric and because we require only one response type for this work, the ON responses from the chip are suppressed.

### 3 The Neurotrophic Model of Plasticity

---

We briefly describe our neurotrophic model of anatomical, activity-dependent, competitive synaptic plasticity and its application to the development of topography. A fuller description, derivation, and discussion of the model and its application to various systems can be found elsewhere (Elliott & Shadbolt, 1998a, 1998b; 1999; Elliott et al., 2001). Our discussion of the model is at first general. We then discuss specific features of the implementation of the model necessitated by its coupling to the neuromorphic chip.

Let letters such as  $i$  and  $j$  label afferent or presynaptic cells and letters such as  $x$  and  $y$  label target or postsynaptic cells; the vector character of these labels is left implicit for the purposes of notational simplicity unless indicated otherwise. Let  $s_{xi}$  denote the number of synapses projected from afferent  $i$  to target  $x$ , and let  $a_i \in [0, 1]$  denote the activity of afferent  $i$ . Then our basic equation for the evolution of  $s_{xi}$  is given by

$$\frac{ds_{xi}}{dt} = \epsilon s_{xi} \left[ \frac{(a + a_i)\rho_i}{\sum_j s_{xj}(a + a_j)\rho_j} \sum_y \Delta_{xy} \left( T_0 + T_1 \frac{\sum_j s_{yj}a_j}{\sum_j s_{yj}} \right) - 1 \right]. \quad (3.1)$$

In this equation,  $T_0$  and  $T_1$  represent, respectively, an activity-independent and a maximum activity-dependent release of NTF by target cells; the parameter  $a$  represents a resting NTF uptake capacity by afferents. These three parameters, through the combination  $c = T_0/(aT_1)$ , determine whether competitive dynamics are present in the model: for values of  $c$  below a calculable threshold, competition occurs, while for values of  $c$  exceeding this threshold, competition breaks down (Elliott & Shadbolt, 1998a). Analysis shows that  $a$  must be neither too large nor too small, so we set  $a = 1$  as in most previous work.  $T_1$  is an overall scale factor for the number of synapses and may be set for reasons of numerical convenience, but without any loss of generality, to  $T_1 = 20$ . Selecting a value of  $T_0 = 0$  will ensure the maximum rate of

map development, so this is what we use. The function  $\Delta_{xy}$  characterizes the diffusion of NTF between target cells. Alternatively, it may be thought of as encoding lateral interactions between target cells, with excitatory lateral interactions (positive  $\Delta_{xy}$ ) enhancing target cell NTF release and inhibitory lateral interactions (negative  $\Delta_{xy}$ ) suppressing NTF release. For simplicity, we take  $\Delta_{xy}$  to be a gaussian with characteristic width  $\sigma = 0.75$ , consistent with previous work (Elliott & Shadbolt, 1999). NTF receptor dynamics are encoded in the function  $\rho_i = \bar{a}_i / \sum_x s_{xi}$ , where  $\bar{a}_i$  denotes the recent average activity of afferent cell  $i$ . It can be shown (Elliott & Shadbolt, 1998a) that if  $\Delta_{xy}$  is nonnegative, then NTF receptor dynamics are required to ensure that each afferent cell always innervates at least one target cell. However, if  $\Delta_{xy}$  possesses a Mexican hat profile, with an excitatory center and inhibitory surround, NTF receptor dynamics can be discarded. Finally,  $\epsilon$  sets the overall rate of plasticity. Here, we set  $\epsilon = 0.02$ , which is sufficiently small to prevent development from being disrupted by too great a dependence on temporally recent stimulation and sufficiently large to allow adequately short run times.

