Real-time mixed-signal electronic circuits for understanding and implementing neural computation

(how to exploit the physics of an imprecise computing substrate)

Giacomo Indiveri

Institute of Neuroinformatics
University of Zurich and ETH Zurich

Fast Machine Learning for Science
ETH Zurich, September 1, 2025

Universität Zürich[UZH]

ETH zürich

Trends for reducing energy in AI computing systems

Conventional approaches

1. Make application specific (lose general purpose flexibility)
2. Quantize parameters (reduce bit precision)
3. Minimize resource usage (reduce accuracy)

Novel approaches

1. Reduce data movement (implement in-memory computing)
2. Reduce clock switching (use asynchronous circuits)
3. Exploit all the physics of the devices (mix analog and digital)

Trends for reducing energy in AI computing systems

Universität Zürich[UZH] ETH zürich

Conventional approaches

1. Make application specific (lose general purpose flexibility)
2. Quantize parameters (reduce bit precision)
3. Minimize resource usage (reduce accuracy)

Novel approaches

1. Reduce data movement (implement in-memory computing)
2. Reduce clock switching (use asynchronous circuits)
3. Exploit all the physics of the devices (mix analog and digital)

Trends for reducing energy in AI computing systems

Conventional approaches

1. Make application specific (lose general purpose flexibility)
2. Quantize parameters (reduce bit precision)
3. Minimize resource usage (reduce accuracy)

Novel approaches

1. Reduce data movement (implement in-memory computing)
2. Reduce clock switching (use asynchronous circuits)
3. Exploit all the physics of the devices (mix analog and digital)

Principles of neural design

- Co-localize memory and computation (local processing, local state variables)
- Maximize fine-grain parallelism (massively parallel arrays of memory and processing)
- Adapt time-scales to data rates (match circuit time constants to signal dynamics)
- Use the “physics of computation” (exploit properties of computing substrate)

Clock speed

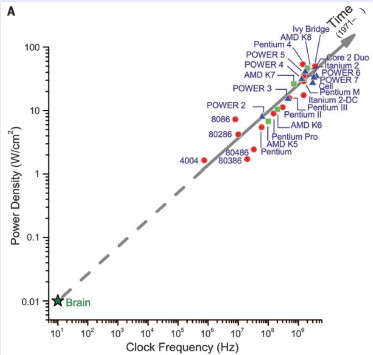

Brains outperform faster computing systems in many sensory processing tasks at lower speeds, with less power.

Computing substrate

- Slow, noisy and variable processing elements.
- Distributed across space (no time multiplexing).
- Local connectivity, small world networks.
- Massively parallel computation.
- Continual always-on learning.
- Real-time spatio-temporal signal processing.

Time represents itself

The brain uses the time evolution of the physical system to implement its computations. Neural circuits compute by exploiting the natural time evolution of their hardware substrate. [Sterling & Laughlin, 2017]

The origins

- Fundamental research
- *Emulation* of neural function
- Subthreshold analog
- Asynchronous digital

1980s

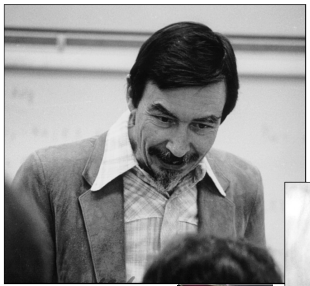

Carver Mead

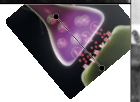

Misha Mahowald

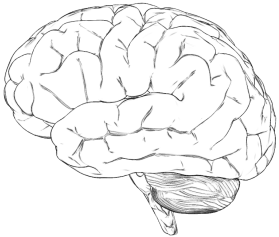

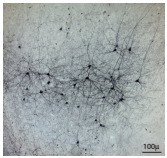

[Nuno da Costa, INI, 2008]

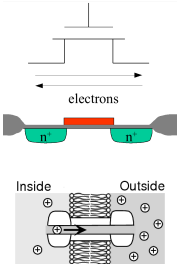

1. Rooted in neuroscience (but covers physics, computer science, microelectronics, . . .)
2. Exploits the physics of electronic devices to emulate the biophysics of neural systems.
3. Lets time represent itself.
4. Explores sensory-processing systems that interact intelligently with the real world.
5. Develops cognitive agents that produce autonomous behavior.

To "*understand by building*", start from the faithful emulation of neural circuits using subthreshold analog circuits.

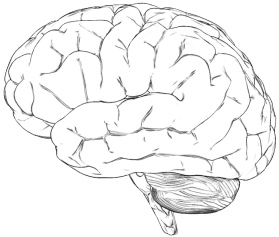

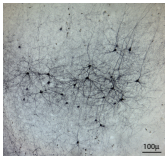

[Nuno da Costa, INI, 2008]

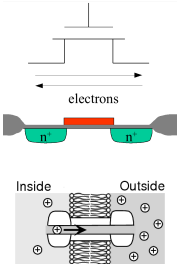

1. Rooted in neuroscience (but covers physics, computer science, microelectronics, . . .)
2. Exploits the physics of electronic devices to emulate the biophysics of neural systems.
3. Lets time represent itself.
4. Explores sensory-processing systems that interact intelligently with the real world.
5. Develops cognitive agents that produce autonomous behavior.

To "*understand by building*", start from the faithful emulation of neural circuits using subthreshold analog circuits.

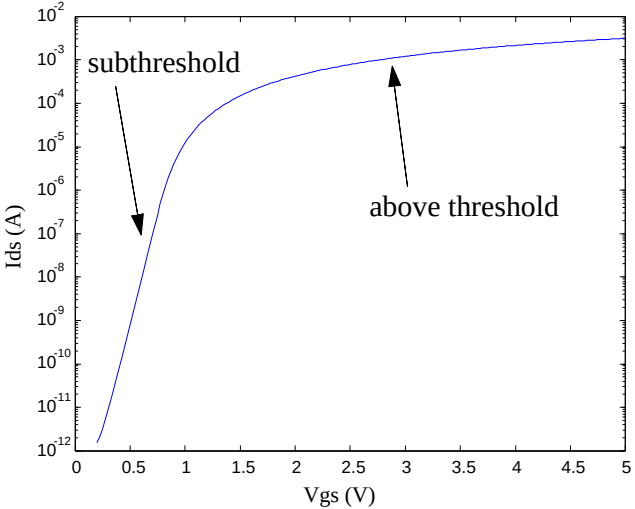
subthreshold
above threshold
Ids (A)
Vgs (V)
10^-2
10^-3
10^-4
10^-5
10^-6
10^-7
10^-8
10^-9
10^-10
10^-11
10^-12
0
0.5
1
1.5
2
2.5
3
3.5
4
4.5
5

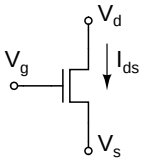
Vd
Vg
Ids
Vs

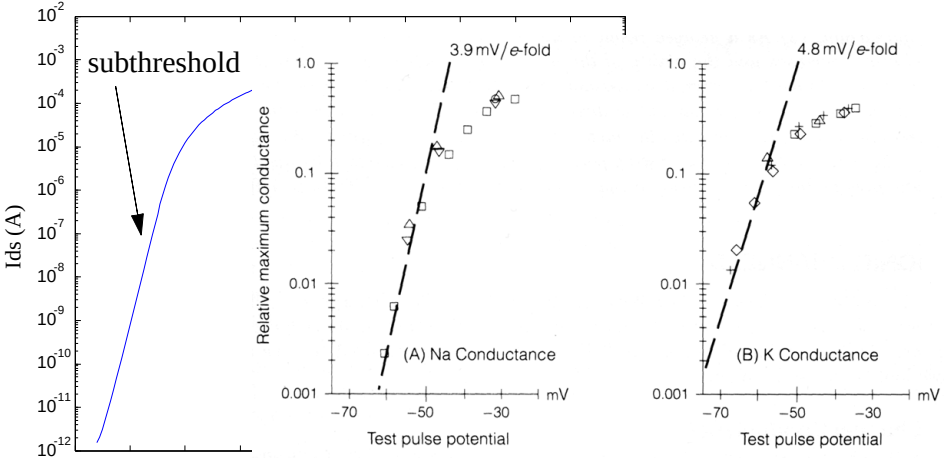

FIGURE 4.6 Exponential current–voltage characteristic of voltage-dependent channels. At high voltages, the fraction of channels that are open approaches unity, causing a saturation of the curves. (*Source:* [Hodgkin et al., 1952b, p. 464].)

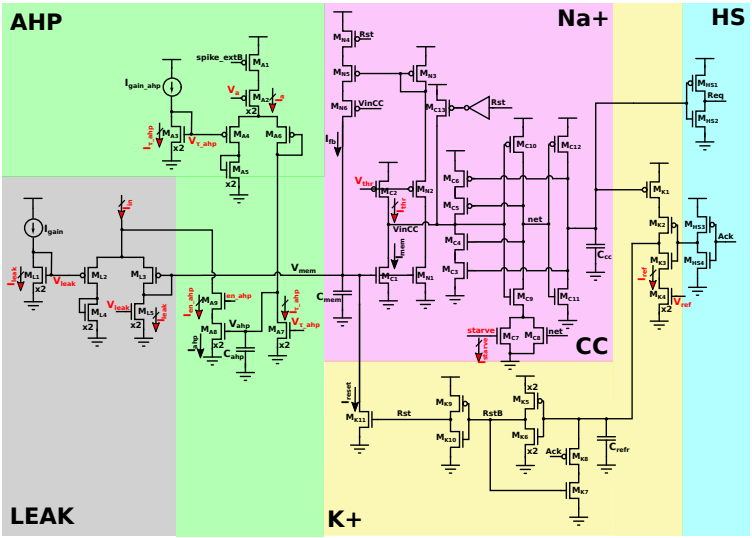

Adaptive Exponential I&F neuron model (beyond HH)

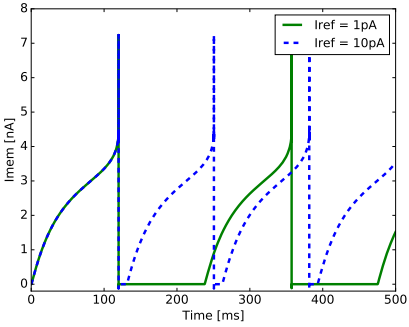

Work	[12]	[19]	[36]	**This work**
Techn.	180 nm CMOS	28 nm CMOS	28 nm FDSOI	22 nm FDSOI
Type	Mixed	Mixed	Mixed	Mixed
V_{dd}	1.8 V	0.7-1 V	1 V	0.8 V
Freq	30 Hz	-	30 Hz	30 Hz
Results	Experimental	Experimental	Simulation	Simulation
En./spike	883 pJ	2.3 nJ-30 nJ	50 pJ	14 pJ

[Brette and Gerstner, 2005, Rubino et al., IEEE TCAS, 2020]

