

UNIVERSITY OF ZURICH AND ETH
ZURICH

MASTER THESIS

**Unsupervised Learning in
Neuromorphic Devices using
Winner-Take-All networks**

Author:
Raphaela Kreiser

Supervisor:
Giacomo Indiveri

*A thesis submitted in fulfillment of the requirements
for the degree of MSc UZH ETH in Neural Systems and
Computation*

*in the
Neuromorphic Cognitive Systems Group
Institute of Neuroinformatics*

July 2016

Declaration of Authorship

I, Raphaela KREISER, declare that this thesis titled, "Unsupervised Learning in Neuromorphic Devices using Winner-Take-All networks" and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Date:

“Quote...”

Dave Barry

UNIVERSITY OF ZURICH AND ETH ZURICH

Abstract

Institute of Neuroinformatics

Master in Neural Systems and Computation

**Unsupervised Learning in Neuromorphic Devices using
Winner-Take-All networks**

by Raphaela KREISER

The Thesis Abstract is written here (and usually kept to just this page). The page is kept centered vertically so can expand into the blank space above the title too...

Acknowledgements

The acknowledgments and the people to thank go here, don't forget to include your project advisor...

Contents

Declaration of Authorship	i
Abstract	iii
Acknowledgements	iv
1 Introduction	1
1.1 Motivation	1
1.2 Thesis Outline	4
2 Architecture and Learning Abilities of the Neuromorphic Processor	5
2.1 Reconfigurable On-line Learning Spiking Neuromorphic Processor	5
2.2 Spike-Based Learning Algorithm	7
2.3 The Silicon Neuron Block	9
2.4 The Long-Term Plasticity Synapse Array	10
2.5 The Short-Term Plasticity Synapse Array	12
2.6 Learning on the Network Level	13
3 Unsupervised Learning on neuromorphic hardware	15
3.1 Winner-Take-All Network	15
3.2 Network Configuration on Chip	19
3.3 Pattern Classification Task	20
3.4 Experimental Results	22
3.4.1 Synapse Dynamics and sensitivity to the parameters .	22
3.4.2 Classification performance	24
3.4.3 Learning four different patterns	30
3.4.4 Discussion	33
4 Sequence Learning	35
4.1 Conceptual model	35
4.1.1 Dynamic Neural Fields	36
4.2 Different behaviour resulting from population sizes, connectivity and WTAs	38
4.3 Hardware realisation	41
4.4 Experimental Results: Learning and forgetting sequences	42
4.4.1 Memorising and forgetting a simple sequence	42
4.4.2 Learning longer sequences	45
4.4.3 Learning repeating items	47

4.4.4	Serial order errors	48
4.4.5	Discussion	49
5	Conclusion	52
A	Setting suitable biases	54
	Bibliography	65

List of Figures

2.1	Micro-photograph of the ROLLS neuromorphic processor (Qiao et al., 2015)	6
2.2	Block diagram of the ROLLS chip architecture.	6
2.3	These two figures show the dynamics of the stochastic learning rule, with a LTP transition on the right. The post-synaptic activity in the top figure is kept at 150 Hz (left) and 250 Hz (right) via a teacher signal, while the pre-synaptic activity presented in the third figure from the top is kept at 50 Hz. The second figure from the top show the changes in calcium concentration in response to the post-synaptic activity, and the bottom figures show the changes in the internal weight variable. Simulation from (Ling, 2015).	9
2.4	Learning circuits in action. The yellow trace shows the position of the weight. High meaning that its value is w_{max} , low meaning that its weight is w_{min} . The red trace shows the state of the calcium variable, which is active low. A down decrease of the red trace means an increase in the post-synaptic calcium concentration. The green trace is the digital down signal, which is active high, and the blue trace is the digital up signal which is active low. Here, a down transition of the weight takes place. After the neuron fires, calcium increases. Because calcium is below θ_2 only a down signal can be initiated, leading to a weight transition from w_{max} to w_{min}	11
2.5	An up transition of the weight is taking place. After three postsynaptic spikes, the calcium post-synaptic concentration increases more and more. Because at the time of a presynaptic spike, the membrane threshold is above θ_{mem} and the calcium concentration is high enough, two up signals occur, leading to a weight transition from w_{min} to w_{max}	11
2.6	Time-plot of activity of a silicon neuron and plastic synapse. The black trace shows the input events, the pink trace the membrane potential and the light blue trace the calcium concentration (being active low). As the calcium concentration increases and the membrane potential is above θ_{mem} the weight increases, as shown as the dark blue trace.	12

3.1	A constant step input is being converted to a Gaussian output by the recurrent connections of neurons in a WTA configuration and because of the neurons mismatch. The left side shows how the step input is converted to a Gaussian output when it is fed to recurrently connected neurons. The right shows the excitatory and inhibitory connections underlying the WTA configuration.	16
3.2	A schematic of the neuron's recurrent connections on the device. Recurrent connections are shown by connections of the neurons soma. The neuon whose soma is coloured red represents the gloabl inhibitory neuron. The input to the neuron is given as two Gaussian curves with different maxima (green curve). The global inhibition and local excitation leads to an amplification of the the higher input and a suppression of the lower input. The resulting output is given as the blue curve.	17
3.3	WTA behaviour on the ROLLS chip	18
3.4	Matrix of recurrent connections of all 256 neurons. One WTA network consists of 63 neurons. The excitatory neurons are connected to their three nearest neighbors and to the global inhibitory neuron. The global inhibitory neuron is connected to all excitatory neurons. Blue represents an inhibitory connection whereas red means excitatory.	19
3.5	Setting the WTA connectivity on the ROLLS chip.	20
3.6	The schematic of the recurrently connected excitatory neurons receiving either input pattern 1 or input pattern 2 on their plastic synapses.	21
3.7	The blue curves are two Gaussians fitted to the probability measures to make LTD transitions due to a 110 Hz and a 10 Hz input frequency respectively. The red curves are two Gaussians fitted to the probability measures to make LTP transitions due to a 110 Hz and a 10 Hz input frequency. These probabilities depend on the postsynaptic firing rate of the neuron.	23
3.8	The synaptic weight matrix after initialization.	25
3.9	Maximal achievable WTA behaviour due to mismatch. Spiking for input pattern one is shown in blue, spiking for input pattern two is shown in red.	25
3.10	Synaptic weight matrix after two trials, i.e. stimulating one time with input pattern one followed by input pattern two, each for 150ms.	26
3.11	Synaptic weight matrix after fourteen trials with different sets of biases.	27

3.12 Raster plot of the output spikes after stimulating with input pattern one (blue) and input pattern two (red) for 1000ms in the testing phase. Only the first impulse of activity is shown. Each pattern was shown 13 times for 150ms during learning. The inhibitory time constant is much smaller than the excitatory time constant.	28
3.13 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).	28
3.14 Firing activity during testing, after having learned the weight matrix shown in Fig.3.11b, using fast excitation and slow inhibition. Neurons quickly excite their neighbors, and as soon as the slower, but very strong inhibition becomes active, it switches off all activity.	29
3.15 Histogram of the selectivity of the neurons, responding to input pattern one (left), input pattern two (right), and to both input patterns (middle).	30
3.16 Activity after having learned the weight matrix shown in Fig.3.11b and being stimulated with input pattern one (blue) and input pattern two (red). The time constants of excitation and inhibition are only slightly different and the stimulus onset is slightly shifted. Input pattern two starts 300 ms later than input pattern one.	31
3.17 The four different inputs used for stimulating the learning synapses. The step signal is given to one fourth of learning synapses. It is a 200 ms Poisson spike train with a frequency of 110 Hz, the other synapses are stimulated with Poisson spike trains of 10 Hz.	32
3.18 Weight matrix of the plastic synapses after learning 4 different input patterns.	32
3.19 Histograms on the selectivity after learning four different input patterns.	33
4.1 The serial order architecture, introduced in Sandamirskaya and Schoener, 2010.	35
4.2 Connectivity matrix, sent to the ROLLS chip to encode the serial order architecture	37
4.3 Activity of the neurons after stimulating the first ordinal group together with the first content group using a continuos WTA configuration with inhibitory population for ordinal neurons.	39
4.4 Plastic weight matrix for learning a sequence with three elements after initilization	40
4.5 External input used for generating and recalling a sequence consisting of five elements.	42
4.6 Learning sequence ABC	43
4.7 Learning sequence CAB	43
4.8 Learning and forgetting a sequence	45
4.9 Learning sequence EABDC	46

4.10	Learning sequence ABDEC	47
4.11	Learning sequence AAC	47
4.12	Common learning and recalling errors.	51
A.1	Silicon neuron schematics. The NMDA block implements a voltage gating mechanism; the LEAK block models the neuron's leak conductance; the spike-frequency adaptation block AHP models the after-hyper-polarizing current effect; the positive-feedback block Na ⁺ models the effect of the Sodium activation and inactivation channels; reset block K ⁺ models the Potassium conductance functionality.	56
A.2	Short-term plasticity synapse array element and dynamics of recurrent synaptic connections. (A) Block diagram of the synapse element. (B) Transistor level schematic diagram of the excitatory and inhibitory pulse-to-current converters.	58
A.3	Post-synaptic learning circuits for evaluating the algorithm's weight update and 'stop-learning' conditions. The DPI circuit MD1-5 integrates the post-synaptic neuron spikes and produces a current proportional to the neuron's Calcium concentration. Three current-mode winner-take-all circuits WTA, WTAUP, and WTADN compare the Calcium concentration current to three set thresholds sl_thmin!, sl_thdn!, and sl_thup!, while the neuron's membrane current is compared to the threshold sl_memthr!.	61
A.4	Long-term plasticity synapse array element. (A) Plastic synapse configuration logic block diagram. (B) Timing diagram for broadcast and recurrent activation modes in one synapse using 4-phase handshaking protocol. Dashed red lines show the sequence between signals. (C) Schematic diagram of the bi-stable weight update and current generator blocks.	64

List of Tables

A.1	Biases of the silicon neuron	55
A.2	Biases of the silicon neuron (continued)	57
A.3	Biases for the recurrent synaptic connections	59
A.4	Biases for the recurrent synaptic connections (continued) . . .	60
A.5	Biases for implementing the learning rule	62
A.6	Biases for the dynamics of the learning synapses and long-term potentiation	63
A.7	Biases for the dynamics of the learning synapses and long-term potentiation (continued)	64

Chapter 1

Introduction

1.1 Motivation

Much research has been done in the field of Machine Learning and reliable learning algorithms for prediction and decision making were developed. Neural networks were used to implement these learning functions. However such algorithms follow von Neumann computer architectures and utilize digital logic. These algorithms are fast and precise, however they underperform routine cognitive tasks that seem to be simple for biological systems. In addition, the power consumption of digital von Neumann computer architectures is immense compared to the one of the human brain and the performance of machine learning algorithms, especially in tasks that involve autonomous realtime interactions with the environment, suffers in the presence of noisy and uncontrolled sensory input (Ambrogio et al., 2016). In contrast to computational and organisational principles like Boolean logic, common in present day computers, the nervous system carries out robust computation using hybrid analog and digital processing elements. Computation is distributed, event-driven, collective and massively parallel, and it makes use of adaptation, selforganization, and learning. Neuromorphic circuits aim to recreate the biophysics and network architectures behind biological neural processes using the physics of silicon. (Indiveri and Liu, 2015) Hybrid digital and analog circuits were constructed that emulate the basic dynamical properties of biological neurons and synapses (Douglas, Mahowald, and Martin, 1901). These circuits have been integrated into very large scale integration (VLSI) devices in order to build real-time sensory systems (Liu and Delbruck, 2010). The resulting compact, real-time, and energy-efficient devices can be used to build neuromorphic behaving systems that have cognitive abilities and it can perform learning tasks. In fact



this hardware harnesses some of the outstanding properties of neuronal computing substrate: the parallel nature of computation, co-location of memory and processing, and low-power event-based computing principle. Not only are neuromorphic systems a tool that can be used to translate knowledge about neuroscience into these systems and create computation similar to the brain, they can also give an understanding about the model limitations leading further research in the right direction. A recent addition to VLSI implementation of neuromorphic systems is the Re-configurable On-Line Learning Spiking (ROLLS) neuromorphic processor. (Qiao et al., 2015) This device has the novel ability to carry out on-chip on-line learning and to configure the neuron connections into the desired neural network. The learning circuits of the chip enable the network to find solutions for tasks that require adaptation to input signals, and that can interact in real-time with its environment.

In this project the ROLLS device is configured in a soft winner-take-all (WTA) network, which has been proposed as a neural model to explain object recognition (Riesenhuber and Poggio, 1999b). The winner-take-all network carries out a hybrid analog and digital computing mechanism. It integrates the strength of local stimuli in an analog fashion but at the same time selectively amplifies the strongest and suppresses the weakest signal. In this work, a WTA architecture is being used to carry out an unsupervised learning task. Unlabeled input data arrives at the plastic synapses of the 256 Integrate and Fire (I&F) neurons coded as 10 Hz or 110 Hz frequencies distributed over 256 input synapses. The aim is that the network finds a hidden structure within the data and learns to classify it accordingly.

In order to show the vast learning abilities that the ROLLS neuromorphic processor can carry out, a neural-dynamic architecture for serial-order memory, routed in neuroscience and in cognitive science was implemented on the ROLLS chip in order to perform a sequence learning task which can be used on embedded systems to perform cognitive tasks in realtime. The proposed neuromorphic cognitive architecture has a high potential for cognitive robotics since it can process large amounts of noisy sensory signals coming from neuromorphic sensors in parallel. Eventually a robot can learn to associate input scenes from a low-power DVS camera using a silicon retina (Lichtsteiner and Delbrück, 2005) with actions, all based on the hierarchical neural network architecture implemented on the ROLLS neuromorphic processor.

The field of behavioral robotics developed neuronal controllers for behaving systems which are currently used for insect-like robots (Manoonpong, Parlitz, and Wörgötter, 2013) and robotic salamanders (Ijspeert et al., 2007). These can adapt to changing environments and work well in settings where reactive behavior is enough. However, they do not scale up to more complex behaviours, which require a representation of the environment. In this project we aim for cognitive robotics that are able to understand complex scenes and that can plan longer sequences accordingly. A dynamical systems approach from the field of cognitive science was used to find an architecture that gives a suitable representation of the environment to enable intelligent (Thelen and Smith, 1994).

Here, behavior of an agent is generated by a dynamical system, defined over behavioural variables that describe the immediate sensory state of the agent and its possible movements. In this framework, the concept of state is spanning continuous behavioural spaces over different perceptual and motor dimensions. Continuous dynamic neural fields are defined over these dimensions that stabilise different attractor states corresponding to the perceived states, decisions, and planned actions. These dynamics can account for working memory and stabilisation of decisions among alternative states (Schoener, 2008), which are routed in theoretical understanding of large-scale neuronal processes (Wilson and Cowan, 1973; Grossberg, 1988; Ermentrout, Galan, and Urban, 2008). This dynamic neural fields approach (Schoener, 2008; Sandamirskaya et al., 2013; Sandamirskaya, 2013a) has been used both to account for visual cognition of humans (Johnson, Spencer, and Schoener, 2008) and to control cognitive robots (Richter, Sandamirskaya, and Schoner, 2012).

Inspired by the dynamic neural fields approach, the 256 neurons of the neuromorphic device were reconfigured to account for different groups of neurons which excite one group and inhibit another. The detailed network will be explained in chapter 4. This cognitive model can explain how humans learn sequences such as phone numbers and car plates, and maintain them in the short-term memory (Henson, 1998). Every number or letter has a position in space and activates its own memory with a certain memory strength. If the strength of the memory is too small the recall is likely to fail.

In this model, ordinal groups of neurons represent action fields, which can be used to carry out an action according to a learned element or simply recall it. Ordinal and memory groups are highly self-excitatory and therefore only



stop firing when they are inhibited by an external stimulus. When an ordinal group is suppressed, its active memory excites the next ordinal group, leading to a sequence of events. At the same time another large group of neurons can be used to represent content. This content can be learned by the ordinal neurons. Associations between content and ordinal/action neurons are learned by an STDP learning rule. The learning dynamics of the neuromorphic device are in accordance with Hebb's postulate: "Cells that fire together, wire together" (Hebb, 1949). Hence, when content neurons fire at the same time as ordinal neurons, the synapses from ordinal neurons to the content neurons will potentiate and an association is being learned. During recall, the ordinal neurons will automatically activate the associated content.

1.2 Thesis Outline

This report is structured as follows. After the introduction of the first section, the second section describes the architecture of the ROLLS neuromorphic processor and gives a short description of the silicon neuron and synapse circuits that are being used. The spike-based learning mechanism will be introduced and the plasticity of the network and how it constructs the classification of different input patterns will be explained. In the second section it will be described how the neuromorphic processor can be used to carry out an unsupervised learning task. The applied methods that lead to the networks ability of performing this task will be introduced. These methods include the winner-take-all configuration, the specification of input patterns, the initialisation of the network as well as how to read out the learned classes from the network. Learning dynamics and classification results will be shown. The fourth chapter introduces the sequence learning task. The conceptual model will be described, different network architectures on chip as well as the final hardware implementation will be explained. Finally the results of learning and forgetting sequences will be shown. After considering serial order errors occurring on the hardware and how these are related to human behaviour, a conclusion of the neuromorphic learnig abilities will be given. The appendix includes the circuits of the ROLLS chip as well as the biases needed to perform the proposed tasks.

Chapter 2

Architecture and Learning Abilities of the Neuromorphic Processor

2.1 Reconfigurable On-line Learning Spiking Neuromorphic Processor

A compact, low-power, artificial neural processing systems having real-time on-line learning abilities, was used in this project to carry out an unsupervised pattern recognition as well as a sequence learning task. The reconfigurable on-line learning Spiking Neuromorphic Processor (ROLLS) designed by Ning Qiao (Qiao et al., 2015) is a full-custom mixed signal VLSI device which comprises neuromorphic learning circuits that can emulate the biophysics of real spiking neurons and dynamic synapses in real-time. This chip comprises 128K synapse circuits and 256 neuron circuits, realised with analogue electronics. These comprise biologically plausible dynamics and bi-stable spike-based plasticity mechanisms necessary for their on-line learning abilities.

The 256 neuron circuits can be reconfigured, thus allowing us to build brain-inspired neural network connectivities. With the right circuit biases the neural networks are able to carry out short-term as well as long-term plasticity mechanisms. In this thesis it will be shown that in the right configuration and with the right set of biases, this device is able to reliably perform unsupervised learning tasks in real-time. Fig. 2.1 shows a micro-photograph of the global architecture of the ROLLS neuromorphic processor. Two distinct synapse arrays can be recognised. One array of 256x256 programmable

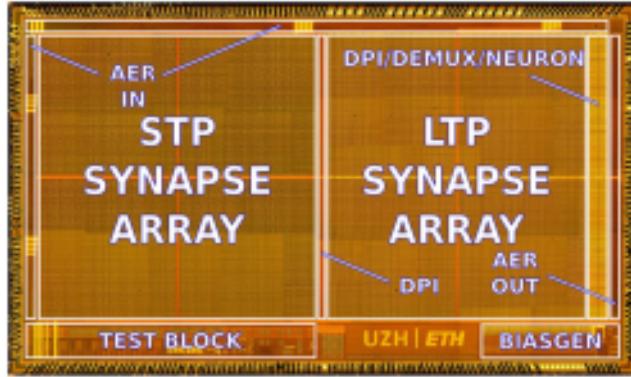


FIGURE 2.1: Micro-photograph of the ROLLS neuromorphic processor (Qiao et al., 2015).

synapses with short-term plasticity (STP) circuits on the left, and one array of 256x256 plastic synapse circuits for modelling long-term plasticity (LTP) mechanisms on the right.

The device also comprises a 256x2 row of linear integrator filters denoted as “virtual synapses”. These can be used to direct external inputs to the neurons. Stimulation of a single virtual synapse has a similar effect as stimulating the target neuron with multiple independent inputs and is therefore an efficient tool that can be used to directly obtain postsynaptic spikes on a desired output frequency. The silicon neurons can be recurrently connected through a programmable logic state of a synaptic matrix. The block-diagram of the chip architecture is shown in Fig. 2.2.

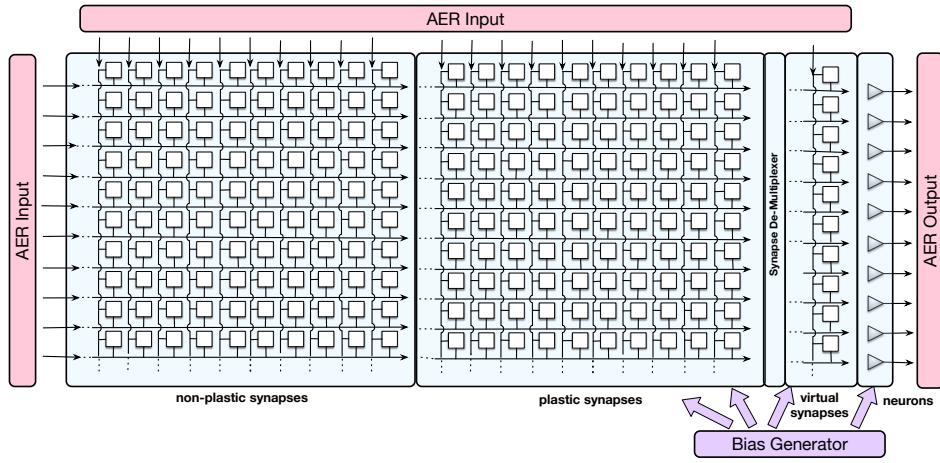


FIGURE 2.2: Block diagram of the ROLLS chip architecture.

