Hebbian Learning in Winner-Take-All Networks

Tristan Robitaille
Leonardo Martinelli
Mohamed Malek Abid

June 2nd, 2025

1 Introduction

Topographic maps are a canonical feature of sensory systems, preserving the spatial relationships of stimuli across neural projections. These maps emerge early in development, even before sensory experience, suggesting internal activity-dependent mechanisms. This isn't to say that there is no input to these neural circuits responsible for computation of sensory information; rather, similar to the Chomskyan and biolinguistic concept of the human faculty for language, this process can be interpreted as developing a structured base framework for vision that enables us to acquire vision more efficiently when we eventually see the external world. Retinal waves, spontaneous bursts of activity sweeping across the retina, are known to guide retinotopic organization via Hebbian plasticity before vision is functional [1, 2].

Theoretical work has shown that Hebbian learning, together with local excitation and global inhibition, can drive the emergence of spatially ordered maps through competitive self-organization [3]. Winner-take-all (WTA) networks operate on this fundamental principle: populations compete through shared inhibitory circuits, with the most strongly driven unit suppressing its neighbors, allowing the network to sharpen input representations and enforce exclusivity [4]. Our group selected this specific project due to our interest in biological learning; we were motivated by the desire to build an on-chip demonstration of developmental mechanisms like synaptogenesis and pruning, guided solely by unsupervised local activity.

Towards this goal, we implemented a biologically inspired Hebbian learning process — featuring separate growth and pruning phases — on a soft WTA architecture running on the Dynap-SE. Our aim was to evaluate whether such a system can form topographic maps and exhibit selective responses to input structure.

2 Methods

2.1 Tuning Neurons on Dynap-SE

The silicon neurons on the Dynap-SE operate in the sub-threshold regime and thus are subject to significant process variation and are sensitive to variables such as temperature and supply voltage. Before investigating learning in Winner-Take-All networks, we aimed to ensure that we were obtaining a neuronal response that roughly matches biological constraints with regards to time constants.

Specifically, by using an oscilloscope and modifying the DC current parameter in short pulses and observing a single neuron from the chip, we tuned the refractory period constant until we observed a refractory period of approximately 5 ms. Furthermore, observing the decay from our neuron, we tuned the membrane time constant ($\tau_{membrane}$) to result in a fall-time between 10–20 ms. Our final TAU1 = (5, 200), deviated from the default (7, 200) that we inherited from previous Winner-Take-All network projects. The magnitude of the DC_P parameter was selected such that the neural activity would be characteristic of a neuron receiving a subthreshold current when fitting with the decay, and above threshold to induce high spiking for the refractory period analysis.

2.2 Winner-Take-All Architecture

We implemented a soft winner-take-all (WTA) network as the output layer, comprising 16 excitatory populations, each consisting of 8 neurons, and a single global inhibitory population composed of 20 neurons. Each excitatory population is connected with AMPA synapses in a recurrent fashion and to its immediate neighboring populations. The first and last population are also connected to each other. All excitatory populations project to the global inhibitory population, which in turn provide uniform inhibitory feedback to all excitatory populations (Figure 1). The excitatory populations are connected to the inputs using 2 to 4 connections (randomly chosen). Small note: Because of the hardware limiting us to 64 inputs per neuron, we implemented a "half-all-to-all" network, where each even input is connected to each even WTA population and each odd input is connected to each odd WTA population. We chose this architecture for two principal reasons. First, the local excitatory connectivity enables the propagation of spatial proximity information across adjacent units, a process crucial for the self-organization of topographic maps [3]. Second, the global inhibition imposes a competitive dynamic that effectively suppresses spurious or diffuse activity, an essential mechanism during the early phases of learning when connectivity is still random and unrefined.

