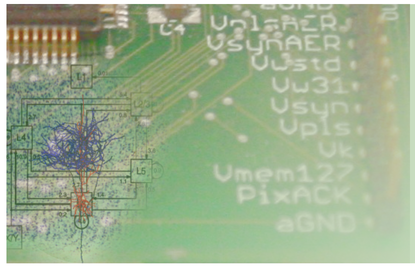

Spike-based synaptic plasticity for on-chip, online learning in real-time

Fabio Stefanini[1], Mattia Rigotti[2], Sadique Sheik[1], Stefano Fusi[2], Giacomo Indiveri[1]

1 Institute of Neuroinformatics (INI), Univeristy of Zurich and ETH Zurich, Zurich Switzerland.
2 Center for Theoretical Neuroscience, College of Physicians and Surgeons, Columbia University, New York, NY, USA

uzh | eth | zürich
Institute of Neuroinformatics
Neuromorphic Cognitive Systems

erc European Research Council

Introduction

The remarkable ability of biological systems to adapt to new, unknown environmental conditions, learn from their experience, and extract relevant information from the world, is yet unparalleled by artificial systems. Software simulations are demanding in terms of computing power and are not suited for robotic applications. Analog neuromorphic hardware is compact and low-power but it is unprecise and inhomogeneous. Here we investigated the properties of neuromorphic hardware related to the intrinsic variability of its components. Inspired by theoretical analysis, methods of computer science and biological evidences, we developed a network that exploits the intrisic variability of the hardware implementation to improve classification performance.

Classification of real-world stimuli

A linear classifier (e.g. a perceptron) is a key element for memory formation in neural networks. We applied our approach to solve the task of classifying hand-written characters from the MNIST database [1].

- Input stimuli are projected into a *high-dimensional space* using a large population of neurons and random connections [2].
- Classes become *linearly separable*.
- The classifier comprises a single neuron receiving inputs through plastic synapses with *stochastic STDP*.

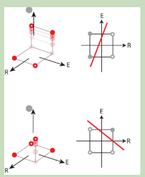

Fig.1 With randomly connected neurons it is possible to solve the non-linear separability of the input classes. The XOR case is shown as an explanatory example.

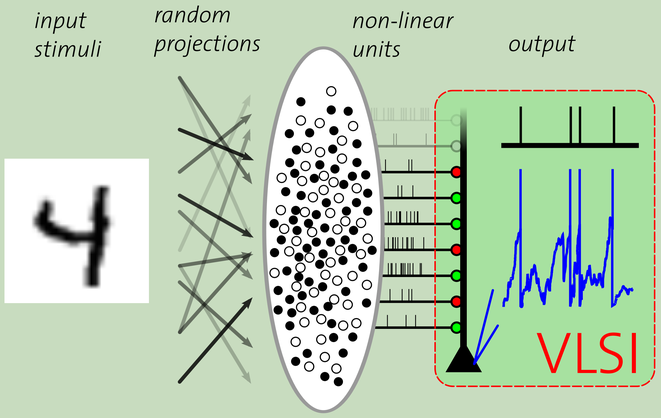

Fig.2 Scheme of the neural network for pre-processing and classification on VLSI.

Neuromorphic hardware setups

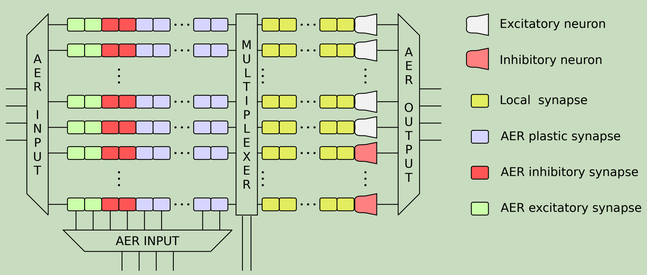

Fig.3 Block diagram of the neuromorphic chip with plastic synapses (CMOS, 0.35um technology, ~10mm²)

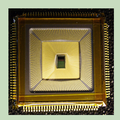

- Compact, low-power, real-time, bio-inspired computation [3]
- Hybrid analog-digital implementation using standard VLSI, CMOS technology
- Circuits for synapses with STDP and Leaky Integrate-and-Fire neuron models

Classification with imprecise hardware

We use neuromorphic hardware in VLSI technolgy to emulate spiking neural networks in real-time. We carried-out simulations to assess the effects of variability of the learning dynamics.