Peripheral analog/digital input-output circuits for both receiving and transmitting spikes in real-time off-chip follow an AER (address-event representation) protocol and can be used to stimulate individual synapses or neurons

on the chip. An on-chip programmable bias generator, optimized for sub-threshold circuits (Delbrück et al., 2010) allows the user to program the properties of the synapses and neurons (such as time constants, pulse widths, etc.), creating networks with different properties and topologies.

The ROLLS was fabricated using a standard 180 nm CMOS 1P6M process. It occupies an areas of 51.4 mm^2 with approximately 12.2 million transistors.

The neural network architecture and the parameters of the neuromorphic core are defined and communicated to the processor via a high-level Python framework, called PyNCS (Stefanini et al., 2014). All silicon synapses process input spikes in real-time as they arrive and silicon neurons transmit the spikes they produce as they are generated. Hence, input data is processed instantaneously and the circuits are operating with time-constants well-matched to behaving systems that interact with the environment in real-world scenarios. The synapse circuits of the processor are able to convert input spikes into output currents which have non-linear dynamics, due to adaptation and learning features.

2.2 Spike-Based Learning Algorithm

An important property of biological synapses is their ability to exhibit different forms of plasticity. Plasticity mechanisms produce long-term changes in the synaptic strength of individual synapses. Different strengths of synapses can help to form memories and to learn features of input stimuli. The mechanism leading to an increase of synaptic strength is denoted as long-term potentiation whereas the mechanism leading to a decrease of synaptic strength is denoted as long-term depression. In the proposed neuromorphic processor, the silicon synapses can model these plasticity mechanisms with a Spike-Time Dependent Plasticity (STDP) rule which has been widely spread in computational neuroscience literature (Abbott and Nelson, 2000). In STDP mechanisms, synaptic weight values are updated based on the timing of presynaptic and postsynaptic spikes. However, STDP rules alone have poor memory retention performance (Billings and Rossum, 2009) and require local mechanisms to learn both spike-time correlations and mean firing rates of input patterns. (Senn, 2002) That is why the ROLLS processor implements a spike-driven synaptic plasticity rule proposed by Brader et al. (Brader, Senn, and Fusi, 2007). This rule has been shown to reproduce

many of the neural dynamics observed in biology because it does not rely on spike-timing alone but rather 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. On long time-scales the weights are bi-stable, meaning that they converge to a bounded high or low state. The learning rule comprises a stochastic weight update mechanism avoiding that all synapses get updated in exactly the same way. This stochasticity can be obtained by simply exploiting the variability in the input spike trains of a Poisson process, and the variability in the post-synaptic neuron's membrane potential. The weight-update is evaluated upon the arrival of each pre-synaptic spike and is given for a synapse i by the following equations:

$$w_i = w_i + \Delta * w^+ \text{ if } V_{mem}(t_{pre}) > \theta_{mem} \text{ and } \theta_1 < Ca(t_{pre}) < \theta_3 \quad (2.1)$$

$$w_i = w_i - \Delta * w^- \text{ if } V_{mem}(t_{pre}) > \theta_{mem} \text{ and } \theta_1 < Ca(t_{pre}) < \theta_2 \quad (2.2)$$

where w_i represents an internal variable that encodes the bi-stable synaptic weight. The terms w^+ and w^- determine the amplitude of the variables increase and decrease respectively, following the update. V_{mem} represents the post-synaptic neuron's membrane potential at the time of the pre-synaptic spikes arrival. θ_{mem} is a threshold term that determines whether the neurons membrane at time of the presynaptic spike is higher than this threshold, in order to increase the weight or to decrease it otherwise. Ca is the term that represents the neuron's Calcium concentration, which is proportional to the neuron's recent spiking activity. θ_1 , θ_2 , and θ_3 are the thresholds that determine in which conditions the weights are allowed to be increased, decreased or will not be updated. The stop-learning conditions are useful for normalising the weights of all synapses afferent to the same neuron. It was shown by (Senn and Fusi, 2005) that these conditions can extend the memory lifetime of recurrent spiking neural networks, as well as increase their capacity. In addition to the weight update rule, the internal variable of the synapse w_i is constantly being driven toward one of the two stable states depending on whether it is above or below a given threshold θ_w

$$\frac{d}{dt}w_i = +C_{drift} \text{ if } w_i > \theta_w \text{ and } w_i < w_{max} \quad (2.3)$$

$$\frac{d}{dt}w_i = -C_{drift} \text{ if } w_i < \theta_w \text{ and } w_i > w_{min} \quad (2.4)$$

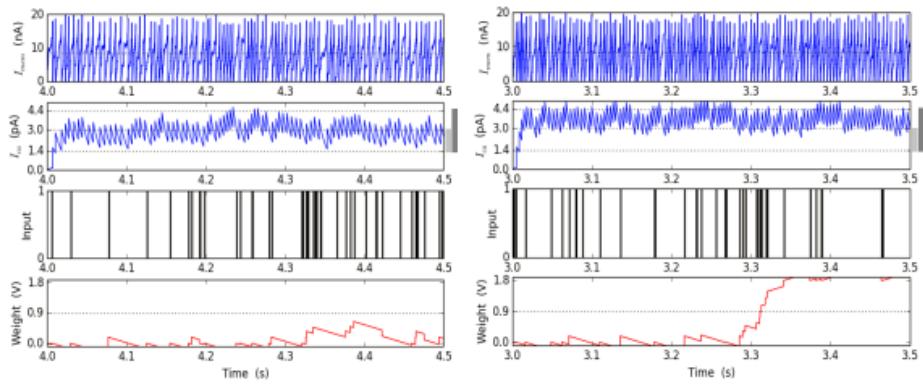


FIGURE 2.3: These two figures show the dynamics of the stochastic learning rule, with a LTP transition on the right. The post-synaptic activity in the top figure is kept at 150 Hz (left) and 250 Hz (right) via a teacher signal, while the pre-synaptic activity presented in the third figure from the top is kept at 50 Hz. The second figure from the top show the changes in calcium concentration in response to the post-synaptic activity, and the bottom figures show the changes in the internal weight variable. Simulation from (Ling, 2015).

where C_{drift} represents the rate at which the synapse is driven to its bounds w_{max} and w_{min} . Finally, the binary weight of the synapse is a thresholded version of the variable w_i :

$$J_i = J_{max}f(w_i, q_J) \quad (2.5)$$

where $f(w_i, q_J)$ is a threshold function with threshold q_J , and J_{max} is the maximum synaptic efficacy. Ning et al. showed that the ROLLS neuromorphic processor can faithfully implement this learning algorithm (Qiao et al., 2015).

2.3 The Silicon Neuron Block

The neuron circuit integrated in the ROLLS chip was derived from the adaptive exponential I&F circuit proposed in Indiveri et al. (Indiveri et al., 2011). It exhibits neural behaviours, such as spike-frequency adaptation, a refractory period mechanism and an adjustable spiking threshold mechanism. The circuit schematics are shown in appendix A. The circuit models the NMDA voltage gating function, the neuron's leak conductance, the effect of Sodium activation and inactivation channels for producing the spike, as well as the effect of the Potassium conductance, resetting the neuron and

implementing a refractory period mechanism. The neuron circuit equations are nearly the same as the ones of the adaptive I&F neuron model, and the repertoire of behaviors that it can reproduce were analysed by Brette and Gerstner (Brette and Gerstner, 2005). The transistor biases are globally tunable parameters which can be set by the on chip Bias Generator (BG). There are 13 tuneable parameters, making it possible to let the neuron produce different kinds of behaviours. See appendix A for more information. The neuron soma circuit and the post-synaptic plasticity circuits integrate the spikes produced by the neuron into a current that models the neuron's Calcium concentration. This current is then being compared to three threshold currents θ_1 , θ_2 , and θ_3 as shown in equation 2.1 and 2.2. At the same time, the neuron's membrane current is being compared to θ_{mem} . The outcome of these comparisons sets the digital 'down' (DN) and 'not up' (NUP) signals, which are buffered and transmitted in parallel to all synapses afferent to the neuron's long-term plasticity array.

2.4 The Long-Term Plasticity Synapse Array

Each of the 256x256 synapse circuits in the long-term plasticity array includes event-based programmable logic circuits for configuring plastic connections of the network, as well as circuits that implement the previously mentioned learning algorithm. The circuits are shown in the appendix A. The bistable state of the synaptic weight can be set by sending an AER event with the matching address and by asserting the configuration signals `set_high` and `set_low`. The heights of the digital up and down signals that arise after comparisons made with the learning rule can be changed in the bias generator. In addition, the up and down drift rate and the threshold θ_{mem} , which determines the threshold of a weight transition can be set, see equation 2.3 and 2.4. Therefore, it can be directly defined how fast a synapse should learn to depress and to potentiate.

Fig. 2.4 and Fig. Fig.2.5 show experimental results that highlight the features of both synapse and neuron learning circuits in action: weight updates are triggered when the pre-synaptic spikes arrive, and when the post-synaptic neuron's Calcium concentration is in the appropriate range. Depending on the value of the Calcium concentration, the digital up and down signal turn on or off. The weight's internal variable is increased or decreased

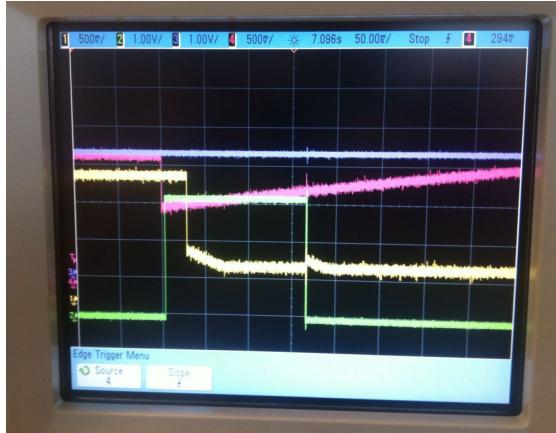


FIGURE 2.4: Learning circuits in action. The yellow trace shows the position of the weight. High meaning that its value is w_{max} , low meaning that its weight is w_{min} . The red trace shows the state of the calcium variable, which is active low. A down decrease of the red trace means an increase in the post-synaptic calcium concentration. The green trace is the digital down signal, which is active high, and the blue trace is the digital up signal which is active low. Here, a down transition of the weight takes place. After the neuron fires, calcium increases. Because calcium is below θ_2 only a down signal can be initiated, leading to a weight transition from w_{max} to w_{min} .

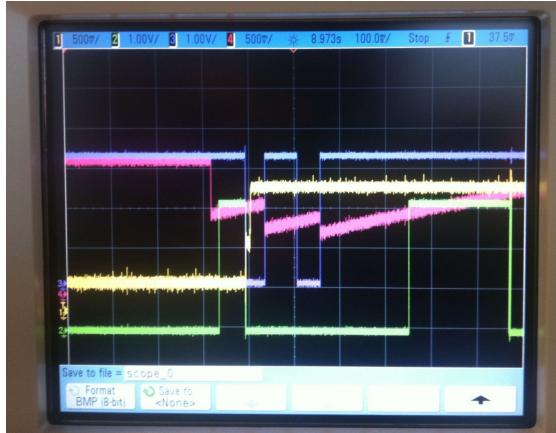


FIGURE 2.5: An up transition of the weight is taking place. After three postsynaptic spikes, the calcium post-synaptic concentration increases more and more. Because at the time of a presynaptic spike, the membrane threshold is above θ_{mem} and the calcium concentration is high enough, two up signals occur, leading to a weight transition from w_{min} to w_{max} .

depending on where the membrane potential is with respect to the membrane threshold. This variable is actively driven to the low or high bounds, according to whether it is below or above the weight threshold θ_w .

Fig. 2.6 shows the data taken from the oscilloscope during potentiation, here

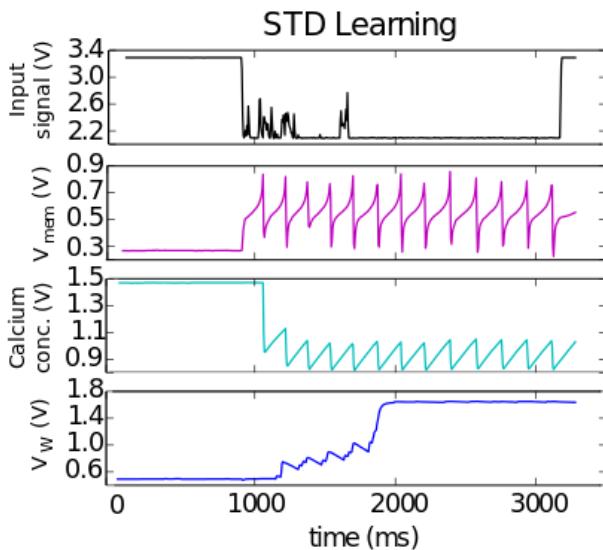


FIGURE 2.6: Time-plot of activity of a silicon neuron and plastic synapse. The black trace shows the input events, the pink trace the membrane potential and the light blue trace the calcium concentration (being active low). As the calcium concentration increases and the membrane potential is above θ_{mem} the weight increases, as shown as the dark blue trace.

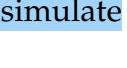
together with the input events (black trace) and the membrane potential (pink trace).

2.5 The Short-Term Plasticity Synapse Array

The circuits of the STP synapses allow the user to program the strength of recurrent synaptic connections to one out of four possible weight values. In addition, it can be specified if a synaptic connection should be excitatory or inhibitory. An excitatory synapse comprises circuits for modelling Short-Term Depression (STD) dynamics (Rasche and Hahnloser, 2001; Boegerhausen, Suter, and Liu, 2003) where the magnitude of the excitatory post-synaptic potential (EPSP) decreases with every input spike, and recovers slowly in absence of inputs. More details about the underlying silicon neuron and synapse circuits can be found in the original paper of the ROLLS chip (Qiao et al., 2015).

2.6 Learning on the Network Level

In the previous section we have described the spike-based learning rule, as well as the single building blocks implemented on the ROLLS chip. Each neuron has three different activation modes. First, direct activation through its virtual synapses. Second, recurrent activation through programmable synapses, which is usually initiated by the firing of other neurons, that are recurrently connected to the given neuron. And third, activation through plastic synapses that have the ability of potentiation and depression. The strength of the recurrent connections is fixed and can be set to one out of four different weights. Configuring the recurrent connections of the neurons, their synaptic strength, and setting the neurons to be either excitatory or inhibitory lets us create different types of neural networks. These neural networks can have different abilities. As an example, it was shown that groups of neurons can create, modify and preserve memories ([Technology](#) and [Tianlu, 2015](#)).

Learning protocols can guide the synaptic modification by making the post-synaptic neuron respond differently to different kinds of input patterns. In an unsupervised learning setting, input patterns   are analysed and classified   on their statistical structure. The neural network is able to discover new and hidden structure without requiring previous training. Unsupervised learning is an interesting topic in the machine learning community and many reliable algorithms, such as clustering, principal component analysis (PCA), estimation maximisation and others have been developed to  simulate unsupervised learning.

However, these algorithms utilise digital logic, fast and precise, but very power-consuming. Neuromorphic circuits that recreate the biophysics and network architectures behind biological neural processes show a better performance in the presence of noise, they are more power-efficient and operate in real-time. Inspired by the recurrent organisation of the neurons in the cerebral cortex, where the interaction between excitatory and inhibitory neurons ensure a well-balanced overall system ([Zhang and Sun, 2011](#)), we use the connectivity scheme of a soft winner-take-all (WTA) for an unsupervised learning task. The computation of a soft max due to the cooperative-competitive behaviour of soft WTA implementations is compatible with evidence from neuroscience. An example is the cooperative feature resolution that is processed by dendrites of the superficial pyramidal neurons in layer

2/3, as well as the output selection mechanism to motor structure that is carried out by pyramidal neurons in layer 5 (Douglas and Martin, 2007).

Chapter 3

Unsupervised Learning on neuromorphic hardware

3.1 Winner-Take-All Network

A nonlinear computational principle that is able to compute decision making was shown to exist in prefrontal cortex (Wang, 2012) and it is referred to as the winner-take-all network (WTA). In this framework neurons compete with each other for activation. In a hard WTA, only one neuron stays active, while all others are being suppressed and finally shut off. This behaviour emerges from the excitatory and inhibitory connections between the neurons. In a WTA network, neurons have excitatory connections with their nearest neighbours, therefore locally enhancing each others activity. At the same time, a single or a small group of inhibitory neurons exists which globally inhibit all excitatory neurons at the same time. This inhibitory population only gets its excitatory input from the excitatory neurons and is therefore only activated when the excitatory neurons are active as well. The strength of the global inhibition can be set in a way to be left with only one ‘winning neuron’, i.e. in a hard WTA setting, or with a group of winning neurons, i.e. in a soft WTA setting. This type of computation has been proposed in hierarchical models of vision (Riesenhuber and Poggio, 1999a), and models of selective attention and recognition (Carpenter and Grossberg, 1987).

Fig. 3.1 shows how a constant input to the neurons is converted to a Gaussian output curve. Some silicon neurons of the ROLLS chip are more excitable than others, due to mismatch. Hence, in a WTA configuration the neurons that fire the most amplify their neighbours, while weaker neurons

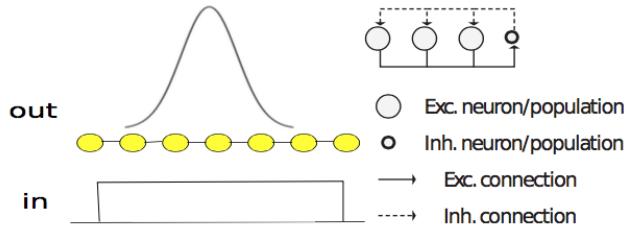


FIGURE 3.1: A constant step input is being converted to a Gaussian output by the recurrent connections of neurons in a WTA configuration and because of the neurons mismatch. The left side shows how the step input is converted to a Gaussian output when it is fed to recurrently connected neurons. The right shows the excitatory and inhibitory connections underlying the WTA configuration.

get further suppressed. At the right of the figure one can see the recurrent connections, that establish this kind of behaviour.

In this project a WTA implementation is used for an unsupervised learning task of pattern recognition. A soft WTA implementation which enables VLSI systems to selectively enhance the contrast between inputs (Chicca et al., 2014) helps the network to autonomously find a hidden structure in the different input patterns. Excitatory neurons were connected via excitatory recurrent connections to themselves resulting in self-excitation, and to their three nearest neighbors. At the same time there is an excitatory all-to-one connection from the excitatory neurons to the one inhibitory neuron, which globally inhibits all the excitatory neurons. With setting the biases of the synapse circuits one can specify the amount of inhibition, resulting in the sustained activity of a desired number of neurons. Fig.3.2 shows a schematic of the recurrently connected neurons.

Fig.3.3 shows a raster plot of the postsynaptic spikes of 63 neurons configured in a WTA network where 20 neurons receive a gaussian input with maximum 110Hz and other 20 neurons receive a gaussian input with maximum 10Hz. As soon as neurons start to fire the inhibitory neuron is active, suppressing the activity of all neurons. Despite the great inhibition from the inhibitory neuron, some neurons that receive a higher input frequency manage to stay active until the input is gone. The neurons that receive lower input frequencies are suppressed immediately and do not manage to become active. Mismatch and high excitation of neighbouring neurons leads to a small shift in the active group compared to the one that actually receives the highest input.

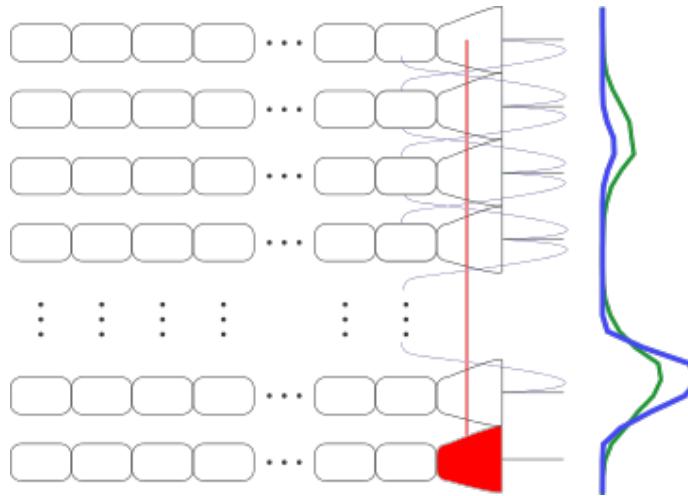
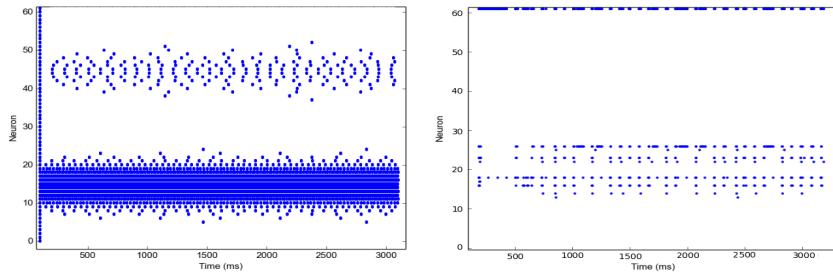


FIGURE 3.2: A schematic of the neuron’s recurrent connections on the device. Recurrent connections are shown by connections of the neurons soma. The neuron whose soma is coloured red represents the gloabl inhibitory neuron. The input to the neuron is given as two Gaussian curves with different maxima (green curve). The global inhibition and local excitation leads to an amplification of the the higher input and a suppression of the lower input. The resulting output is given as the blue curve.

