University of Pennsylvania
ScholarlyCommons

Departmental Papers (BE)

Department of Bioengineering

September 2005

Balancing Guidance Range and Strength Optimizes Self-Organization by Silicon Growth Cones

Brian Taba
University of Pennsylvania

Kwabena Boahen
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/be_papers

Recommended Citation

Taba, B., & Boahen, K. (2005). Balancing Guidance Range and Strength Optimizes Self-Organization by Silicon Growth Cones. Retrieved from https://repository.upenn.edu/be_papers/71

Postprint version. Published in *Lecture Notes in Computer Science*, Volume 3697, Artificial Neural Networks: Formal Models and Their Applications-ICANN 2005, 2005, pages 1027-1034.
Publisher URL: <http://dx.doi.org/10.1007/11550907>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/be_papers/71
For more information, please contact repository@pobox.upenn.edu.

Balancing Guidance Range and Strength Optimizes Self-Organization by Silicon Growth Cones

Abstract

We characterize the first hardware implementation of a self-organizing map algorithm based on axon migration. A population of silicon growth cones automatically wires a topographic mapping by migrating toward sources of a diffusible guidance signal that is released by postsynaptic activity. We varied the diffusion radius of this signal, trading strength for range. Best performance is achieved by balancing signal strength against signal range.

Comments

Postprint version. Published in *Lecture Notes in Computer Science*, Volume 3697, Artificial Neural Networks: Formal Models and Their Applications-ICANN 2005, 2005, pages 1027-1034.

Publisher URL: <http://dx.doi.org/10.1007/11550907>

Balancing Guidance Range and Strength Optimizes Self-Organization by Silicon Growth Cones

Brian Taba and Kwabena Boahen

University of Pennsylvania, Philadelphia PA 19104, USA,
{btaba, boahen}@seas.upenn.edu,
WWW home page: <http://www.neuroengineering.upenn.edu/boahen>

Abstract. We characterize the first hardware implementation of a self-organizing map algorithm based on axon migration. A population of silicon growth cones automatically wires a topographic mapping by migrating toward sources of a diffusible guidance signal that is released by postsynaptic activity. We varied the diffusion radius of this signal, trading strength for range. Best performance is achieved by balancing signal strength against signal range.

1 Introduction

Neuromorphic engineers seek to migrate the computational efficiency of neurobiological systems into engineering applications by building silicon chips that faithfully reproduce neural function. For example, silicon retinæ now emulate up to thirteen different cell types to encode distinct stimulus properties in four types of spiking ganglion cell output [1]. This level of detail is possible because the biological retina is relatively accessible experimentally. However, no comparable circuit description exists for higher order processing centers.

Where circuit details are lacking, a viable alternative is to leverage recent rapid progress in developmental neuroscience to design systems that can self-organize their own connectivity. During development, neurons wire themselves into their mature networks by sprouting axonal and dendritic precursors called *neurites*. Each neurite is tipped by a sensory structure called a *growth cone* that uses local chemical cues to guide the elongating neurite. Growth cones move by continually extending and retracting finger-like appendages called *filopodia* whose dynamics are biased by diffusible ligands [2].

Adopting this developmental approach, we previously described the first self-organizing map chip that is based on neurite outgrowth [3]. Our chip's neuromorphic cells are equipped with growth cone circuits that enable them to wire themselves into a mature network automatically, without an explicit blueprint.

Although analogous self-organizing algorithms have been successfully implemented digitally, previous analog implementations required high precision components that are expensive in chip area (e.g., [4],[5]). By contrast, neurobiological

Fig. 1. Neurotropic attraction model. **a.** A stimulus coactivates a contiguous patch of source cells α , which fire spikes down their axons to induce their growth cones $c(\alpha)$ to excite nearby target cells x . **b.** Active target cells release neurotrophin $n(x)$ into the extracellular medium. **c.** Neurotrophin spreads laterally before being bound by active growth cones, which measure the direction of the local gradient, indicated by the arrow. **d.** Active growth cones climb the gradient by displacing other growth cones.

systems achieve tolerable performance with low precision components. By mimicking their growth cones, our neuromorphic approach realizes the low power benefit of analog implementation without the associated cost in chip area.

