University of Zurich[UZH]

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Institute of Neuroinformatics (INI)

Building blocks for learning and inference in neuromorphic systems

Dr. Michael Pfeiffer

pfeiffer@ini.phys.ethz.ch

Institute of Neuroinformatics

University of Zurich and ETH Zurich, Switzerland

NICE Workshop, Albuquerque NM, February 26th 2014

Presenting work and ideas by (among others) …


Giacomo Indiveri


Tobi Delbruck


Shih-Chii Liu


Rodney Douglas


Emre Neftci
(now UCSD)


Kevan Martin


Peter O'Connor
(now Braincorp)


Danny Neil


Ueli Rutishauser
(Caltech)


Matthew Cook


Elisabetta Chicca
(now Univ. Bielefeld)


Jonathan Binas

Why investigate spike-based computation?

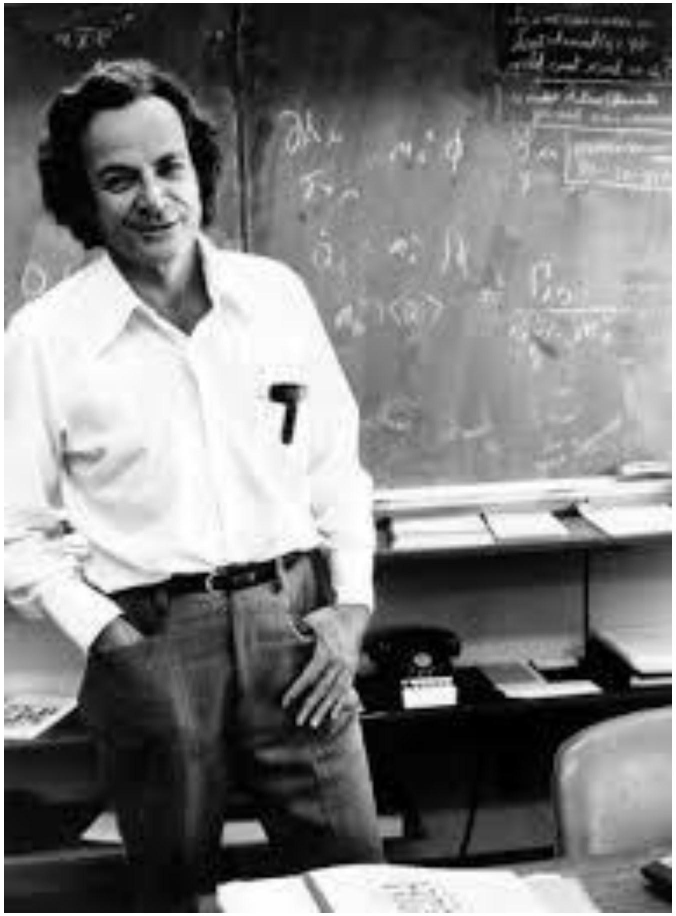

Richard Feynman (1918-1988)

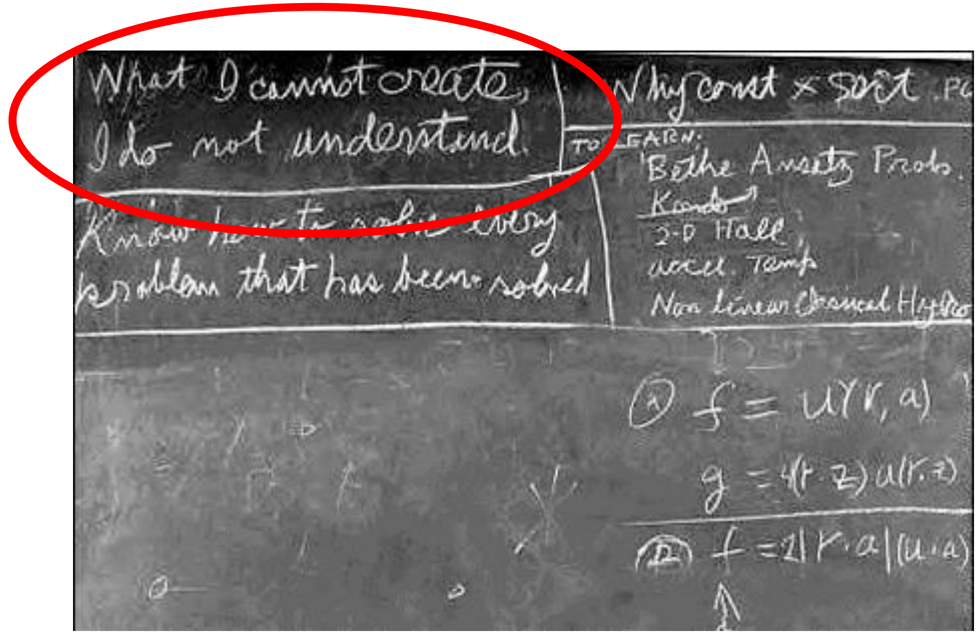

And this is how we create cognitive behavior in neuromorphic systems:
a) ...