To apply this model to the development of a topographic mapping between an afferent structure and a target structure, we suppose that the afferent and target cells are separately organized into two square arrays of cells. Although not necessary, for convenience we take these arrays to be of the same size, so that the number of target cells equals the number of afferent cells. Were the topographic mapping between these two sheets of cells perfect, each afferent cell would innervate precisely one target cell—the target cell that would be in register with it were the two sheets superimposed without distortion. Such a mapping ensures that neighborhood relations on the afferent sheet are preserved on the target sheet and vice versa and are expected to develop from an initially unrefined mapping in which these relations are not at first respected. We establish an initial unrefined pattern of connectivity between the afferent and target sheets by following the method used by Goodhill (1993). Let  $d_{\max}$  denote the maximum distance (in units of cell spacing) between any pair of cells in the target sheet. For a given afferent cell, let  $d$  denote the distance between any randomly selected target cell and the target cell to which the afferent cell would project were topography perfect. Then the number of synapses between the afferent cell and the randomly selected target cell is set to be proportional to

$$\beta \left( 1 - \frac{d}{d_{\max}} \right) + (1 - \beta)n, \quad (3.2)$$

where  $n \in [0, 1]$  is a randomly selected number for each such pair of target and afferent cells. The parameter  $\beta \in [0, 1]$  defines the topographic bias in the initial projections (Goodhill, 1993). A value of  $\beta = 0$  gives no initial topographic bias, with entirely random connections between the two sheets. A value of  $\beta = 1$  gives maximum bias, so that an afferent cell projects max-

imally to its topographically preferred target cell, with projections falling linearly with increasing distance away from this cell. Consistent with previous work, we usually set  $\beta = 0.5$  (Elliott & Shadbolt, 1999); we will also examine the impact of decreasing the value of  $\beta$  on our results.

In order to depict in graphical form the topographic representation of the afferent sheet on the target sheet, we use the following method. First, for each target cell, we calculate the center of mass (CoM) of afferent projections to that cell, so that we evaluate

$$\vec{M}_x^{\text{CoM}} = \frac{\sum_{\vec{i}} \vec{i} s_{\vec{x}\vec{i}}}{\sum_{\vec{i}} s_{\vec{x}\vec{i}}}, \quad (3.3)$$

where we have restored the vectorial nature of the afferent and target cell labels for clarity. The positions defined by these vectors are then plotted in space and connected by lines according to the nearest-neighbor relations on the target sheet: the positions corresponding to  $\vec{M}_x^{\text{CoM}}$  and  $\vec{M}_y^{\text{CoM}}$  are connected by a line if, and only if, cells  $\vec{x}$  and  $\vec{y}$  are nearest neighbors on the target sheet; this is repeated for all such pairs. Increasing deviations of the resulting pattern of lines away from a regular square grid of lines represent increasing disruption of the topographic representation of the afferent sheet on the target sheet.

We now turn to a discussion of the features of the implementation of the model that are particular to its coupling to the neuromorphic chip. The silicon retina is a square array of  $16 \times 16$  pixels. Hence, the array of target cells is also  $16 \times 16$ . In free running mode, we discard data corresponding to a zero address on the bus and so ignore the (0,0) pixel on the retina. Correspondingly, we discard the (0,0) neuron in the target sheet. This is achieved by removing all projections from the (0,0) afferent pixel to the target sheet and all projections to the (0,0) target cell from the afferent sheet.

Our plasticity model is not a spiking model, but rather uses an instantaneous firing rate,  $a_i \in [0, 1]$ . In principle, we could capture one spike at a time from the chip and process this single spike according to equation 3.1. However, single spikes processed in this way by our model would lack any spatial or temporal information about the structure of the optical stimuli presented to the chip, and consequently topographic refinement would not occur. We therefore capture, or "bin," several spikes before simultaneously processing all of them according to equation 3.1. In bursting mode, it is possible that each pixel could spike several times during a binning read, thus necessitating a somewhat ad hoc procedure for converting spike numbers with variable maxima per bin into instantaneous rates. In nonbursting mode, for a typical moving edge, each stimulated pixel spikes precisely once. In this mode, therefore,  $a_i$  conveniently takes only binary values, so that  $a_i \in \{0, 1\}$ , without the need for any arbitrariness in the conversion to an instantaneous rate.