We previously described the design of the Neurotrope1 chip and its ability to self-organize topography [3]. Here, we extend those results by illustrating how topographic self-organization depends on the range of the diffusible growth cone guidance signal. Section 2 introduces the learning algorithm, Section 3 describes the neuromorphic implementation, and Section 4 presents chip measurements from a topographic self-organization task. We find that performance on this task is optimized by balancing signal strength with signal range.

2 Neurotropic Axon Guidance

In our implementation, an active growth cone’s filopodia bind a diffusible guidance signal called *neurotrophin* that is released by active target cells (Fig. 1). The growth cone measures the local concentration gradient by comparing the rates at which its filopodia accumulate neurotrophin, and maximizes its neurotrophin uptake by climbing the measured gradient. Since filopodial neurotrophin binding is gated by presynaptic activity and neurotrophin release is gated by postsynaptic activity, neurotropic gradient ascent implements a Hebbian update rule under which cells that fire at the same time wire to the same place.

This algorithm for the self-organization of connections from one layer of neurons to another is formally described as follows. Source cells occupy nodes of a regular two-dimensional (2D) lattice, while growth cones and target cells occupy nodes of separate regular 2D lattices that are interleaved. Nodes are indexed by their positions in their respective layers, labeled by Greek letters in the source layer (e.g. $\alpha \in \mathbb{Z}^2$) and by Roman letters in the target layer (e.g. $x, c \in \mathbb{Z}^2$).

Target cell x fires at a rate $a_{TC}(x)$ that is proportional to its excitation:

$$a_{TC}(x) = \sum_{\alpha} a_{SC}(\alpha) A(c(\alpha) - x)$$

Fig. 2. Axon remapping. **a.** Cell bodies tag their spikes with their source layer addresses, which are decoded through the forward lookup table into the target layer addresses of their growth cones. **b.** Migrating growth cones decode their target layer addresses through the reverse lookup table into the source layer addresses of their cell bodies, which index their entries in the forward lookup table. **c.** Growth cones swap locations by modifying their four entries in the forward and reverse lookup tables.

where $a_{SC}(\alpha)$ is the activity of source cell α , and $A(c(\alpha) - x)$ is the branch density of the excitatory arbor elaborated by the growth cone-tipped axon trunk projected by source cell α to target site $c(\alpha)$. $A(c(\alpha) - x)$ decreases with target layer separation $\|c(\alpha) - x\|$.

Active target cells release neurotrophin, which spreads laterally until consumed by constitutive scavenger processes, establishing a concentration

$$n(x) = \sum_y a_{TC}(y) N(x - y)$$

that sums contributions from all active target cells y , weighted by a spreading function $N(x - y)$ that decreases with $\|x - y\|$.

An active growth cone located at $c(\alpha)$ uses a local winner-take-all function to identify the node

$$c'(\alpha) = \arg \max_{x \in \mathcal{C}(c(\alpha))} n(x)$$

that contains the most neurotrophin during the growth cone's activity, where $\mathcal{C}(x)$ contains x and its nearest neighbors. Upon identifying $c'(\alpha)$, the growth cone swaps nodes with the growth cone currently occupying $c'(\alpha)$, moving the entire axon arbor with it, thereby increasing its neurotrophic uptake while maintaining a constant axon density. Growth cones initiate swaps independently, at a rate $\lambda(\alpha) \propto a_{SC}(\alpha) \max_{x \in \mathcal{C}(c(\alpha))} n(x)$. Swaps are serviced asynchronously, in the order in which they arrive.

Software simulation of similar equations yields self-organized topography when driven with appropriate presynaptic correlations [6]. In this paper, we probe the effect of varying the width of $N(x - y)$ in a hardware implementation.