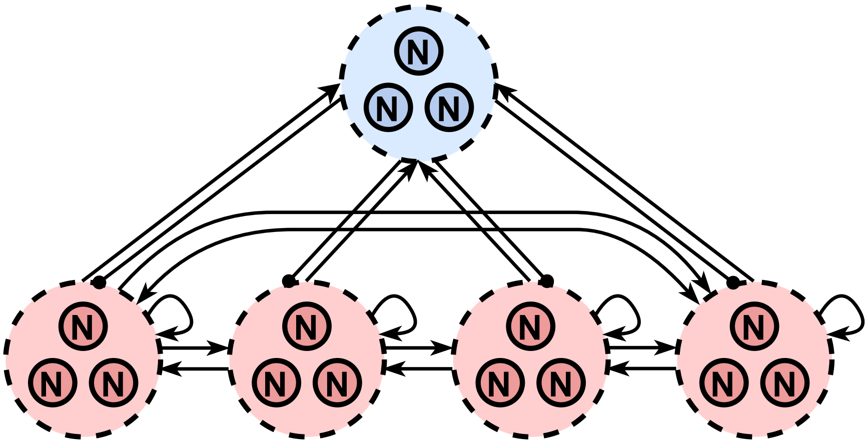

Figure 1: **Winner-Take-All Network.** A schematic representation of the network is reported with 4 populations instead of the actual 16, showing local and recurrent connectivity between excitatory populations and one-to-one connectivity between excitatory and inhibitory populations.

2.3 Traveling Wave Input to the Network

For the input layer, we considered an array of 16 spike generators in order to match the dimensionality of the output layer. The two layers are connected via a connectivity matrix indicating the number of parallel connections (AMPA synapses) between each spike generator and excitatory population in the WTA. Spike generators are sending Poisson trains of activity to the WTA layer. We defined a Gaussian spatial profile of activation, so the spike generator corresponding to the peak of the Gaussian will have maximum firing rate of 100 Hz. Thus, the firing rate r_i of each input neuron i was computed as:

$$r_i = A \cdot \exp\left(-\frac{(x_i - \mu(t))^2}{2\sigma^2}\right)$$

where A is the peak rate (e.g., 100 Hz), x_i is the neuron's position index, $\mu(t)$ is a time-dependent center that drifts across the input layer, and σ is the width controlling spread. To emulate the correlated activity waves observed in the retina, we imposed drift dynamics on the Gaussian input, thus having it travel across spike generators (Figure 2).

2.4 Hebbian Learning with Growth and Pruning Phases

We finally implemented a learning algorithm with the computer in the loop. Following the theoretical results [3], we chose an unsupervised rate-based Hebbian learning. Instead of imposing a global normalization procedure at each iteration we modified the classic Hebbian learning by distinguishing two different phases. A growth phase, emulating synaptogenesis, where the number of connections between the two layers just increases based on the correlation between the activities followed by a pruning phase, emulating synaptic pruning, where spurious connections are eliminated in the absence of correlation. During the growing phase, the number of connections increases by 2 if the product of the normalized input and output is

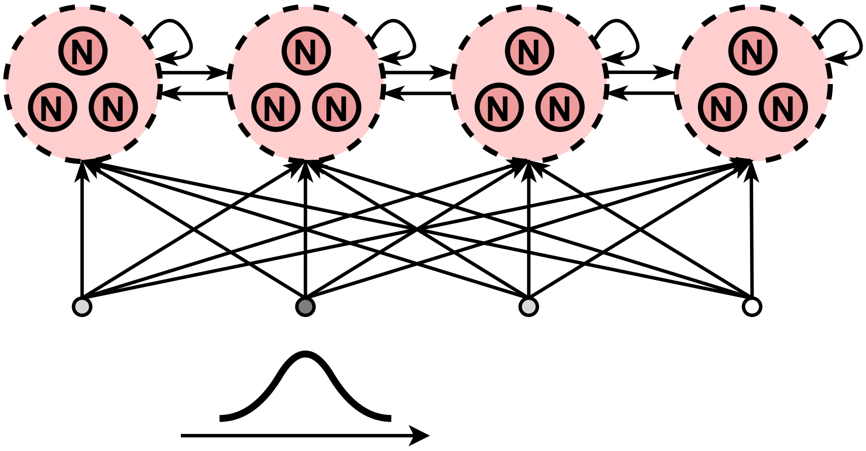

Figure 2: **Input Layer Connectivity.** The lower dots represent spike generators connected to the excitatory populations. Spike generators are activated according to a wave-like dynamics with a traveling Gaussian activation shape.

above 0.85. During the pruning phase, the number of connections decreases by 1 if the product of the normalized input and output is below 0.5.