Synapse circuits

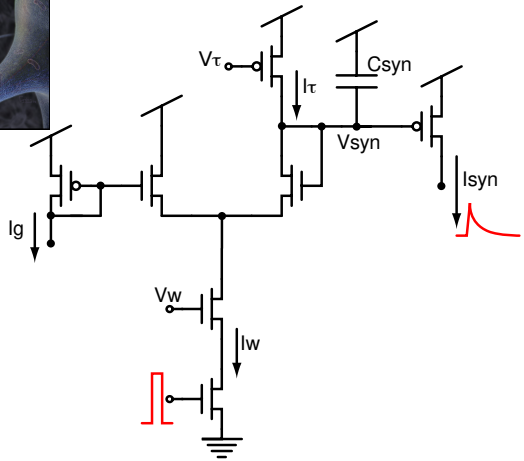

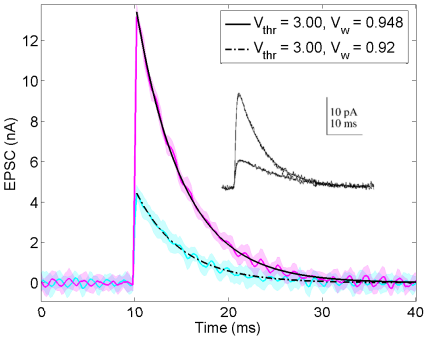

$$\tau \frac{d}{dt} I_{syn} + I_{syn} = \frac{I_g I_w}{I_\tau}$$

$$\tau = \frac{C U_T}{\kappa I_\tau}$$

[Bartolozzi, Indiveri, NECO 2007]

Synapse circuits

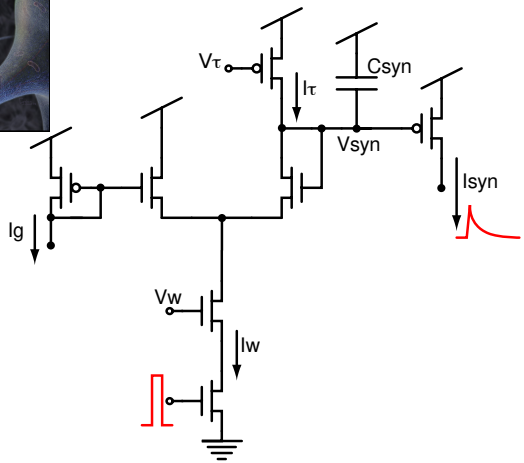

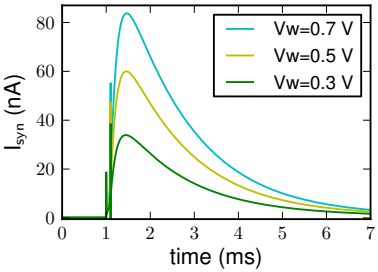

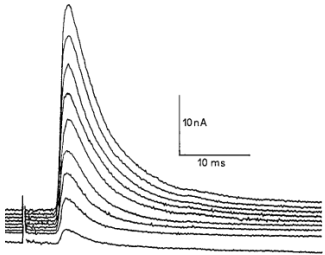

[Sumislawska et al., ISCAS 2016]

Synapse circuits

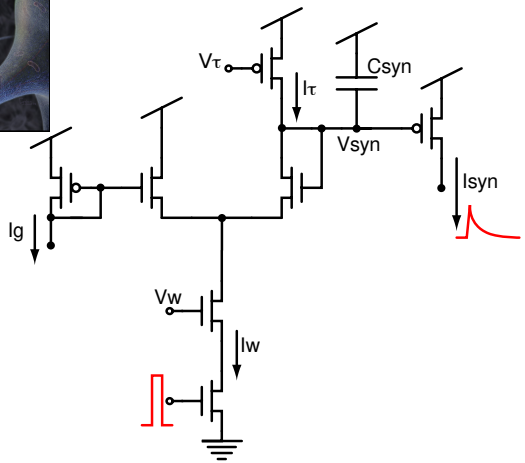

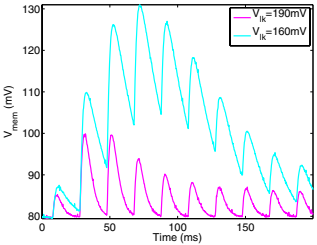

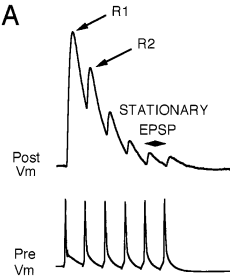

[C. Rache & Hahnloser, 2001] [M. Boegerhausen et al., 2003]

Mixed-signal neuromorphic chips

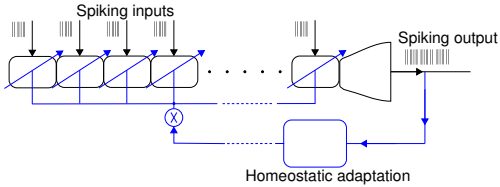

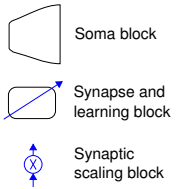

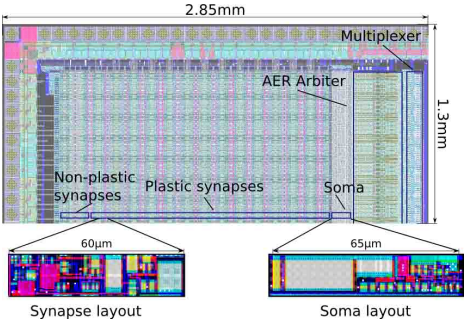

- Spiking neural networks (SNNs)
- Analog subthreshold circuits.
- Slow temporal, non-linear dynamics.
- Massively parallel operation.
- Compatible with memristive devices
- Inhomogeneous, imprecise, and noisy.

Mixed-signal neuromorphic chips

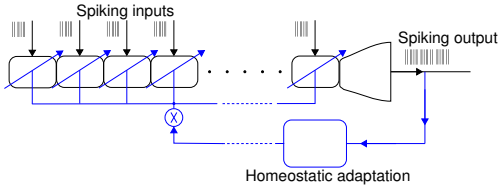

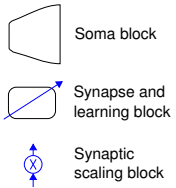

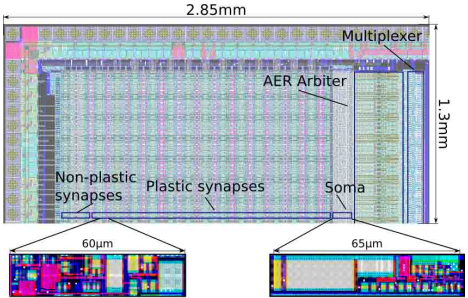

- Spiking neural networks (SNNs)
- Analog subthreshold circuits.
- Slow temporal, non-linear dynamics.
- Massively parallel operation.
- Compatible with memristive devices
- Inhomogeneous, imprecise, and noisy.
- Fast asynchronous digital circuits for routing spikes.
- Reprogrammable network topology

Background

We are building physical, real-time, signal processing systems for real-world sensory data.

Why SNNs? Why spikes?

Background

We are building physical, real-time, signal processing systems for real-world sensory data.

Requirements

1. Robust communication of analog signals across long distances through noisy channels.
2. Local processing, multi-core architectures and distributed computing.
3. Low power and low-latency.

Why SNNs? Why spikes?

Background

We are building physical, real-time, signal processing systems for real-world sensory data.

Requirements

1. Robust communication of analog signals across long distances through noisy channels.
2. Local processing, multi-core architectures and distributed computing.
3. Low power and low-latency.

Optimal solution for communication and computation

- The optimal method that minimizes bandwidth and power consumption for achieving this goal, under these constraints, is pulse-frequency modulation. [A. Mortara et al., 1995, K. Boahen, 1998]
- "*Counter to intuition, computing with spikes can be extremely efficient on neuromorphic hardware even when the problem being solved is mathematically formulated in terms of activity rates.*" [M. Davies, Intel, 2019]

Also real neurons are diverse and inhomogeneous

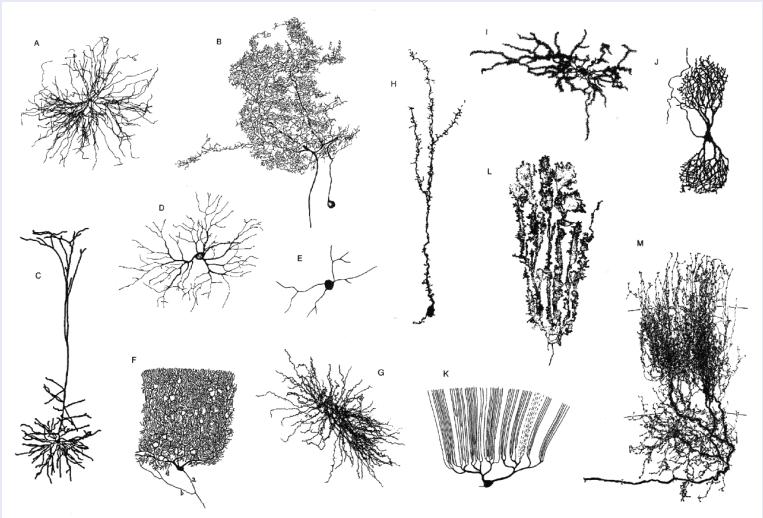

(adapted from [B. Mel, 1994])

Device mismatch effects

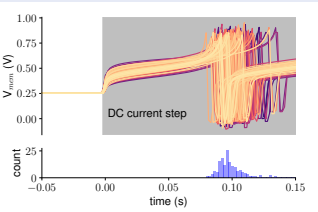

Time-to-first spike measured across 256 different neurons

Real and silicon neurons

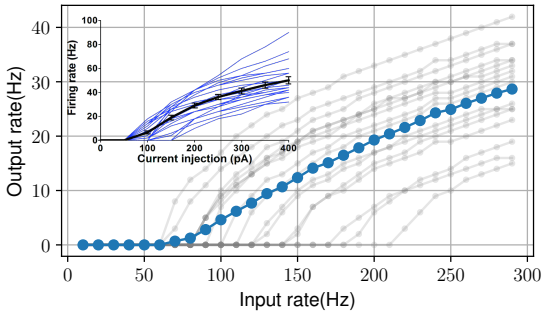

Device mismatch effects

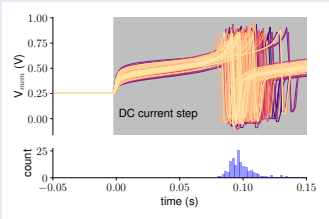

Time-to-first spike measured across 256 different neurons

How to cope with mismatch?

- Use populations of neurons and average over space and time
- Employ negative feedback, adaptation, and learning mechanisms

Device mismatch effects

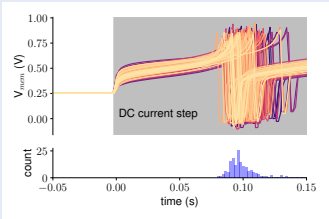

Time-to-first spike measured across 256 different neurons

How to cope with mismatch?