3 Neuromorphic Implementation

We implemented this model in a full custom VLSI chip that interleaves a 24×40 array of growth cone circuits with a 24×20 array of target cell circuits and a neurotrophin spreading network. The Neurotrope1 chip was fabricated through MOSIS using the TSMC $0.35\mu\text{m}$ process, and is 11.5 mm^2 in area. Axons are

Fig. 3. a. Neurotrophe1 system. **b.** Neurotrophe1 cell mosaic. The neurotrophin spreading network (grey) is interleaved with the array of target cell circuits (TC). Each growth cone circuit (GC) occupies one node of the spreading network and extends filopodia to the three adjacent nodes, expressing neurotrophin receptors (black) at all four nodes.

implemented as entries in an off-chip lookup table that are updated by Neurotrophe1 activity, as described in Subsection 3.1. Subsection 3.2 explains how Neurotrophe1 computes these updates from neurotrophin in the spreading network and Subsection 3.3 explains how the spreading network circuit shapes $N(x - y)$.

3.1 Axon Remapping

Chips in our system exchange spikes encoded in the address-event representation (AER) [7], an asynchronous protocol that pools spikes from all the cells on the same chip onto a shared data link. AER tags each spike with the address of its originating cell body for transmission off-chip onto the data link. Each spike is filtered through a *forward lookup table* that translates the source layer address of its origin into the target layer address of its destination (Fig. 2a). The receiving chip uses the delivered address to route the spike to the appropriate target.

An axon is remapped to a new target site simply by updating its entry in the lookup table. Updates are requested by Neurotrophe1 growth cone circuits and communicated as address-events to a Ubicom ip2022 microcontroller for processing. Each update request identifies a pair of adjacent growth cones whose target layer addresses are to be swapped. These addresses are translated through a *reverse lookup table* that decodes target layer growth cone addresses into the source layer cell body addresses that index the forward lookup table (Fig. 2b). Axons migrate by modifying their entries in each lookup table (Fig. 2c).

Both lookup tables are stored in a random access memory (RAM) chip (Fig. 3a). The ip2022 processes the axon updates computed by the growth cone circuits and overwrites the appropriate RAM cells. The ip2022 also supports a universal serial bus (USB) over which a computer can write to and read from the RAM. Any AER-compliant device can implement the source cell population; in this paper, we simulate the source layer with a second ip2022.

3.2 Axon Updates

Axon updates are computed by Neurotrophe1 using the growth cone circuits described in [3]. Each growth cone occupies one node of the neurotrophin spread-

Fig. 4. a. Neurotrophin spreading circuit. **b.** SPICE simulation of node voltages V_r during a neurotrophin release pulse, as a function of distance r from the release site. $V_{spread} = 2.0V$.

ing network and extends filopodia to the three adjacent nodes, expressing neurotrophin receptors at all four sites (Fig. 3b). Neurotrophin is represented as charge in the spreading network. During a presynaptic spike, receptor circuits tap charge from their nodes of the spreading network and store it on separate capacitors. The first capacitor voltage to cross a threshold triggers an update request and resets all the growth cone’s capacitors. If the receptor at the growth cone body won the race to threshold, no action is required and the request is dropped. If one of the filopodial receptors won, the growth cone transmits a request off-chip to swap places with the growth cone currently occupying the winning node.

Gradient measurements are noisy, so we require the ip2022 to process multiple requests from the same pair of adjacent growth cone circuits before actually executing the swap. We maintain a running count in the RAM of accumulated requests for each pair, and only execute a swap after its count exceeds a preprogrammed threshold. An executed swap resets the counts of five affected growth cone pairings among the two growth cone circuits and their four neighbors. The effect is to screen out spurious update requests and brake growth cone velocity.

3.3 Neurotrophin Spreading Circuit

The neurotrophin spreading function $N(x - y)$ is shaped by the transistor circuit in Fig. 4a, which implements a neighborhood function similar to [8]. Transistors M1 and M2 are gated by a facilitation circuit that only allows charge to be injected into the spreading network during a burst of postsynaptic activity, since $\sim tb$ requires several consecutive spikes to bring it low enough to open M1, although $\sim ts$ allows individual spikes to open M2. Similarly, a growth cone can only sample charge during bursts of presynaptic activity through transistors M3-5. (The sampled current I_{NT} is limited by bias V_b .)