Figure 1: The development of the topographic mapping between the neuromorphic chip's retinal sheet and a target sheet of neurons, with the binning at each iteration of 32 spikes. Each map is generated as discussed in the text and represents the state of topography at the iteration number indicated immediately above it.

Because the retinal array is  $16 \times 16$ , a horizontally or vertically moving edge generates successive sequences of 16 spikes, the spikes being generated by the line of pixels on which the edge instantaneously falls. We therefore bin integral multiples of 16 spikes during the data acquisition process, before processing the spike data using equation 3.1. In fact, we will show results from the binning of 16, 32, and 64 spikes.

#### 4 Results

---

We now turn to a presentation of our results. A run typically consists of 10,000 to 20,000 iterations, each iteration corresponding to the capturing and processing, via equation 3.1, of a fixed number of spikes from the chip. The number of spikes captured per iteration influences the final quality of the topographic mapping from the retinal sheet to target sheet and affects the overall rate of topographic error reduction. We consider the binning of 16, 32, and 64 spikes and perform several complete runs for each bin size. For each bin size, all runs are qualitatively very similar. We therefore characterize the results for each bin size by randomly selecting a representative run from each case.

For the binning of 32 spikes, we show, in Figure 1, the refinement of topography over the period of development. By approximately 10,000 it-



erations, topographic development is largely complete, with no significant further changes occurring with increased iteration numbers. Initially, the topographic map is tightly folded and disrupted. For example, for target cells near the edges of the sheet, considerable afferent input is derived initially from the opposite edge of the afferent sheet. Hence, the center of mass of afferent projections at first is skewed toward the center of the afferent sheet. As development proceeds, the map unfolds, with the centers of mass drifting in the direction of perfect projections, although nonuniformities in the emerging topographic grid remain. By about 6000 iterations, the map is nearing its final structure. Edge effects remain even at 10,000 iterations, although these would likely be reduced by greatly increasing the number of iterations. The "hole" at the top left corner of the map corresponds to the missing (0,0) afferent and target cells. Although the afferent input is derived from a real-world device and not simulated, these results compare very favorably with those obtained from simulated retinal inputs (Goodhill, 1993; Elliott & Shadbolt, 1999). Corresponding to Figure 1, in Figure 2 we show the development of the receptive field of a typical target cell near the center of the target sheet. The development of this cell's receptive field shows that, as with the development of the overall topographic map, the receptive field is nearly mature at 6000 iterations.

In Figure 3, we show the final topographic maps obtained from binning 16, 32, and 64 spikes, where the data for 32 spikes are identical to those used to generate Figures 1 and 2; these maps represent 17,500, 10,000, and 10,000 iterations of development, respectively. There is little qualitative difference between the maps for the binning of 32 or 64 spikes; both are nearly, although not quite, perfect. For 16 spikes, however, the map, although well developed, does not exhibit the quality of the maps for larger numbers of spikes. This is because the captured activity patterns represent only the leading edge of a horizontally or vertically moving bar and lack any structure in a direction perpendicular to the orientation of the bar's edge. A very similar result would be obtained with simulated retinal data in which just single lines of retinal cells are activated across the retina.

Figure 4 shows the change in the overall topographic error in the maps as development proceeds for the 16, 32, and 64 spike data sets shown in Figure 3. At each iteration, the topographic error is calculated simply as the average over all target cells of the distance (in units of cell spacings) between a target cell's center of mass input and the afferent cell that would project to the target cell were topography perfect. The greatest rate of error reduction is achieved with the binning of 64 spikes. Nevertheless, although 32 spikes result in a slower decrease in topographic error, the final errors in the mature states for both 32 and 64 spikes are nearly identical. The residual error in these systems is due largely to the edge effects previously discussed. With 16 spikes, the error reduces and stabilizes by around 10,000 iterations but remains higher than the final error for 32 or 64 spikes. In addition to edge effects, this higher error reflects the slightly greater dis-




Figure 2: The development of the receptive field of a target cell near the center of the target sheet for the same system as shown in Figure 1. Each square represents an afferent cell's input to the target cell, with the gray-scale value of the square indicating the number of synapses projected by that afferent cell. A white square ("Min") denotes no synapses, a black square ("Max") denotes the largest number of synapses projected by an afferent to the target cell, and shades of gray interpolate between these extremes.