Building Blocks…

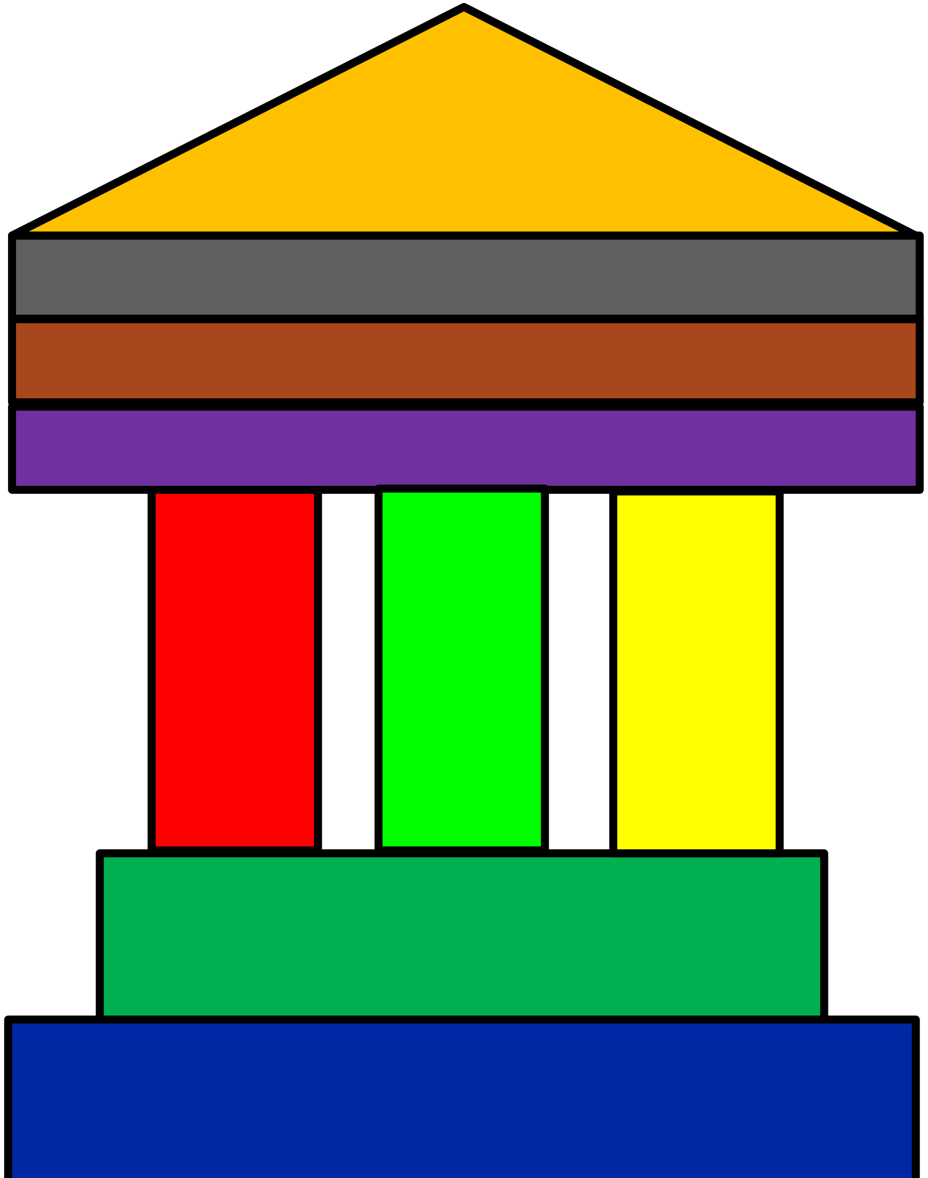

Building Blocks…

Relevant for real-world problems

Scalable and adaptable

Mapping onto neuromorphic hardware

Brain-inspired

Efficient input representations

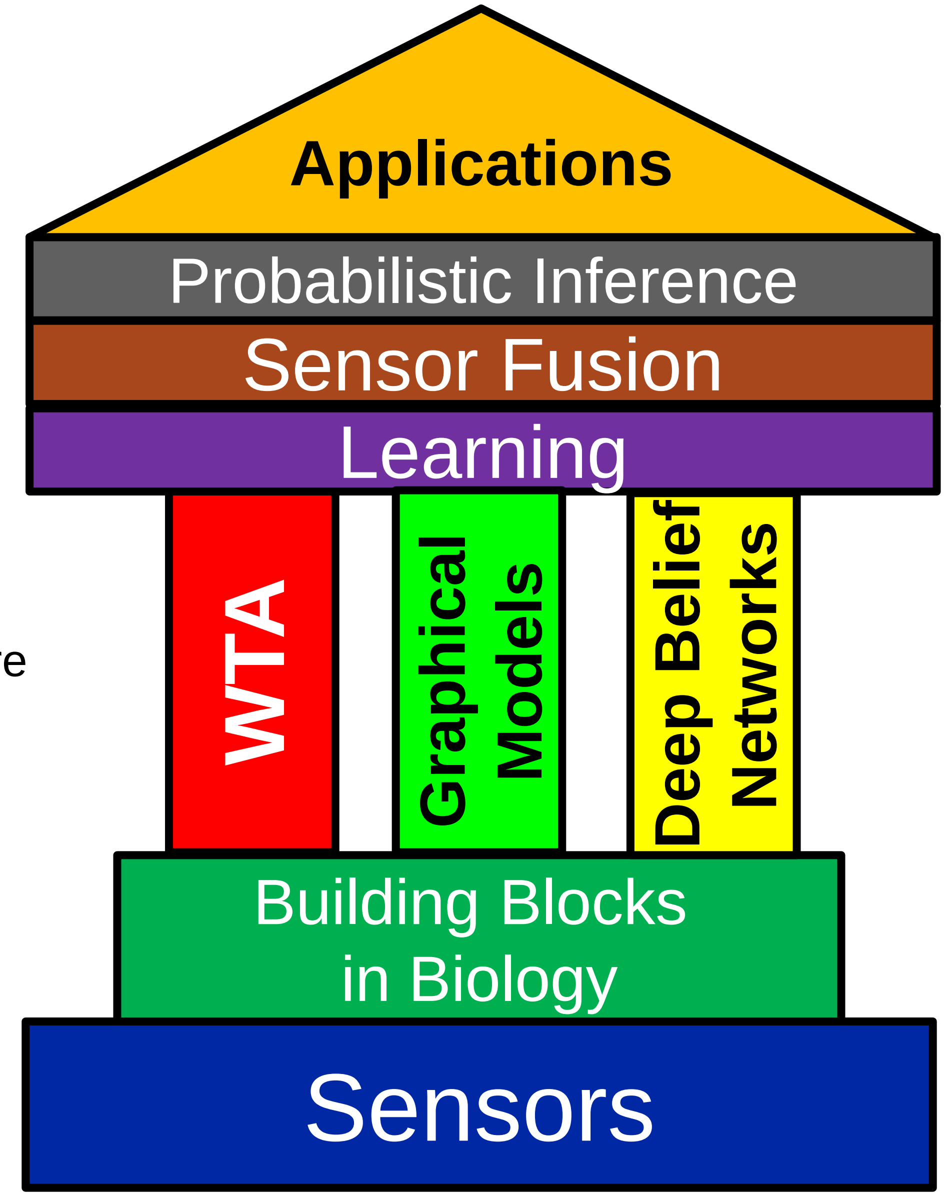

Building Blocks for Neuromorphic Engineering

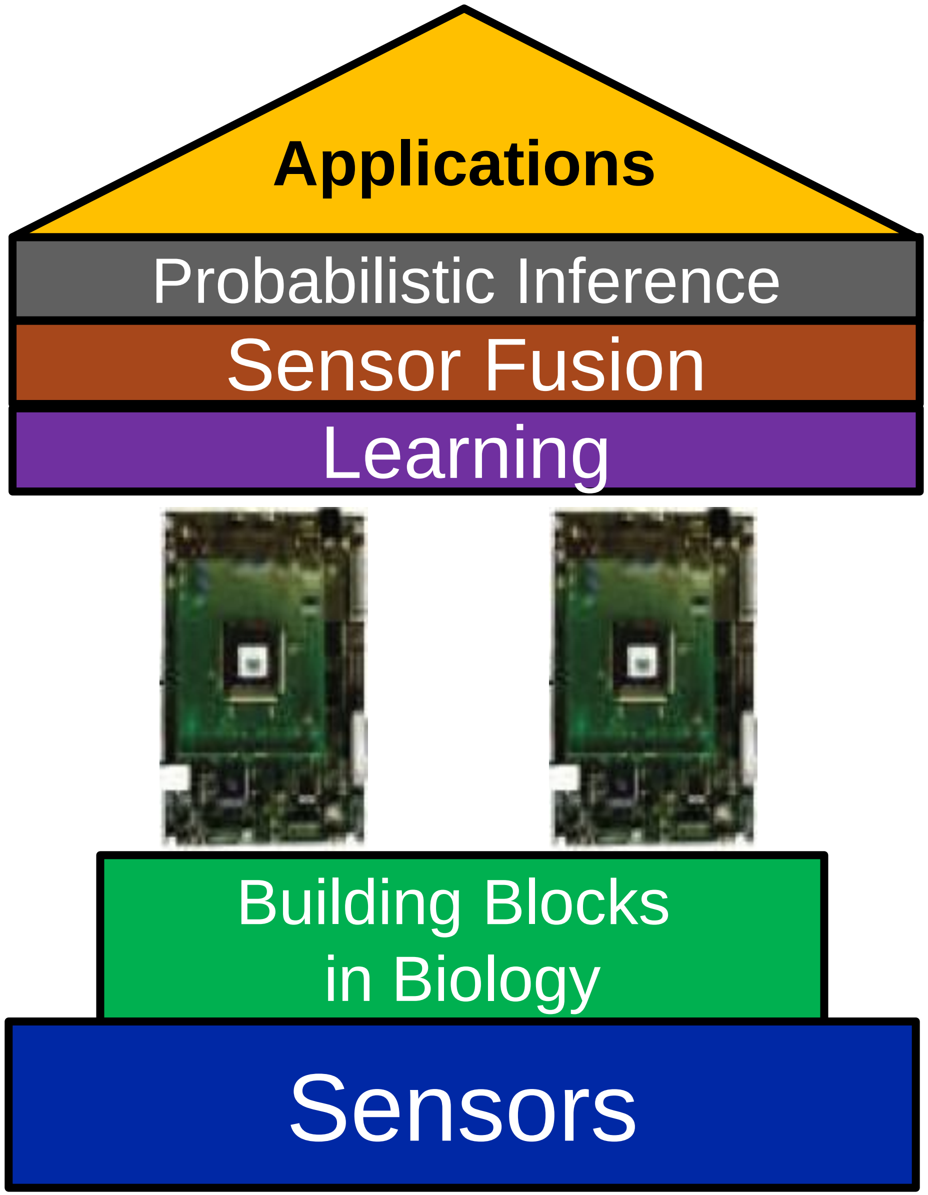

Mapping onto neuromorphic hardware

Building Blocks…

Sensors

Conventional vs. Event-based Sensors

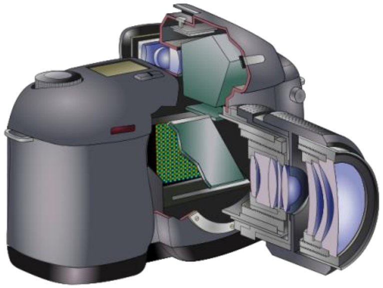

Conventional camera:

- Shoots still images or sequences of frames at constant frame rate
- Same high resolution over the entire image
- Every pixel behaves similarly
- Massive amounts of data
- Mostly redundant data for processing sequences / videos