Between spikes, charge spreads laterally through the unnumbered pFETs until shunted to ground through one of the unnumbered nFETs located at each lattice node. The two gate biases V_{spread} and V_{shunt} control the distance to which

Fig. 5. Topographic self-organization. **a.** Source layer maps generated from random initial projection at time steps n . Source cells are colored by the target layer locations of their growth cones. **b.** Source layer maps generated from perfect initial projection.

charge can spread from its injection site. V_{shunt} gates the shunt transistor hosted by each lattice node. This transistor enters saturation as the node voltage rises and becomes a constant current sink. Incoming current in excess of this sink charges the node capacitance until the node voltage exceeds V_{spread} , allowing the remaining current to flow outward to as many nodes as needed to sink all the injected current. Larger shunt currents sink the injected current closer to the release site, so increasing V_{shunt} reduces the spreading range (Fig. 4b).

4 Topographic Self-Organization

To induce the growth cone population to self-organize a topographic image of the source layer, we drive them with bursts of presynaptic spikes from contiguous patches of coactive source cells. Each patch consists of a randomly selected source cell and its three immediately adjacent neighbors, a presynaptic activity correlation kernel with sufficient structure to instruct topographic ordering [9].

We trace the evolution of the growth cone population by sampling the contents of the forward lookup table every five minutes, an interval long enough to present each of the 24×20 possible patches about once per sample. Starting from a random projection at $n = 0$ (Fig. 5a), small chunks of local topography are visible by $n = 20$. These local crystals eventually merge into the larger, more global clusters observed at $n = 200$. A similar endstate is reached by a perfect initial projection as it relaxes to a more sustainable topographic level (Fig. 5b).

We evaluate performance quantitatively by defining an *order parameter* $\Phi^{(n)}$ whose value measures the relative topography at a source cell's location at time n . One definition of topography is that adjacent source cells extend axons to adjacent target sites, so we choose $\Phi^{(n)}$ for a given source cell to be the average target layer distance separating its growth cone from those of its three nearest source layer neighbors in the n th sample. (In a perfectly topographic projection, $\Phi^{(n)} = 1$.) The population average $\langle \Phi^{(n)} \rangle$ converges to similar values from both random and perfect initial projections (Fig. 6a), so sustainable topography is not limited by the initial conditions, but by some intrinsic property of the system. This intrinsic limit depends on the neurotrophin spreading range, which we control with the shunt bias V_{shunt} . At equilibrium, $\langle \Phi^{(n)} \rangle$ is minimized by an intermediate value of V_{shunt} that corresponds to an optimal spreading range (Fig. 6b).

To investigate this optimal spreading range, we examine the probability $P^{(n)}(\Phi)$ that a pair of growth cones projected by adjacent source cells is separated in the target layer by a distance Φ in the n th axon projection sample.

Fig. 6. **a.** Order parameter Φ evolution from random (black) and perfect (grey) initial projections. **b.** Average equilibrium order parameter $\langle \Phi^{(100)} \rangle$ dependence on neurotrophin spreading range, as controlled by V_{shunt} . **c.** Instantaneous order parameter probability $P^{(n)}(\Phi)$ measured during evolution from random initial projection. Grey: $P^{(0)}(\Phi)$; black: $P^{(200)}(\Phi)$.

We construct $P^{(n)}(\Phi)$ from the relative frequency with which each value of Φ is observed within the n th sample of the population ensemble of growth cone positions downloaded from the RAM. Perfect growth cone guidance elicits a $P^{(\infty)}(\Phi)$ that is 1 at $\Phi = 1$ and 0 elsewhere. Unguided growth cones are distributed in proportion to the number of target sites located a distance Φ from a given attractor (grey in Fig. 6c). This number initially increases as $2\pi\Phi$, but falls to zero at large Φ , since growth cone separations cannot exceed the finite array dimensions. The actual distribution achieved by our physical system lies somewhere between these two extremes (black in Fig. 6c).