Figure 3: Final states of the topographic mappings between the neuromorphic chip's retinal array and a target sheet for three different levels of spike binning per iteration. The number of spikes per bin used to generate a map is indicated immediately above the map.




Figure 4: The change in the overall topographic error in the development of the topographic maps, calculated as discussed in the text, plotted against iteration number, for the same three data sets used to generate Figure 3. The number attached to each curve indicates the number of spikes per bin used to generate the data.

ruption in the maps generated by the binning of fewer spikes, as seen in Figure 3.

In terms of maximizing the rate of topographic error reduction, for the parameter values used in the simulations of our plasticity model, bins of 64 spikes appear to be approximately optimal. If we bin larger numbers of spikes per iteration (96 or 128 spikes, for example; data not shown), we find that the error curves initially follow very closely that for 64 spikes in Figure 4 but eventually depart from it, leaving higher residual errors. This behavior reflects the way in which the spike bin size affects the correlations in the afferent activity patterns experienced by the developing neuronal network. Once afferent correlations (and hence spike bin size) exceed some threshold, competition between spatially neighboring afferent cells begins to break down or becomes tempered, although competition between more distant afferent cells remains unchanged. Hence, a gross, overall retinotopy continues to emerge rapidly, but topographic refinement at a fine-grained level does not occur because nearby afferents' activities are too well correlated. For a sufficiently large number of spikes per bin, no error reduction will occur at all, corresponding to a complete breakdown of competitive dynamics.

For bins of 32 spikes and 10,000 iterations in each case, we examine the dependence of the final maps on the parameter  $\beta$ , which determines the topographic bias in the initial, unrefined mapping between the afferent and




Figure 5: Final states of the topographic mappings between the neuromorphic chip's retinal array and a target sheet for three different values of the initial topographic bias parameter  $\beta$ . The value of  $\beta$  used to generate a map is indicated immediately above the map. In all cases, 32 spikes are binned per iteration.

target sheets. All the data above are generated for a value  $\beta = 0.5$ , which falls midway between no initial bias,  $\beta = 0.0$ , giving a completely random initial map, and maximal bias  $\beta = 1.0$ , giving a well-ordered initial map. Figure 5 shows representative examples of the maps generated with smaller values of  $\beta$ , so that the initial topographic bias is smaller. For  $\beta = 0.4$  and  $\beta = 0.3$  (data not shown), the final maps are essentially identical to those generated with  $\beta = 0.5$ . For  $\beta = 0.2$ , we see a well-ordered map, although distortions particularly around the edges of the map are beginning to appear. For  $\beta = 0.1$ , these edge distortions become more significant, but the center of the map remains well structured. For  $\beta = 0.0$ , corresponding to an initially completely random map, a few small patches of locally refined topography emerge, but there is no global order to the map. Figure 6 shows the change in the overall topographic error in these maps during development for these three smaller values of  $\beta$ . The topographically refined patches for  $\beta = 0.0$  explain why the topographic error actually increases over time in this case: the absence of global order means that synaptic connections may be removed from areas of lower topographic error in favor of higher correlations in areas of greater topographic error.

Finally, we examine the robustness of our results to noise. For all data presented above, spontaneous firing of the silicon retina has been completely suppressed, so that pixels generate a spike only in response to a moving edge. Turning up the spontaneous firing rate introduces noise into the captured activity patterns, allowing us to determine how robust our results are to such noise. We bin 64 rather than 32 spikes in order to allow a reasonable number of both edge and spontaneous events to be caught during each capture phase, and set  $\beta$  to its usual value,  $\beta = 0.5$ . In Figure 7, we show typical examples of captured events, in response to a horizontally moving vertical edge. The percentage figure associated with each set of data gives an estimate of the amount of noise in the data, calculated simply as the ratio of the number of spontaneous to total spike events per bin. Figure 8 shows




Figure 6: The change in the overall topographic error in the development of the topographic maps, calculated as discussed in the text, plotted against iteration number, for the same three data sets used to generate Figure 5. The number attached to each curve indicates the value of  $\beta$  used to generate the data.