DVS – The silicon retina

Tobi Delbruck: siliconretina.ini.uzh.ch

DVS128 Sensor

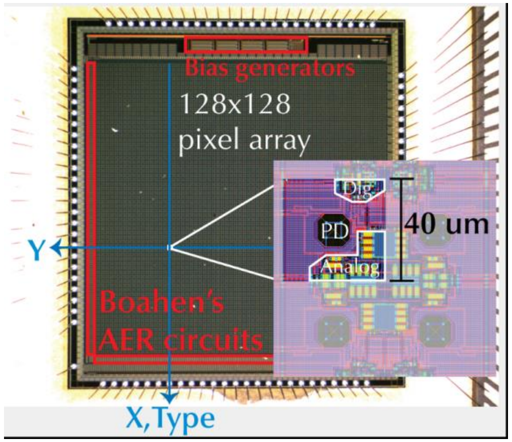

128x128 Pixel DVS Chip

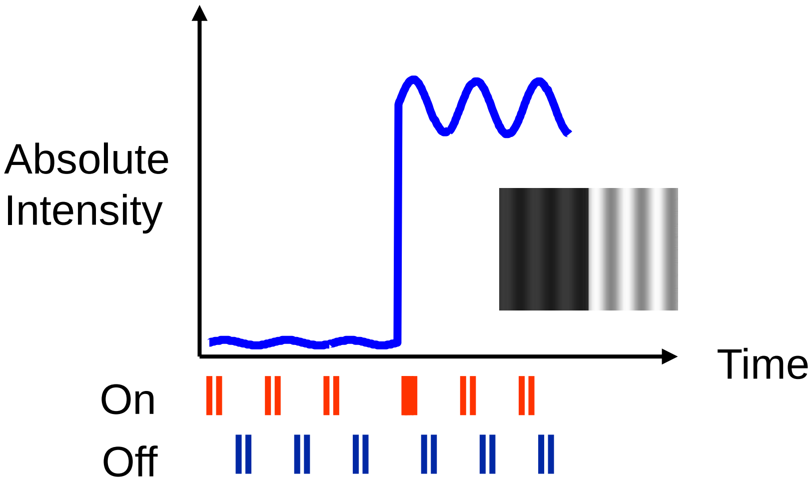

Function of a single pixel

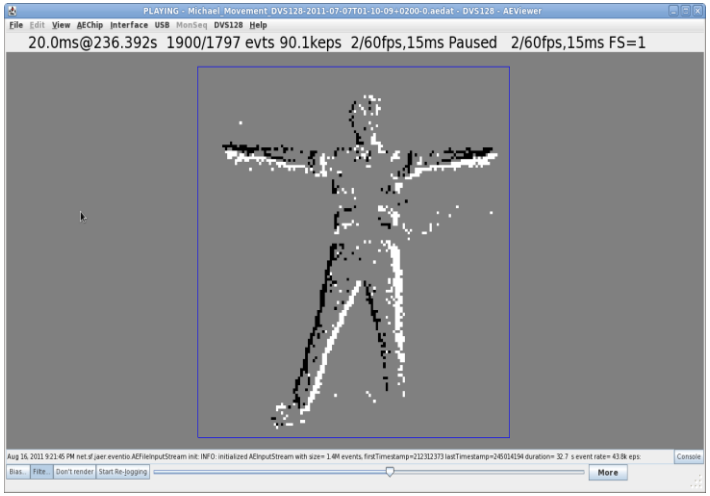

Activity of all pixels

Advantages of DVS

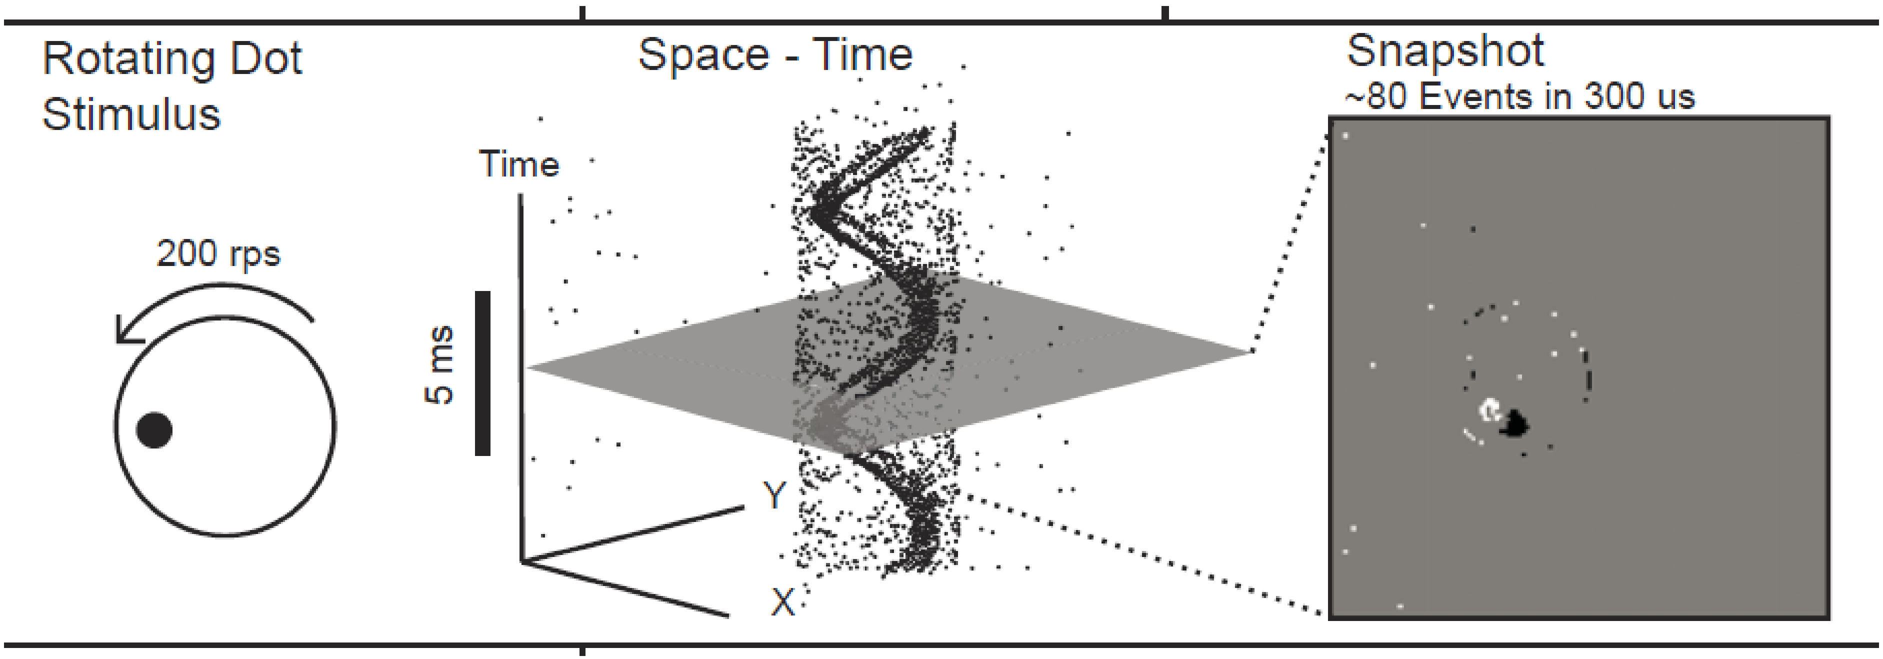

+ High Speed and precise time resolution

+ Low data rate

+ Ideally suited for real-time tracking, robotics, …

- Low spatial resolution

- No intensity measurement

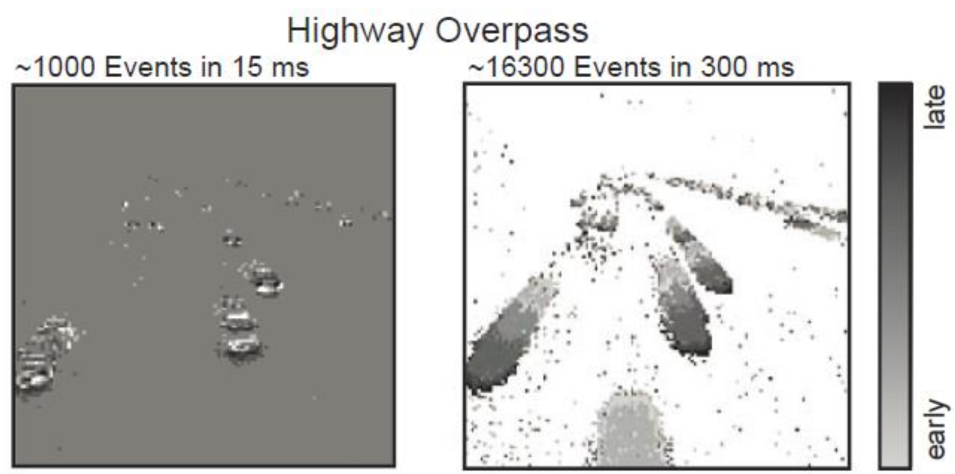