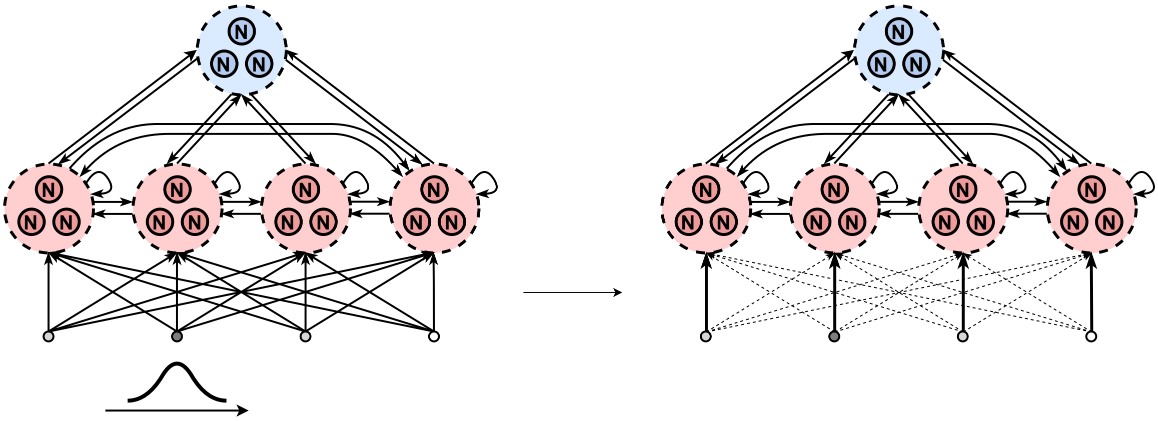

Figure 3: **Learning Topographic Maps.** A schematic representation of the full pipeline. On the left the network with traveling wave inputs with initial random connectivity. On the right the network after learning with strengthened topographic connectivity.

2.5 WTA Discriminative Properties

We tested the WTA network across two axes of variation: input amplitude, a and position, μ. In each trial, two Gaussian bumps were used as input Figures 4 and 5 summarize network responses in a series of experiments exploring different input amplitudes and spatial separations (a and μ, respectively, for each bump).

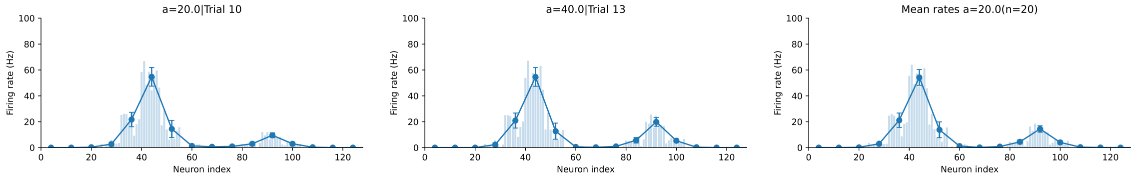

Figure 4: **Amplitude Sweep.** Representative firing rates for varying a_2, amplitude of the second bump, given a fixed $a_1 = 100$.

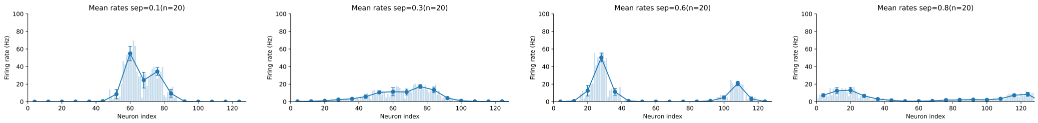

Figure 5: **Spatial Separation Sweep.** Spatial input separations affect the network's discriminative performance as we vary μ_1, μ_2 from $(0.4, 0.5) \rightarrow (0.1, 0.9)$ Note that since the first and last excitatory populations share lateral connectivity, this circular structure causes spatially adjacent inputs near the boundaries to interact as if they were neighboring in the center.