- Use populations of neurons and average over space and time
- Employ negative feedback, adaptation, and learning mechanisms

Choosing bit resolution

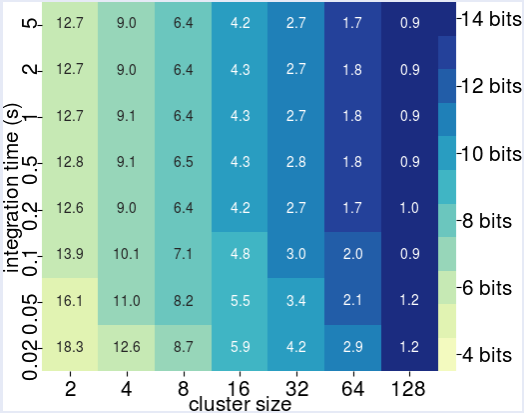

Coefficient of variation and Equivalent Number of Bits (ENOB)

[Zendrikov et al., 2023]

Robust representations via population coding

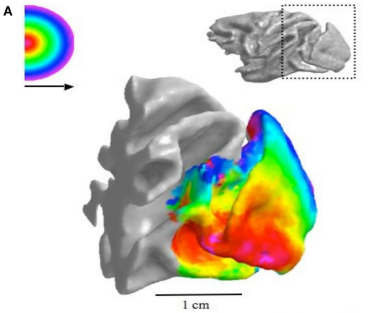

⟵ Retinotopic and orientation maps representing the preference of neurons in the visual cortex for the location and orientation of a stimulus on the visual field.

Robust representations via population coding

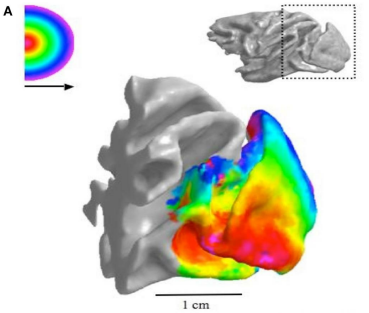

⟵ Retinotopic and orientation maps representing the preference of neurons in the visual cortex for the location and orientation of a stimulus on the visual field.

Orientation tuning: ⟶ Non-human primate response to moving bars (top); Neuromorphic processor response to flashing bars (bottom)

Feature tuning via populations of neurons

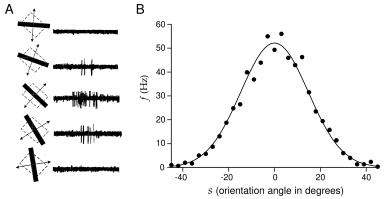

[Dayan & Abbott, 2005]

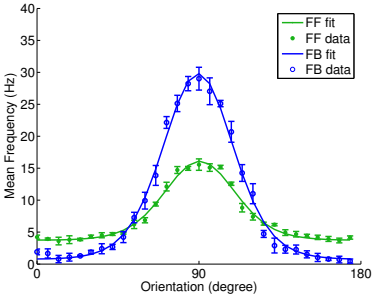

[Chicca et al., 2007]

Added values of population coding

False myth		Reality
• Neural responses are slow.	$\Longrightarrow$	**Populations** of noisy neurons have very fast response times. [T. Tchumatchenko et al., 2011]
• Neurons need to fire at high firing rates to achieve high precision.	$\Longrightarrow$	Sparse neural **population** activity (in space *and* time) can represent signals with high accuracy. [Denève at al., 2017]
• Neurons need to be accurate to propagate precise information across layers.	$\Longrightarrow$	To propagate signals across multiple SNN layers, it is necessary to use inhomogeneous **populations** of neurons. [M.C.W. van Rossum, et al., 2002]

Using **populations** of mismatched silicon neurons in mixed-signal neuromorphic processors has pros and cons:

Costs	Benefits
Increased area	Lower latency Lower power consumption Increased robustness Increased fault-tolerance

Ensemble and stochastic learning can exploit variability of inhomogeneous synapses.

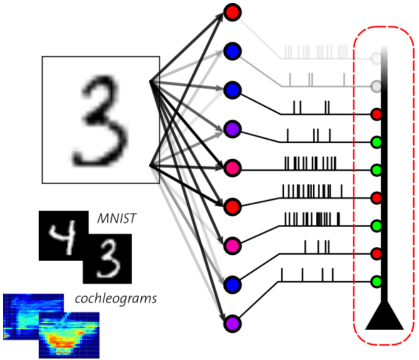

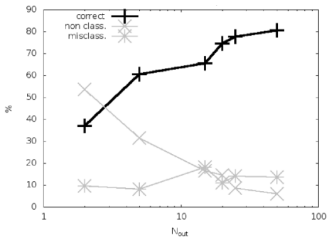

MNIST	deep/CNN (Hinton et al. 2012)	98.4%
	random + bistable synapses	$\sim$ 85%
	random + bistable synapses + (mod. protocol)	$\sim$ 96%
TIMIT	deep/CNN (Hinton et al. 2012)	77%
	VLSI cochlea + bistable synapses	$\sim$ 60%

On-line bagging techniques

AdaBoost theorem: $1-\text{error}(H_{final}) \geq 1 - e^{-2\gamma^2 N}$ [Y. Freund And R. E. Schapire, 1995]

Robust on-chip learning

- Bi-stable synapses with STDP circuits

  [Indiveri et al, 2006]

- Spike-driven synaptic plasticity with stop-learning

  [Mitra et al, 2009,Qiao et al., 2015]

- Error-propagation with local learning

  [Cartiglia et al., 2020]

- Dendritic Hebbian synaptic plasticity with stop-learning

  [Rubino et al, 2023]

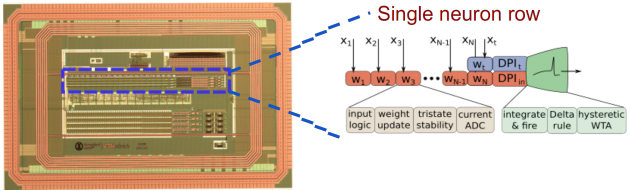

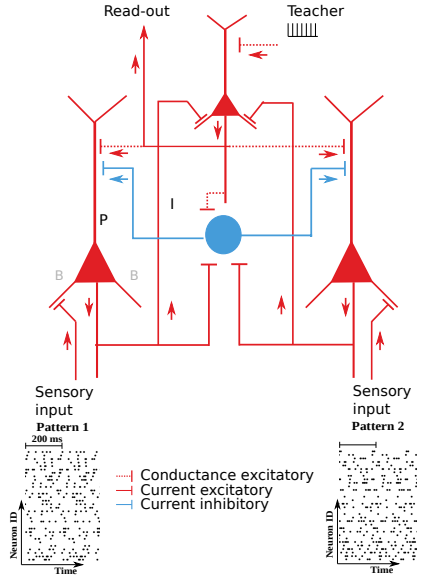

Attractor networks on neuromorphic chips

Synaptic matrix

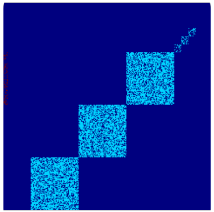

Inhibitory neurons

Excitatory neurons

Excitatory neurons

Excitatory neurons

Output spikes

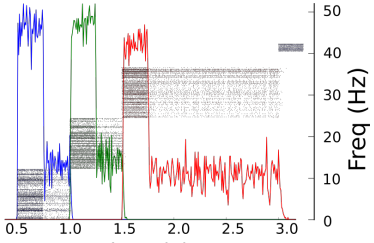

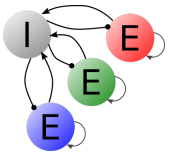

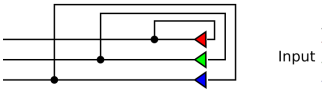

Input

Input

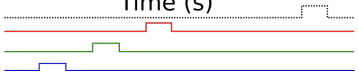

Methods and tools:

- Mean Field Theory
- Effective Response Function
- Self consistency condition

[M. Giulioni et al., 2012, M. Giulioni et al., 2015]

$$\tau \frac{d}{dt}\mu = -\mu + W\nu_{in} - \beta$$

$$\nu_{out} = \Phi(\mu(\nu_{in}), \sigma)$$

$$\nu_{out} = \nu_{in} = \nu$$

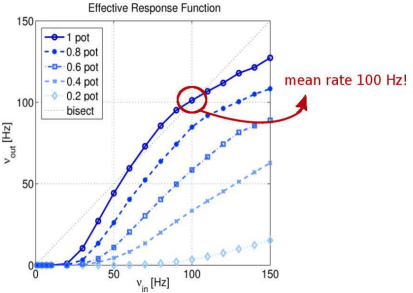

Soft Winner-Take-All networks

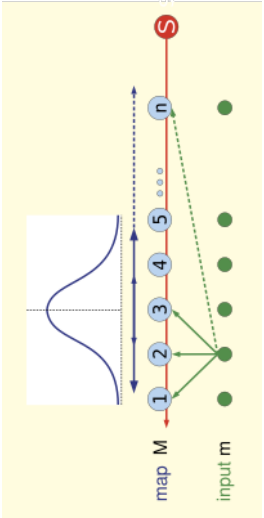

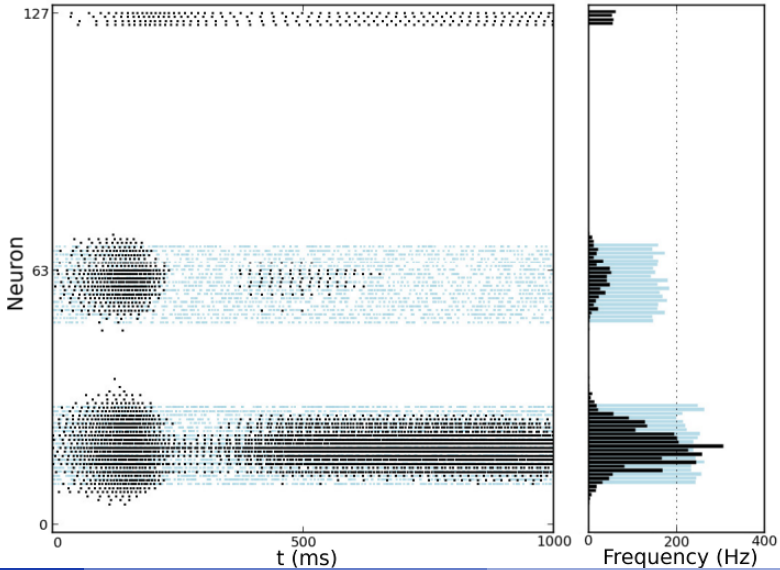

Soft Winner-Take-All drifting

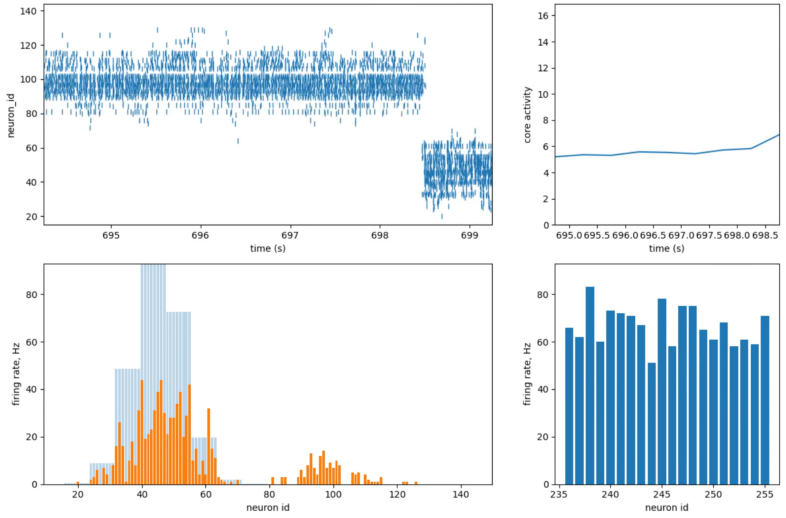

[Zendrikov et al., 2023]

