

On-chip unsupervised learning in Winner-Take-All networks of spiking neurons

Raphaela Kreiser, Timoleon Moraitis, Yulia Sandamirskaya, Giacomo Indiveri
Institute of Neuroinformatics
University of Zurich and ETH Zurich
Zurich, Switzerland
Email: raphaela.kreiser@uzh.ch

Abstract—The ability to learn re-occurring patterns in real-time sensory inputs in an unsupervised way is a key feature of neural networks that can enable them to carry out complex tasks directly, or to simplify the learning process of subsequent layers in powerful deep network configurations. Dedicated neuromorphic computing electronic systems can implement low-power real-time neural network inference engines. However, unsupervised on-line learning in these systems remains an open challenge. In this paper, we demonstrate spike-based unsupervised learning in a neuromorphic hardware device that has ultra low-power spiking neuron circuits and on-chip plasticity synapse circuits implemented with analog electronics. We configure populations of silicon neurons in a soft winner-take-all (WTA) network configuration, which enables them to learn the classification of different spike-rate input patterns in an unsupervised manner. We demonstrate the ability of this neuromorphic hardware to perform unsupervised learning of the pattern classification in real-time, and characterize its robustness as a function of network parameters and network configuration.

I. INTRODUCTION

Unsupervised learning is a powerful learning mechanism, used by biological systems when learning to classify objects in categories based on their similar properties. In machine learning, when data is unlabeled or no teacher to provide rewards is available, unsupervised learning can be used to infer hidden structures in the data [1], [2]. The self-organizing map (SOM) is an algorithm based on a competitive artificial neural network (ANN), originally introduced by Kohonen [3]. After training a SOM based on similarity between the already stored and new input vectors, clusters of neuronal nodes emerge in an unsupervised way that correspond to different data classes. This principle has been realised in a simulated spiking neural network for pattern classification [4] and the capacity of a neuromorphic device to store patterns in attractor-states, shaped by synaptic connections, was analyzed in [5]. Here, we introduce a simplified hardware neural architecture that uses a similar to SOM principle, which implements unsupervised clustering in real-time neuromorphic hardware. The neuromorphic hardware we use in this work comprises hybrid analog/digital circuits that directly emulate the biophysics of the biological processes they model [6] so that computation is distributed, event-driven, and massively parallel. In order to achieve robust computation, given the noisy and variable nature of analog circuits, adaptation, self-organisation, and learning are of paramount importance.

In this work, we show how the silicon neurons of the Reconfigurable On Line Learning Spiking (ROLLS) neuromorphic processor [7], configured in a soft winner-take-all (WTA) topology, can perform unsupervised pattern recognition. We use unlabelled Poisson spike trains with different frequencies as exemplary patterns that stimulate plastic synapses of the silicon neurons. The plastic synapses either depress or potentiate depending on different pre- and postsynaptic firing rates. After learning, different groups of neurons respond to different input patterns because of varying potentiated and depressed synapses. Neurons within a group form local clusters and fire with a similar postsynaptic frequency when stimulated with one of the patterns, thus demonstrating unsupervised clustering.

We proceed with a brief introduction to the ROLLs neuromorphic processor and the WTA architecture that performs unsupervised learning. We then present the pattern classification task and the evaluation of the architecture.

II. NEUROMORPHIC PROCESSOR

The ROLLs neuromorphic processor [7] is a full-custom mixed signal analog/digital VLSI device. It comprises analog neuromorphic circuits that emulate the biophysics of spiking neurons and dynamic synapses in real-time and asynchronous digital circuits that manage the transmission and routing of spikes, using the Address Event Representation (AER). The chip comprises a column of 1×256 neurons, an array of 256×256 non-plastic programmable synapses, an array of 256×256 plastic synapses, and 8×256 time-multiplexed “virtual” synapses that can be used to provide direct external input to neurons.

A block-diagram of the chip architecture is shown in Fig. 1. Peripheral input-output AER circuits for both receiving and transmitting off-chip spikes in real time can be used to stimulate individual synapses or neurons on the chip. Silicon synapses process spikes as they arrive, and produce output currents with biologically plausible temporal dynamics. Silicon neurons integrate these currents to generate and transmit spikes in real time. The on-chip programmable bias generator [8], allows the user to program the properties of the synapses and neurons (time constants, pulse widths, etc.). The network topology can be configured using the PyNCS high-level Python software framework [9].