2.6 Learning Topographic Maps

We trained the network using the Hebbian-like learning rule described above by performing 3 growing phases followed by 3 pruning phases and recorded the number of connections between the input and the WTA populations as a connection matrix, as shown in Figure 6. We recorded the learning process and compiled a video showing the evolution of the connection matrix, available at this URL. As can be seen in the video, a diagonal of strong connections emerges in the connection matrix, which corresponds to the 1-to-1 connections between the input and the WTA populations. It can also be seen that some connections outside of the 1-to-1 diagonal are strengthened, especially for WTA populations 8 and 9.

3 Discussion

Our results demonstrate that Hebbian learning with distinct growth and pruning phases can produce coherent topographic mappings in a soft WTA network. After sufficient training, a clear one-to-one alignment emerges between input and output units, with almost all of the WTA excitatory clusters/populations specializing for specific spatial input positions. This supports earlier computational models by Willshaw and von der Malsburg that proposed competition plus correlation suffices for map formation [3].

However, one key observation is that some inherent variability was present in our silicon neurons. Some populations, notably 8 and 9, formed disproportionately strong connections across many input positions. These "columns" in the connection matrix likely reflect hyperexcitable neurons that dominated due to slightly lower leak or higher bias currents. [5].

Furthermore, despite successful learning of the desired mapping, our network struggled with competitive discrimination. In experiments with two spatially distinct Gaussian inputs, the weaker bump was not reliably suppressed. Ideally, a winner-take-all network should attenuate less dominant inputs. Here, however, both bumps were preserved in output activity, sometimes with the weaker input having higher gain. This suggests that the network acts more as a soft selector than a strict WTA, blurring input distinctions.

We hypothesized that increasing inhibitory strength might sharpen discrimination. However, increasing the connectivity of excitatory-to-inhibitory (e_i) links in the network yielded no noticeable improvement in ameliorating the issues with discrimination between two bumps in certain configurations, namely those with a small difference in amplitude or position . Thus we currently believe that issues with discrimination may stem from broader connectivity interactions, or potentially from mismatched excitatory loop strengths. It is also possible that additional tuning of lateral inhibition or other connectivity may be adequate for sharper selectivity.

Future directions include scaling to 2D structures, which would be a better model of our biological inspiration, integrating STDP to account for pre-post differences in spikes, and leveraging more advanced neuromorphic platforms such as Dynap-SE2 to explore the potential for leveraging dendritic computation to enhance our learning strategy.