Auditory Cortex hierarchy
(from sounds to words)

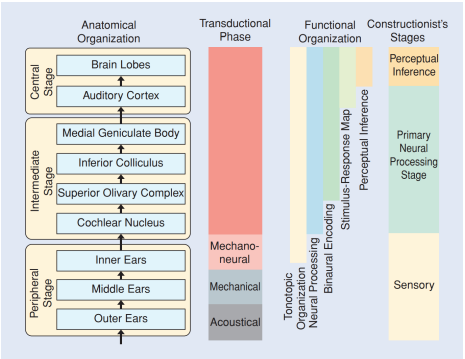

[R. Munkong & B.-H. Juang, 2008]

From signal processing to computing

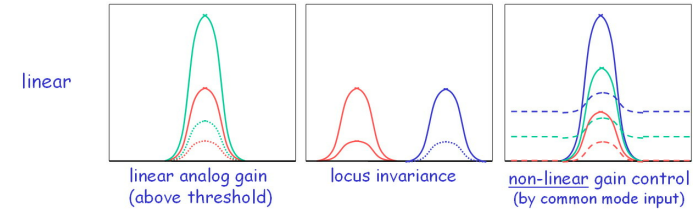

[R.J. Douglas & K.A.C. Martin, 2013]

Auditory Cortex hierarchy
(from sounds to words)

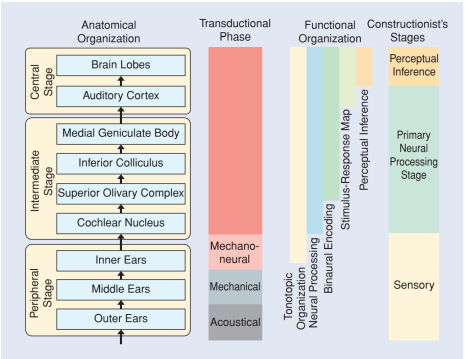

[R. Munkong & B.-H. Juang, 2008]

From signal processing to computing

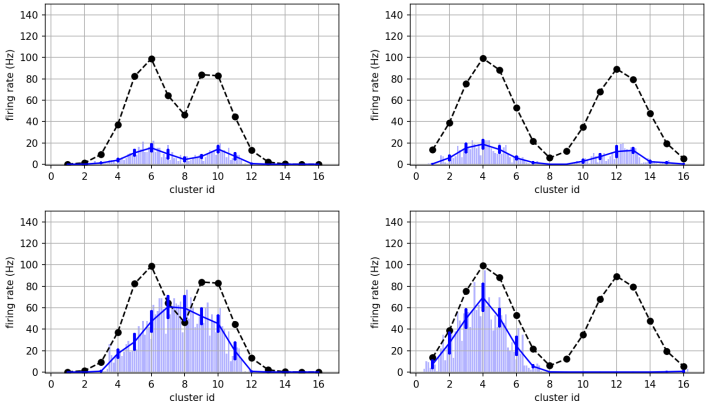

Auditory Cortex hierarchy

(from sounds to words)

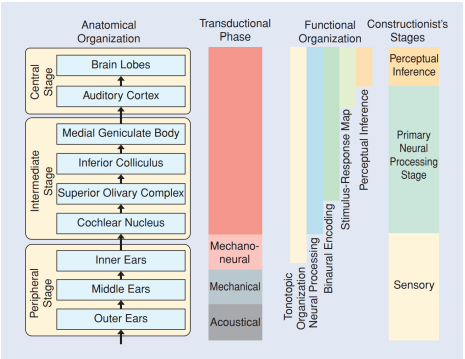

[R. Munkong & B.-H. Juang, 2008]

From signal processing to computing

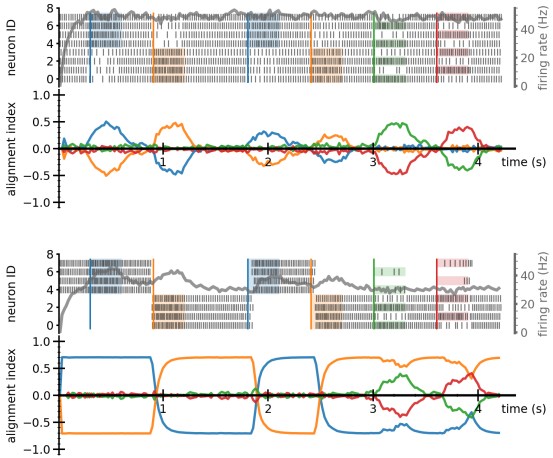

[Zendrikov et al., 2025]

Neuromorphic vs conventional processors

Pros

- Low-power (< 1 mW)
- Low latency
- ...

Cons

- High area
- High variability, noisy
- Low(er) accuracy

Neuromorphic vs conventional processors

Pros

- Low-power (< 1 mW)
- Low latency
- ...

Cons

- High area
- High variability, noisy
- Low(er) accuracy

What are they good for?

- Closed-loop sensory-motor processing
- Multi-modal sensory fusion
- Always-on on-line learning

What are they bad at?

- High precision number crunching
- High accuracy pattern recognition
- Batch processing of large data sets

Neuromorphic vs conventional processors

Pros

- Low-power (< 1 mW)
- Low latency
- ...

Cons

- High area
- High variability, noisy
- Low(er) accuracy

What are they good for?

- Closed-loop sensory-motor processing
- Multi-modal sensory fusion
- Always-on on-line learning

What are they bad at?

- High precision number crunching
- High accuracy pattern recognition
- Batch processing of large data sets

Open challenges

- How to do reliable computation using a noisy and heterogeneous computing substrate (check).
- How to "program" this computational substrate, combining *computational primitives* (NSMs).

Application examples: extreme edge computing

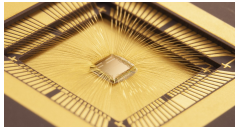

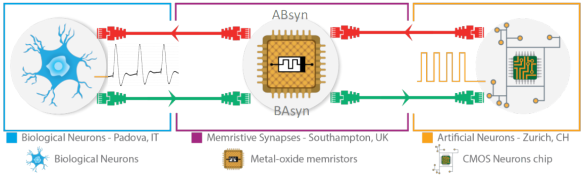

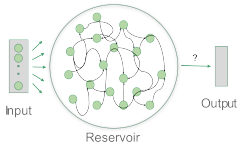

- On-line ECG anomaly and heart rate detection [De Luca et al., 2025,Bauer et al., 2019]
- High-Frequency Oscillation (HFO) detection [Sharifhazileh, Burelo et al., 2021]
- Adaptive pace-maker with neuromorphic CPG network [Abu-Hassan et al., 2019]
- On-line classification of EMG signals [Donati et al., 2019]
- Neuromorphic pattern generation circuits for bioelectronic medicine [Donati et al., 2021]
- Closed-loop bidirectional brain machine interfaces with in rats and cell-cultures [Boi et al., 2016] [Serb et al. 2020]
- Zebra-finch “Bird’s Own Song” classification [Corradi et al., 2015]
- Instantaneous stereo depth estimation with a stereo-vision setup [Risi et al., 2021]
- On-line detection of vibration anomalies using balanced spiking neural networks [Dennler et al., 2021]
- Closed-loop obstacle avoidance on roving robot [Milde et al. 2017]
- Closed-loop robot head position control with a neuromorphic processor [Zhao et al., 2020]

Application examples: extreme edge computing

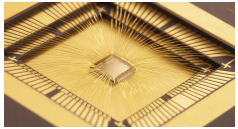

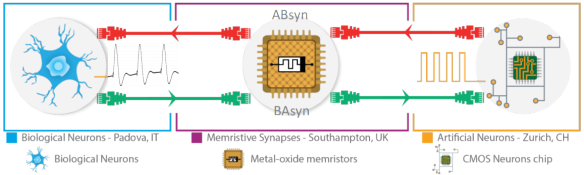

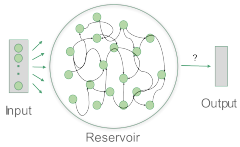

- On-line ECG anomaly and heart rate detection [De Luca et al., 2025,Bauer et al., 2019]
- High-Frequency Oscillation (HFO) detection [Sharifhazileh, Burelo et al., 2021]
- Adaptive pace-maker with neuromorphic CPG network [Abu-Hassan et al., 2019]
- On-line classification of EMG signals [Donati et al., 2019]
- Neuromorphic pattern generation circuits for bioelectronic medicine [Donati et al., 2021]
- Closed-loop bidirectional brain machine interfaces with in rats and cell-cultures [Boi et al., 2016] [Serb et al. 2020]
- Zebra-finch “Bird’s Own Song” classification [Corradi et al., 2015]
- Instantaneous stereo depth estimation with a stereo-vision setup [Risi et al., 2021]
- On-line detection of vibration anomalies using balanced spiking neural networks [Dennler et al., 2021]
- Closed-loop obstacle avoidance on roving robot [Milde et al. 2017]
- Closed-loop robot head position control with a neuromorphic processor [Zhao et al., 2020]

On-line classification of EMG signals

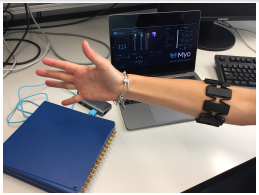

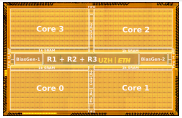

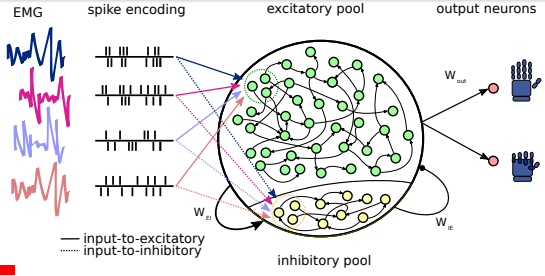

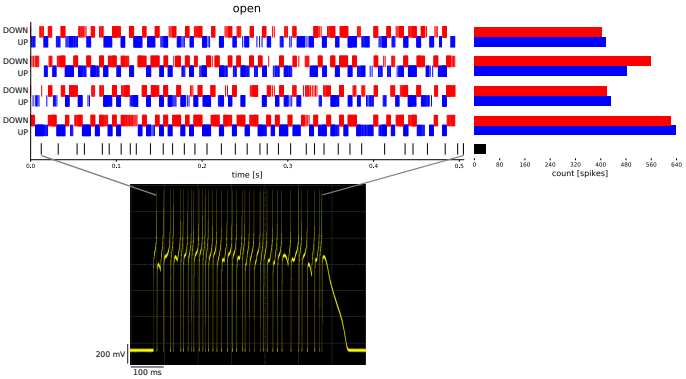

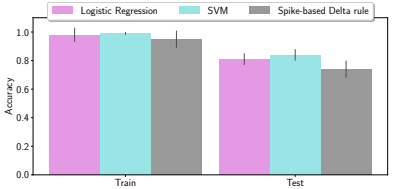

[Donati et al., 2019]