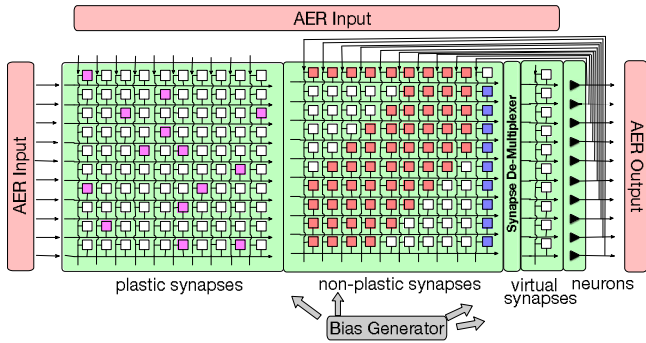


Fig. 1. Block diagram of the ROLLS chip architecture. Triangles on the right represent *silicon neurons*, squares *synapses*, organized in three arrays: *plastic*, *non-plastic*, and *virtual*. The *AER blocks* manage the input and output traffic of spikes, the *bias generator* allows to program different parameter settings of the analog circuits.

A. On Chip Spike-Based Learning

In the ROLLS neuromorphic processor, the plastic silicon synapses emulate a spike-driven synaptic plasticity rule that updates the synaptic weights according to the timing of the pre-synaptic spike, the state of the post-synaptic neuron’s membrane potential, and its recent spiking activity [10]. While weights are updated with gradual analog changes with every spike, on long time-scales they converge to one of two possible analog high/low states, i.e., they are bi-stable. The weight-update is evaluated upon the arrival of each pre-synaptic spike. It is positive if at the time of the pre-synaptic spike the membrane potential of the post-synaptic neuron is above a given threshold, θ_{mem} , while the calcium variable of that neuron (that represents the average activity of the post-synaptic neuron) is in a range set by thresholds θ_{min} and θ_{max} . If the membrane potential of the postsynaptic neuron is below θ_{mem} , and the calcium variable is in the range $[\theta_{min}, \theta_2]$, then the synapse is depressed. No update happens under other conditions. The stop-learning conditions, determined by the thresholds θ_{min} and θ_{max} , are useful for normalising the weights of all synapses afferent to the same neuron.

In addition to the event-driven weight update rule, the weight of the synapse, w_i , is constantly driven toward one of the two stable states (w_{min} and w_{max}) with a low slew-rate, depending on whether it is above or below a threshold θ_w .

B. Silicon Neuron Block

The neuron circuit integrated in the ROLLS chip exhibits biologically realistic neuronal behaviors, such as spike-frequency adaptation, adjustable refractory period, and spiking threshold. The neuron equations derived from the circuit closely resemble those of the adaptive exponential I&F neuron model [11]. Thirteen globally tunable parameters enable the neuron to produce a range of different firing behaviours.

III. WTA NETWORK

In a soft WTA network, neurons amplify their local activity by being excitatorily connected to their nearest neighbors and

globally inhibit each other by being all connected to a common pool of inhibitory neurons, which inhibits all excitatory neurons back. As a consequence, the network “selects” the local group receiving the strongest excitation and suppresses the activity of all other neurons in the network via the global inhibition. This type of soft-max computation has been proposed in hierarchical models of vision [12] and in models of selective attention and recognition [13]. Similar connectivity patterns were analysed in a continuum model of dynamic neural fields (DNFs) to account for self-organisation and pattern formation in cortical and thalamic nervous tissue [14] and to develop neural-dynamic cognitive architectures [15]. Soft WTA hardware networks have also been shown to support the extraction of hidden structure in input patterns by selectively enhancing the contrast between inputs [16]. Importantly, soft WTA networks of silicon neurons have been shown to be instrumental in reducing the effect of device mismatch of transistor, due to fabrication imperfections, and in producing robust network behavior by stabilizing activity of the most active neurons. When stimulating all neurons with the same input, the WTA selects neurons with the highest postsynaptic average activity. Device mismatch will cause different input patterns to select different winning neuron populations, laying the foundation for the unsupervised learning.

IV. PATTERN CLASSIFICATION

Here we analyze an unsupervised pattern recognition task in which the input patterns consist of Poisson spike trains that stimulate the 256 plastic synapses of every neuron with different mean frequencies. We tuned the circuit biases to initiate long-term depression (LTD) of plastic synapses when stimulated with low frequencies (10 Hz) and long-term potentiation (LTP) when stimulated with high frequencies (110 Hz). We configured the chip to implement four soft-WTA networks, as described in Fig. 2b. To assess the effect of device mismatch, we examined robustness and differences across the four networks behaviors. Each network comprises 62 excitatory neurons and a single global inhibitory neuron. Excitatory neurons have excitatory connections to themselves and to their six (three per side) nearest neighbors. The inhibitory neuron receives excitatory input from all excitatory neurons and inhibits them back. The recurrent connections that form the soft-WTA network are realized using the ROLLS non-plastic programmable synapses. Fig. 2b shows the recurrent connectivity matrix of all four WTA networks instantiated on the ROLLS chip. The plastic synapses, which are stimulated by the input patterns, are initialized randomly with a 5% probability of being in the potentiated state (see Fig. 2a, and schematically in Fig.1, left).