In this setting of unsupervised learning, it is desired that neurons modify their weights according to different postsynaptic spiking frequencies that were initiated by different input patterns. After the stimulus is gone, and a weight update was performed, the neurons should stop firing, so that a new input stimulus can be presented and the firing for the previous input does not interfere with the firing for the next input pattern. The time constant for the excitatory connections must be very high and the time constant of inhibitory connections must be very low, so that excitation is slow whereas inhibition is fast. Thus, the excitatory neurons can be quickly inhibited. This inhibition results in suppressed activity which in turn leads to a smaller excitation of the inhibitory neuron. Finally, the inhibitory neuron does no longer receive input and is shut off. A new input pattern can be presented. The time of the stimulation as well as the strength of inhibition should be just enough to learn LTP and LTD at the winning neurons. A too long stimulation will result in the forgetting of previously learned patterns because neurons will quickly synchronise their postsynaptic firing frequency due to the recurrent connections. More and more neurons will learn the same input, and the increasing postsynaptic firing rate will lead to



(A) Input spiketrain consisting of two gaussian bumps used to postsynaptic spikes. The 10Hz stimulate 63 neurons in an WTA gaussian bump is clearly super-configuration. The stronger pressed, i.e. the neurons receiving gaussian bump (centered around 10Hz) remain silent, while the neuron 15) has a maximum frequency of 110Hz whereas the weaker one (centered around 110Hz) stay active until the end of stimulating. (B) Raster plot of the resulting configuration. The inhibitory neuron (number neuron 45) has a maximum frequency of 10Hz. 5% of plastic synapses are initialized to be randomly potentiated. Input is given to all 256 plastic synapses of each neuron.

FIGURE 3.3: WTA behaviour on the ROLLS chip

potentiation on all synapses. Instead, only a few synapses should potentiate, in order to account for high and low frequencies of the input. The high inhibition results in sparse firing and not many neurons are able to perform learning transitions. This sparse firing is needed, in order to be able to learn many different patterns. The firing of too many neurons would also result in all of them learning the first pattern, leaving only a few neurons that can potentially learn the next pattern. The solution to arrive at local groups of neurons, that encode the same pattern is to present the different input patterns over many trials. The potentiated weights of the previously winning neurons lead to increased activity so that also the activity of their neighbors is being increased. Thus a broader group of neurons fires for a specific input.



3.2 Network Configuration on Chip

In this project 252 silicon neurons were used to implement four WTA networks. Four WTA's were chosen in order to examine robustness and differences in the network behaviours due to mismatch in the underlying circuits. Each WTA comprises 62 excitatory neurons and a single global inhibitory neuron. Excitatory neurons are self-excitatory and have excitatory connections to their three nearest neighbours, all of which are having the same connection strength. The inhibitory neuron receives excitatory input from all excitatory neurons and sends out the same amount of inhibition to all excitatory neurons. These connections were made via programmable synapses and stay fixed during learning.

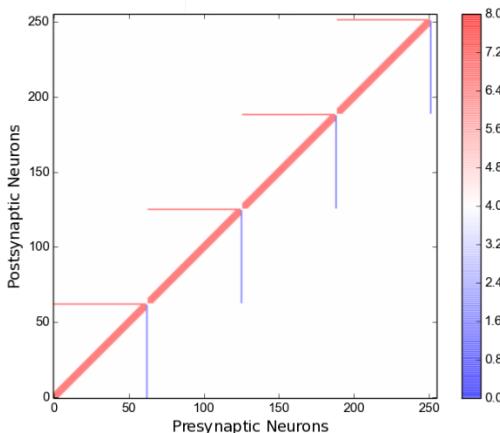


FIGURE 3.4: Matrix of recurrent connections of all 256 neurons. One WTA network consists of 63 neurons. The excitatory neurons are connected to their three nearest neighbors and to the global inhibitory neuron. The global inhibitory neuron is connected to all excitatory neurons. Blue represents an inhibitory connection whereas red means excitatory.

Fig.3.4 shows the recurrent connection matrix of the programmable synapses and Fig.3.5 shows a clipping on how this connectivity is set on the ROLLS chip. External frequency input patterns are used to stimulate the plastic synapses, so that neurons can learn an association between input frequency and the exact synapse which is stimulated by this frequency. To make firing activity possible, we initialised 5% of the plastic synapses to be randomly potentiated. Thus, we only need a high gain of the plastic synapses, but do not require to artificially initiate postsynaptic spikes on the neuron. This random initialisation of weights is supported by biological models.



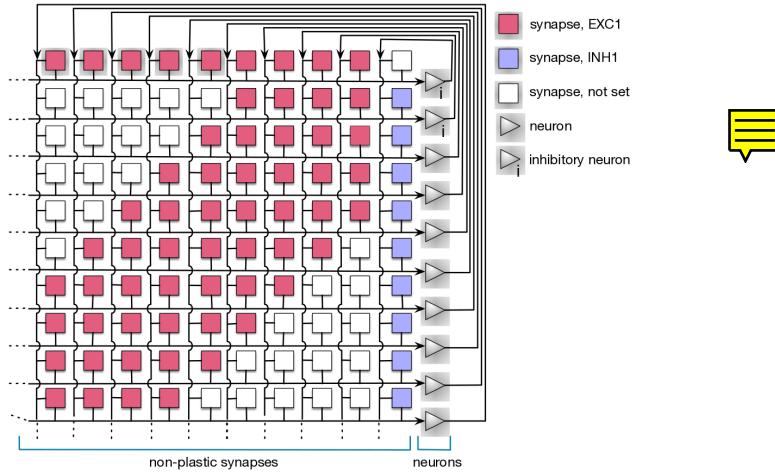


FIGURE 3.5: Setting the WTA connectivity on the ROLLS chip.

Song et.al suggest that STDP, acting without hypothetical constraints on randomly potentiated synapses, can reproduce the remapping observed in adult cortex (Song and Abbott, 2001). This behaviour can also be seen in the postsynaptic firing of our network. In the first trials, where every neuron was stimulated through the plastic synapses with the same input, the random weights lead to almost equally distributed firing rates across all the neurons. However, mismatch in the underlying circuits makes some neurons fire more than others, and the WTA connections further enhance the difference of these firing rates. Thus the neurons with a high firing rate potentiate, whereas the ones with a low firing rate depress. After many trials structure emerges and we are left with only a few groups of neurons that respond to a given input. In the section 3.4 it will be explained that also the input frequency determines the probability of LTD and LTP transitions. This is especially needed for pattern recognition tasks, where the patterns are made of different input frequencies.



3.3 Pattern Classification Task

To demonstrate the mechanism and the performance of unsupervised learning on the ROLLS chip, a pattern classification task was applied. We differentiate between two different input patterns. Input patterns are encoded as sets of Poisson spike trains that stimulate the neuron's plastic synapses with different mean frequencies. The neuron's postsynaptic firing rate represents the classification of the input pattern.

The two input patterns are shown in Fig.3.6. All neurons are stimulated on

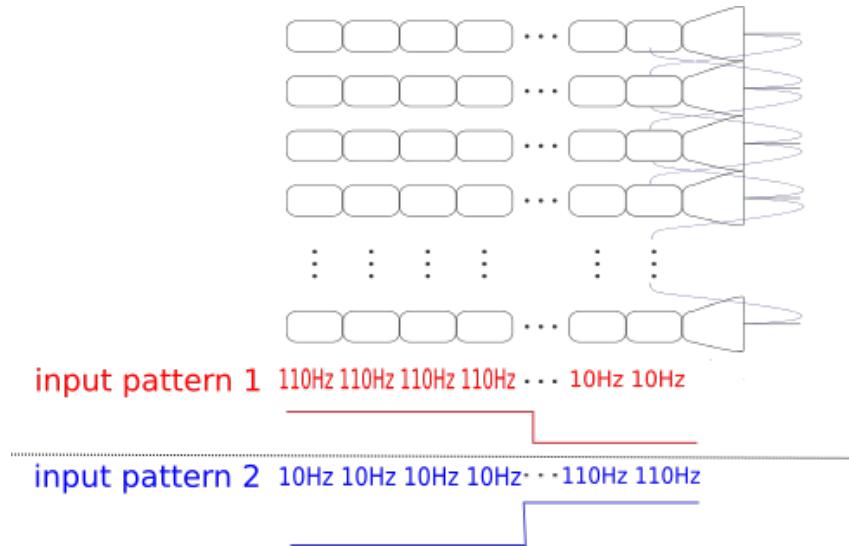


FIGURE 3.6: The schematic of the recurrently connected excitatory neurons receiving either input pattern 1 or input pattern 2 on their plastic synapses.

their plastic synapses with the same input pattern. The goal is that some neurons learn to fire for a given input while other neurons learn to fire for another input pattern depending on the input pattern with which the neurons are stimulated. The neurons implicitly learn to encode different features of the input, which the neurons can use later on, to infer about the similarity of patterns. There can be 256 input features encoded by the 256 plastic synapses of each neuron. Similar to cortical neurons, every silicon neuron on the chip is not exactly the same (Goris, Movshon, and Simoncelli, 2014). This variability, which is mainly due to mismatch in the circuits, makes the neurons fire in slightly different ranges although the input stimuli are the same for each neuron. The recurrent connections that were used to obtain the WTA behaviour enhance this variability, resulting in local groups of neurons firing at a similar frequency range. The biases that implement the learning rule were set in such a way that the synapses receiving 110 Hz potentiate and the synapses receiving 10 Hz depress their weights when the neuron is firing at a specified postsynaptic rate. The neurons that are below a certain postsynaptic firing can not make a weight transition due to the θ_{min} threshold. If the postsynaptic firing of these neurons is so weak that the calcium concentration does not exceed this threshold, a weight transition cannot take place. These neurons keep their initial random weights, so that they can still respond to the next input pattern. After stimulating the plastic synapses of all excitatory neurons for 150 ms with input pattern one, input pattern two is used for stimulating, again all excitatory neurons

for 150 ms. The neurons that learned input pattern one will not respond to input pattern two because some of their synapses that previously received 10 Hz depressed. Others still have randomly potentiated weights at their synapses, which are now receiving 110 Hz input. These might be the new ?winning? neurons. Thus, different winners are found for different input patterns. After stimulating multiple times, neighbouring neurons encode the same input pattern, resembling mapping structures, like retinotopy that was found in the visual cortex (Engel, Glover, and Wandell, 1997).



3.4 Experimental Results

3.4.1 Synapse Dynamics and sensitivity to the parameters

The goal of the pattern classification task was, to show that neurons closely connected in a neighbourhood learn to encode a specific pattern while another group encodes another pattern without requiring a teacher signal. The circuit parameters underlying the learning rule, the recurrent connectivity and the soma are working together to achieve this behaviour. Different input frequencies require different parameters for learning, since the input directly influences the postsynaptic spiking frequency as well as the presynaptic spiking frequency. The biases where set in a way that LTP transitions at the synapses are likely for a 110 Hz input rate if the postsynaptic spike frequency was greater than 8 Hz and LTD transitions are likely if the input frequency was 10 Hz and the postsynaptic rate was smaller than 30 Hz. The goal was to maximise the LTD and LTP transitions for these input frequencies, but they can certainly be changed if the networks has to learn different values of input frequencies.

Fig.3.7 shows the average probability of potentiation and depression of each neuron depending on its input and output frequency. To arrive at these curves, the recurrent connections of the neurons were removed and all plastic weights where set to w_{min} . The input frequency to the plastic synapses was swept from 10 Hz to 120 Hz in steps of 10 Hz, and for each of these input frequencies the neurons were at the same time stimulated via virtual synapses with frequencies that where swept from 200 Hz to 600 Hz in steps of 100 Hz to obtain different postsynaptic spiking frequencies. For the measurement of LTP, the postsynaptic frequency of each neuron was recorded. Then it was counted how many of each neurons 256 synapses potentiated to

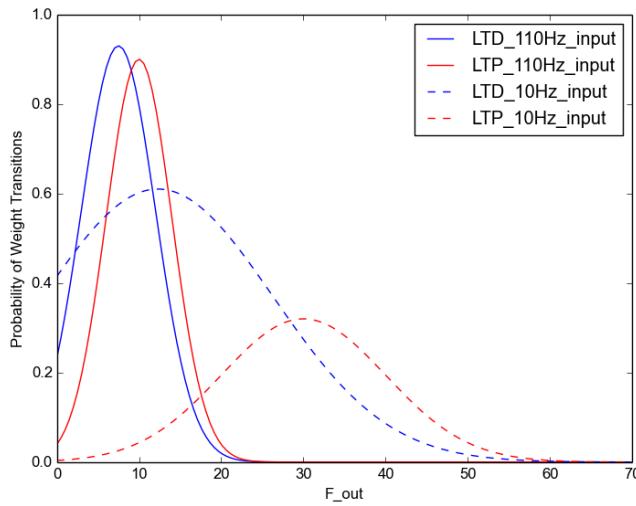


FIGURE 3.7: The blue curves are two Gaussians fitted to the probability measures to make LTD transitions due to a 110 Hz and a 10 Hz input frequency respectively. The red curves are two Gaussians fitted to the probability measures to make LTP transitions due to a 110 Hz and a 10 Hz input frequency. These probabilities depend on the postsynaptic firing rate of the neuron.

obtain the potentiation probability. The same was done for the LTD learning curves, with the only difference that the weights were initially set to w_{max} and it was recorded how many of the 256 learning synapses of each neuron depressed. The learning behaviour was achieved by finding the right biases of the circuits. A description on how to find the optimal biases can be found in appendix A.

In the learning curves from Fig.3.7 it becomes clear that there is only a very small range for LTP to be more likely than LTD. Higher postsynaptic frequencies are generally more likely to occur after many preceding presynaptic spikes, therefore higher postsynaptic frequencies together with a 110 Hz input are more likely to lead to LTP than to LTD. Since the plastic synapses are used for stimulation in the learning task many presynaptic spikes will automatically lead to a higher postsynaptic frequency followed by LTP, and a smaller presynaptic frequency will directly lead to a smaller postsynaptic frequency, followed by LTD.

The small range where LTP is more likely than LTD is due to the high inhibition in the WTA network. The inhibition has to be high in order to only find a few winning neurons where learning is allowed to take place. In addition,

the same neuron has to perform LTD and LTP transitions for the same postsynaptic frequency, only by differentiating the input to the plastic synapses. This only happens in a very small postsynaptic range since the postsynaptic firing rates are normalised by the WTA network and also because every neuron receives the same input. However, the inhibition and the WTA behaviour help to filter out exactly these few neurons spiking in the optimal range to underly both LTP for high input frequencies and LTD for low input frequencies. These transitions occur at a postsynaptic frequency around 10 Hz. The curves in Fig.3.7 show that at 10 Hz postsynaptic frequency, there is a maximum for LTD transitions for 10 Hz input and a maximum for LTP transition for 110 Hz input. The red LTP curves for 10 Hz and 110 Hz input have minimal overlap. This is very import to avoid that neurons learn potentiate for both input frequencies. 10 Hz input frequencies require a very high output frequency to make LTP transitions. This is avoided in the WTA network, because the global inhibition makes suppresses activity making high output frequencies impossible.

LTP is generally less likely to occur for small input frequencies. Its maximum value for low input frequencies is just 0.3 whereas the maximum for LTD under the condition of a high input frequency is about 0.9. These specific conditions that lead to LTP and LTD on different synapses but on the same neuron limit the postsynaptic firing range that can be used. At the same time, it is desired to make the WTA behavior as strong as possible to obtain only a few winning neurons for a given input pattern. A strong WTA has high inhibition, because high excitation directly increases the inhibition. Therefore postsynaptic firing rates are low. The highly restricted postsynaptic frequency range is not a problem since it makes learning only depend on the input frequency, the current state of the weight and its gain. The postsynaptic rate is normalised directly by these factors, it cannot become too high because of the inhibition that is initiated by the overall activity of the neurons. The high gain of the the potentiated weights will make neighbouring neurons spike at a similar frequency, so that they are able to learn the same input pattern.

3.4.2 Classification performance

In this section learning and classification of the two input patterns will be evaluated on four different WTA networks.

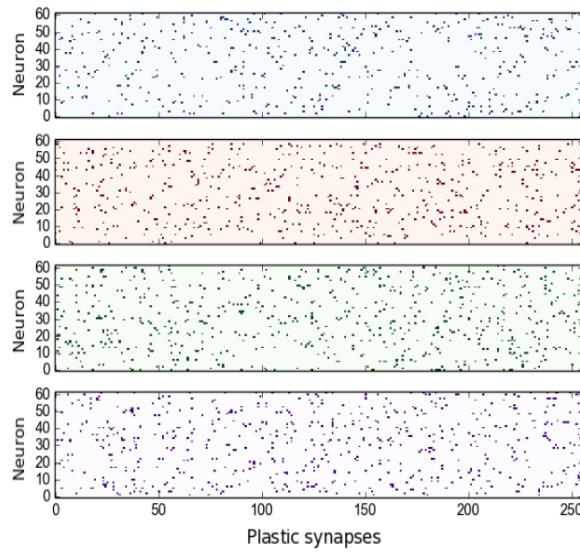


FIGURE 3.8: The synaptic weight matrix after initialization.

Fig.3.8 shows the initial plastic synaptic weight matrix with 5% potentiated synapses.

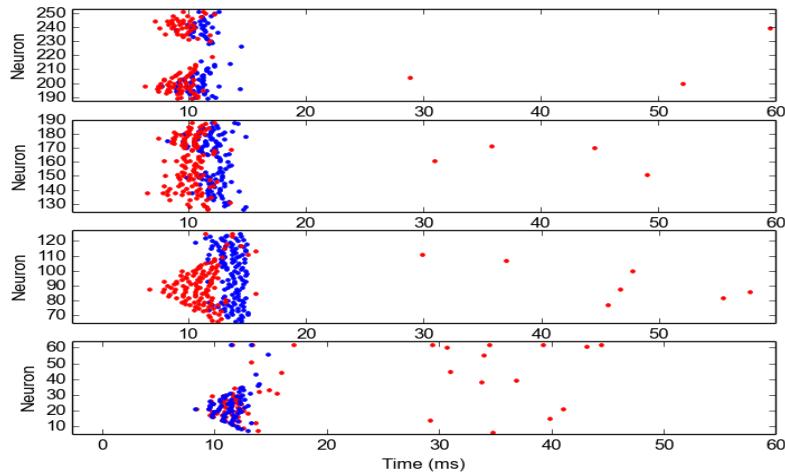


FIGURE 3.9: Maximal achievable WTA behaviour due to mismatch. Spiking for input pattern one is shown in blue, spiking for input pattern two is shown in red.

Fig.3.9 shows the postsynaptic spikes after presenting input pattern one (blue) and after presenting input pattern two (red). In this scenario some neurons are clearly more excitable due to mismatch. They excite their neighbors until the excitation to the inhibitory neuron is so strong that all activity is being switched off. Inhibition is almost immediate. Since the input pattern is used to stimulate for 150ms firing activity occurs again after the

inhibition has decreased to zero. Only the first impulse of activity is shown here.

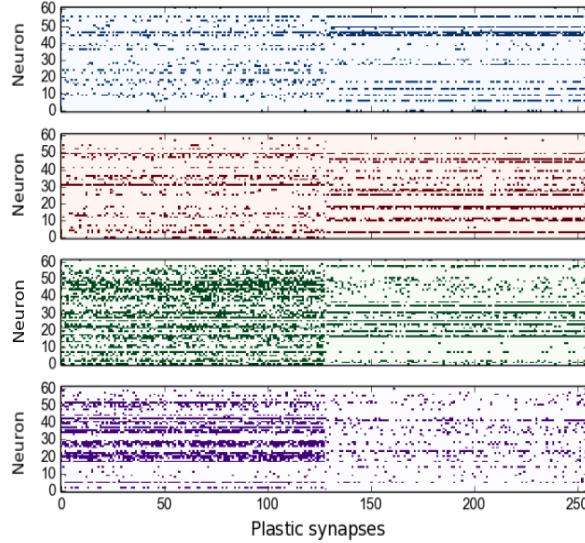


FIGURE 3.10: Synaptic weight matrix after two trials, i.e. stimulating one time with input pattern one followed by input pattern two, each for 150ms.

Fig.3.10 shows the synaptic weight matrix after these two trials, i.e. stimulating one time with input pattern one followed by input pattern 2, each for 150ms.