Physical computation for “edge intelligence”

We are now entering the era of *neuromorphic intelligence* in which dedicated cognitive “chiplets” will be used to provide intelligence to a multitude of extreme edge-computing devices

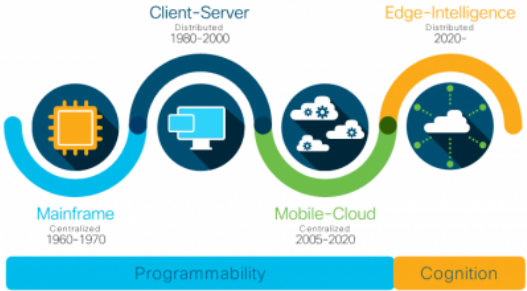

SynSense

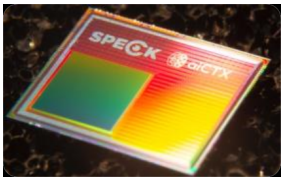

- Health monitoring
- Wearable sensors
- Environmental sensing
- Industrial monitoring
- Intelligent machine vision
- Consumer applications

nature
International journal of science

EDITORIAL · 06 FEBRUARY 2018

Big data needs a hardware revolution

Artificial intelligence is driving the next wave of innovations in the semiconductor industry.

Software companies make headlines but research on computer hardware could bring bigger rewards. Credit: Morris MacMatzen/Getty

- Conventional AI increasing power requirements are unsustainable.
- New emerging memory technologies will benefit from massively parallel processing architectures.
- Neuroscience and machine learning are uncovering powerful and robust neural processing methods.
- Hardware implementations of spiking neural networks and sparse event-based sensory-processing systems are starting to show their advantages.
- This is the perfect time to follow the “neuromorphic computing and engineering” approach for starting a hardware revolution.

nature
International journal of science

EDITORIAL · 06 FEBRUARY 2018

Big data needs a hardware revolution

Artificial intelligence is driving the next wave of innovations in the semiconductor industry.

Software companies make headlines but research on computer hardware could bring bigger rewards. Credit: Morris MacMatzen/Getty

- Conventional AI increasing power requirements are unsustainable.
- New emerging memory technologies will benefit from massively parallel processing architectures.
- Neuroscience and machine learning are uncovering powerful and robust neural processing methods.
- Hardware implementations of spiking neural networks and sparse event-based sensory-processing systems are starting to show their advantages.
- This is the perfect time to follow the “neuromorphic computing and engineering” approach for starting a hardware revolution.

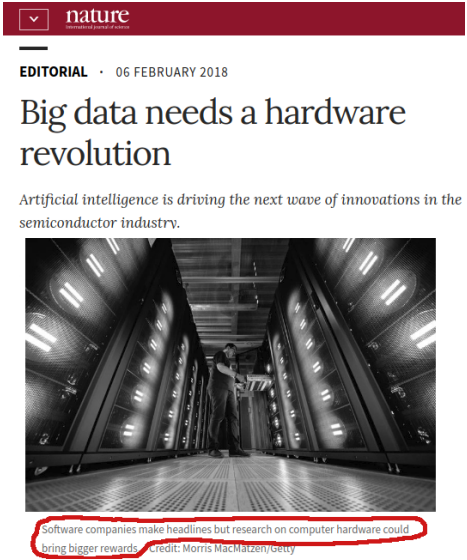

nature

EDITORIAL · 06 FEBRUARY 2018

Big data needs a hardware revolution

Artificial intelligence is driving the next wave of innovations in the semiconductor industry.

Software companies make headlines but research on computer hardware could bring bigger rewards. Credit: Morris MacMatzen/Getty

- Conventional AI increasing power requirements are unsustainable.
- New emerging memory technologies will benefit from massively parallel processing architectures.
- Neuroscience and machine learning are uncovering powerful and robust neural processing methods.
- Hardware implementations of spiking neural networks and sparse event-based sensory-processing systems are starting to show their advantages.
- This is the perfect time to follow the “neuromorphic computing and engineering” approach for starting a hardware revolution.

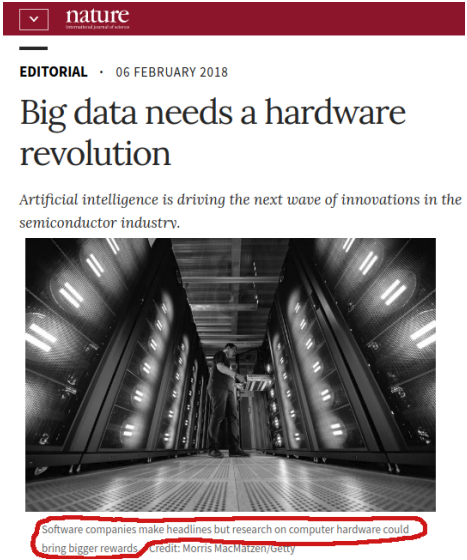

nature

EDITORIAL · 06 FEBRUARY 2018

Big data needs a hardware revolution

Artificial intelligence is driving the next wave of innovations in the semiconductor industry.

Software companies make headlines but research on computer hardware could bring bigger rewards. Credit: Morris MacMatzen/Getty

- Conventional AI increasing power requirements are unsustainable.
- New emerging memory technologies will benefit from massively parallel processing architectures.
- Neuroscience and machine learning are uncovering powerful and robust neural processing methods.
- Hardware implementations of spiking neural networks and sparse event-based sensory-processing systems are starting to show their advantages.
- This is the perfect time to follow the “neuromorphic computing and engineering” approach for starting a hardware revolution.

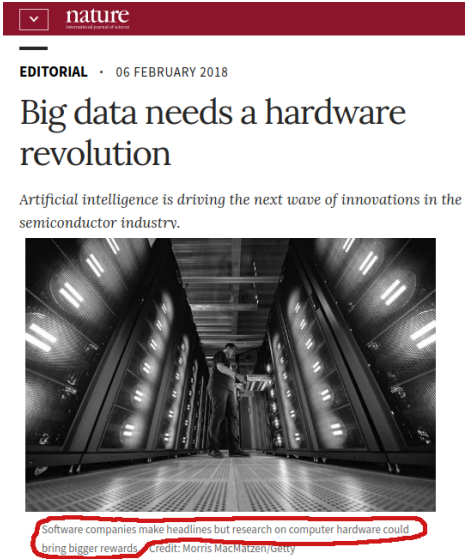

nature

EDITORIAL · 06 FEBRUARY 2018

Big data needs a hardware revolution

Artificial intelligence is driving the next wave of innovations in the semiconductor industry.

Software companies make headlines but research on computer hardware could bring bigger rewards. Credit: Morris MacMatzen/Getty

- Conventional AI increasing power requirements are unsustainable.
- New emerging memory technologies will benefit from massively parallel processing architectures.
- Neuroscience and machine learning are uncovering powerful and robust neural processing methods.
- Hardware implementations of spiking neural networks and sparse event-based sensory-processing systems are starting to show their advantages.
- This is the perfect time to follow the “neuromorphic computing and engineering” approach for starting a hardware revolution.

Acknowledgments

institute of neuroinformatics

Universität Zürich[UZH] ETH zürich

- Elisa Donati
- Chiara De Luca
- Sapta Ghosh
- Chenxi Wen
- Dmitrii Zendrikov
- Junren Chen
- Zhe Su
- Farah Baracat
- Patrick Bösch
- Giovanni Camisa
- Olympia Gallou
- Leonardo Martinelli
- Maryada
- Qi Shen
- Tianshu Shen
- Shyam Narayanan

erc

ICT H2020

FNSNF Swiss National Science Foundation

Thank you for your attention

Backup slides

NEUROMORPHIC
Computing and Engineering

OPEN ACCESS

No APCs in 2024

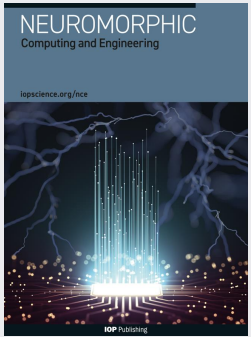

A multidisciplinary, open access journal devoted to the application and development of neuromorphic computing, devices, and systems in advancing new scientific discovery and realising emerging new technologies.

Editor-in-Chief
Giacomo Indiveri University of Zurich, Switzerland

Indexed in Scopus and Web of Science
IMPACT FACTOR COMING IN JUNE **2024**

IOP Publishing

Iopscience.org/nce
nce@ioppublishing.org

@IOPneuromorphic

http://capocaccia.cc/

- Interdisciplinary, international, diverse
- Morning lectures, afternoon hands-on work-groups
- Active and lively discussions (no powerpoint)
- Concrete results, establishment of long-term collaborations

Capo Caccia, Sardinia, Italy. April 28 - May 11, 2024

Take home message for neuromorphic intelligence

Academic/basic research

- Study real brains, start from small neural circuits/systems
- Take into account all properties of neurons and synapses
- Focus on fundamental problems (ignore incremental benchmarks)
- "There's plenty of room at the bottom" (large scale is not all)

Applied/industrial research

- Choose a specific problem to solve that is not being solved yet
- Consider it's requirements in it's entirety, from end to end
- Be open to using the best of all possible approaches (analog and digital)
- Build the full ecosystem for your solution (devices, software, users)

Early access:

Bottom-Up and Top-Down Approaches for the Design of Neuromorphic Processing Systems: Tradeoffs and Synergies Between Natural and Artificial Intelligence, Frenkel and Indiveri, Proceedings IEEE, 2023.