A. Frequency dependence of LTD and LTP transitions

Fig. 3 shows that the average probability of potentiation (LTP) and depression (LTD) of the plastic synapses depends (For the given circuit parameters) on the input and output spike frequencies. These curves were obtained using the following protocol: for the LTP probability plot, we set all plastic weights

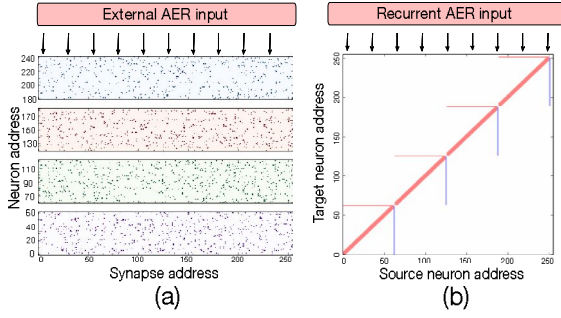


Fig. 2. Connectivity on the ROLLS chip. (a) Four sets of *plastic synapses*, connecting inputs to the four WTAs; 5% of the synapses are initialised to be randomly potentiated. (b) Matrix of the *non-plastic* recurrent connections of all 256 neurons. Four WTA network contain 63 excitatory and one inhibitory neurons each. *Blue* is an inhibitory, and *red* an excitatory connection.

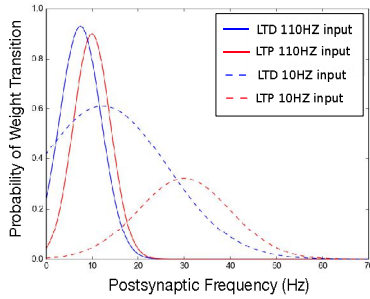


Fig. 3. The blue curves are Gaussian fits to the measured probability of *LTD transitions* and the red curves – to the measured probability of *LTP transitions* at input frequency of 110 Hz (solid) and 10 Hz (dotted). The transition probabilities depend on the postsynaptic firing rate of the neuron.

to w_{\min} and all non-plastic connections to zero. We sweep the input frequency to the plastic synapses from 10 Hz to 120 Hz in steps of 10 Hz, and for each of these conditions we stimulated the post-synaptic neurons via virtual synapses with input frequencies that were swept from 200 Hz to 600 Hz in steps of 100 Hz to obtain different postsynaptic firing rates. For each neuron, we counted how many of its 256 plastic synapses switch state from low (depressed, w_{\min}) to high (potentiated, w_{\max}). Similarly, for the LTD plots, we initialized the synaptic weights to w_{\max} , applied the same stimulation protocol as for the LTP, and counted the number LTD transitions per neuron.

For the classification task, we use input patterns that consist of a combination of the 110 Hz and 10 Hz mean frequencies. To achieve best performance, a maximum for LTP at 110 Hz and for LTD at 10 Hz input frequency is desired to reside at the same postsynaptic frequency. This way the synapses of an active neuron undergo both LTP and LTD transitions only depending on the input frequency. The learning curves from Fig. 3 show that these two maxima occur for a postsynaptic frequency of 10-13 Hz. We have configured the WTA network to suppress highly active neurons by the global inhibition to the point where they fire in this postsynaptic frequency range. Local excitation in the WTA network makes neighboring neurons fire at a similar frequency, so they will tend to learn the same input patterns. Neurons with a low output are suppressed even more by the global inhibition and do not undergo weight transitions. They retain the random weights

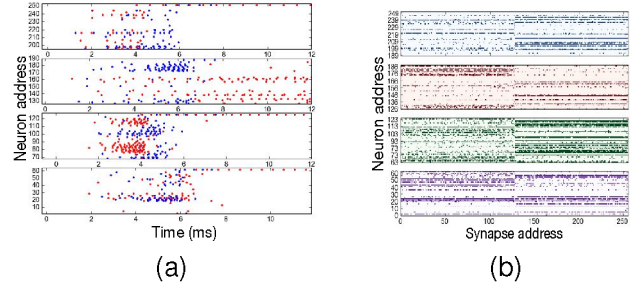


Fig. 4. (a) Postsynaptic spikes when stimulating with *input pattern one* (blue) and *input pattern two* (red) after 14 learning trials. Spikes of each stimulation were superimposed in order to highlight the different firing behavior for different input patterns. (b) Synaptic weight matrix after stimulating fourteen times with each pattern.

and can potentially be recruited for different input patterns.