Fig.3.11a shows the synaptic weight matrix after fourteen learning trials. Whereas after only one stimulation with input pattern one and two, the weights are mostly potentiated for the first pattern, the network finds an equilibrium to equally potentiate synapses for pattern one and for pattern two. The high gain of plastic synapses leads to the activation of more neighbouring neurons and so to the emergence of clusters. However, after approximately fourteen trials training should be stopped since more neighbors are being activated and thus potentiated. This leads to the fact, that the latest pattern potentiates more and more neurons, leaving no capacity to learn another pattern. Fig.3.11a also shows that the high mismatch of neurons leading to different activity also leads to different learning behaviour. While the WTA shown in green perfectly learns to differentiate between two patterns, the other three WTA's are much less active and therefore are much slower in learning the two patterns. Slightly modified biases, which make the excitability of the neurons higher lead to the weight matrix shown

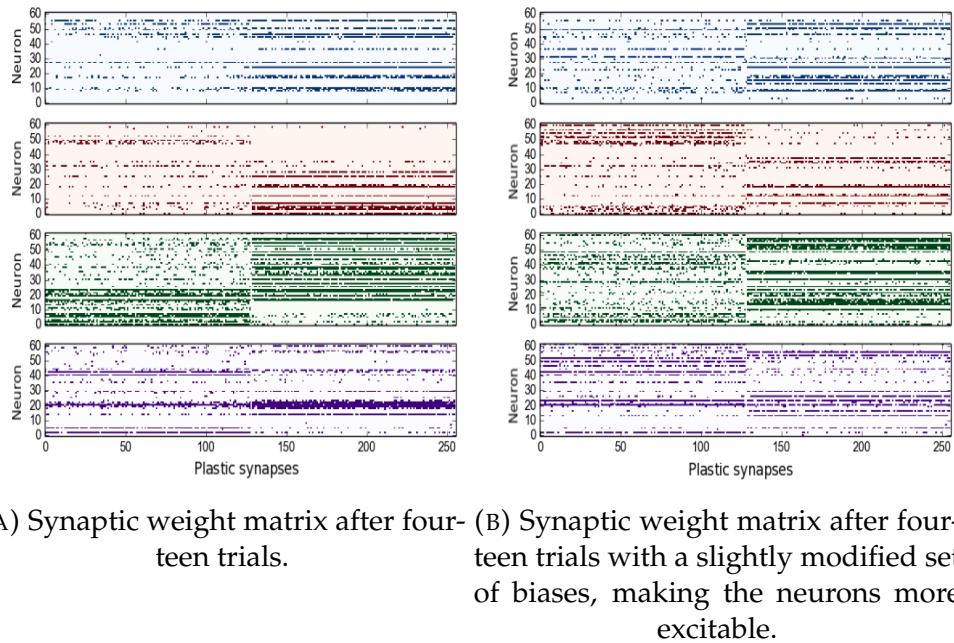


FIGURE 3.11: Synaptic weight matrix after fourteen trials with different sets of biases.

in Fig.3.11b. The WTAs show a more similar behaviour, there are approximately two clusters for each WTA that correspond to input pattern 1 and two clusters that correspond to input pattern 2.

There is a tradeoff between optimal WTA behaviour and the number of learned input patterns. In the optimal WTA, clusters of neurons learn to potentiate for one input immediately. It turns out that these neurons are the most excitable neurons of the circuit. Hence, the next input has much less neurons available for learning and these are also less excitable. To obtain an equal number of neurons to learn the different patterns, the cluster sizes become smaller and firing rates in the testing phase become weaker. Another important feature of the learning is that broader groups of neurons should potentiate over time, after presenting the input patterns many times. This leads to a stable solution of learned inputs and amounts to an equal number of neurons learning each pattern. This behaviour is obtained by changing the parameters responsible for the weight gain. Potentiated weights lead to an increased firing of the neuron when the input frequency to the potentiated weights is high. This in turn leads to an increased firing of the neighbors, due to the WTA configuration, which again leads to learning in these neighbouring neurons.

To test if the neurons, which learned a given input pattern, fire only for

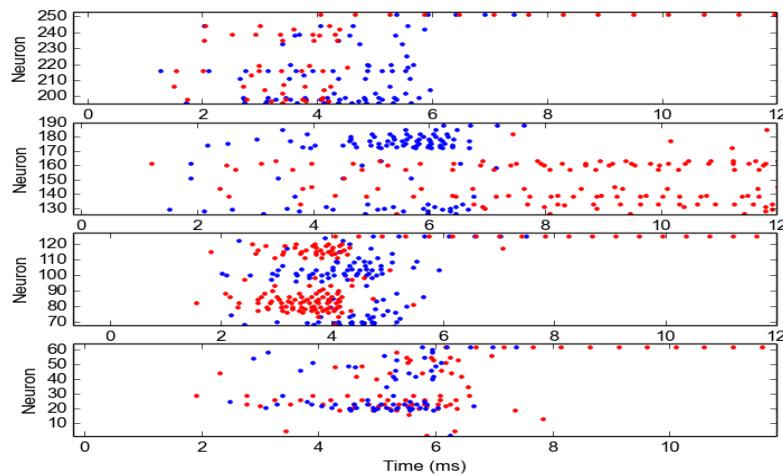


FIGURE 3.12: Raster plot of the output spikes after stimulating with input pattern one (blue) and input pattern two (red) for 1000ms in the testing phase. Only the first impulse of activity is shown. Each pattern was shown 13 times for 150ms during learning. The inhibitory time constant is much smaller than the excitatory time constant.

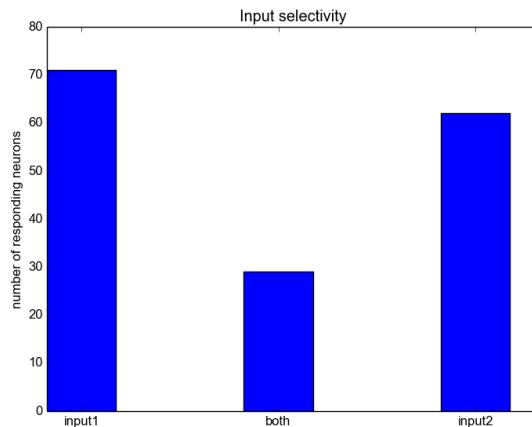


FIGURE 3.13: Histogram of the selectivity of the neurons: number of neurons responding to input pattern one (left), input pattern two (right), and to both input patterns (middle).

this exact pattern, θ_{min} was set to its maximum to stop the learning process. In addition, the inhibition was decreased to allow higher output frequencies, so to see a clearer difference between specialised and non-specialised neurons. After learning the weight matrix shown in Fig. 3.11b and stimulating the neurons with input pattern one and two, the resulting raster plot is shown in Fig. 3.12.

The learned weights were used to study the dynamics of the WTA using different time constants. Throughout the experiments, we chose the time constant for inhibition to be much smaller than the time constant of excitation. Evidence was found within superficial layers of neocortex (Rutishauser, Slootine, and Douglas, 2012) where inhibitory time constants are usually smaller due to smaller cell bodies of inhibitory neurons (McCormick et al., 1985; Koch, Rapp, and Segev, 1996). Smaller inhibition is also used to arrive at Fig. 3.12.

Fig. 3.13 shows a histogram of the selectivity of neurons. It was compared which neurons are only active for input one, which ones are only active for input two and which neurons became active for both input patterns. The selectivity is improved when clusters become larger and better separable in space for different input patterns. The achieved clustering behaviour is comparable to state-of-the-art machine learning algorithms based on DBSCAN, where points that are closely packed together in space are clustered together(Ester et al., 1996).

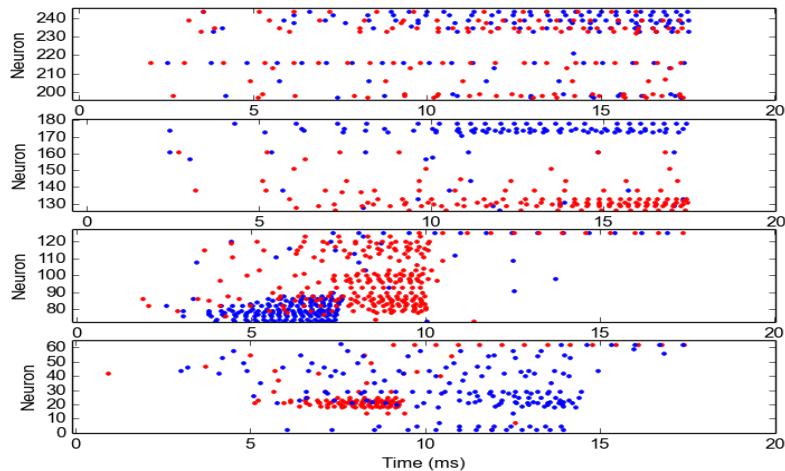


FIGURE 3.14: Firing activity during testing, after having learned the weight matrix shown in Fig.3.11b, using fast excitation and slow inhibition. Neurons quickly excite their neighbors, and as soon as the slower, but very strong inhibition becomes active, it switches off all activity.

Another scenario is shown in Fig. 3.14. Here, excitation is much faster than inhibition. Therefore, neurons have more time to excite their neighbors until being switched off. The firing activity after learning the weight matrix of

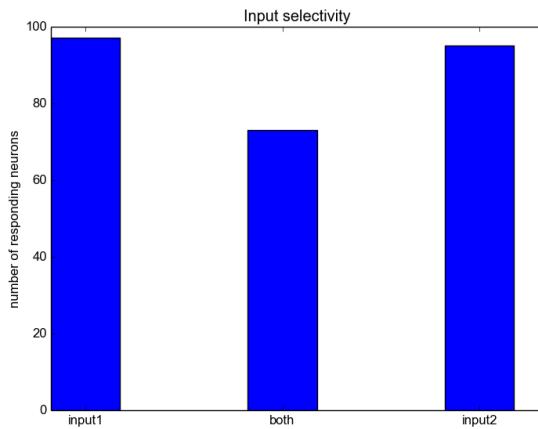


FIGURE 3.15: Histogram of the selectivity of the neurons, responding to input pattern one (left), input pattern two (right), and to both input patterns (middle).

Fig. 3.11b is shown in Fig. 3.14. Fig. 3.15 shows the corresponding selectivity histogram. Using fast excitation and slow inhibition leads to more density in activity, winners for each pattern can be seen more clearly. However, this also leads to more overlapping activity, meaning that many neurons fire for the same pattern due to the longer time they have to excite neighbors which do not belong to the learned pattern. This is another argument, why excitation should not be faster than inhibition during the learning phase. In fact, this would mean that the same winner would be found for every pattern and different clusters would not slowly evolve over time, but rather the same clusters would be found for both patterns after the first learning trials.

Similar time constants are not appropriate for this kind of learning. Making excitation only slightly faster than inhibition leads to oscillatory behaviour. If the inhibition strength is set to strong enough biases, activity can eventually be switched off. Only the most strongly activated neurons fire in the last oscillation. This behaviour is shown in Fig. 3.16, where blue spikes correspond to the activity after being stimulated with pattern one, and red spikes correspond to the activity after being stimulated with input pattern two. Although, most neurons again fire for both patterns, it can be recognized that some groups fire more for one input pattern than for another.

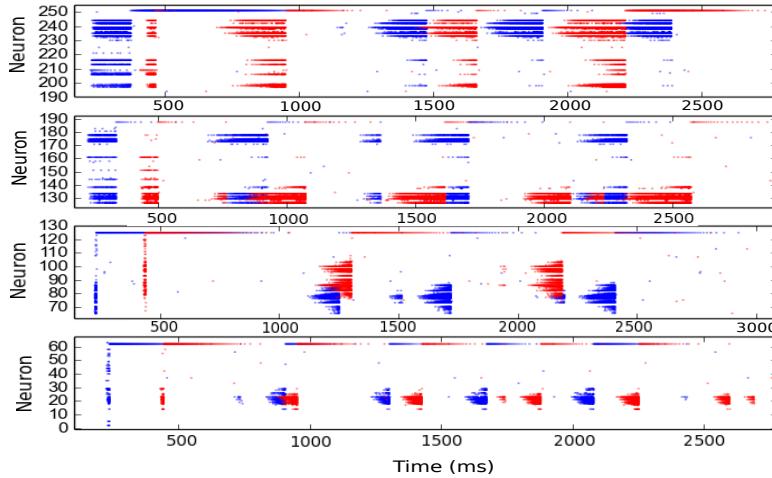


FIGURE 3.16: Activity after having learned the weight matrix shown in Fig.3.11b and being stimulated with input pattern one (blue) and input pattern two (red). The time constants of excitation and inhibition are only slightly different and the stimulus onset is slightly shifted. Input pattern two starts 300 ms later than input pattern one.

3.4.3 Learning four different patterns

In the following it will be shown that each WTA comprising 63 neurons, is able to learn to differentiate between more than two different patterns in an unsupervised manner. Patterns consist of high and low frequencies given to different synapses. Here we used again frequencies of 10 Hz and 110 Hz. Each neuron has 256 learning synapses, meaning that there are infinitely many possible patterns that could be learned. However, the number of neurons is limited and mismatch leads to a variability in learning and firing activity, which worsens the performance. However, we can show that a WTA comprising only 63 neurons can at least learn to differentiate between four different input patterns. The four different input patterns are shown in Fig. 3.17. The learned synaptic weight matrices after alternately stimulating for 200ms with each pattern for five times is shown in Fig. 3.18 on the left. After stimulating with each pattern for ten times, the weight matrix changed as shown on the right of Fig. 3.18. Next, the biases of the circuit were slightly modified to stop the learning process and to enhance the excitability. The same input patterns were used for stimulating the neurons for 1000ms.

Fig.3.19 shows histograms of the neurons selectivity to four input patterns. Mismatch leads to the fact that highly excitable neurons are active for all

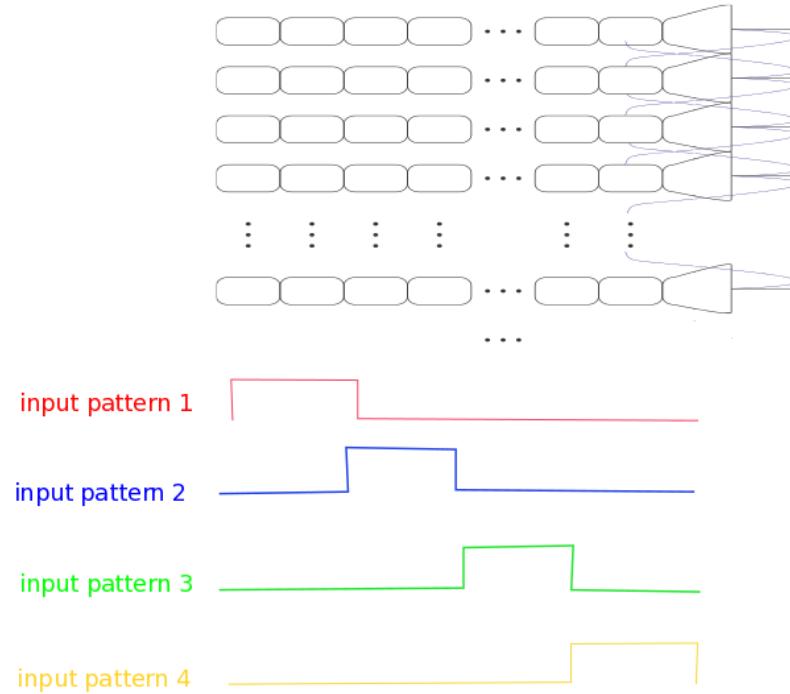
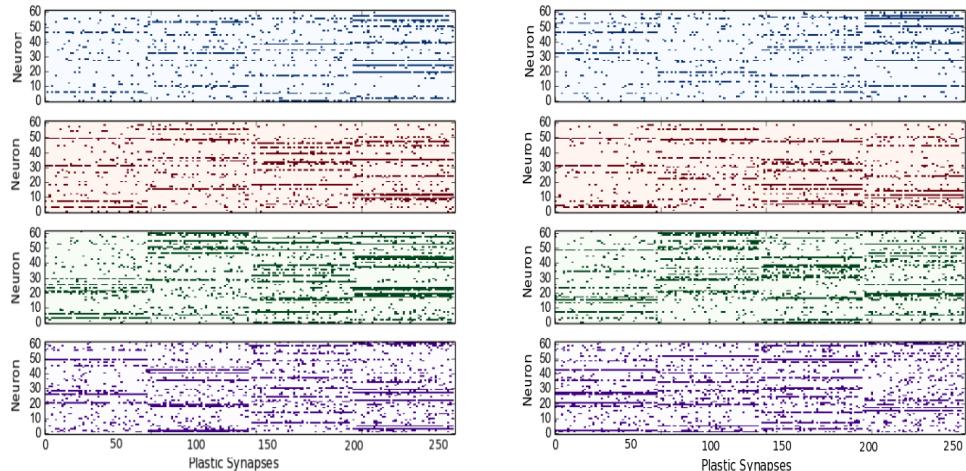


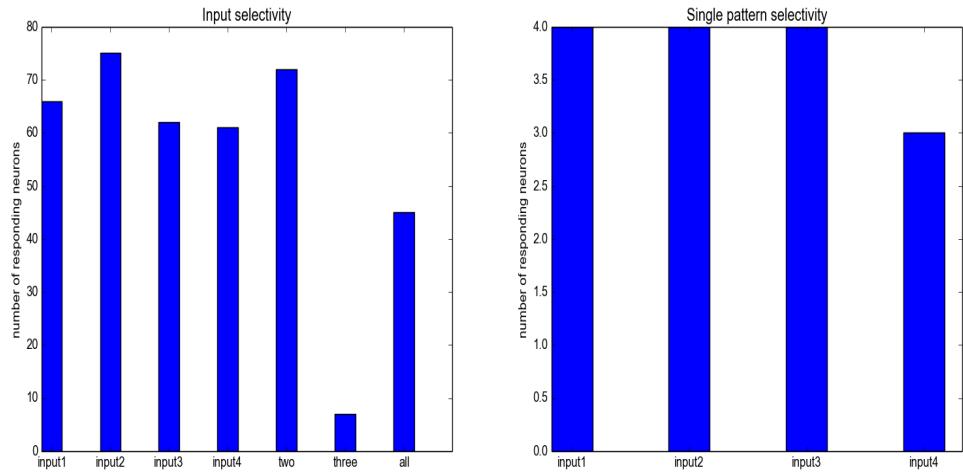
FIGURE 3.17: The four different inputs used for stimulating the learning synapses. The step signal is given to one fourth of learning synapses. It is a 200 ms Poisson spike train with a frequency of 110 Hz, the other synapses are stimulated with Poisson spike trains of 10 Hz.



(A) Plastic weights after 4 different input patterns were alternately used to stimulate the synapses. Each pattern was presented five times for 200 ms.

(B) Plastic weights after 4 different input patterns were alternately used to stimulate the synapses. Each pattern was presented ten times for 200ms.

FIGURE 3.18: Weight matrix of the plastic synapses after learning 4 different input patterns.



(A) Number of neurons that fire for each combination of input patterns. (B) Number of neurons that show perfect selectivity to only one input pattern.

FIGURE 3.19: Histograms on the selectivity after learning four different input patterns.

patterns. However, Fig.3.19b shows that there are a few neurons selective only for one pattern.

3.4.4 Discussion

In this chapter we have shown that the silicon neurons of the ROLLS chip can be configured in a WTA network helping the neurons to find a hidden structure in input patterns. Dynamics were discussed on the single neuron, learning and network level and they were optimised by setting the transistor biases of the circuits in order to solve an unsupervised pattern recognition task. With this knowledge questions arise to what extend the network is able to code for similarity. Similar spiking activity of different groups of neurons is initiated when stimulated with similar patterns. Further experiments will lead to clarity to what extend this similarity coding can be used. In addition, the ROLLS chip can be used for pattern classification on datasets where white pixels are coded as high frequencies and black pixels as low frequencies. Common datasets can be used to evaluate the accuracy. Furthermore, larger WTA networks can be used in order to learn more different patterns as the dataset sizes or the data complexity increases. Another outlook of this work is to build hierarchies. Different



WTA networks that learn different features can be coupled so that eventually neurons in **high-level hierarchies** become able to learn more complex features, just like it is the case in the visual system (Van Essen, Anderson, and Felleman, 1992).

In comparison to common machine learning clustering algorithms like k-means or DBSCAN, this model does not require previous knowledge of the data, such as the number of clusters or the epsilon radius and the minimum number of points required to form a dense region, like in DBSCAN (Ester et al., 1996). Clusters are learned via an unsupervised STDP learning rule by spiking neurons in a WTA configuration. Another interesting property that arises from our model is that one pattern is learned by more than one cluster of neurons. This overdetermined feature leads to robust recognition, even if one group of neurons fail, the pattern will still be recognized by another group. The more patterns are being learned, the less clusters are available for learning each pattern. Thus, the pattern classification becomes vulnerable for many patterns. In order to perform unsupervised pattern classification for many patterns, more silicon neurons are needed. As the technology of building neuromorphic devices progresses, it might soon be possible to perform robust and reliable unsupervised machine learning tasks directly on hardware. This would contribute to a fast and low-power solution, being able to perform complex machine learning tasks.

Looking at each single WTA it can be seen that the learning behaviour between the WTA's is quite different. This is a result of the random initial weights as well as of the mismatch of the transistors in silicon synapses and in the silicon neuron. However, in this unsupervised learning mechanism we make use of the slightly different properties of the circuits. Instead of obtaining precise behaviour of each building block, we use the neural network as a whole to arrive at a reliable and robust result. In addition, the network is robust in the sense that the failing of one building block will not have any influence on its overall learning capability.