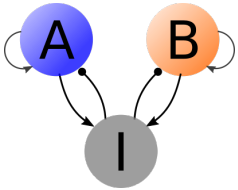

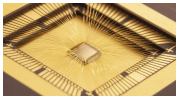

[R. Krause et al., 2021]

[E. Donati et al., 2021]

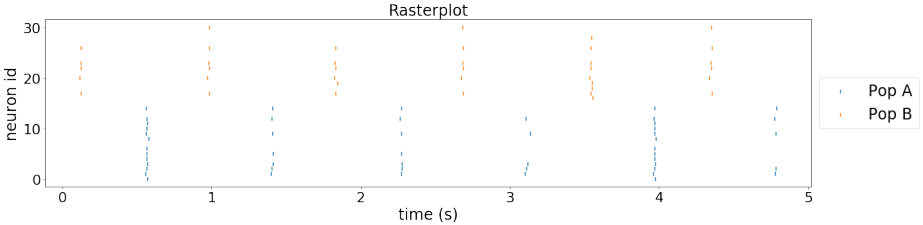

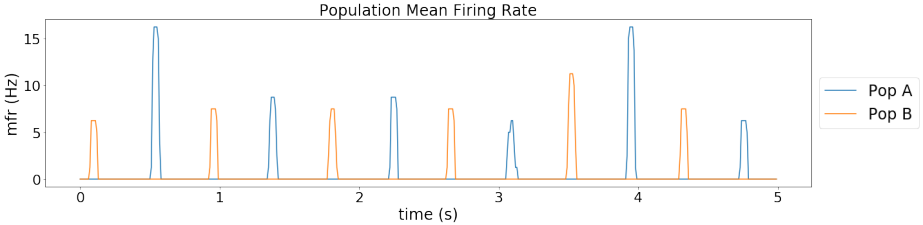

Neural oscillators and Central Pattern Generators

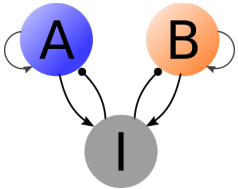

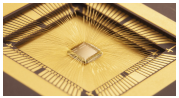

[R. Krause et al., 2021]

[E. Donati et al., 2021]

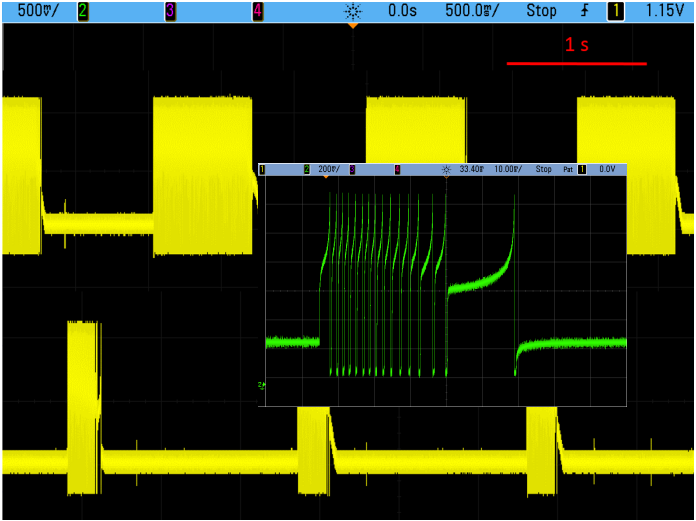

Excitation-Inhibition balanced networks

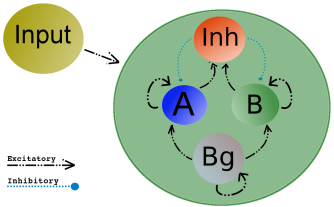

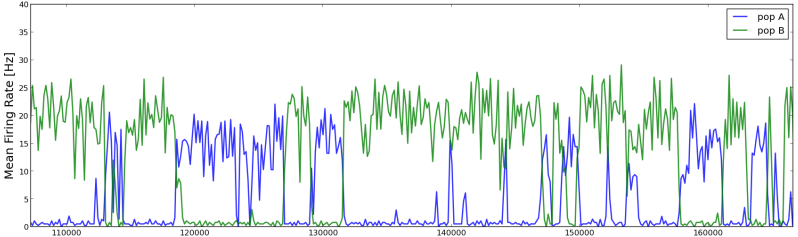

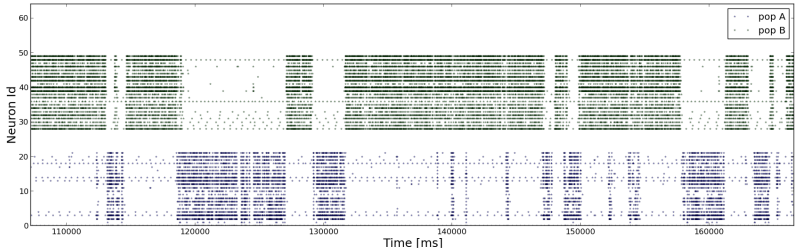

The canonical micro-circuit as a soft-WTA

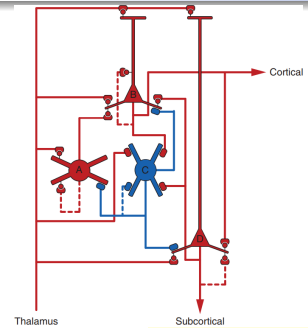

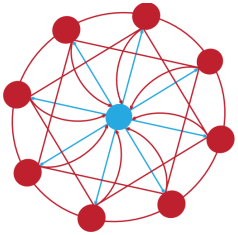

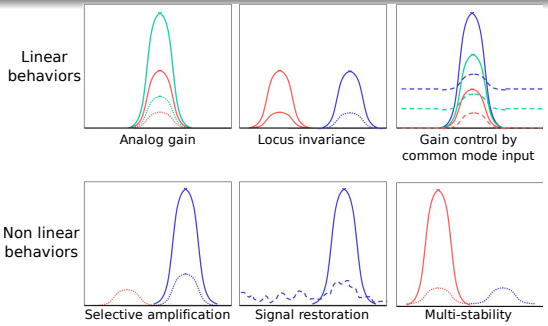

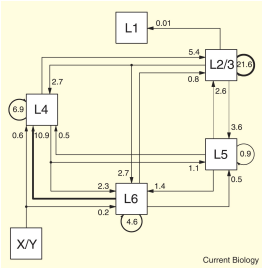

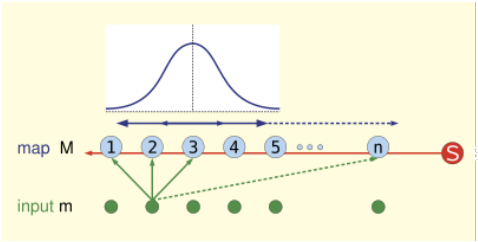

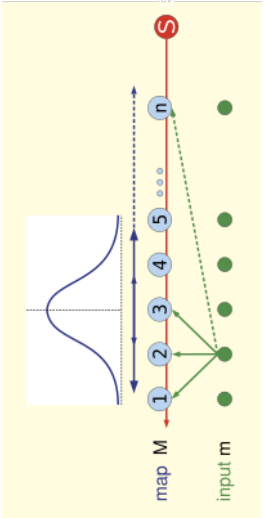

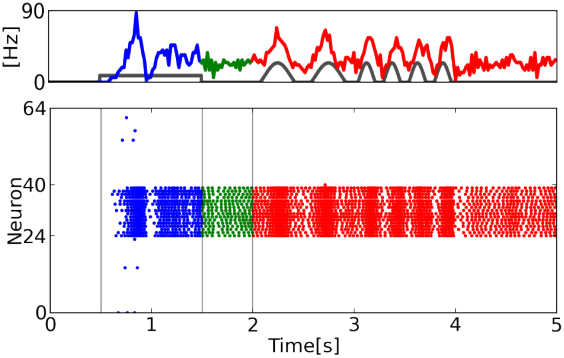

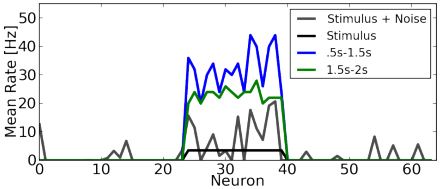

[Neftci et al., 2013]

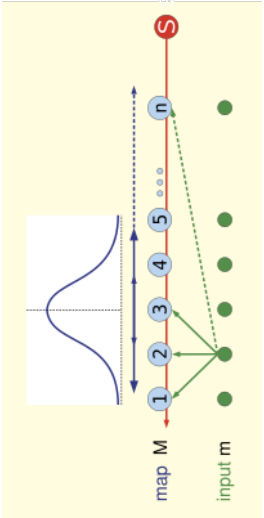
S
n
5
4
3
2
1
map M
input m

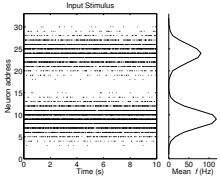
Input Stimulus
Neuron address
Time (s)
Mean f (Hz)

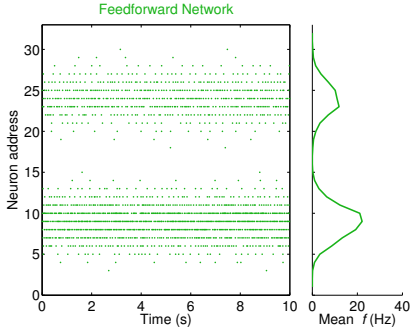
Feedforward Network
Neuron address
Time (s)
Mean f (Hz)

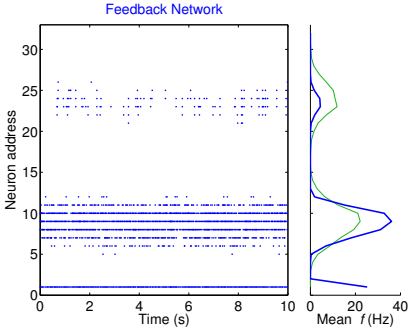
Feedback Network
Neuron address
Time (s)
Mean f (Hz)

Soft Winner-Take-All networks

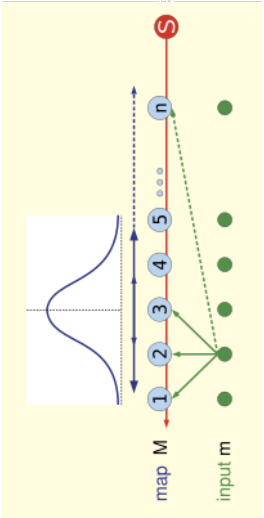

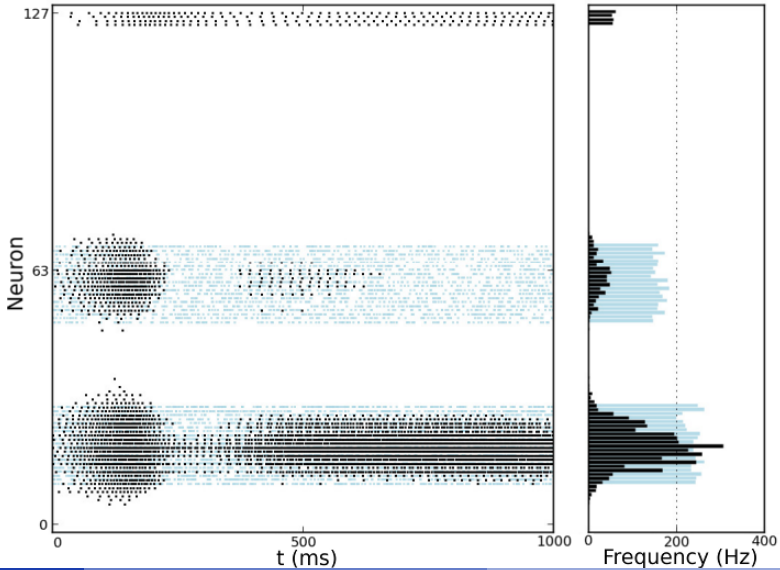

Soft Winner-Take-All drifting

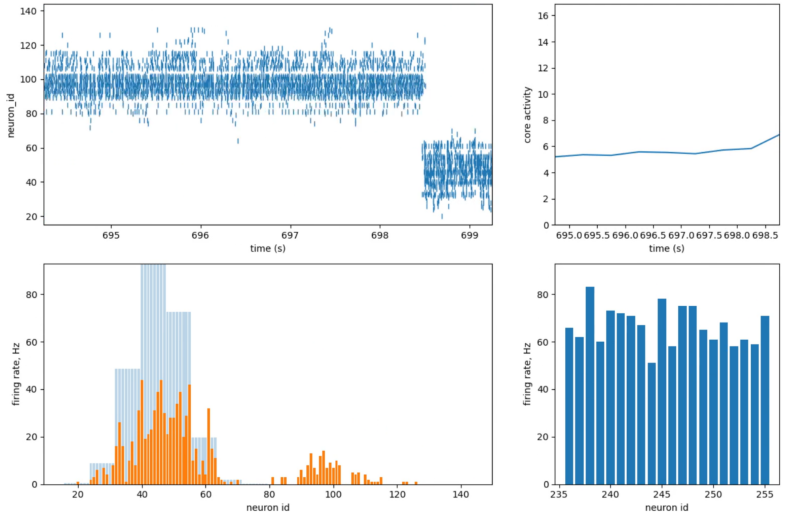

[Zendrikov et al., 2023]