Figure 7: Representative examples of captured spikes (64 spikes binned per iteration) in response to a moving edge in the presence of spontaneous firing from the chip's pixels. Each  $16 \times 16$  array of squares represents the array of pixels on the chip, a black square denoting a pixel that has spiked during the binning period and a white one a pixel that has not spiked. The percentage figure above each array gives an estimate of the level of spontaneous firing for that data set.

the maps generated in the presence of spontaneous activity from the chip after 10,000 iterations. For a spontaneous firing rate of  $\sim 40\%$  (so that  $\sim 40\%$  of active pixels are not associated with an edge), the final map, although exhibiting slight distortions, is quite good. The topographic error remaining in this map is  $\sim 0.5$  (in units of cell spacings), which compares very favorably to the results above despite the high level of noise. Even for a spontaneous firing rate of  $\sim 50\%$ , the final map still possesses well-developed structure,




Figure 8: Final states of the topographic mappings between the neuromorphic chip's retinal array and a target sheet for three different levels of spontaneous activity on the chip. The spontaneous activity level present in generating a map is indicated immediately above the map.




Figure 9: The change in the overall topographic error in the development of the topographic maps, calculated as discussed in the text, plotted against iteration number, for the same three data sets used to generate Figure 8. The number attached to each curve indicates the spontaneous activity level used to generate the data.

although the distortions are much larger, giving a remaining topographic error of  $\sim 1.0$ . For  $\sim 60\%$  spontaneous firing, however, the final map has barely unfolded, the remaining topographic error being  $\sim 4.5$ , close to its initial value. Figure 9 shows the change in the overall topographic error in these maps during development for the data sets shown in Figure 8.

## 5 Discussion

---

We have shown that a silicon retina, responding to a moving bar generated on an LCD display, can be used, in conjunction with a biologically inspired neurotrophic model of neuronal development, to drive the topographic refinement of afferent projections to a sheet of target neurons and establish a nearly perfect one-to-one mapping between afferent and target cells. The emergence of a refined topographic map occurs robustly, largely independent of the size of spike bins provided that the binned data contain sufficient structure in a direction perpendicular to the bar's edge. However, larger spike bins do lead to greater rates of topographic error reduction, although the final errors associated with the mature maps do not exhibit much variation as a function of spike bin size.

For our choice of parameters, bins of 64 spikes give roughly optimal performance, resulting in the greatest rate of topographic error reduction. Larger numbers of spikes per bin begin to correlate the activities of nearby afferents too strongly, resulting in a breakdown or tempering of competitive dynamics between spatially neighboring afferent cells above some threshold in the strength of afferent correlations. The precise location of this threshold depends in part on the function  $\Delta_{xy}$ , which characterizes the diffusion of NTF between target cells. Narrowing  $\Delta_{xy}$  increases the threshold but weakens the coupling between target cells so that there is less tendency for the receptive fields of nearby target cells to be similar. Overall, therefore, while narrowing  $\Delta_{xy}$  will increase this threshold, it will also tend to disrupt the emergence of topography.

Our results, generated from a real-world device that captures some of the key features of retinal processing, compare very favorably with results obtained by simulating afferent input (Elliott & Shadbolt, 1999). This is despite the intrinsic noise present in all real-world devices—for example, the occasional failure of some pixels to respond to the stimulus and the temporal jitter of the responses. Indeed, in many respects, the system presented here performs better than earlier simulations. In some earlier work (Elliott & Shadbolt, 1999), we considered afferent input to consist of randomly activated pixels smeared by a gaussian smoothing function (cf. Goodhill, 1993). For the same learning rate used here, that system required well in excess of 100,000 iterations in order to generate a refined topographic map (see figure 6 in Elliott & Shadbolt, 1999). In contrast, our study here requires only on the order of 10,000 iterations to produce a mature map when coupled to the silicon retina. It may be argued that our stimulus (a moving bar) is rather unnatural. However, similar differences between the rates of development in systems with and without preprocessing of the raw optical image are also seen in robotic studies (unpublished observations). Retinal preprocessing of the raw optical image therefore appears, from a developmental point of view, to be of considerable advantage, allowing maps to achieve maturity more quickly than otherwise possible.