B. Classification performance

As a proof of concept, the following initial experiment was chosen. We used two populations of neurons that produce patterns of Poisson spike trains with the same rates, but that project to distinct orthogonal sets of synapses (e.g., the first population projects to the first 128 synapses, and the second population projects to the other 128 sets of synapses). Specifically, every neuron in a soft-WTA network is stimulated with the same input pattern, which consist of either 110 Hz Poisson spike trains on the plastic synapses 0-128, and 10 Hz Poisson spike trains on synapses 129-255 (pattern one), or the other way around (pattern two).

The neurons that learned input pattern one have less chance to fire for input pattern two, as some of the synapses which had previously received 10 Hz have been depressed. Neurons that did not fire did not depress their randomly potentiated synapses, therefore increasing the probability to fire for the next pattern.

After repeatedly being stimulated by the two patterns, groups of neighbouring neurons change their plastic synapses to be more responsive to one of the patterns. Fig. 4b shows the synaptic weight matrix after fourteen learning trials. The network found an equilibrium, in which a similar amount of synapses are potentiated for each pattern.

The difference among the weight matrices of the four WTA networks highlights the effect of mismatch of silicon neurons. The circuit biases used here aimed to achieve most similar behavior and successful classification of patterns in all WTAs.

After the silicon neurons learned the weight matrix shown in Fig. 4b, a stimulation with input pattern one (blue) and input pattern two (red) resulted in the postsynaptic spikes shown in Fig. 4a. To emphasize timing of activity of every neuron, only the first spikes are shown here. When stimulating with a given pattern, the postsynaptic firing rates of different neurons can be used for the classification.

To test how well the neurons differentiate between the patterns, their activity was analyzed by forming a histogram of




Fig. 5. Histogram of the selectivity of the neurons: number of neurons responding to *input pattern one* (left), *input pattern two* (right), and to *both input patterns* (middle).

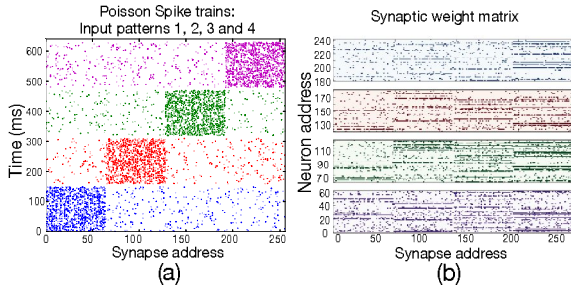


Fig. 6. (a) Four different input patterns that stimulate the plastic synapses. $\frac{1}{4}$ of the synapses receive 150 ms Poisson spike trains with 110 Hz, and the remaining synapses with 10 Hz mean frequency. (b) Plastic synaptic weight matrix after stimulating with each pattern five times.

the selectivity of neurons, shown in Fig. 5. We can observe that majority of neurons became selective to one of the patterns.

The number of patterns that can be successfully differentiated increases linearly with the number of neurons comprising a WTA network. We show that a WTA with 63 neurons can learn four different patterns, however by suffering loss in the selectivity to individual inputs.

Patterns again consist of Poisson spike trains with high (110Hz) and low (10Hz) mean frequencies, with $\frac{1}{4}$ of plastic synapses receiving the high frequency input. Fig. 6(a) shows the four different input patterns used.

Fig. 6b shows the weight matrix of plastic synapses after stimulating with each pattern five times. Further training does not improve clustering, because θ_{max} prevents learning for high postsynaptic frequencies. This stop-learning is necessary to avoid saturation of neurons' activity and overwriting of the learned patterns with the most recent one.

A decrease in performance is indicated by only four neurons being selective to input pattern one, two, and three, while for pattern four only three neurons are exclusively selective.

V. CONCLUSION

In this work, we evaluated the ability of neuromorphic hardware to perform unsupervised learning by training WTA networks to learn and recognize different frequency patterns. Such frequency patterns can encode features of the objects and the neuromorphic chip can potentially learn to classify objects in an unsupervised manner. During learning, clusters form as an effect of the WTA configuration and a Hebbian-learning weight update. The results show that one pattern is learned by more than one group of neurons making the network robust

to failure of a single group. A limitation of our neuromorphic hardware is the limited number of neurons. Thus, only a few patterns can reliably be learned. Increasing the number of neurons would result in better selectivity performance.