Chapter 4

Sequence Learning

4.1 Conceptual model

A neural-dynamic architecture for serial order working and long-term memory was used to implement sequence learning on the ROLLS chip. It is based on dynamical neural field theory (Sandamirskaya and Schoener, 2010; Duran and Sandamirskaya, 2012) and its architecture is shown in Fig. 4.1.

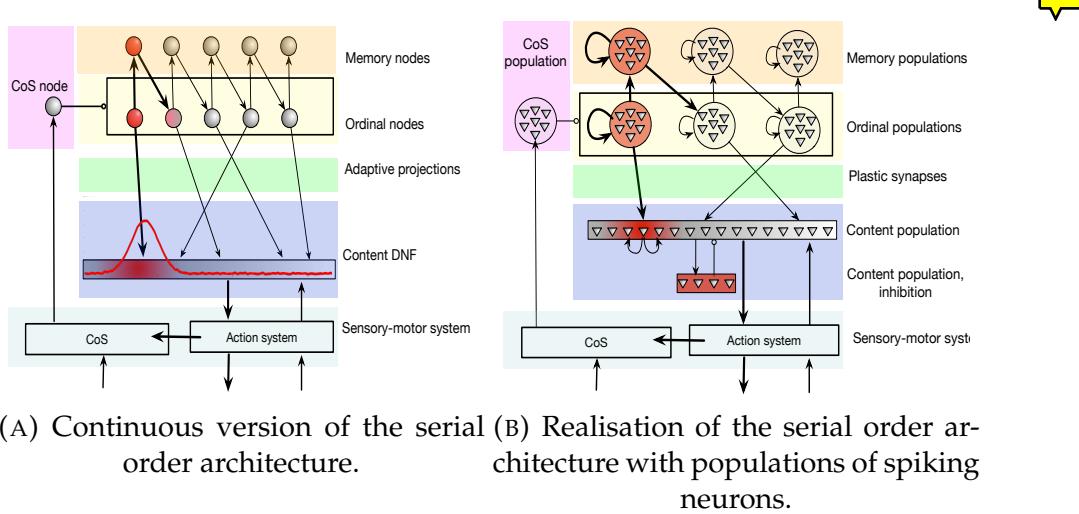


FIGURE 4.1: The serial order architecture, introduced in Sandamirskaya and Schoener, 2010.

Groups of ordinal neurons represent position in a sequence. They have all-to-all synaptic connections to content neurons in a dynamical neural field. These content neurons represent the perceptual state or action that has to be stored in the given location (modeled as ordinal groups). The content DNF is connected to the action system of the agent and sets attractors to generate behavior. The synaptic weights are initially set to zero but the synapses of active ordinal neurons to active content neurons will potentiate when these

neurons are active at the same time. The start of a sequence is triggered by external excitatory input to the first ordinal group. At the same time an external excitatory signal is given to some content neurons. The content neurons are in a WTA configuration so that only a group of neurons stays active while input noise will not lead to the firing of neurons. The ordinal groups are in a WTA configuration as well, to make sure that only one ordinal group learns one teacher signal at a time. The ordinal group then activates its memory group via excitable connections. Memory groups are firing independently of each other and they remain active for the whole cycle due to their high self-excitation. This memory constantly remembers that its element has already been active. At the same time each active memory group slightly inhibits its element, to make sure that the excitation from the most recent memory to the next item is stronger than the excitation from a very previous memory which already activated an item. A condition of satisfaction system detects when each action or perceptual state has reached the intended outcome (encoded in the coupling between the action and the CoS system, (Richter, Sandamirskaya, and Schoner, 2012)). Then the CoS node triggers a sequential transition by inhibiting the ordinal nodes.

The position of the content neurons can be used for encoding. For each active ordinal group, some content neurons will be stimulated in order to learn an association. After the complete learning phase, an external input is send to an inhibitory group of neurons to switch off all activity. The recall is then being triggered by an excitatory stimulation of the first ordinal group. This model is inspired by neuronal findings about serial order encoding in the cortex (Carpenter, Georgopoulos, and Pellizzer, 1999; Beiser and Houk, 1998) and behavioral findings on serial order errors (Henson, 1998).

Fig. 4.2 shows how the serial order architecture can be implemented on the ROLLS chip. Fig. 4.2a shows the architecture for a sequence where three elements are learned and Fig. 4.2b where a sequence of five elements is learned. In the figures, a group of 100 (or 70) neurons in the upper right corner forms the content DNF which is implemented by a WTA network as described in section 3.1.

4.1.1 Dynamic Neural Fields

In this work, we use the conceptual and mathematical framework of Dynamic Neural Fields to realise a cognitive architecture on neuromorphic

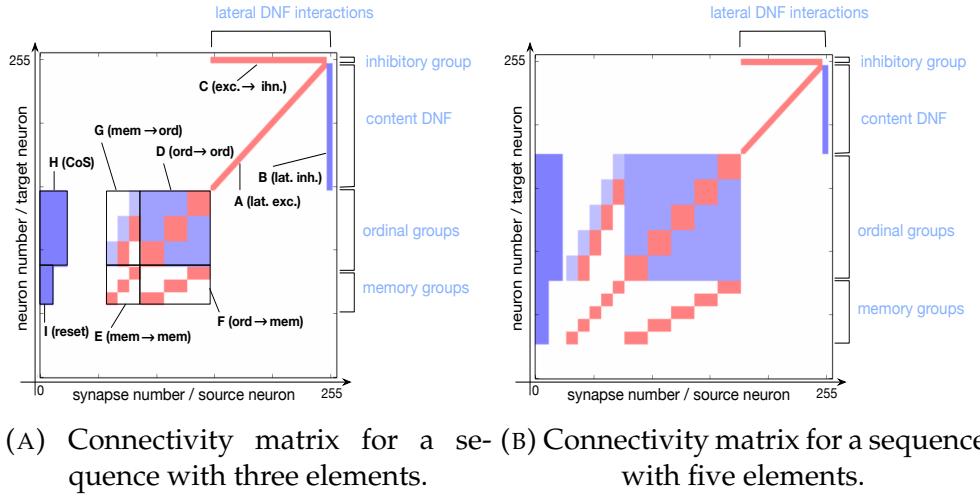


FIGURE 4.2: Connectivity matrix, sent to the ROLLS chip to encode the serial order architecture

hardware. The analytical studies were carried out by Yulia Sandamirsakya (Sandamirsakya, 2013b). In a Dynamic Neural Field (DNF) a continuous activation function following the dynamical system equations, accounts for the activation of a neuronal population, Eq. (4.1):

$$\begin{aligned} \tau \dot{u}(x, t) = & -u(x, t) + h \\ & + \int f(u(x', t)) \omega(x, x') dx' + S(x, t). \end{aligned} \quad (4.1)$$

Where $u(x, t)$ is the activation function at time t of a DNF. The parameter space x describes the state of the system and $-h$ is a negative resting level, setting values of $u(x, t)$ to be below zero. This can be regarded as an output threshold that keeps noise from initiating activation when no external input is provided. $S(x, t)$ is the external input and $f(u)$ is a sigmoidal non-linearity, Eq. (4.2) transforms the output of the DNF to be a sigmoidal. The output is zero for negative values of $u(x, t)$, positive for positive $u(x, t)$ and saturating for large values of $u(x, t)$:

$$f(u(x, t)) = (1 + e^{-\beta u(x, t)})^{-1}. \quad (4.2)$$

$\omega(x, x')$ is the interaction kernel that determines connectivity between positions x and x' on the DNF. Typically, the interaction kernel has a “Mexican

hat" shape with a short-range excitation and a long-range inhibition implementing a soft winner-take-all connectivity pattern, Eq. (4.3):

$$\omega(x, x') = c_{\text{exc}} e^{-\frac{(x-x')^2}{2\sigma_{\text{exc}}^2}} - c_{\text{inh}} e^{-\frac{(x-x')^2}{2\sigma_{\text{inh}}^2}}. \quad (4.3)$$

When the DNF receives a bi-modal input with two gaussian bumps, lateral interactions lead to the strengthening of one and suppression of the other input bump. The dynamics are the same as in the WTA network described in section 3.1 and the output of the DNF, $f(u(x, t))$, is a single stabilised activity peak.

These neuronal dynamics can be related to perceptual and motor parameters measured in behavioral experiments with humans or animals.

It has been shown that DNF architectures can be used in robotic applications (Richter, Sandamirskaya, and Schoner, 2012), such as tracking in a robotic table-top scenario along with object recognition and object pose estimation (Faubel and Schoner, 2009).

In the following it will be shown that a DNF architecture for serial order (Sandamirskaya and Schoener, 2010) can be implemented on neuromorphic hardware.

4.2 Different behaviour resulting from population sizes, connectivity and WTAs

The proposed model for sequence learning can be realised by using different types of winner-take-all configurations and by using discrete or continuous groups of neurons. The different configurations have different advantages and disadvantages which will be discussed here. The most robust configuration can be seen in Fig.4.2. Here the WTA network of the content neurons (as described in section 3.1), is different from the one of the ordinal neurons. The WTA of the ordinal neurons uses excitation among neighbouring neurons and inhibition among distant neurons. In this way, different weights of excitation and inhibition can be used in the two different WTA's to obtain different firing behaviour, i.e. content neurons stop firing immediately when not stimulated, whereas ordinal groups keep firing

until they are suppressed by an external inhibitory signal. Studies in neuroscience found evidence for WTA's with global inhibition similar to the one described in section 3.1 (Pouille et al., 2009). Taking inspiration from real neural networks, a WTA with global inhibition is used to implement the content DNF. However, using this type of WTA also for the ordinal groups degrades the networks ability to learn sequences. This is because biases for recurrent connections are limited, making it impossible to achieve different firing behaviours in ordinal and in content groups while using the same type of connectivity and the same biases. Another possible implementation of the serial order architecture would be to use continuous neurons in the ordinal neurons. Here not only time, which is coded in the content neurons, but also space, which is coded in the ordinal neurons would be continuous. This implementation was tested, but it has the disadvantage of spontaneous switching to the next ordinal element as can be seen in Fig. 4.3. It is not pos-

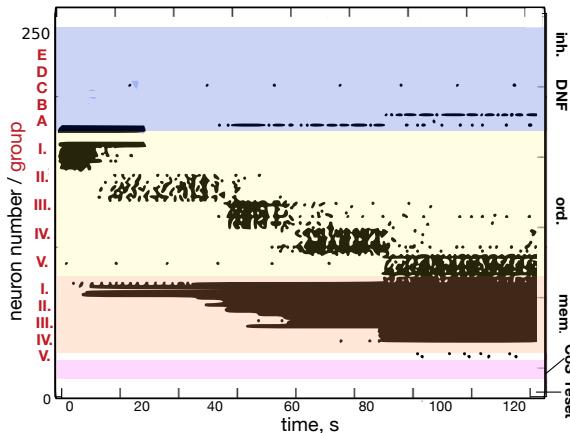
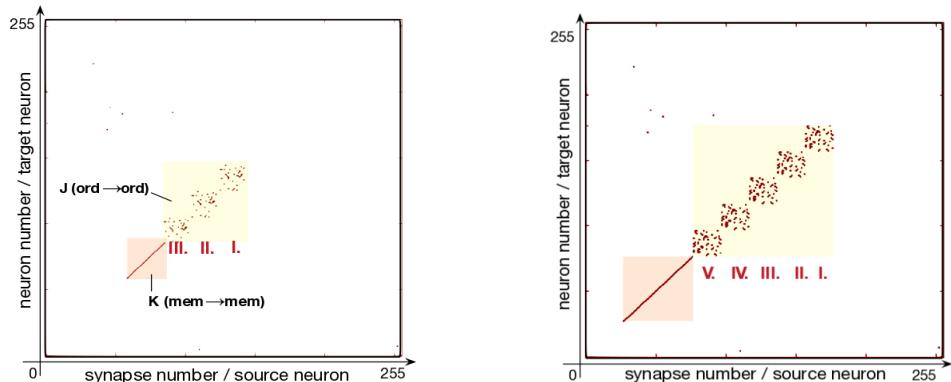


FIGURE 4.3: Activity of the neurons after stimulating the first ordinal group together with the first content group using a continuos WTA configuration with inhibitory population for ordinal neurons.

sible for the ordinal neurons to have clearly defined groups that keep firing until being suppressed by an external signal. This is again due to the fact that the same set of biases cannot account for different behaviours, i.e. in ordinal and content neurons. Automatic sequence learning would be possible by finding out the time that is needed to switch to the next element and stimulating the content neurons accordingly. The spontaneous switching cannot be avoided since there is no parameter space in which the inhibition of the ordinal WTA is stronger than the inhibition from the active memory.



(A) Plastic weight matrix for learning a sequence with three elements after initialization
(B) Plastic weight matrix for a sequence with five elements

FIGURE 4.4: Plastic weight matrix for learning a sequence with three elements after initialization

A suitable parameter space can be found for the configuration in Fig. 4.2. A main feature of this configuration is that memory neurons are highly self-exitable to remain active, even without input from the ordinal neurons. A memory group has all-to-all excitatory connections, with the excitatory weight set to a high value. However, the strength of this self-excitaiton is not enough to account for sustained activity by itself. As a solution, potentiated plastic synapses which have a high gain, can be used to enhance the self-excitation of these memory neurons. These plastic self-connections can be seen on the diagonals in the lower left corner in Fig. 4.4. At the same time memory neurons slightly inhibit the ordinal neurons. The ordinal neurons should be inhibited less by their memory than by their WTA configuration to avoid spontaneous switching of ordinal groups. Hence, a weaker inhibitory weight is used for the all-to-all inhibition from neurons in a memory group to neurons in an ordinal group, than the larger inhibitory weight that is required for inhibition among ordinal groups. The strength of inhibition is shown by the shade of the blue colour in Fig.4.2., where darker blue means stronger inhibition. Excitation between ordinal and memory groups is set to the maximum weight, as well as the excitation and inhibition of the content WTA with its global inhibitory group. To allow ordinal groups to sustain activity until they are being externally depressed, plastic potentiated weights are required. 10% to 30% of randomly plastic potentiated weights are used, Fig. 4.4. The amount of self-excitation that is needed depends on the characteristics and current state of the chip, such as humidity and

temperature. To generalise the parameter setting to other chips, memory groups have to be more active than ordinal groups to make sure that these are not only firing due to the input from ordinal neurons. For all neurons it is important that inhibition is immediate, to immediately suppress noise and obtain WTA behaviour. Excitation on the other hand needs to be strong but slow, so that it is carried out over time and activity can be sustained. To implement this model on the ROLLS chip, we configured the connectivity of neurons as shown in Fig.4.2.

4.3 Hardware realisation

In order to implement this model on the ROLLS chip, we configured the connectivity of neurons as described in the previous section. The size of the neurons groups varies depending on the length of the sequence. The minimal population size is 10 neurons per group and the longest sequence that can be learned consists of five elements. The length of the sequence is limited by the number of silicon neurons on the chip. A device with more neurons would be able to robustly learn longer sequences. In this implementation, each ordinal group consists of 20 neurons, each memory group of 10 neurons and the number of content neurons varies from 70 to 100 neurons depending on the length of the sequence. For longer sequences the neurons available to encode content is reduced in order to keep enough neurons for a stable serial order architecture. 12 neurons are used for external inhibition of ordinal neurons, accounting for the condition of satisfaction (CoS) signal. This leads to activation of the next ordinal group, initiated by the excitable connections from the active memory of the previous element to the next ordinal group. Another 12 inhibitory neurons are used to switch off all activity after learning or recall, so that the next trial can be started. This group is used to reset the neurons. The external input signal for a sequence of 5 elements is shown in Fig.4.5. All the neurons belonging to the first ordinal group are stimulated for 6000ms with 200Hz via virtual synapses. A signal to the content neurons consists of a high Gaussian signal with the maximum being 900Hz and a standard deviation of 5. All other neurons receive random values between 0 and 10Hz. The content neurons are stimulated for 6000ms. The inhibitory signal suppressing the ordinal neurons comes into action by stimulating a group of inhibitory neurons for 500ms with 1000 Hz.

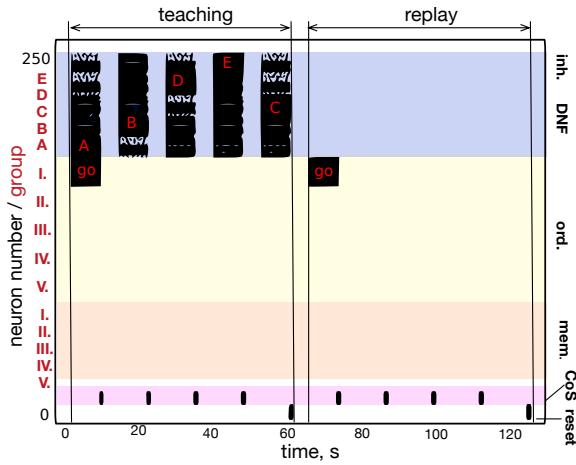


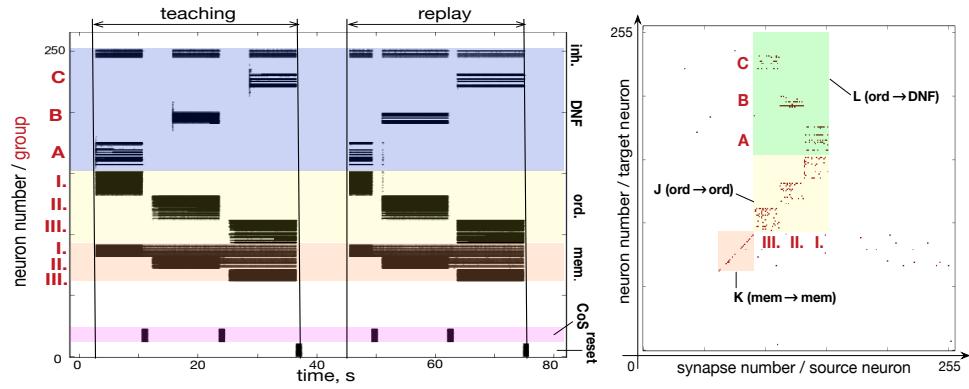
FIGURE 4.5: External input used for generating and recalling a sequence consisting of five elements.

The output frequency should be low in order to make the system biologically plausible and because the bandwidth with the AER circuits of the chip is limited. If this bandwidth is exceeded, then the server will drop events and the spikes that you get at the output are only a subset of the ones that are really there. Output frequencies from 10Hz to 70Hz were recorded in this task.

4.4 Experimental Results: Learning and forgetting sequences

4.4.1 Memorising and forgetting a simple sequence

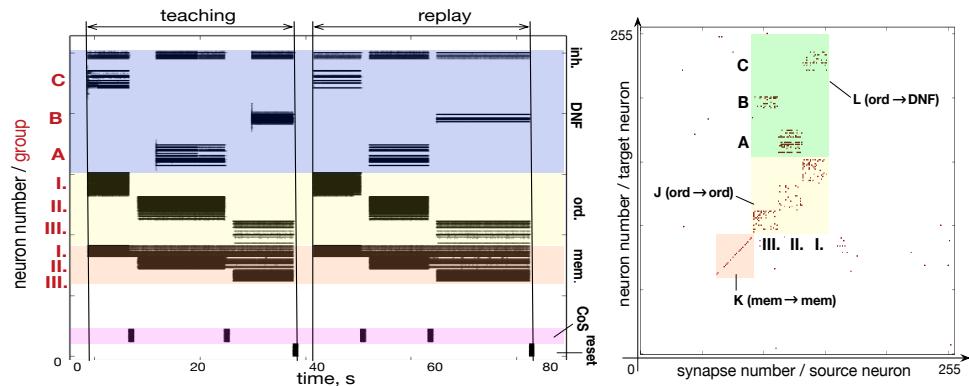
Simple sequences consisting of three elements were successfully memorised and recalled on neuromorphic hardware. The left part of Fig. 4.6a shows the activity of neurons when the teacher signal ABC is presented. The right part shows the activity after only the first ordinal group is being activated. Hence, the activation of content neurons is carried out autonomously due to learned connections from ordinal to content neurons. The WTA configuration in the content neurons together with the mismatch of silicon neurons leads to different firing rates, leading to different probabilities of potentiation at the synapses. Therefore, the potentiated synapses for each element as well as their activity during recall look slightly different. Fig. 4.6b shows the matrix of plastic synapses after sequence ABC was being learned.



(A) Neuron activity during learning (left) and (B) Weights after learning the sequence ABC

FIGURE 4.6: Learning sequence ABC

Fig. 4.7a shows the activity of neurons during learning and recalling of sequence CAB. Here, the content of A was presented for a 9000ms during learning, whereas C and B where presented for 6000ms. The proposed model can learn and recall items for different time windows. A switch in the ordinal group, which leads to learning of a new item and turning on the corresponding memory is only carried out when the activity of ordinal neurons is being suppressed by an inhibitory signal coming from the CoS group. Learning one item for a longer time than other items leads to more potentiation between this item and its ordinal group. Hence, the item that is learned longer is consolidated more in its representation in the ordinal group. Fig. 4.7b shows the corresponding weight matrix, with A being slightly more potentiated.