Binary variables V0,V1: V0≠V1

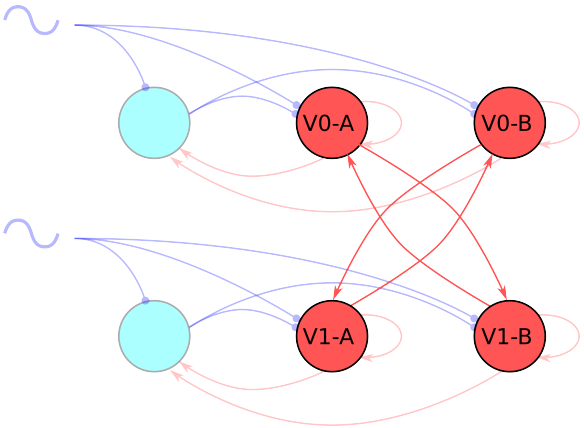

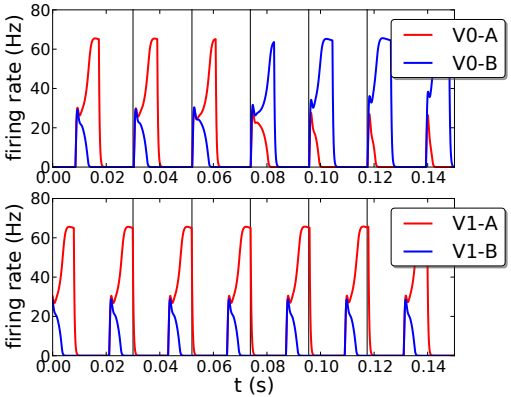

In absence of external input (evidence), the network settles to the lowest energy state (all constraints satisfied).

[Mostafa et al., 2015]

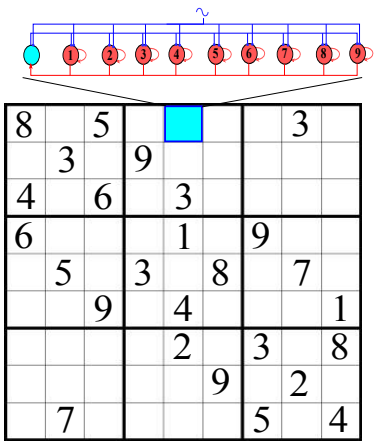

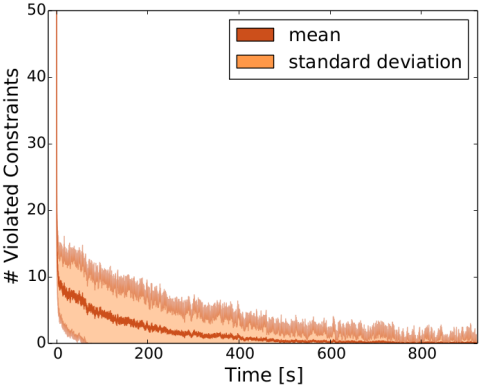

Can be applied to all Boolean satisfiability problems, such as graph coloring problem, SAT, etc.

[Mostafa et al., 2015]

Exploiting the device mismatch in the neuron's refractory period.

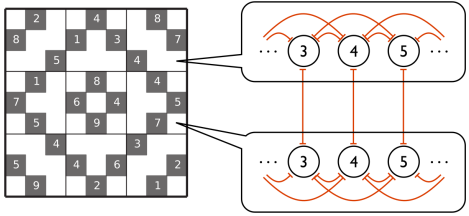

[Bias et al., 2016]

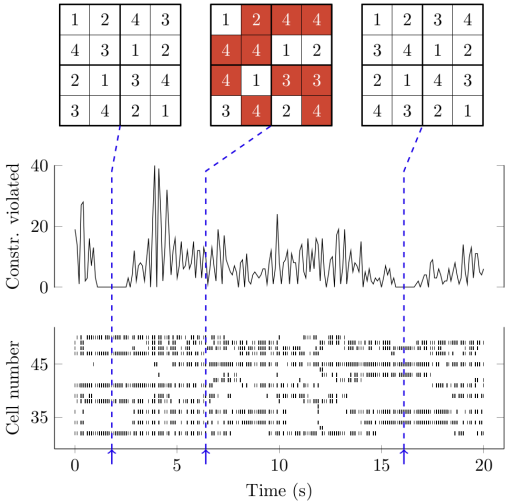

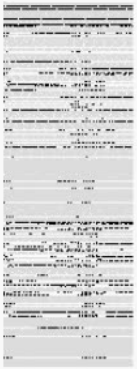

Coupling multiple WTAs to process variables

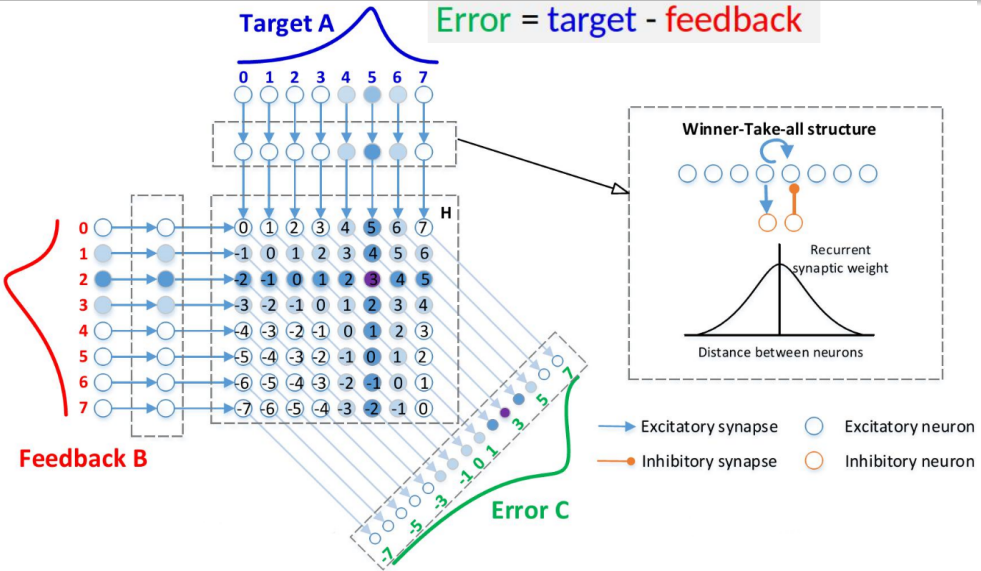

[J. Zhao et al., 2020]

Coupling multiple WTAs to process variables

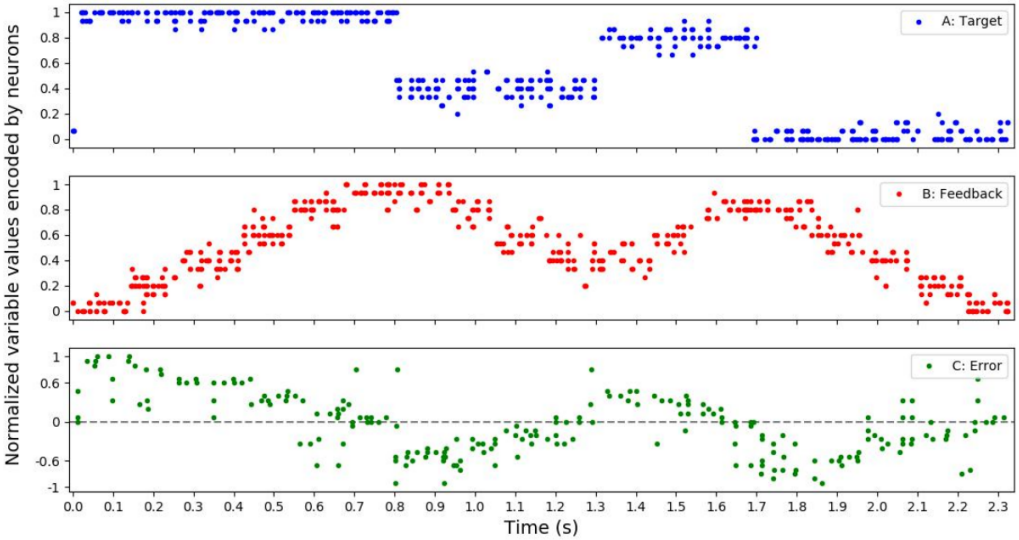

16 neurons

16 neurons

31 neurons

Finite State Machines vs Neural State Machines

A finite-state machine (FSM) is a mathematical model of computation used to design both computer programs and sequential logic circuits. It is conceived as an abstract machine that can be in one of a finite number of states. [Wikipedia]

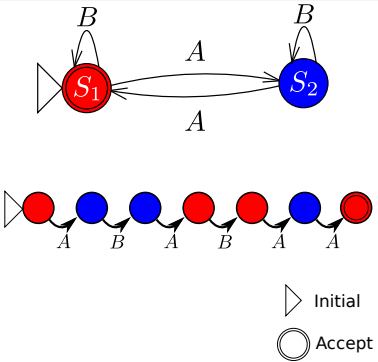

- Recognizes regular expression $B^*[AB^*A]^*$

[Minsky, 1967]