Spontaneous activity is a key feature of real nervous systems that allows, for example, the encoding of negative quantities by mechanisms that reduce spontaneous activity. Although we have not considered such encoding here, any model of neuronal development must be robust to the presence of spontaneous activity. By turning up the level of spontaneous activity on the silicon retina, we have shown that reasonable maps are generated even in the presence of high levels of spontaneous firing, with up to approximately half of all spike events not being stimulus related.

The used retina chip proves to be a suitable input sensor for the plasticity model. It provides the possibility of varying such parameters as refractory period and spontaneous firing rate of the afferent neurons and therefore allows us to evaluate the performance of the model under different input conditions. In the nonbursting mode, the retina chip implements very sparse coding, which makes the topographic refinement process more efficient. If, in addition, spontaneous firing is suppressed, the output signal is robust and reproducible with an error rate (number of missed and spurious events divided by total number of events) of approximately 5%.

## 6 Conclusion

---

The bar-stimulated silicon retina coupled to a neurotrophic model is a comparatively simple system that has been used here to demonstrate, for the first time, that a biologically inspired model of neuronal plasticity and a neuromorphic chip can operate successfully in unison. Future work might see the mounting of the chip on a mobile robotic platform for more realistic patterns of stimulation, the activation of both the ON and the OFF pathways in the chip and an examination of their segregation on some target structure, and using two such chips in a binocular configuration to examine the development of structures such as ocular dominance columns.

## Acknowledgments

---

T.E. thanks the Royal Society for the support of a University Research Fellowship. J.K. was supported in part by the Swiss National Foundation Research SPP grant. We thank David Lawrence of the Institute of Neuroinformatics for his invaluable help with interfacing the chip to the PC. We thank the Telluride 2000 Neuromorphic Engineering Workshop for having facilitated this work.

## References

---

- Boahen, K. (1999). Retinomorphic chips that see quadruple images. In *Proc. 7th Int. Conf. on Microelectronics for Neural, Fuzzy, and Bio-Inspired Systems* (pp. 12–20). Los Alamitos, CA: IEEE Computer Society Press.
- Boahen, K. A. (2000). Point-to-point connectivity between neuromorphic chips using address events. *IEEE Trans. Circuits and Systems II*, 47, 416–434.