The optimal spreading range balances a neurotrophin release site's ability to hold growth cones with its ability to attract them (Fig. 7a). For short spreading ranges, $P^{(100)}(\Phi)$ resembles a random distribution except for a small peak at low Φ that captures growth cones that manage to fall within each other's detection horizon. Increasing the spreading range allows growth cones to lure coactive peers from greater distances, siphoning $P^{(100)}(\Phi)$ into a peak at low Φ . However, as the spreading range approaches the array size, the ability of larger attraction basins to rope in more distant growth cones is outweighed by the inability of captured growth cones to localize within the basin. Consequently, for long spreading ranges, $P^{(100)}(\Phi)$ broadens and shifts toward the random distribution.

We dissociate the attraction and confinement aspects of guidance by tracking the topographic evolution of the 25% of growth cones that are furthest from and closest to their topographic neighbors (black dashed lines in Fig. 7b). Proximate growth cones do better at short spreading ranges, but distant growth cones do better at long spreading ranges. The optimal spreading range improves the performance of both proximate and distant growth cones.

5 Conclusions

We characterized the first hardware implementation of a self-organizing map algorithm based on axon migration. We varied the range of the neurotropic signal that guides the silicon growth cone population to automatically wire a topographic map when driven by correlated activity. Long neurotrophin spreading ranges attract distant growth cones but cannot hold them to their targets, while

Fig. 7. Optimal neurotrophin spreading range. Left: short ($V_{\text{shunt}} = 0.40V$); middle: medium ($0.30V$); right: long ($0.20V$). **a.** Equilibrium distribution $P^{(100)}(\Phi)$. Dashed line indicates $\langle \Phi^{(100)} \rangle$. **b.** Evolution of $\langle \Phi^{(n)} \rangle$ within full growth cone population (black solid line) and growth cone subpopulations with the 25% highest and lowest $\Phi^{(n)}$ values (black dashed lines). Grey lines plot the corresponding $\langle \Phi \rangle$ for a random distribution.

short neurotrophin spreading ranges hold growth cones to nearby targets but cannot rescue distant growth cones. This tradeoff between recovery and confinement implies that future systems should address these aspects separately, perhaps through other developmental mechanisms like synaptogenesis, to consolidate accurately placed growth cones, and pruning, to eliminate outliers.

Acknowledgment

This work was supported by the David and Lucille Packard Foundation and the NSF/BITS program (EIA0130822). B.T. received support from the Dolores Zohrab Liebmann Foundation.

References

1. Zaghloul, K.A., Boahen, K.: Optic nerve signals in a neuromorphic chip I: Outer and inner retina models. *IEEE Trans. Bio-Med. Eng.* **51**:4 (2004) 657–666
2. Huber, A.B., Kolodkin, A.L., Ginty, D.D., Cloutier, J.-F.: Signaling at the growth cone: ligand-receptor complexes and the control of axon growth and guidance. *Annu. Rev. Neurosci.* **26** (2003) 509–563
3. Taba, B., Boahen, K.: Topographic map formation by silicon growth cones. *Advances in Neural Information Processing Systems* **15** (MIT Press, Cambridge, eds. Becker, S., Thrun, S., Obermayer, K.) 1163–1170
4. Melton, M., Phan, T., Reeves, D.S., Van den Bout, D.E.: The TInMANN VLSI chip. *IEEE Trans. Neural Networks* **3** (1992) 375–384
5. Porrmann, M., Witkowski, U., Ruckert, U.: A massively parallel architecture for self-organizing feature maps. *IEEE Trans. Neural Networks* **14** (2003) 1110–1121
6. Lam, S.Y.M., Shi, B.E., Boahen, K.A.: Self-organized cortical map formation by guiding connections. *Proc. 2005 IEEE Int. Symp. Circ. & Sys.* (in press)
7. Boahen, K.: Point-to-point connectivity between neuromorphic chips using address-events. *IEEE Trans. Circ. & Sys. II* **47** (2000) 416–434
8. Heim, P., Hochet, B., Vittoz, E.A.: Generation of learning neighbourhood in Kohonen feature maps by means of simple nonlinear network. *Electronics Letters* **27** (1991) 275–277
9. Miller, K.: A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between on- and off-center inputs. *J. Neurosci.* **14** (1994) 409–441