Single WTA

Inhibitory neurons

Excitatory neurons

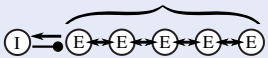

Global Inhibition

Nearest-N Excitation

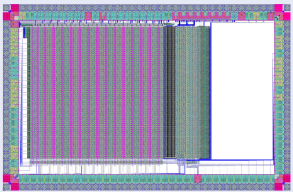

Coupled WTAs

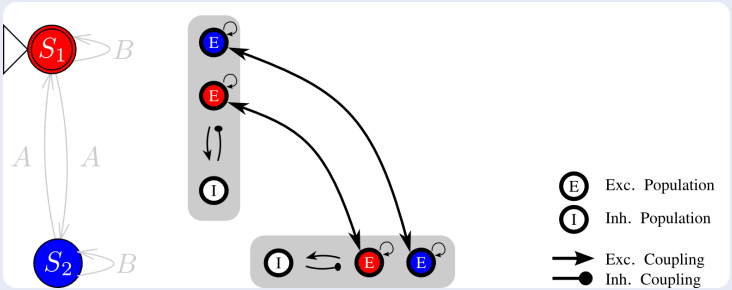

[R. Rutishauser & R.J. Douglas, 2009, R. Rutishauser et al., 2011, E. Neftci et al., 2013]

Single WTA

Inhibitory neurons

Excitatory neurons

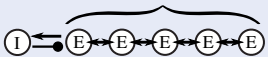

Global Inhibition

Nearest-N Excitation

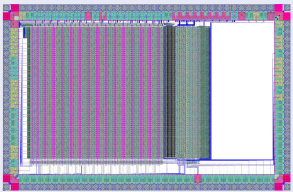

Coupled WTAs

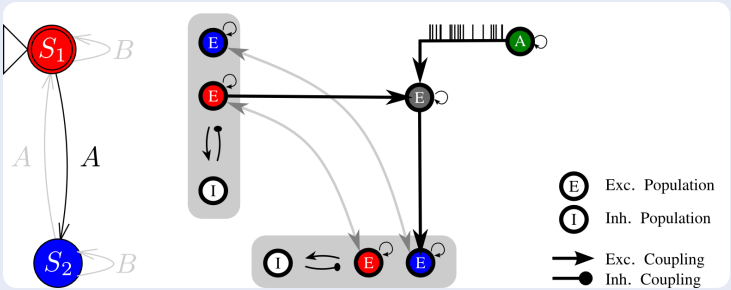

[R. Rutishauser & R.J. Douglas, 2009, R. Rutishauser et al., 2011, E. Neftci et al., 2013]

Single WTA

Inhibitory neurons

Excitatory neurons

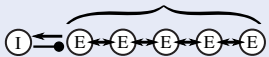

Global Inhibition

Nearest-N Excitation

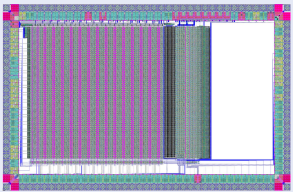

Coupled WTAs

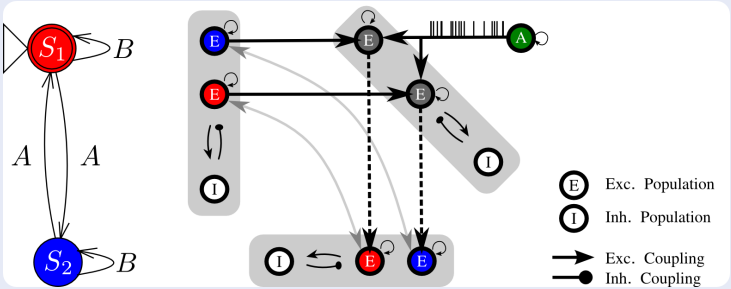

[R. Rutishauser & R.J. Douglas, 2009, R. Rutishauser et al., 2011, E. Neftci et al., 2013]

Neural State Machines: experimental results

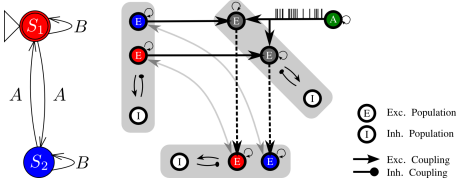

Vision sensor

Linked image not found

Multi-Neuron Chips

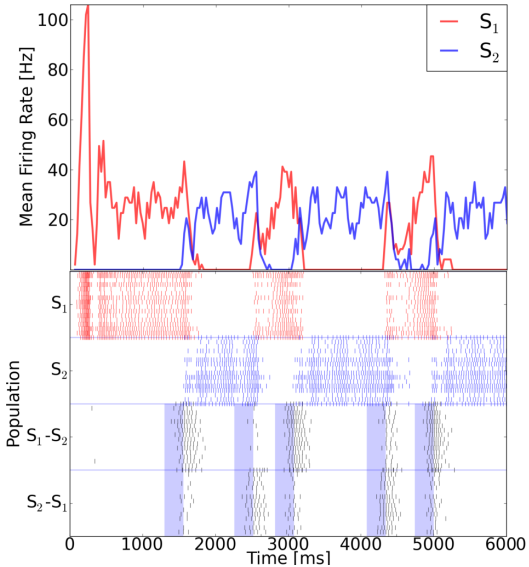

Synthesizing neuromorphic cognitive system

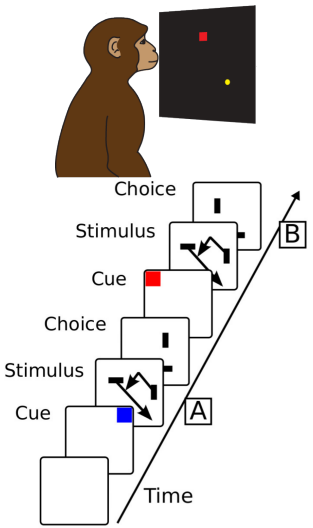

(a) State Machine

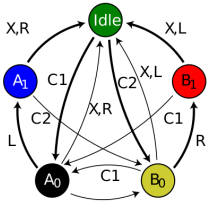

(b) Network Architecture

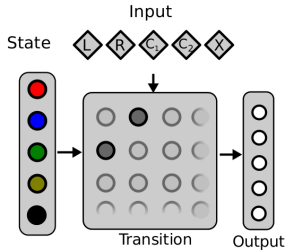

(c) Example Run

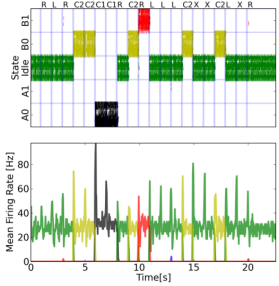

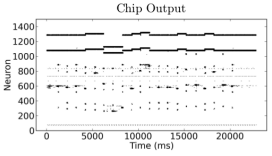

[E. Neftci et al., 2013]

Synthesizing neuromorphic cognitive systems

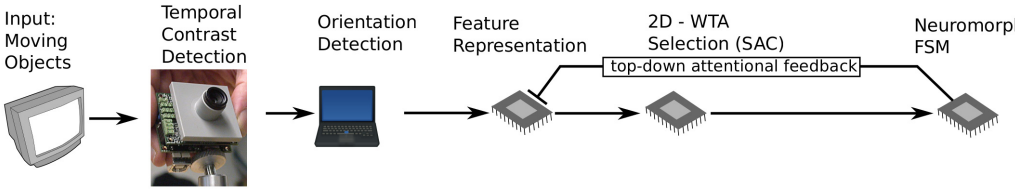

Neuromorphic processors for sensory processing

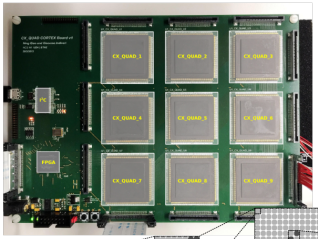

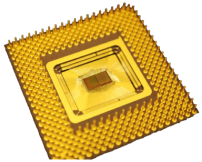

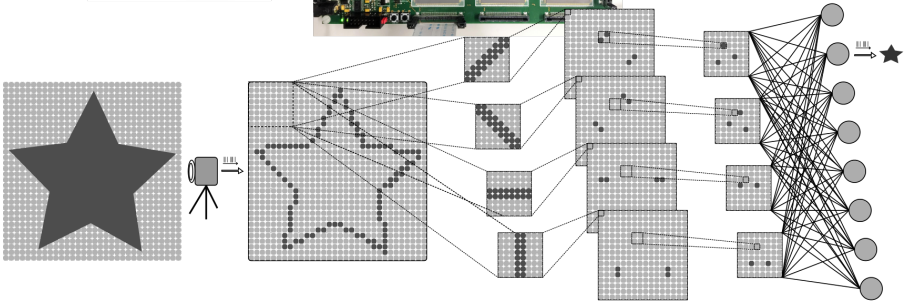

[Indiveri et al., IEDM 2015]

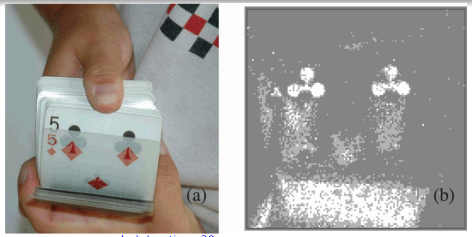
(a)
(b)
symbol duration=30ms
input layer
feature layer
pooling layer
output layer
td1=8ms
td2=8ms
td3=10ms

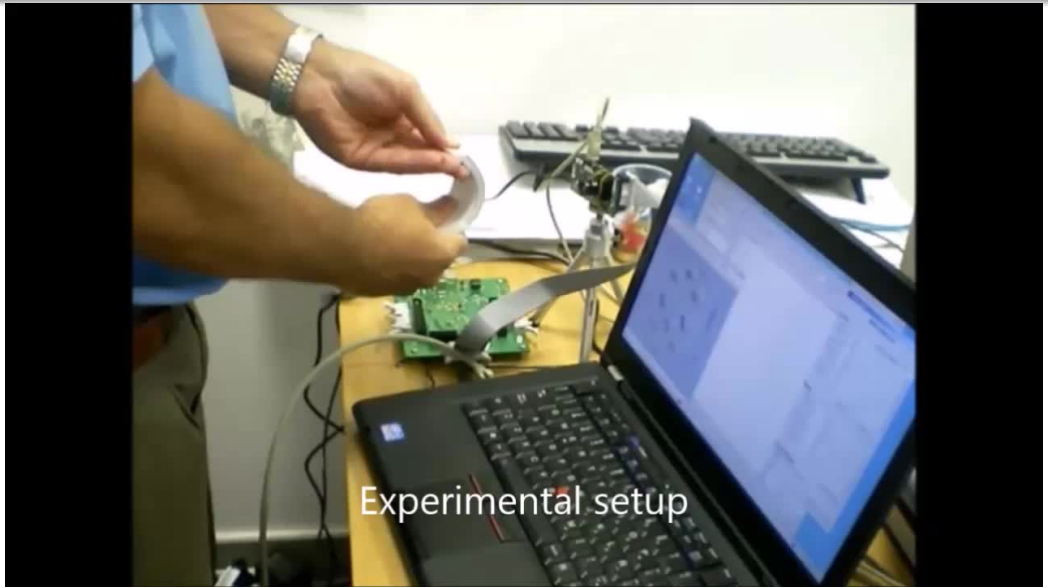
Experimental setup

On-chip learning implementation