Lichtsteiner et al. 2006

Silicon Cochlea: Shih-Chii Liu

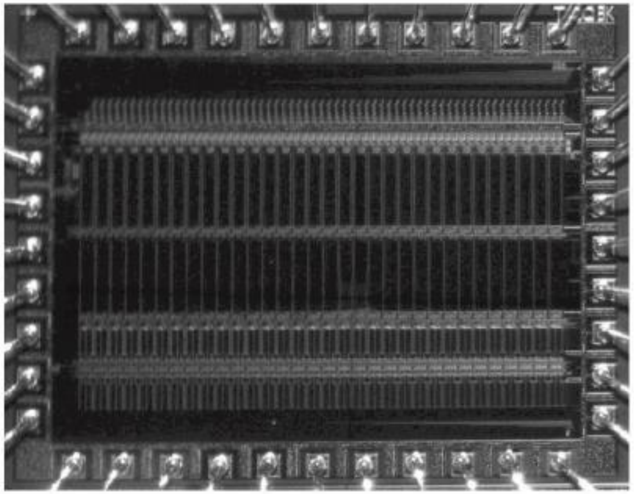

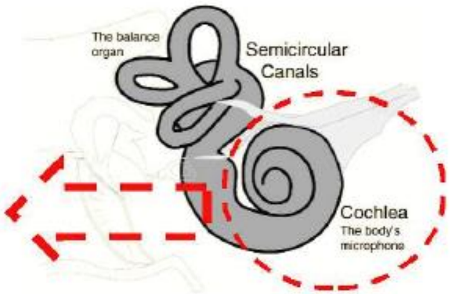

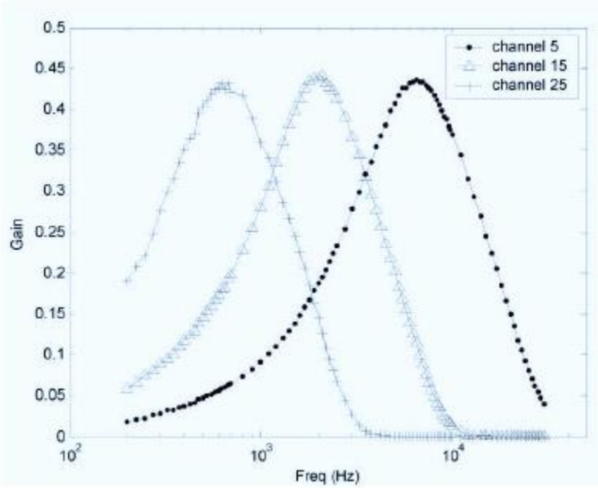

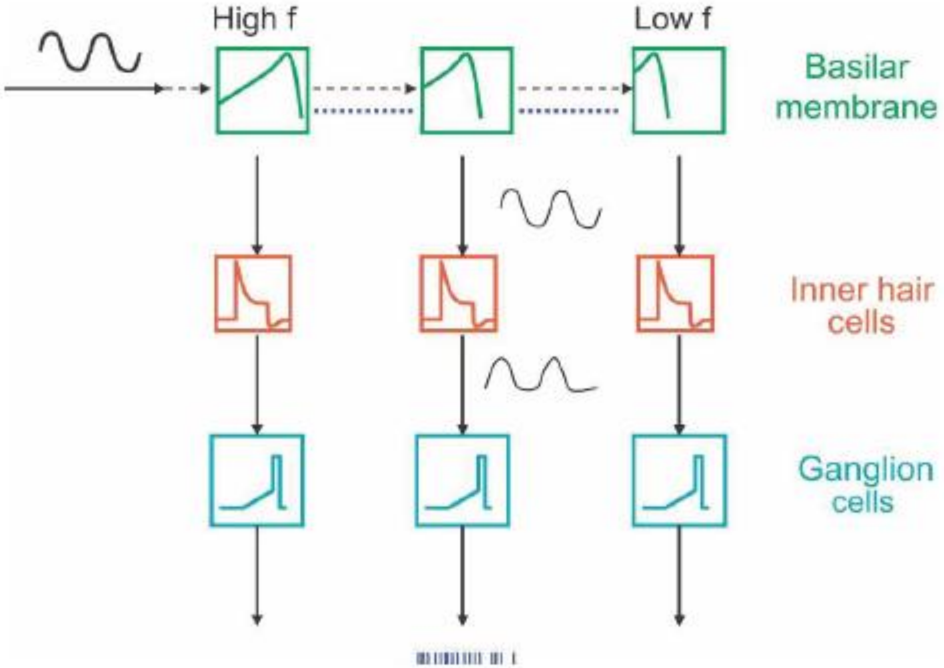

Challenges of Event-based Sensory Processing

1. Standard computer vision or machine learning approaches do not work on **spatio-temporal spike patterns**
2. Working on **asynchronous** sequences is different than on still images
3. Problem of **sensor fusion**
4. Challenge and opportunity of **real-time scenarios**

Our goals:

1. Relate **event-based learning and computation** to established machine learning and inference mechanisms
2. Make event-based algorithms suitable for neuromorphic hardware to run in **real-time and with low energy consumption**
3. Understanding **neural computation** in biology better by creating functional bio-inspired silicon solutions

Application: Gesture Recognition

- Remote-free control of devices with arm and hand gestures
- Collaboration with Jun-Haeng Lee, Hyunsurk Ryu, et al. (SAIT)