- The stochastic plasticity *slows-down* the learning, only after multiple presentations the input pattern is stored.[4]
- A slow-learning process maximizes memory capacity and prevents overfitting at the cost of learning time.
- Several *weak* classifiers can be combined to improve performances and learning speed. [5]

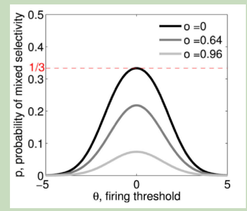

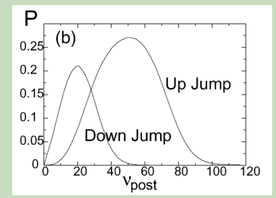

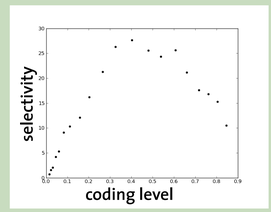

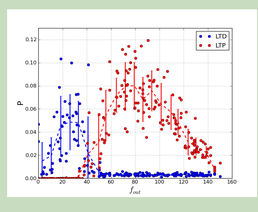

We carried out non-spiking simulations of the network obtaining 96% generalization performance. The performance of the single classifiers is comparable to their hardware, spike-based implementation and suggests the viability of the hardware solutions. The imprecise, analog circuits of our hardware are therefore compatible with machine-learning and neural-network theories that explain how variability in the classifier responses and simple pre-processing can be used to solve typical machine-learning tasks.

Fig.3 Comparison between theory, simulations and hardware. Top-left: theoretical analysis of randomly connected neurons (LTU) showing the advantage of dense representations for mixed selectivity [2]. Bottom-left: peak selectivity is reached for dense representations in neuromorhpic hardware (coding level *f=0.5*). Top-right: Probabilities for potentiation and depression in a simulation of stochastic plasticity [6]. Bottom-right: Transition probabilities measured from the neuromorphic hardware.

Some examples of correctly labeled data

Some examples of wrongly labeled data

labeled as: 4 6 0

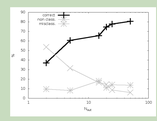

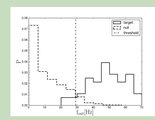

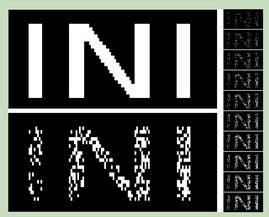

Fig.4 A toy example of the stochastic plasticity at work. *Middle*: the "INI" pattern is converted into trains of spikes and stimulates 3472 VLSI synapses. *Right*: the evolution of synaptic connectivity (white, potentiated synapse; black, depressed). *Bottom*: resulting synaptic levels. *Top-Left*: Combinations of classifiers improve performances of the weak hardware classification. *Bottom-Left*: Output histograms of the perceptron after learning. The neuron is considered selective for the target class.

References

[1] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11):2278–2324, 1998.
[2] M. Rigotti, D.D. Ben Dayan Rubin, X.-J. Wang, and S. Fusi. Internal representation of taskrules by recurrent dynamics: the importance of the diversity of neural responses. Frontiers in Computational Neuroscience, 4(0), 2010.
[3] E. Chicca, F. Stefanini and G. Indiveri. Neuromorphic electronic circuits for building autonomous cognitive systems. Proceedings of the IEEE (Submitted).
[4] W. Senn and S. Fusi. Slow stochastic learning with global inhibition: a biological solution to the binary perceptron problem. Neurocomputing. 58–60:321-326, 2004.
[5] L.Breiman. Bagging predictors. Machine Learning. 24 (2): 123--140, 1996.
[6] J. Brader, W. Senn, and S. Fusi. Learning real world stimuli in a neural network with spike-driven synaptic dynamics. Neural Computation, 19:2881–2912, 2007.

Acknowledgments

This work was supported by the EU ERC Grant "neuroP" (257219).