(A) Neuron activity during learning (left) and (B) Weights after learning the sequence CAB

FIGURE 4.7: Learning sequence CAB

Another ability of this model is the forgetting of previously learned sequences when learning a new sequence. When presented with a new teacher sequence, the winner-take-all connectivity of the content neurons leads to amplification of the activity that is due to the external teacher signal. At the same time the inhibitory group globally inhibits the content neurons, leading to suppressed activity. Since the activity of externally stimulated neurons is greater than the activity of recalled neurons through plastic connections, the recalled activity is largely reduced and eventually switched off completely. The new sequence is being learned in the previously described manner, while the old sequence will slowly be forgotten. The mechanism of forgetting comes about by the STDP learning rule implemented on the device. The activity of ordinal neurons represents presynaptic spiking and these neurons potentiate their synapses to content neurons in case there is a postsynaptic spike on the level of content neurons. On the other hand, if presynaptic spikes occur without being followed by a postsynaptic spike, the synaptic weight is being decreased and eventually gets depressed to its initial minimal state. Fig. 4.8a shows the firing activity when sequence ACB is being learned and recalled. Fig. 4.8b shows the corresponding weight matrix after learning. After learning sequence ACB, another sequence, namely BAC is presented to the content neurons, Fig. 4.8c. The synaptic weight matrix after one trial is shown in Fig. 4.8d. The activity plot shows that the newly presented elements are slowly learned while the old elements are still recalled, because forgetting takes more time than learning. Fig. 4.8e shows the activity of the neurons after stimulating with BAC for four times. The sequence BAC is recalled but traces of ACB still remain. Comparing the weight matrices, one can see that the sequence BAC is much more potentiated than ACB in the weight matrix, Fig. 4.8f. The weights of ACB depressed for the neurons that were not active during learning and recall. These two figures show that now BAC is consolidated much more in memory than is ACB. With the strength of the inhibitory group of the content WTA one can regulate the speed of forgetting. The more the previously learned signal is suppressed, the quicker it will be forgotten. The speed of learning new sequences can be regulated by changing the values of the differential pair integrator (DPI) circuit parameters of the virtual synapses. If the external input is able to initiate more activity, more postsynaptic spikes will follow the ‘presynaptic spikes’ that occur in the ordinal neurons. Hence more synapses will become potentiated.

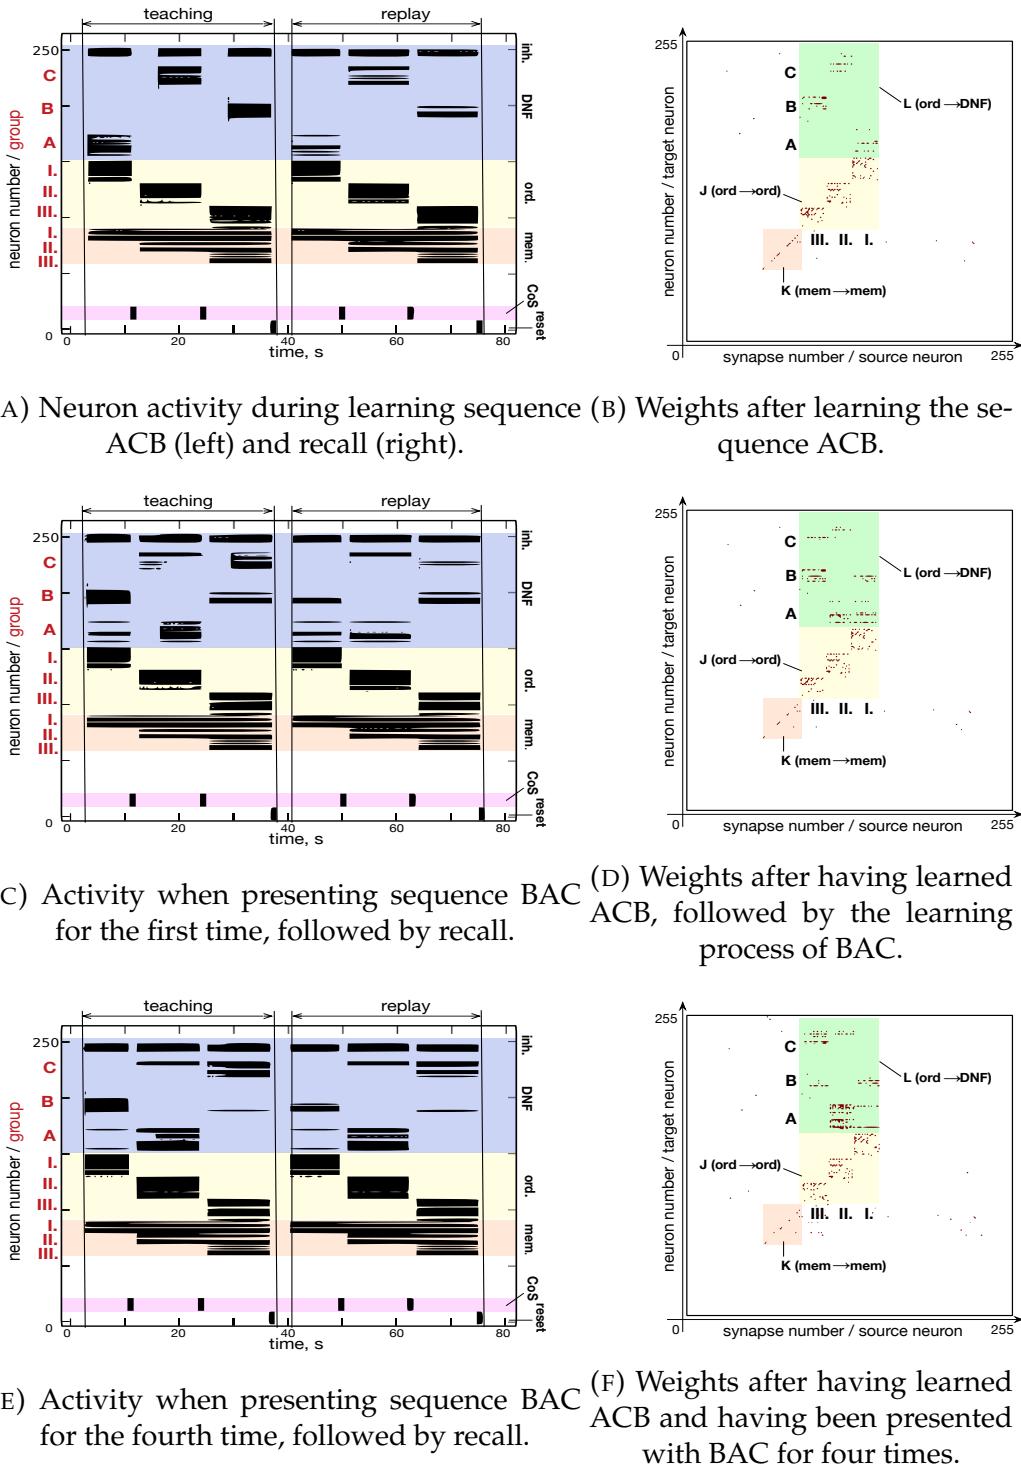
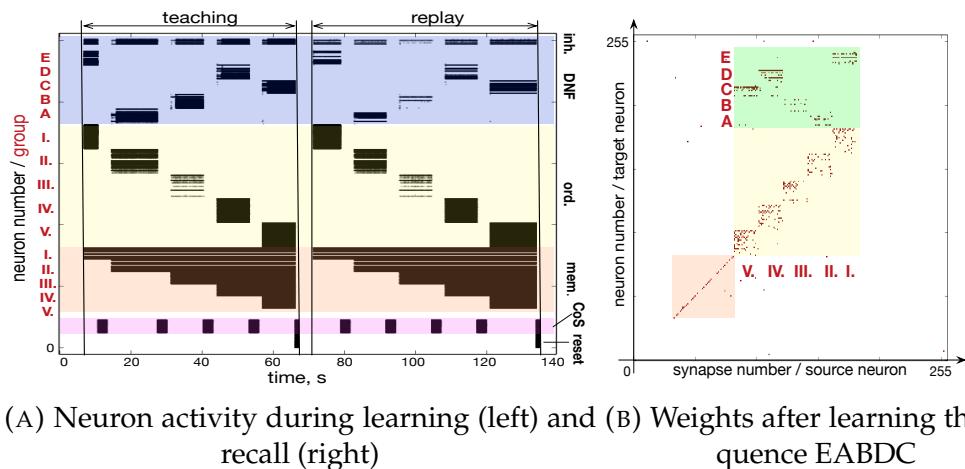


FIGURE 4.8: Learning and forgetting a sequence

4.4.2 Learning longer sequences

In theory the model is able to learn sequences of different lengths and for different sizes of neuron populations. A change in the population sizes

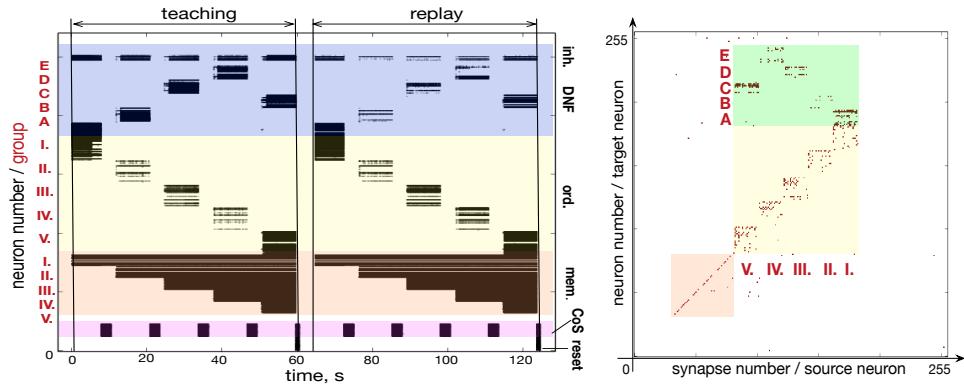
leads to different behaviour of the groups. An increase in size for example leads to increased activity, which could result in spontaneous switching of states, whereas smaller populations, leading to less activity, result in a failure to activate the next memory or the next ordinal group. Thus, the biases of the silicon neuron and synapse circuits have to be tuned accordingly to again obtain the correct behaviour. Another limiting factor is the size of the neuromorphic device. With 256 neurons it is not possible to learn sequences longer than 5 elements. Populations have to be large enough in order to compensate the mismatch of single neurons so that the networks behaviour remains robust. Sequences of five elements can be learned on the ROLLS chip. The population size of the content neurons was decreased to 70 neurons, while all other populations sizes were not changed. Fig. 4.9a shows the activity when a sequence EABDC is being learned (left) and recalled (right). Fig. 4.9b shows the corresponding weight matrix after learning. It can be seen that some groups of neurons are much more excitable



(A) Neuron activity during learning (left) and (B) Weights after learning the sequence EABDC

FIGURE 4.9: Learning sequence EABDC

than others due to the mismatch on the chip. Here, the third ordinal group is clearly weaker than the other groups making it difficult to find the right biases that lead to correct behaviour of the whole network, e.g. increasing the recurrent excitation leads to stronger activity in this group but at the same time the other groups become even stronger leading to spontaneous switches between the groups. Although the resulting variability of how strong an item is consolidated is not optimal, the same behaviour occurs in humans when remembering a sequence. For a long sequence, some items seem to be stronger in our memory than others. Another example of a sequence with five elements together with the learned weight matrix is shown in Fig.4.10. Here ABDEC was learned. Again the activity of ordinal groups



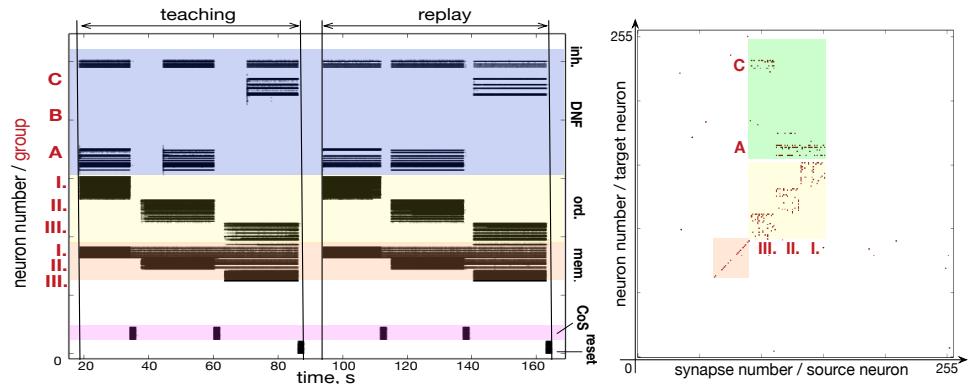
(A) Neuron activity during learning (left) and (B) Weights after learning the sequence ABDEC

FIGURE 4.10: Learning sequence ABDEC

is different, which is not only due to mismatch but also due to the randomly initialised 30% of potentiated synapses of ordinal neurons. However, the chosen biases can cope with these slight changes and still result in the right overall behaviour.

4.4.3 Learning repeating items

The model does not show any difficulties in learning repeating items as can be seen in Fig.4.11. Learning in ordinal to content neurons is not influenced



(A) Neuron activity during learning (left) and (B) Weights after learning the sequence AAC

FIGURE 4.11: Learning sequence AAC

by previously learned items since there are no excitable or plastic connections in between ordinal groups. Therefore each ordinal group can learn its content independently.

4.4.4 Serial order errors

Common errors that are made by the serial order architecture are spontaneous switching of ordinal groups or a failure to switch. The first comes about if excitatory recurrent connections are too strong or the excitation time constant is too small, making excitation of the next state faster than the inhibition. The inhibition that suppresses new activity due to WTA connectivity has to be very strong to only allow switching if all the ordinal neurons activity is externally suppressed. A failure to switch the ordinal group occurs when the external inhibition is not strong and long enough to completely switch off any activity of the previous ordinal group. It can also happen due to mismatch, when the excitation to the previous ordinal group is stronger than the excitation to the next ordinal group. When the slight inhibition from memory groups to their ordinal groups is not strong enough, a previous ordinal group is being switched on again. However, the inhibition from memory groups to their ordinal groups should not be too strong, because high inhibition from memory to ordinal together with high excitation from memory to the next ordinal will lead to spontaneous switching. Hence, the network is very prone to changes in the strength and time constants of excitation and inhibition. Failure to switch from the third to the fourth ordinal group is shown in Fig. 4.12a. Here the network recalled the third element twice instead of recalling the fourth one. Learning was still done correctly since the failure to switch only occurred in the recall.

In the next example, the network failed to switch to the fourth ordinal group. Thus, the third group was presented with three different teacher signals. It is interesting to see that this ordinal group merges together the different teacher signals. This is because recall of the previous item occurs at the same time as the presence of the next teacher signal. Since recall and teacher activity is nearby in space, the WTAs excitability of neighbors leads to merging of these signals instead of suppression. The recalled content is a combination of all three teacher signals. The spiking activity together with the plastic weight matrix is shown in Fig. 4.12b.

An example of spontaneous switching is shown in Fig. 4.12c. As soon as the fourth ordinal group is activated, it activates its memory which in turn excites the fifth ordinal group. The excitation to the fifth ordinal group is much stronger than the self-excitation from the fourth ordinal group leading to an immediate switch. Therefore, the fifth ordinal group is being presented with

two different teacher signals. It learns the first teacher signal, but after being presented with the second one, it potentiates its synapses to this signal and forgets the previous one. This can be seen in the recall. Only the last teacher signal is being recalled by the fifth ordinal group.

A similar case is shown in Fig. 4.12d, with the difference that the spontaneous switch only occurs during recall. The content signals were learned correctly, but the switch from the third ordinal group to the fifth leads to the fact that the fourth content element is not being remembered. This is similar to common errors made by humans, who have learned a long sequence but forgot to recall a specific item, which is stored somewhere in the middle of the sequence (Henson, 1998). At the same time, this example shows, that even when using the same set of biases and recurrent connections, the chip shows high variability as different transitions occur during training and during recall.

Errors of this model always involve neighbouring items. Either the next element overwrites the previous one, if the ordinal groups do not manage to switch, or some items are only quickly learned and partly overwritten by the next item due to spontaneous switching. This leads to confusion during recall, which one is the right element to recall. If the error only occurs during recall, either a previous items is recalled again or an item is being left out due to spontaneous switching. The errors are more likely to occur for longer sequences and they mostly appear towards the end. This is in line with common errors that humans make when learning a long sequence.

4.4.5 Discussion

This project showed that the serial order architecture can account for sequence learning which can be successfully implemented in neuromorphic hardware. With this implementation, we can study different sequence learning algorithms and test different cognitive architectures by changing the recurrent connections or the learning rule. To arrive at an optimal solution for this architecture, it was shown that a discrete representation of ordinal groups performs much better than a continuos representation. In addition it was discovered that potentiated plastic synapses are needed to maintain activity. A current limitation of this implementation is that a change in the population sizes leads to different behaviour of the neuron groups and the biases of the circuits have to be retuned. As an example, larger groups of

neurons lead to stronger self-excitation which could result in spontaneous switching. Whereas smaller population sizes lead to less activity, which could result in a failure to switch to the next state. Furthermore, the ROLLS chips themselves are highly variable due to mismatch. Therefore the implementation of the same architecture on a different chip has to be followed by a retuning of biases. In further studies, we aim on finding rules on how the biases and plastic synapses for self-excitation have to be changed when changing population sizes. Decreasing population sizes would make it possible to learn even longer sequences on the chip. However the length of the sequence is also limited by the 256 neurons on the ROLLS chip and populations should not be decreased too much in order to be able to compensate the mismatch of single neurons. Research is currently being done to fabricate chips with more than 256 configurable neurons and these chips would be able to learn sequences with much more than five elements.

In the future, different approaches that are inspired by biological neural substrate can be used for implementation in order to find a model that best represents the computation carried out by the brain and that at the same time leads to fast, low-power and reliable computation on a neuromorphic device.

The viability of a sequence learning model is usually examined by comparing the errors committed by the model with those made by humans (Henson, 1998). Common mistakes observed in nature involve swapping of neighboring elements but a single error in a sequence does not mean that its following items cannot be recalled anymore, as it is claimed in chaining theory (Wickelgren, 1965). Our sequence learning model commits more errors as sequences get longer. These always involve neighboring items. However a single error is not always followed by errors to recall the next items. Instead, after a single spontaneous switch or a failure to switch, the network still has the ability to correctly recall the following items.



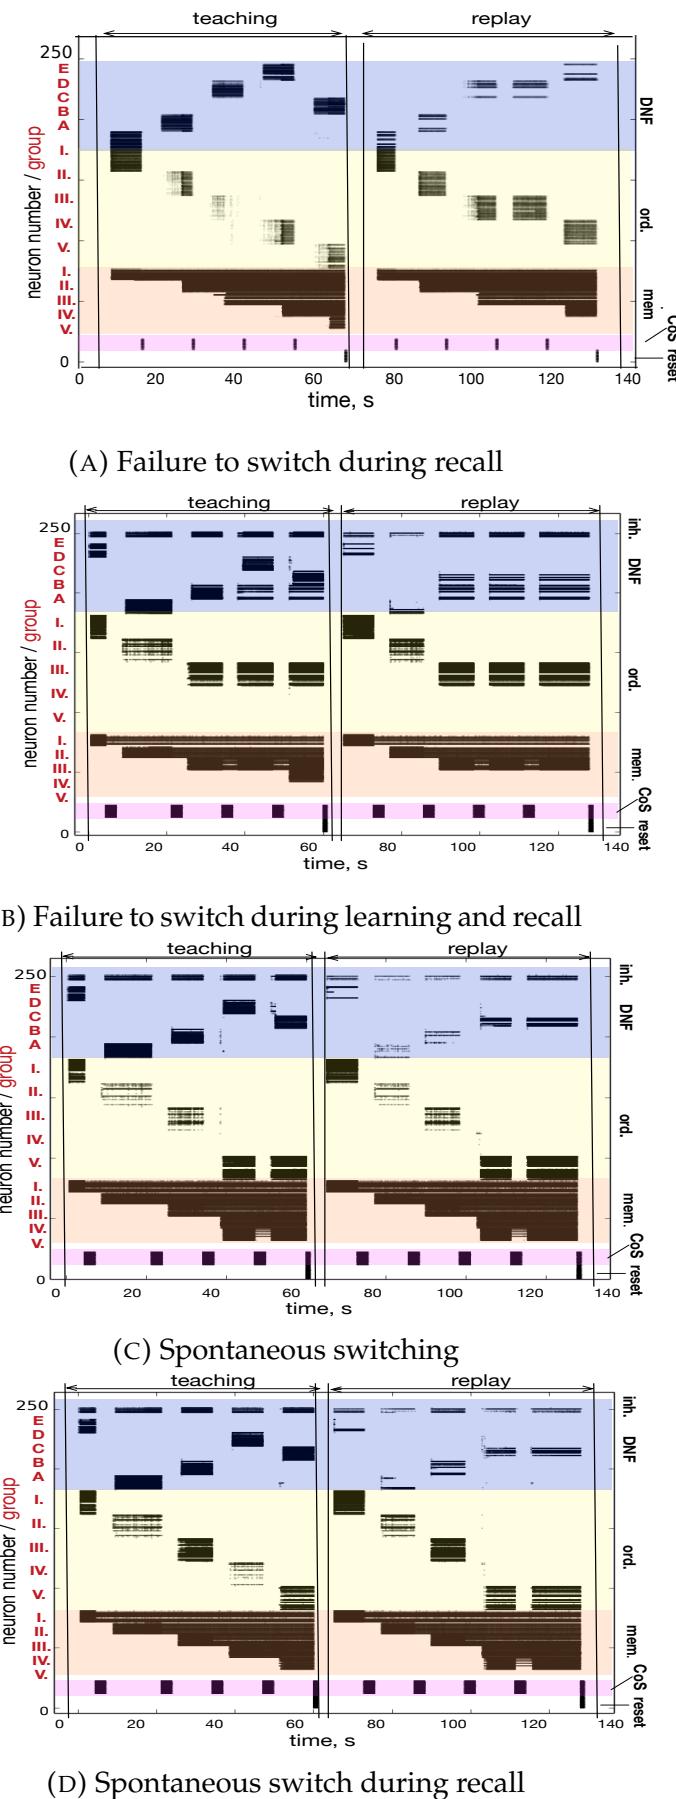


FIGURE 4.12: Common learning and recalling errors.