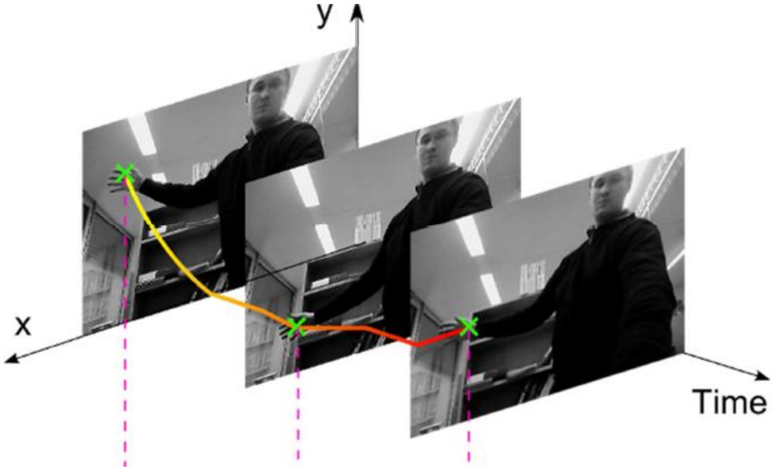

e.g. Kinect

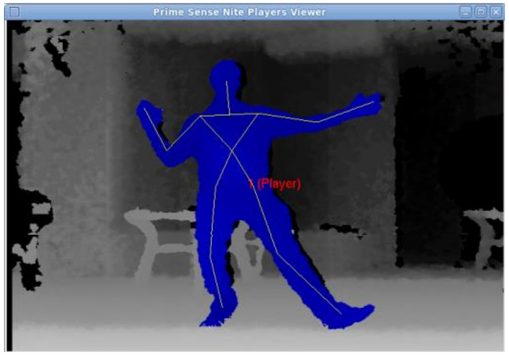

Stereo-DVS-based Gesture Recognition

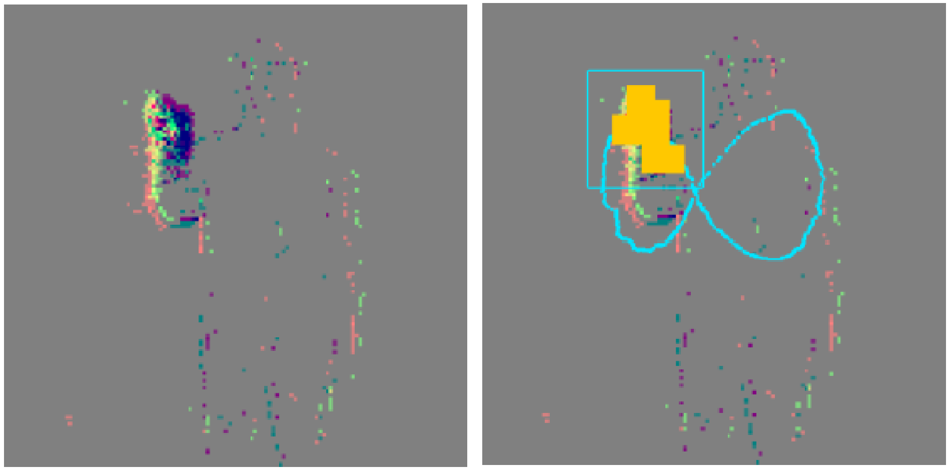

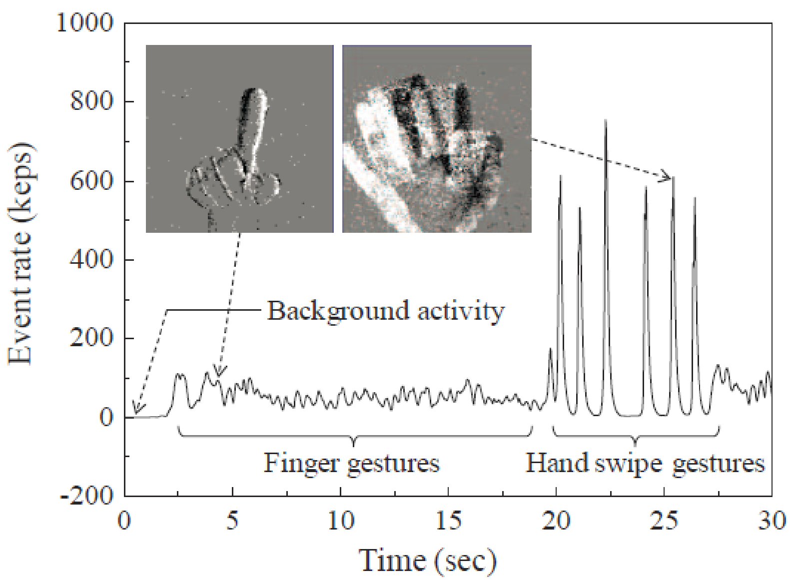

Motion detection is already performed at sensor level

Stereo-DVS setup reduces noise

Simple **clustering** of events for tracking

Data reduction for post-processing

Use data-dependent input rate for **segmentation** of gestures

HMM-based classification of trajectories

- ~97% recognition
- Neglible latency (~20 ms)

[Lee et al. ISCAS 2012, ICIP 2012, IEEE TNNLS 2014]