ACKNOWLEDGMENTS

This project was funded by the UZH grant FK-16-106, EU H2020-MSCA-IF-2015 grant 707373 ECogNet, and EU ERC-2010-StG 20091028 grant 257219 NeuroP.

REFERENCES

- [1] M. Ester, H. P. Kriegel, J. Sander, and X. Xu, "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," *Second Int. Conf. Knowl. Discov. Data Min.*, pp. 226–231, 1996. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.20.2930>
- [2] G. E. Hinton, S. Osindero, and Y. W. Teh, "A fast learning algorithm for deep belief nets," *Neural Comput.*, vol. 18, no. 7, pp. 1527–1554, 2006.
- [3] T. Kohonen, "Analysis of a simple self-organizing process," *Biol. Cybern.*, vol. 44, no. 2, pp. 135–140, 1982.
- [4] P. Diehl and M. Cook, "Unsupervised learning of digit recognition using spike-timing-dependent plasticity," *Frontiers in Computational Neuroscience*, vol. 9, no. August, p. 99, 2015. [Online]. Available: <http://journal.frontiersin.org/article/10.3389/fncom.2015.00099>
- [5] M. Giulioni, F. Corradi, V. Dante, and P. Del Giudice, "Real time unsupervised learning of visual stimuli in neuromorphic VLSI systems," *Nature Publishing Group*, pp. 1–10, 2015. [Online]. Available: <http://dx.doi.org/10.1038/srep14730>
- [6] G. Indiveri, B. Linares-Barranco, T. J. Hamilton, A. van Schaik, R. Etienne-Cummings, T. Delbruck, S. C. Liu, P. Dudek, P. Höfliger, S. Renaud, J. Schemmel, G. Cauwenberghs, J. Arthur, K. Hynna, F. Folowosele, S. Saighi, T. Serrano-Gotarredona, J. Wijekoon, Y. Wang, and K. Boahen, "Neuromorphic silicon neuron circuits," *Front. Neurosci.*, vol. 5, no. May, pp. 1–23, 2011.
- [7] N. Qiao, H. Mostafa, F. Corradi, M. Osswald, F. Stefanini, D. Sumislawska, and G. Indiveri, "A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128K synapses," *Front. Neurosci.*, vol. 9, no. APR, pp. 1–17, 2015.
- [8] T. Delbruck, R. Berner, P. Lichtsteiner, and C. Dualibe, "32-bit configurable bias current generator with sub-off-current capability," in *ISCAS 2010 - 2010 IEEE Int. Symp. Circuits Syst. Nano-Bio Circuit Fabr. Syst.*, 2010, pp. 1647–1650.
- [9] F. Stefanini, E. O. Neftci, S. Sheik, and G. Indiveri, "PyNCS: a microkernel for high-level definition and configuration of neuromorphic electronic systems," *Front. Neuroinform.*, vol. 8, no. 5, p. 73, 2014. [Online]. Available: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4152885&tool=pmcentrez&rendertype=abstract>
- [10] J. M. Brader, W. Senn, and S. Fusi, "Learning Real-World Stimuli in a Neural Network with Spike-Driven Synaptic Dynamics," *Neural Comput. Massachusetts Inst. Technol.*, vol. 19, pp. 2881–2912, 2007.
- [11] R. Brette and W. Gerstner, "Adaptive exponential integrate-and-fire model as an effective description of neuronal activity," *J. Neurophysiol.*, vol. 94, no. 5, pp. 3637–42, 2005. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/16014787>
- [12] M. Riesenhuber and T. Poggio, "Hierarchical models of object recognition in cortex," *Nat. Neurosci.*, vol. 2, no. 11, pp. 1019–25, 1999. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/10526343>
- [13] G. A. Carpenter and S. Grossberg, "A massively parallel architecture for a self-organizing neural pattern recognition machine," *Comput. Vision, Graph. Image Process.*, vol. 37, no. 1, pp. 54–115, 1987.
- [14] H. R. Wilson and J. D. Cowan, "A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue," *Kybernetik*, vol. 13, no. 2, pp. 55–80, 1973.
- [15] Y. Sandamirskaya, "Dynamic neural fields as a step toward cognitive neuromorphic architectures," *Front. Neurosci.*, vol. 7, no. January, p. 276, 2013. [Online]. Available: <http://journal.frontiersin.org/article/10.3389/fnins.2013.00276/abstract>
- [16] E. Chicca, F. Stefanini, C. Bartolozzi, and G. Indiveri, "Neuromorphic electronic circuits for building autonomous cognitive systems," *Proc. IEEE*, vol. 102, no. 9, pp. 1367–1388, 2014.