Chapter 5

Conclusion

In this work, we have shown the learning abilities that can be carried out by neuromorphic hardware. More precisely, we have shown unsupervised learning experiments on the ROLLS chip, which comprises 256 configurable silicon neurons together with virtual, recurrent and plastic synapses. These plastic synapses show learning abilities which have been previously studied in a supervised learning task (Technology and Tianlu, 2015). In this project it was shown that the ROLLS chip is not only successful in supervised learning, but also in unsupervised learning. The biases can be changed so that they are able to initiate depression and potentiation depending on the input frequencies at different synapses. Thus, different patterns of input frequencies can initiate different postsynaptic firing of the neurons. A WTA configuration helps the network to find local clusters of neurons which learn one specific pattern while other clusters learn another pattern. The WTA configuration achieves this by taking advantage of the neurons mismatch, which leads to the fact that some neurons are more active than others. The WTA configuration can amplify the most active groups and suppress the weaker ones, thus enhancing the contrast of the neurons activity. Depression of synapses that are stimulated with low frequencies is a very important feature in the pattern recognition task. This is because different neurons have to learn different inputs. However, the most active neurons tend to fire for both patterns, while weak neurons do not fire for any pattern. The gain of potentiated neurons should be very high in order to lead to different postsynaptic activity, depending on the frequency with which they are stimulated and on the current state of the synaptic weight (potentiated or depressed). The number of patterns that can be learned is limited by the number of neurons on the chip. Here it was shown that four patterns can be learned. However, clusters become smaller when more patterns have to be learned and the classification accuracy of postsynaptic spiking degrades with more

patterns. A solution would be to fabricate neuromorphic chips with more than 256 neurons. This would also be a solution for learning long sequences (with more than five elements) in the sequence learning task which was carried out and explained in chapter 4. With the implementation of the serial order architecture on the ROLLS chip, it was shown that building hierarchical neural networks on neuromorphic hardware lets the chip successfully learn sequences and hence autonomously perform a cognitive task when interfaced with a robot. Furthermore, the ROLLS chip is not only able to account for different learning behaviours that were inspired by the computation carried out in cortex, but it can also be used for studying the plausibility of different cognitive and neuronal models. Here, it was shown that a serial order architecture inspired by Dynamical Neural field theory leads to the desired behaviour of learning and recalling sequences with different length. The next step is to implement this model on another ROLLS chip that has a different setup which can more easily be interfaced with a robot. After retuning the biases for this chip, the robot should learn to associate different positions of visual input coming from a DVS with a movement that it has to make. With this, the behavioral loop can be closed. In future studies the different learning abilities of the chip can be used to implement different neural networks that eventually lead to low-power, biologically-inspired solutions for machine learning tasks, and cognitive robotics.

Appendix A

Setting suitable biases

The behaviour of the chip's neurons and synapses can be modified by changing the biases of the transistors. This is a challenging task because of highly nonlinear and coupled dynamics of the circuits. The circuits on the chips are very sensitive to small changes in the biases and the challenge lies in finding comparably small ranges in which the desired behaviours, e.g. firing rate, WTA behaviour, and learning are in accordance. Because of mismatch, different power supplies, temperature and different setups of the chip, biases have to be retuned for each device. In the following I will show the circuits together with the description of the biases which can be modified. In addition, the values of the biases that were used for the unsupervised learning task and the values for sequence learning are listed in the tables. These biases work well for the device that was used in the experiment, but they will have to be retuned for different devices. In the tables I will give a broad overview which biases most likely have to be changed and what are their purpose in the different experiments. These tables can be used as a guidance if the experiments need to be carried out on different chips in order to speed up the bias tuning.

TABLE A.1: Biases of the silicon neuron

Parameter	Description	Notes	Pattern Recogn.	Sequence Learning
IF_RST	Controls the reset threshold current		1pA	1pA
IF_DC	Controls a constant current injected to all neurons	Small injected current helps neurons to be more excitable, can be set to ground in these experiments to avoid spontaneous firing because of self-excitation	1pA	1pA
IF_BUF	Buffer for oscilloscope		10.9nA	10.9nA
IF_ATHR	Neuron's adaptive threshold current	Small, to avoid fast decrease in activity	1pA	1pA
IF_RFR1	Controls the duration of the neuron's refractory period	Small voltage, e.g. high current, to make fast presynaptic firing possible, leads to increased probability of learning	49.7nA	1.5nA
IF_RFR2	Controls the duration of some specified neuron's refractory period	Same as above, neurons can be set to have different refractory periods	48.2nA	1.5nA
IF_AHW	Controls the adaptaiton time constant	No adaptation required, therefore set to smallest current	1pA	1pA
IIF_AHTAU	Neuron's adaptive Tau		5.8nA	7.4nA

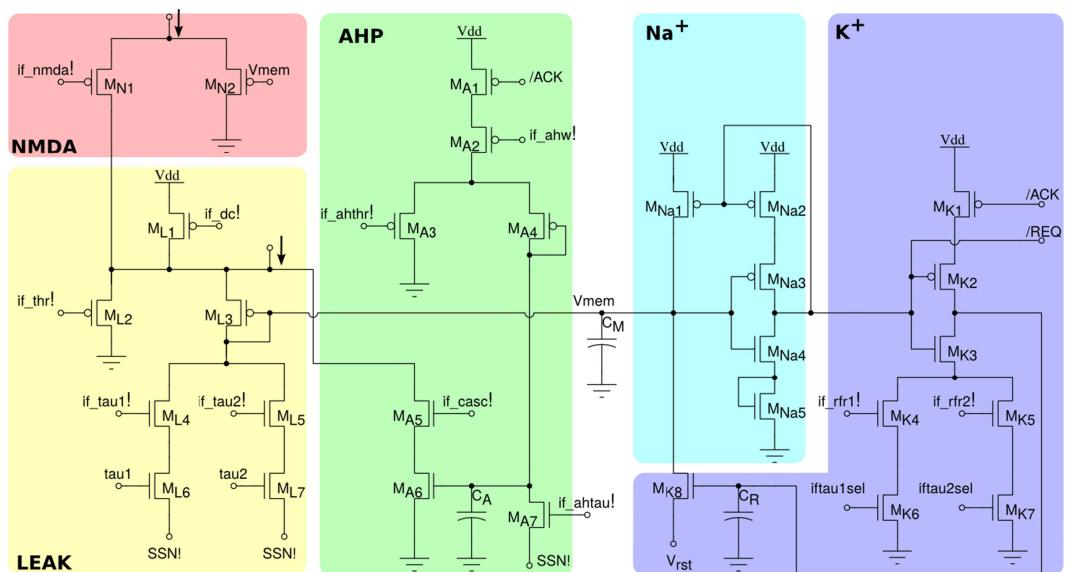


FIGURE A.1: Silicon neuron schematics. The NMDA block implements a voltage gating mechanism; the LEAK block models the neuron's leak conductance; the spike-frequency adaptation block AHP models the after-hyper-polarizing current effect; the positive-feedback block Na^+ models the effect of the Sodium activation and inactivation channels; reset block K^+ models the Potassium conductance functionality.

TABLE A.2: Biases of the silicon neuron (continued)

Parameter	Description	Notes	Pattern Recogn.	Sequence Learning
IIF_TAU2	Time constant of the neurons	Small for unsupervised learning, make specified neurons maximally excitable; increase for sequence learning to avoid spontaneous self-excitation	1.5pA	22.6pA
IF_TAU1	Time constant for specified neurons	Can be used if some neurons do not stop firing, set neurons to Tau1 and increase this current	3.6pA	24uA
IF_NMDA	Controls the sensitivity of the neurons		1.4pA	1pA
IF_CASC	Controls the cascode current		1pA	1pA
IF_THR	Controls the neuron's firing threshold	small to make neurons maximally excitable	17.8pA	280nA

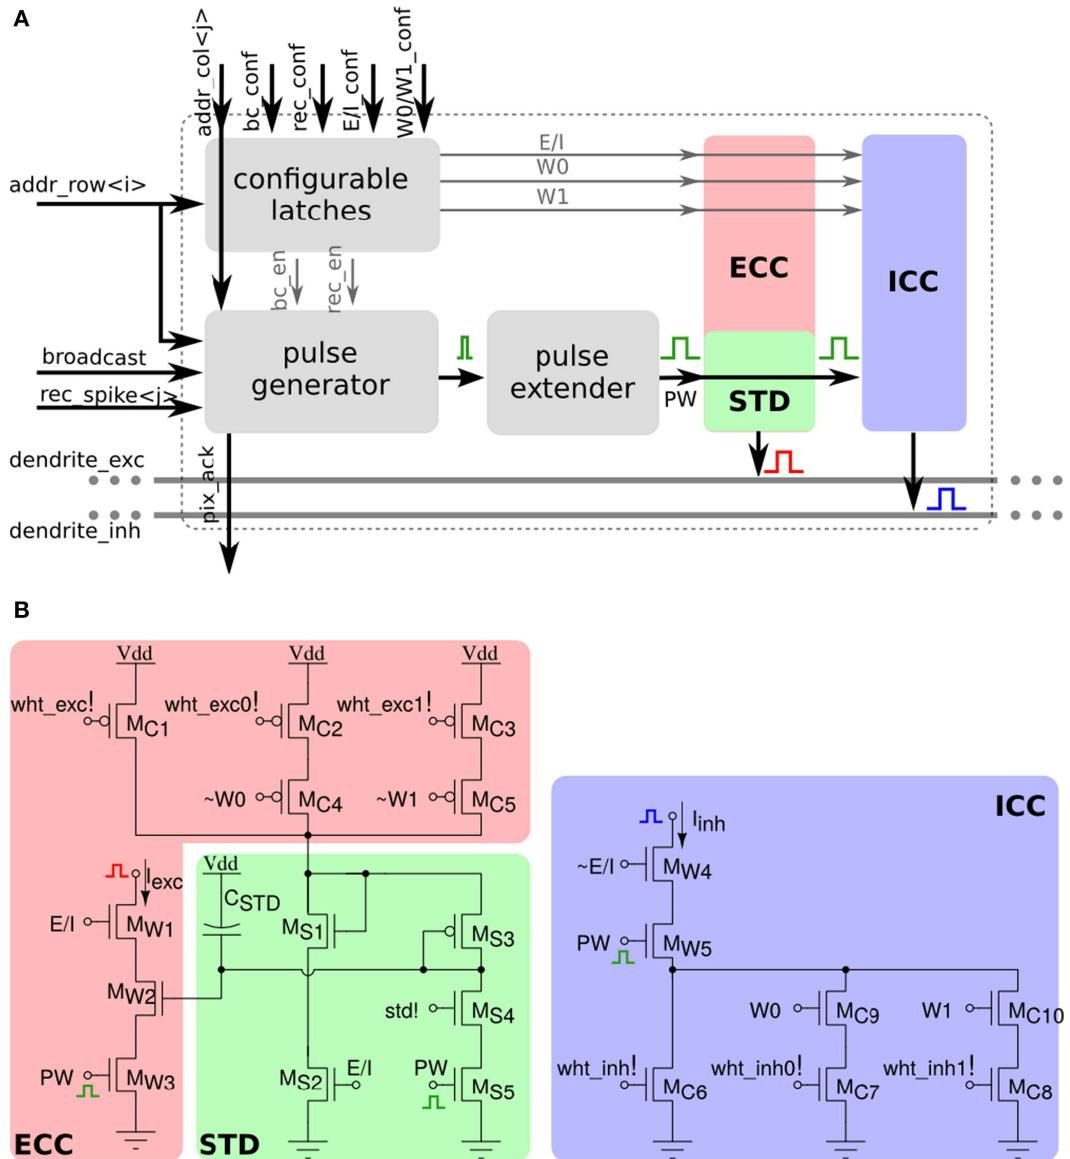


FIGURE A.2: Short-term plasticity synapse array element and dynamics of recurrent synaptic connections. (A) Block diagram of the synapse element. (B) Transistor level schematic diagram of the excitatory and inhibitory pulse-to-current converters.

TABLE A.3: Biases for the recurrent synaptic connections

Parameter	Description	Notes	Pattern Recogn.	Sequence Learning
WHT_STD	Controls the magnitude of short term depression	Minimum, as STD is not used	1pA	1pA
WHT_INH0	Controls the magnitude of the inhibitory current component which is injected into the excitatory DPI when set to value 1	High to achieve maximal WTA behaviour, e.g. suppression of low frequency firing	24uA	8.7nA
PWLK	Controls the pulse width of the synaptic current	Increase to achieve maximal strength of recurrent connections	1.1nA	448.4pA
WHT_INH1	Controls the magnitude of the excitatory current component which is injected into the excitatory DPI when set to value 2	High to achieve maximal WTA behaviour, e.g. suppression of low frequency firing	24uA	199.4nA
WHT_EXC	Controls the magnitude of the excitatory weight independent when set to value 0	High to achieve maximal WTA behaviour, e.g. excitation of neighbouring neurons	665.2nA	110.9nA
WHT_EXC1	Controls the magnitude of the excitatory current component which is injected into the excitatory DPI when set to value 2	High to achieve maximal WTA behaviour, e.g. excitation of neighbouring neurons	2.6uA	1.4uA

TABLE A.4: Biases for the recurrent synaptic connections (continued)

Parameter	Description	Notes	Pattern Recogn.	Sequence Learning
WHT_EXC0	Controls the magnitude of the excitatory current which is injected into the excitatory DPI when set to value 1	High to achieve maximal WTA behaviour, e.g. excitation of neighbouring neurons	1.3uA	291.2nA
WHT_INH	Controls the magnitude of the excitatory weight when set to value 0	High to achieve maximal WTA behaviour	24uA	3.8nA
DPIE_THR	Controls the threshold of excitatory synapses	Not too high, sequence learning: avoid spontaneous switching; pattern recogn.: let inhibition switch off activity	146.1nA	33.1nA
DPIE_TAU	Controls the time constant of inhibitory synapses	High for pattern recogn.: slow excitation, inhibitory neuron can shut off activity, lower for seq. learning: needed for memory	2.5nA	402.2pA
DPII_TAU	Controls the time constant of excitatory synapses	Small for fast inhibition; higher to make it slower for retaining memory	2.2pA	1pA
DPII_THR	Controls the threshold of inhibitory synapses	High for strong inhibition, let only a few neurons learn in pattern recogn., lower for sequence learning, for more activity	5.3uA	262.6nA

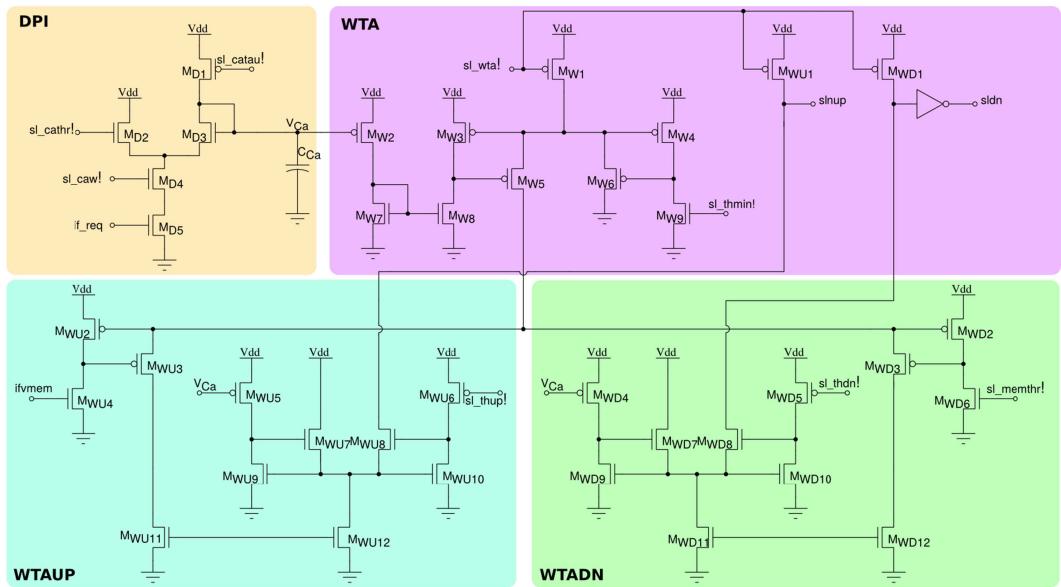


FIGURE A.3: Post-synaptic learning circuits for evaluating the algorithm's weight update and 'stop-learning' conditions. The DPI circuit MD1-5 integrates the post-synaptic neuron spikes and produces a current proportional to the neuron's Calcium concentration. Three current-mode winner-take-all circuits WTA, WTAUP, and WTADN compare the Calcium concentration current to three set thresholds $sl_thmin!$, $sl_thdn!$, and $sl_thup!$, while the neuron's membrane current is compared to the threshold $sl_memthr!$.

TABLE A.5: Biases for implementing the learning rule

Parameter	Description	Notes	Pattern Recogn.	Sequence Learning
SL_THDN	Controls the threshold in the learning rule, sets the limit of LTD	Set higher to obtain LTD for higher frequencies	2.4uA	2.4uA
SL_THUP	Controls the threshold in the learning rule, setting the limit of LTP	Higher than SL_THDN, to obtain LTP for higher firing frequencies	24uA	24uA
SL_MEMTHR	Controls the membrane threshold	High to decrease the probability of pre-, postspike, leading to LTP	121.6nA	121.6nA
SL_BUF	Reset buffer bias of the postsynaptic learning circuit		151.2nA	115.2nA
SL_THMIN	Sets the minimal threshold for learning transitions	High to avoid learning in some neurons; keep random weights for learning at a later stage	49.2nA	49.2nA
SL_WTA	Compares the probability of LTD, LTP and stop learning		115.2nA	115.2nA
SL_CATHR	Controls how high the calcium variable has to be in order to make learning transitions	High to avoid learning in some neurons to keep random weights for learning at a later stage	24uA	24uA
SL_CATAU	Controls the calcium variable's time constant	Low, to get a fast increase of calcium	4.9uA	4.9pA
SL_CAW	Controls the calcium variable's weight	High for large increase in calcium after only a few postsyn. spikes	24uA	24uA

TABLE A.6: Biases for the dynamics of the learning synapses
and long-term potentiation

Parameter	Description	Notes	Pattern Recogn.	Sequence Learning
PA_WDRIFTDN	Controls how much the weight is passively drifting towards w_{min} during the resting state	low, to avoid passive transitions	5.7pA	15pA
PA_WDRIFTUP	Controls how much the weight is drifting towards w_{max} during the resting state	Low, to avoid passive transitions	7.7pA	29.7pA
PA_DELTAUP	Controls the amplitude of a digital up jump	High, same value as PA_DELTADN, so that only a few spikes lead to LTP	92.7nA	1.4nA
PA_DELTADN	Controls the amplitude of a digital down jump	High, same value as PA_DELTAUP	92.7nA	1.4nA
PA_WHIDN	Sets the strength of the weight gain	Maximum for maximal gain of the potentiated weights	19.8uA	6.8uA
PA_WTHR	Controls the threshold of plastic weight	High to obtain more LTD than LTP due to drift	3.1uA	13.9uA
PA_WDRIFT	Controls the strength of passive drift rates	Low, to avoid passive transitions	40.6pA	164.8pA
PA_PWLK	Controls the pulse width of the synaptic current of learning synapses	Increase to make neurons with potentiated weights fire more	1.1nA	448.4pA
PDPI_BUF	Reset buffer of the learning synapse	85.1nA	85.1nA	85.1nA

TABLE A.7: Biases for the dynamics of the learning synapses and long-term potentiation (continued)

Parameter	Description	Notes	Pattern Recog.	Sequence Learning
PDPI_TAU	Controls the time constant of plastic synapses	Low, to make plastic synapses fast	478.2pA	727.8pA
PDPI_THR	Controls the threshold of plastic synapses	High, to increase gain of the plastic weights	136.5nA	44.1nA

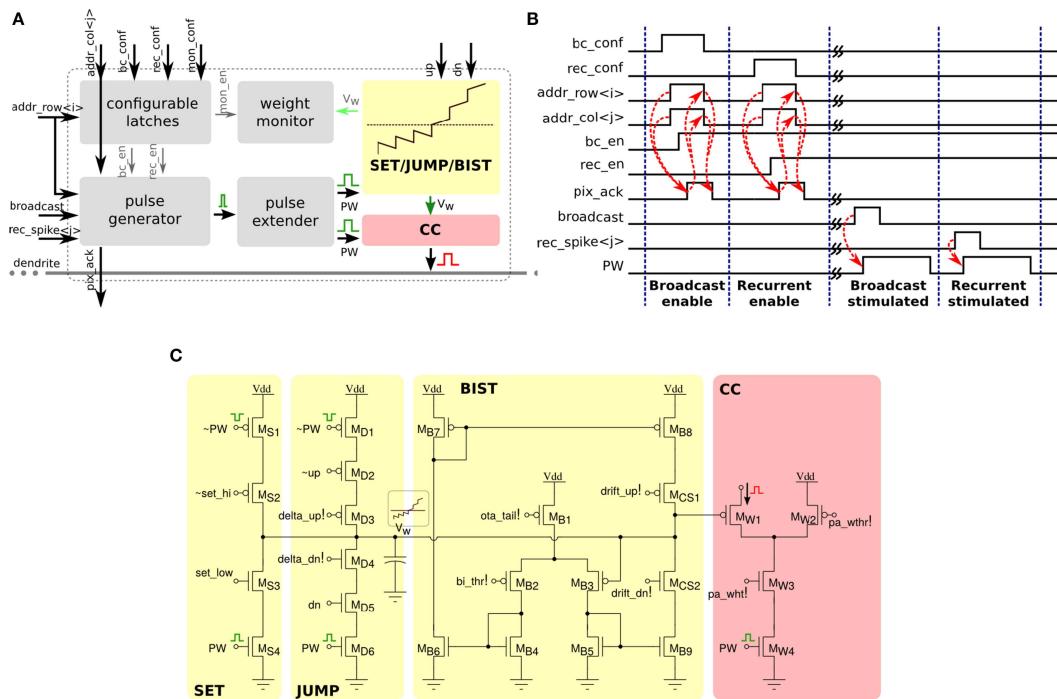


FIGURE A.4: Long-term plasticity synapse array element. (A) Plastic synapse configuration logic block diagram. (B) Timing diagram for broadcast and recurrent activation modes in one synapse using 4-phase handshaking protocol. Dashed red lines show the sequence between signals. (C) Schematic diagram of the bi-stable weight update and current generator blocks.