- Deiss, S. R., Douglas, R. J., & Whatley, A. M. (1998). A pulse-coded communications infrastructure for neuromorphic systems. In W. Maass & C. M. Bishop (Eds.), *Pulsed neural networks* (pp. 157–178). Cambridge, MA: MIT Press.
- Delbrück, T., & Mead, C. A. (1994). Analog VLSI adaptive logarithmic wide-dynamic range photoreceptor. In *Proc. 1994 IEEE Int. Symp. on Circuits and Systems* (pp. 339–342). Piscataway, NJ: IEEE Press.
- Elliott, T., Maddison, A. C., & Shadbolt, N. R. (2001). Competitive anatomical and physiological plasticity: A neurotrophic bridge. *Biological Cybernetics*, *84*, 13–22.
- Elliott, T., & Shadbolt, N. R. (1998a). Competition for neurotrophic factors: Mathematical analysis. *Neural Computation*, *10*, 1939–1981.
- Elliott, T., & Shadbolt, N. R. (1998b). Competition for neurotrophic factors: Ocular dominance columns. *Journal of Neuroscience*, *18*, 5850–5858.
- Elliott, T., & Shadbolt, N. R. (1999). A neurotrophic model of the development of the retinogeniculocortical pathway induced by spontaneous retinal waves. *Journal of Neuroscience*, *19*, 7951–7970.
- Etienne-Cummings, R., & Van der Spiegel, J. (1996). Neuromorphic vision sensors. *Sensors and Actuators A, SNA056*, 19–29.
- Fukushima, K., Yamaguchi, Y., Yasuda, M., & Nagata, S. (1970). An electronic model of the retina. *Proc. IEEE*, *58*, 1950–1951.
- Goodhill, G. J. (1993). Topography and ocular dominance: A model exploring positive correlations. *Biological Cybernetics*, *69*, 109–118.
- Higgins, C., & Koch, C. (2000). A modular multi-chip neuromorphic architecture for real-time visual processing. *Journal of Analog Integrated Circuits and Signal Processing*, *26*, 195–211.
- Indiveri, G., Mürer, R., & Kramer, J. (2001). Active vision using an analog VLSI model of selective attention. *IEEE Trans. Circuits and Systems II*, *48*, 492–500.
- Indiveri, G., Whatley, A. M., & Kramer, J. (1999). Reconfigurable neuromorphic VLSI multi-chip system applied to visual motion computation. In *Proc. 7th Int. Conf. on Microelectronics for Neural, Fuzzy, and Bio-Inspired Systems* (pp. 37–44). Los Alamitos, CA: IEEE Computer Society Press.
- Kramer, J. (2002a). *An integrated optical transient sensor*. Submitted for publication.
- Kramer, J. (2002b). An on/off transient imager with event-driven, asynchronous read-out. In *Proc. 2002 IEEE International Symposium on Circuits and Systems* vol. II (pp. 165–168). Piscataway, NJ: IEEE Computer Society Press.
- Lazzaro, J., Wawrzynek, J., Mahowald, M., Sivilotti, M., & Gillespie, D. (1993). Silicon auditory processors as computer peripherals. *IEEE Trans. Neural Networks*, *4*, 523–528.
- Liu, S.-C., Kramer, J., Indiveri, G., Delbrück, T., Burg, T., & Douglas, R. J. (2001). Orientation-selective aVLSI spiking neurons. *Neural Networks*, *14*, 629–643.
- Mahowald, M. A. (1991). Silicon retina with adaptive photoreceptors. *Proc. SPIE*, *1473*, 52–58.
- Mahowald, M. A. (1994). *An analog VLSI system for stereoscopic vision*. Norwell, MA: Kluwer.
- McAllister, A. K., Katz, L. C., & Lo, D. C. (1999). Neurotrophins and synaptic plasticity. *Annual Review of Neuroscience*, *22*, 295–318.
- Mead, C. (1989). *Analog VLSI and neural systems*. Reading, MA: Addison-Wesley.

- Meyer, R. L. (1983). Tetrodotoxin inhibits the formation of refined retinotopography in goldfish. *Developmental Brain Research*, 6, 293–298.
- Moini, A. (1999). *Vision chips*. Norwell, MA: Kluwer.
- Mortara, A., & Vittoz, E. A. (1994). A communication architecture tailored for analog VLSI neural networks: Intrinsic performance and limitations. *IEEE Trans. Neural Networks*, 5, 459–466.
- Schmidt, J. T., & Edwards, D. L. (1983). Activity sharpens the map during the regeneration of the retinotectal projection in goldfish. *Brain Research*, 209, 29–39.
- Schmidt, J. T., & Eisele, L. E. (1985). Stroboscopic illumination and dark rearing block the sharpening of the regenerated retinotectal map in goldfish. *Neuroscience*, 14, 535–546.
- Sirosh, J., & Miikkulainen, R. (1997). Topographic receptive fields and patterned lateral interactions in a self-organizing model of the primary visual cortex. *Neural Computation*, 9, 577–594.
- Venier, P., Mortara, A., Arreguit, X., & Vittoz, E. A. (1997). An integrated cortical layer for orientation enhancement. *IEEE J. Solid-State Circuits*, 32, 177–186.

---

Received October 23, 2001; accepted April 18, 2002.