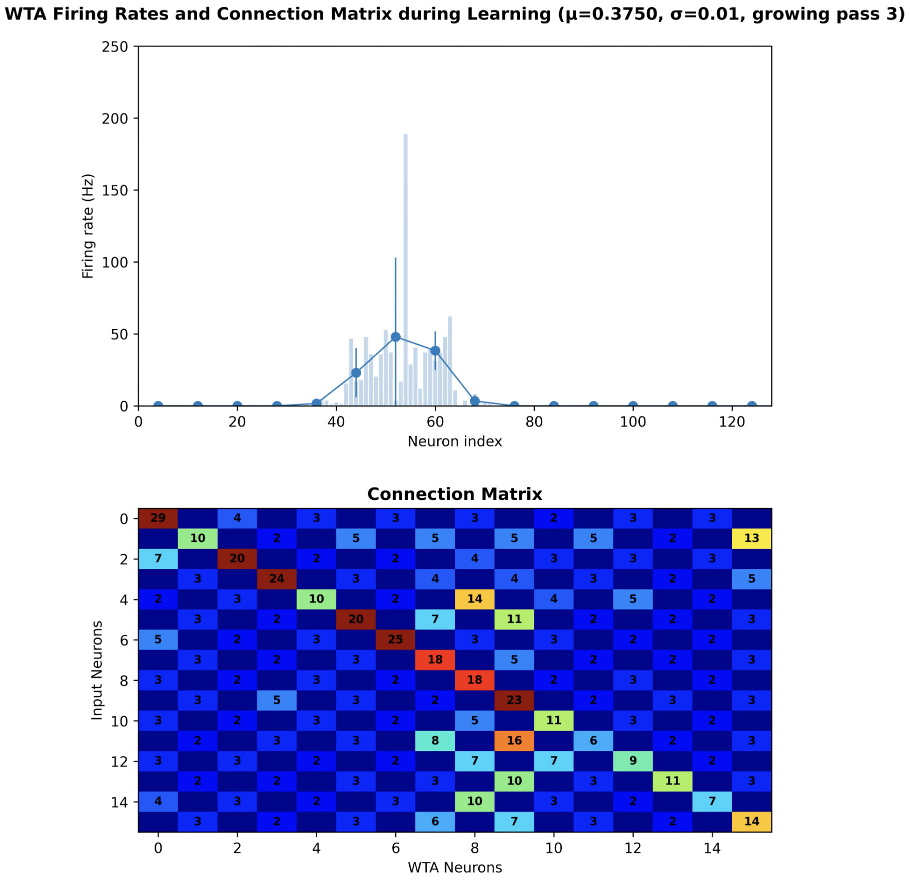

	0		2		4		6		8		10		12		14	
0	29		4		3		3		3		2		3		3	
		10		2		5		5		5		5		2		13
2	7		20		2		2		4		3		3		3	
		3		24		3		4		4		3		2		5
4	2		3		10		2		14		4		5		2	
		3		2		20		7		11		2		2		3
6	5		2		3		25		3		3		2		2	
		3		2		3		18		5		2		2		3
8	3		2		3		2		18		2		3		2	
		3		5		3		2		23		2		3		3
10	3		2		3		2		5		11		3		2	
		2		3		3		8		16		6		2		3
12	3		3		2		2		7		7		9		2	
		2		2		3		3		10		3		11		3
14	4		3		2		3		10		3		2		7	
		3		2		3		6		7		3		2		14

Figure 6: **Hebbian Learning.** A strong diagonal, representing the desired 1-to-1 topographic map emerges during Hebbian learning. Link to video.

4 Conclusion

We demonstrated that our biologically inspired Hebbian learning rule can, in a process incorporating distinct growth and pruning phases, lead to the learning of a self-organized topographic map in a soft WTA network. This was achieved on neuromorphic hardware (Dynap-SE) using structured input patterns that mimic developmental retinal waves. The result is a one-to-one alignment between input and output populations, consistent with our hypothesis and original goal.

However, while the network effectively learned spatial mapping, it failed to robustly suppress weaker competing inputs, highlighting a gap in our soft WTA network's behavior. We believe that if altering the excitatory-inhibitory population connectivity parameters is not sufficient to resolve this limitation, perhaps instead future attempts to combine the topographic learning with a better soft WTA could incorporate more adaptive lateral or inhibitory plasticity to enforce true winner selection.

Overall, this work contributes a concrete demonstration of structural self-organization on analog hardware. It lies at the intersection of systems neuroscience and neuromorphic computing, and opens the door to future extensions that incorporate spike timing, different spatial structures, and other forms of adaptation and learning.

References

[1] Evelyne Sernagor and Vandana Mehta. The role of early neural activity in the maturation of turtle retinal function. *The Journal of Anatomy*, 199(4):375–383, 2001.

[2] Carla J Shatz. Emergence of order in visual system development. *Journal of Physiology-Paris*, 90(3-4):141–150, 1996.

[3] D. J. Willshaw, C. Von Der Malsburg, and Hugh Christopher Longuet-Higgins. How patterned neural connections can be set up by self-organization. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 194(1117):431–445, 1976.

[4] Yanqing Chen. Mechanisms of winner-take-all and group selection in neuronal spiking networks. *Frontiers in computational neuroscience*, 11:20, 2017.

[5] Martin Boerlin, Tobi Delbruck, and Kynan Eng. Getting to know your neighbors: unsupervised learning of topography from real-world, event-based input. *Neural computation*, 21(1):216–238, 2009.