Bibliography

- Abbott, L F and S B Nelson (2000). "Synaptic plasticity: taming the beast." In: *Nature neuroscience* 3.november, pp. 1178–1183. ISSN: 1097-6256. DOI: [10.1038/81453](https://doi.org/10.1038/81453).
- Ambrogio, Stefano et al. (2016). "Unsupervised learning by spike timing dependent plasticity in phase change memory (PCM) synapses". In: *Frontiers in Neuroscience* 10.56, pp. 1–12. ISSN: 1662-453X. DOI: [10.3389/fnins.2016.00056](https://doi.org/10.3389/fnins.2016.00056). URL: http://www.frontiersin.org/neuromorphic{_engineering/10.3389/fnins.2016.00056/abstract.
- Beiser, D G and J C Houk (1998). "Model of cortical-basal ganglionic processing: encoding the serial order of sensory events." In: *Journal of Neurophysiology* 79.6, pp. 3168–3188. ISSN: 0022-3077.
- Billings, Guy and Mark C W van Rossum (2009). "Memory retention and spike-timing-dependent plasticity." In: *Journal of neurophysiology* 101.6, pp. 2775–2788. ISSN: 0022-3077. DOI: [10.1152/jn.91007.2008](https://doi.org/10.1152/jn.91007.2008).
- Boegerhausen, Malte, Pascal Suter, and Shih-Chii Liu (2003). "Modeling short-term synaptic depression in silicon." In: *Neural computation* 15.2, pp. 331–348. ISSN: 0899-7667. DOI: [10.1162/089976603762552942](https://doi.org/10.1162/089976603762552942).
- Brader, Joseph M, Walter Senn, and Stefano Fusi (2007). "Learning Real-World Stimuli in a Neural Network with Spike-Driven Synaptic Dynamics". In: *Neural Computation Massachusetts Institute of Technology* 19, pp. 2881–2912. ISSN: 0899-7667. DOI: [10.1162/neco.2007.19.11.2881](https://doi.org/10.1162/neco.2007.19.11.2881).
- Brette, Romain and Wulfram Gerstner (2005). "Adaptive exponential integrate-and-fire model as an effective description of neuronal activity." In: *Journal of neurophysiology* 94.5, pp. 3637–42. ISSN: 0022-3077. DOI: [10.1152/jn.00686.2005](https://doi.org/10.1152/jn.00686.2005). URL: <http://www.ncbi.nlm.nih.gov/pubmed/16014787>.
- Carpenter, a F, a P Georgopoulos, and G Pellizzer (1999). "Motor cortical encoding of serial order in a context-recall task." In: *Science (New York, N.Y.)* 283.March, pp. 1752–1757. ISSN: 00368075. DOI: [10.1126/science.283.5408.1752](https://doi.org/10.1126/science.283.5408.1752). URL: <http://www.ini.uzh.ch/{~}peterk/>

- Lectures / MotorSystem / WRAPUP / science{_}283{_}1752.pdf.*
- Carpenter, Gail A. and Stephen Grossberg (1987). "A massively parallel architecture for a self-organizing neural pattern recognition machine". In: *Computer Vision, Graphics and Image Processing* 37.1, pp. 54–115. ISSN: 0734189X. DOI: [10.1016/S0734-189X\(87\)80014-2](https://doi.org/10.1016/S0734-189X(87)80014-2).
- Chicca, Elisabetta et al. (2014). "Neuromorphic electronic circuits for building autonomous cognitive systems". In: *Proceedings of the IEEE* 102.9, pp. 1367–1388. ISSN: 00189219. DOI: [10.1109/JPROC.2014.2313954](https://doi.org/10.1109/JPROC.2014.2313954). arXiv: [1403.6428](https://arxiv.org/abs/1403.6428).
- Delbruck, Tobi et al. (2010). "32-bit configurable bias current generator with sub-off-current capability". In: *ISCAS 2010 - 2010 IEEE International Symposium on Circuits and Systems: Nano-Bio Circuit Fabrics and Systems*, pp. 1647–1650. ISBN: 9781424453085. DOI: [10.1109/ISCAS.2010.5537475](https://doi.org/10.1109/ISCAS.2010.5537475).
- Douglas, R.J., M.a. Mahowald, and K.a.C. Martin (1991). "Hybrid analog-digital architectures for neuromorphic systems". In: *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)* 3, pp. 1848–1853. DOI: [10.1109/ICNN.1994.374439](https://doi.org/10.1109/ICNN.1994.374439). URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=374439>.
- Douglas, Rodney J. and K. A C Martin (2007). "Recurrent neuronal circuits in the neocortex". In: *Current Biology* 17.13, pp. 496–500. ISSN: 09609822. DOI: [10.1016/j.cub.2007.04.024](https://doi.org/10.1016/j.cub.2007.04.024).
- Duran, Boris and Yulia Sandamirskaya (2012). "Neural dynamics of hierarchically organized sequences: A robotic implementation". In: *IEEE-RAS International Conference on Humanoid Robots*, pp. 357–362. ISSN: 21640572. DOI: [10.1109/HUMANOIDS.2012.6651544](https://doi.org/10.1109/HUMANOIDS.2012.6651544).
- Engel, Stephen A., Gary H. Glover, and Brian A. Wandell (1997). "Retinotopic organization in human visual cortex and the spatial precision of functional MRI". In: *Cerebral Cortex* 7.2, pp. 181–192. ISSN: 10473211. DOI: [10.1093/cercor/7.2.181](https://doi.org/10.1093/cercor/7.2.181).
- Ermentrout, G. Bard, Roberto F. Galan, and Nathaniel N. Urban (2008). *Reliability, synchrony and noise*. DOI: [10.1016/j.tins.2008.06.002](https://doi.org/10.1016/j.tins.2008.06.002).
- Ester, Martin et al. (1996). "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise". In: *Second International Conference on Knowledge Discovery and Data Mining*, pp. 226–231. ISSN: 09758887. DOI: [10.1.1.71.1980](https://doi.org/10.1.1.71.1980). arXiv: [10.1.1.71.1980](https://arxiv.org/abs/10.1.1.71.1980). URL: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.20.2930>.

- Faubel, Christian and Gregor Schoner (2009). "A neuro-dynamic architecture for one shot learning of objects that uses both bottom-up recognition and top-down prediction". In: *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2009*, pp. 3162–3169. ISBN: 9781424438044. DOI: [10.1109/IROS.2009.5354380](https://doi.org/10.1109/IROS.2009.5354380).
- Goris, Robbe L T, J Anthony Movshon, and Eero P Simoncelli (2014). "Partitioning neuronal variability." In: *Nature neuroscience* 17.6, pp. 858–65. ISSN: 1546-1726. DOI: [10.1038/nn.3711](https://doi.org/10.1038/nn.3711). arXiv: [NIHMS150003](https://arxiv.org/abs/1503.0003). URL: <http://dx.doi.org/10.1038/nn.3711>.
- Grossberg, Stephen (1988). *Nonlinear neural networks: Principles, mechanisms, and architectures*. DOI: [10.1016/0893-6080\(88\)90021-4](https://doi.org/10.1016/0893-6080(88)90021-4).
- Hebb, Donald O (1949). "The first stage of perception: growth of the assembly". In: *The Organization of Behavior* 4, pp. 60–78. ISSN: 03010082. DOI: [10.1016/0301-0082\(84\)90021-2](https://doi.org/10.1016/0301-0082(84)90021-2). URL: papers2://publication/uuid/E774CC14-246F-481D-85E7-1EA1A8642798.
- Henson, R. N.a (1998). "Short-term memory for serial order: the Start-End Model." In: *Cognitive psychology* 36.2, pp. 73–137. ISSN: 0010-0285. DOI: [10.1006/cogp.1998.0685](https://doi.org/10.1006/cogp.1998.0685).
- Ijspeert, Auke Jan et al. (2007). "From swimming to walking with a salamander robot driven by a spinal cord model." In: *Science (New York, N.Y.)* 315.5817, pp. 1416–1420. ISSN: 0036-8075. DOI: [10.1126/science.1138353](https://doi.org/10.1126/science.1138353).
- Indiveri, Giacomo and Shih Chii Liu (2015). "Memory and Information Processing in Neuromorphic Systems". In: *Proceedings of the IEEE* 103.8, pp. 1379–1397. ISSN: 00189219. DOI: [10.1109/JPROC.2015.2444094](https://doi.org/10.1109/JPROC.2015.2444094). arXiv: [arXiv:1506.03264v1](https://arxiv.org/abs/1506.03264v1).
- Indiveri, Giacomo et al. (2011). "Neuromorphic silicon neuron circuits". In: *Frontiers in Neuroscience* 5.MAY, pp. 1–23. ISSN: 16624548. DOI: [10.3389/fnins.2011.00073](https://doi.org/10.3389/fnins.2011.00073).
- Johnson, Jeffrey S., John P. Spencer, and Gregor Schoener (2008). "Moving to higher ground: The dynamic field theory and the dynamics of visual cognition". In: *New Ideas in Psychology* 26.2, pp. 227–251. ISSN: 0732118X. DOI: [10.1016/j.newideapsych.2007.07.007](https://doi.org/10.1016/j.newideapsych.2007.07.007).
- Koch, C, M Rapp, and I Segev (1996). "A brief history of time (constants)." In: *Cerebral cortex (New York, N.Y. : 1991)* 6.2, pp. 93–101. ISSN: 1047-3211. DOI: [8670642](https://doi.org/10.1007/BF02262501). URL: [http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=8670642&retmode=xml](https://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=8670642&retmode=xml)

- ref{\&}cmd=prlinks\$\backslash npapers3://publication/uuid/8C9D1E9B-3B21-4255-8190-4D78966263F2.
- Lichtsteiner, P. and T. Delbrück (2005). "A 64??64 AER logarithmic temporal derivative silicon retina". In: *2005 PhD Research in Microelectronics and Electronics - Proceedings of the Conference II.1*, pp. 406–409. DOI: [10.1109/RME.2005.1542972](https://doi.org/10.1109/RME.2005.1542972).
- Ling, Doris (2015). "Behavioral Simulations of Spike-based Learning VLSI Circuits". In:
- Liu, Shih Chii and Tobi Delbrück (2010). "Neuromorphic sensory systems". In: *Current Opinion in Neurobiology* 20.3, pp. 288–295. ISSN: 09594388. DOI: [10.1016/j.conb.2010.03.007](https://doi.org/10.1016/j.conb.2010.03.007). URL: <http://dx.doi.org/10.1016/j.conb.2010.03.007>.
- Manoonpong, Poramate, Ulrich Parlitz, and Florentin Wörgötter (2013). "Neural control and adaptive neural forward models for insect-like, energy-efficient, and adaptable locomotion of walking machines." In: *Frontiers in neural circuits* 7.February, p. 12. ISSN: 1662-5110. DOI: [10.3389/fncir.2013.00012](https://doi.org/10.3389/fncir.2013.00012). URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/3570936/>{\&}tool=pmcentrez{\&}rendertype=abstract.
- McCormick, D a et al. (1985). "Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex." In: *Journal of neurophysiology* 54.4, pp. 782–806. ISSN: 0022-3077.
- Pouille, F. et al. (2009). "Input normalization by global feedforward inhibition expands cortical dynamic range." In: *Nature neuroscience* 12.12, pp. 1577–85. ISSN: 1546-1726. DOI: [10.1038/nn.2441](https://doi.org/10.1038/nn.2441). URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/19881502/>.
- Qiao, Ning et al. (2015). "A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128K synapses". In: *Frontiers in Neuroscience* 9.APR, pp. 1–17. ISSN: 1662453X. DOI: [10.3389/fnins.2015.00141](https://doi.org/10.3389/fnins.2015.00141).
- Rasche, Christoph and Richard H R Hahnloser (2001). "Silicon synaptic depression". In: *Biological Cybernetics* 84.1, pp. 57–62. ISSN: 03401200. DOI: [10.1007/s004220170004](https://doi.org/10.1007/s004220170004).
- Richter, Mathis, Yulia Sandamirskaya, and Gregor Schoner (2012). "A robotic architecture for action selection and behavioral organization inspired by human cognition". In: *IEEE International Conference on Intelligent Robots and Systems*, pp. 2457–2464. ISBN: 9781467317375. DOI: [10.1109/IROS.2012.6386153](https://doi.org/10.1109/IROS.2012.6386153).

- Riesenhuber, M. and T. Poggio (1999a). "Hierarchical models of object recognition in cortex." In: *Nature neuroscience* 2.11, pp. 1019–25. ISSN: 1097-6256. DOI: [10.1038/14819](https://doi.org/10.1038/14819). URL: <http://www.ncbi.nlm.nih.gov/pubmed/10526343>.
- Riesenhuber, Maximilian and Tomaso Poggio (1999b). "Hierarchical models of object recognition in cortex." In: *Nature neuroscience* 2.11, pp. 1019–25. ISSN: 1097-6256. DOI: [10.1038/14819](https://doi.org/10.1038/14819). URL: <http://www.ncbi.nlm.nih.gov/pubmed/10526343>.
- Rutishauser, Ueli, Jean-Jacques Slotine, and Rodney J Douglas (2012). "Competition through selective inhibitory synchrony." In: *Neural computation* 24.8, pp. 2033–52. ISSN: 1530-888X. DOI: [10.1162/NECO_a_00304](https://doi.org/10.1162/NECO_a_00304). arXiv: [1201.2845](https://arxiv.org/abs/1201.2845). URL: <http://www.ncbi.nlm.nih.gov/pubmed/22509969>.
- Sandamirskaya, Yulia (2013a). "Dynamic neural fields as a step toward cognitive neuromorphic architectures." In: *Frontiers in neuroscience* 7, p. 276. ISSN: 1662-4548. DOI: [10.3389/fnins.2013.00276](https://doi.org/10.3389/fnins.2013.00276). URL: <http://journal.frontiersin.org/article/10.3389/fnins.2013.00276/abstract>.
- (2013b). "Dynamic neural fields as a step toward cognitive neuromorphic architectures." In: *Frontiers in neuroscience* 7.January, p. 276. ISSN: 1662-4548. DOI: [10.3389/fnins.2013.00276](https://doi.org/10.3389/fnins.2013.00276). URL: <http://journal.frontiersin.org/article/10.3389/fnins.2013.00276/abstract>.
- Sandamirskaya, Yulia and Gregor Schoener (2010). "An embodied account of serial order: How instabilities drive sequence generation". In: *Neural Networks* 23.10, pp. 1164–1179. ISSN: 08936080. DOI: [10.1016/j.neunet.2010.07.012](https://doi.org/10.1016/j.neunet.2010.07.012).
- Sandamirskaya, Yulia et al. (2013). "Using Dynamic Field Theory to extend the embodiment stance toward higher cognition". In: *New Ideas in Psychology* 31.3, pp. 322–339. ISSN: 0732118X. DOI: [10.1016/j.newideapsych.2013.01.002](https://doi.org/10.1016/j.newideapsych.2013.01.002).
- Schoener, Gregor (2008). "Dynamical systems approaches to cognition". In: *The Cambridge handbook of computational psychology*, pp. 101–126. ISSN: 02732297. DOI: [10.1016/j.dr.2005.10.002](https://doi.org/10.1016/j.dr.2005.10.002).
- Senn, Walter (2002). "Beyond spike timing: The role of nonlinear plasticity and unreliable synapses". In: *Biological Cybernetics* 87.5-6, pp. 344–355. ISSN: 03401200. DOI: [10.1007/s00422-002-0350-1](https://doi.org/10.1007/s00422-002-0350-1).

- Senn, Walter and Stefano Fusi (2005). "Learning only when necessary: better memories of correlated patterns in networks with bounded synapses." In: *Neural computation* 17.10, pp. 2106–2138. ISSN: 0899-7667. DOI: [10.1162/0899766054615644](https://doi.org/10.1162/0899766054615644).
- Song, Sen and L. F. Abbott (2001). "Cortical development and remapping through spike timing-dependent plasticity". In: *Neuron* 32.2, pp. 339–350. ISSN: 08966273. DOI: [10.1016/S0896-6273\(01\)00451-2](https://doi.org/10.1016/S0896-6273(01)00451-2).
- Stefanini, Fabio et al. (2014). "PyNCS: a microkernel for high-level definition and configuration of neuromorphic electronic systems." In: *Frontiers in neuroinformatics* 8.5, p. 73. ISSN: 1662-5196. DOI: [10.3389/fninf.2014.00073](https://doi.org/10.3389/fninf.2014.00073). URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4152885/> tool=pmcentrez&rendertype=abstract.
- Technology, Information and Electrical Engineering Tianlu (2015). "Semester Project Towards Competitive Learning in Neuromorphic Devices Department of Information Technology and Electrical Engineering". In: July.
- Thelen, E and L B Smith (1994). *A dynamic systems approach to the development of cognition and action*. Vol. 512, p. 376. ISBN: 0262200953. DOI: [10.1162/jocn.1995.7.4.512](https://doi.org/10.1162/jocn.1995.7.4.512).
- Van Essen, D C, C H Anderson, and D J Felleman (1992). "Information processing in the primate visual system: an integrated systems perspective." In: *Science (New York, N.Y.)* 255.5043, pp. 419–423. ISSN: 0036-8075. DOI: [10.1126/science.1734518](https://doi.org/10.1126/science.1734518). URL: http://www.ncbi.nlm.nih.gov/sites/entrez?Db=pubmed&DbFrom=pubmed&Cmd=Link&LinkName=pubmed_pubmed&LinkReadableName=RelatedArticles&IdsFromResult=1734518&ordinalpos=3&itool=EntrezSystem2.PEntrez.Pubmed.Pubmed_RVDocSum.
- Wang, Xiao Jing (2012). "Neural dynamics and circuit mechanisms of decision-making". In: *Current Opinion in Neurobiology* 22.6, pp. 1039–1046. ISSN: 18736882. DOI: [10.1016/j.conb.2012.08.006](https://doi.org/10.1016/j.conb.2012.08.006). arXiv: [NIHMS150003](https://arxiv.org/abs/1500.03). URL: <http://dx.doi.org/10.1016/j.conb.2012.08.006>.
- Wickelgren, Wayne A. (1965). "Short-term memory for phonemically similar lists". In: *The American Journal of Psychology* 78.4, pp. 567–574. URL: <http://www.jstor.org/stable/1420917>.
- Wilson, H R and J D Cowan (1973). "A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue." In: *Kybernetik* 13.2, pp. 55–80. ISSN: 0340-1200. DOI: [10.1007/BF00288786](https://doi.org/10.1007/BF00288786).

- Zhang, Zhi and Qian Quan Sun (2011). "The Balance Between Excitation And Inhibition And Functional Sensory Processing In The Somatosensory Cortex". In: *International Review of Neurobiology* 97, pp. 305–332. ISSN: 00747742. DOI: [10.1016/B978-0-12-385198-7.00012-6](https://doi.org/10.1016/B978-0-12-385198-7.00012-6).