Building Blocks…

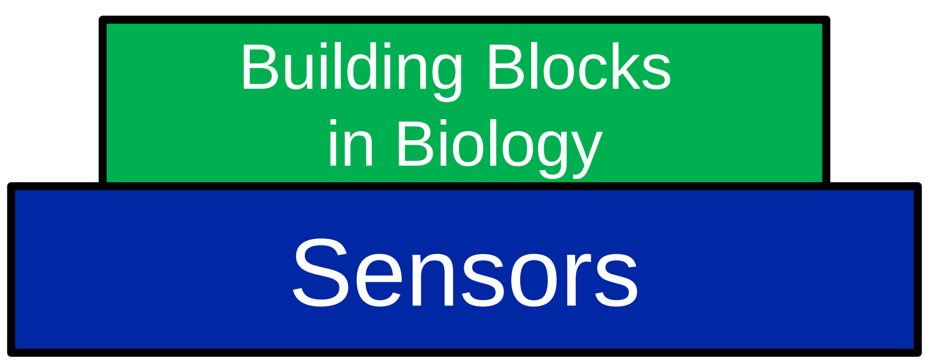

Canonical microcircuits in the cortex

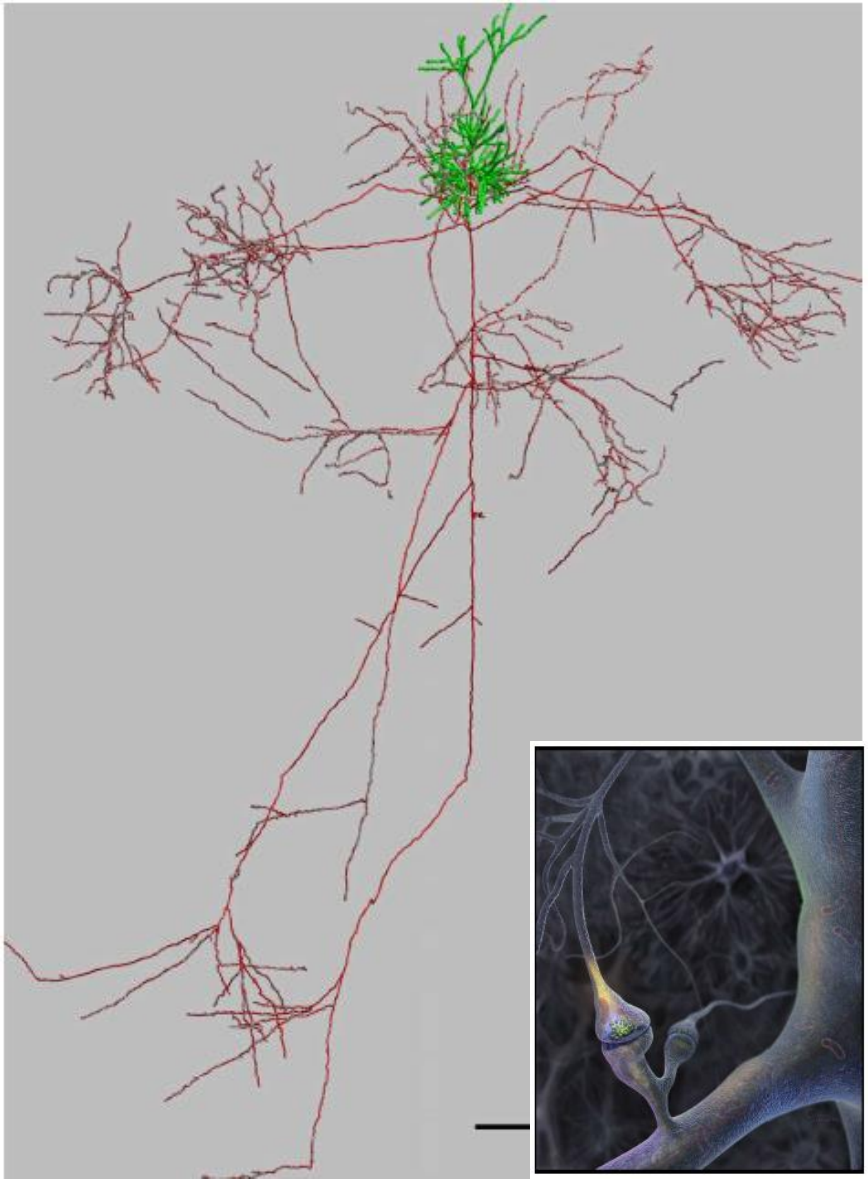

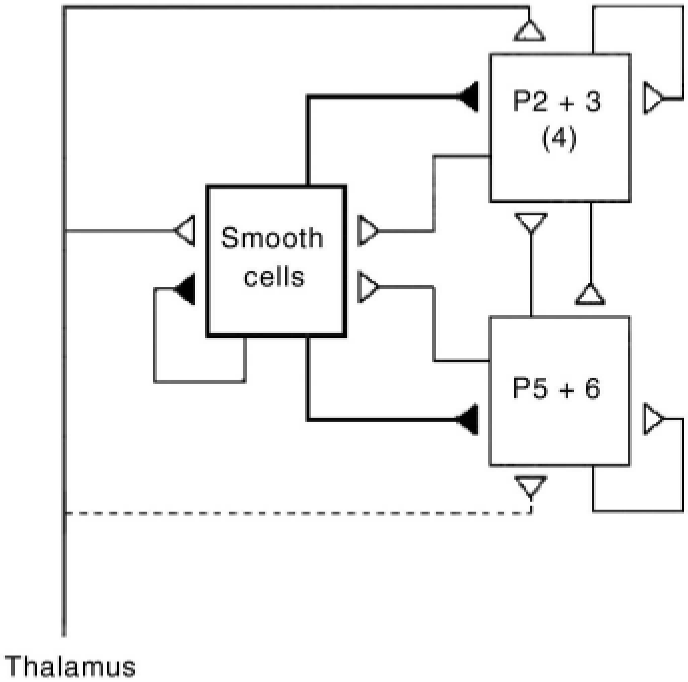

[Douglas and Martin, 1991] suggested **Winner-take-all architecture** for canonical circuits of cat V1 (based on anatomy and electrophysiology)

Building Blocks…

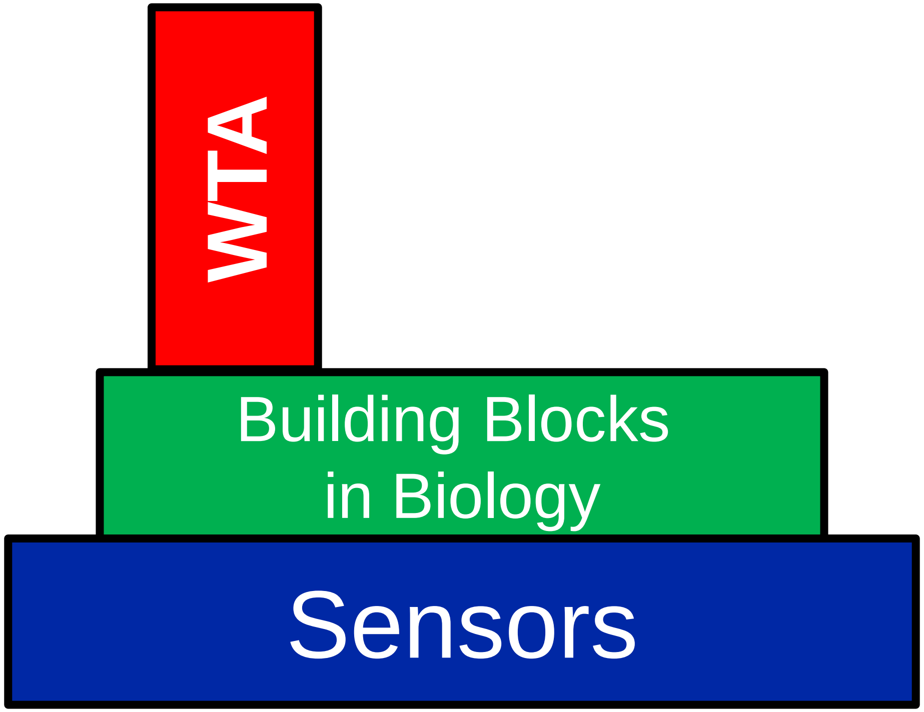

sWTA networks as computational modules

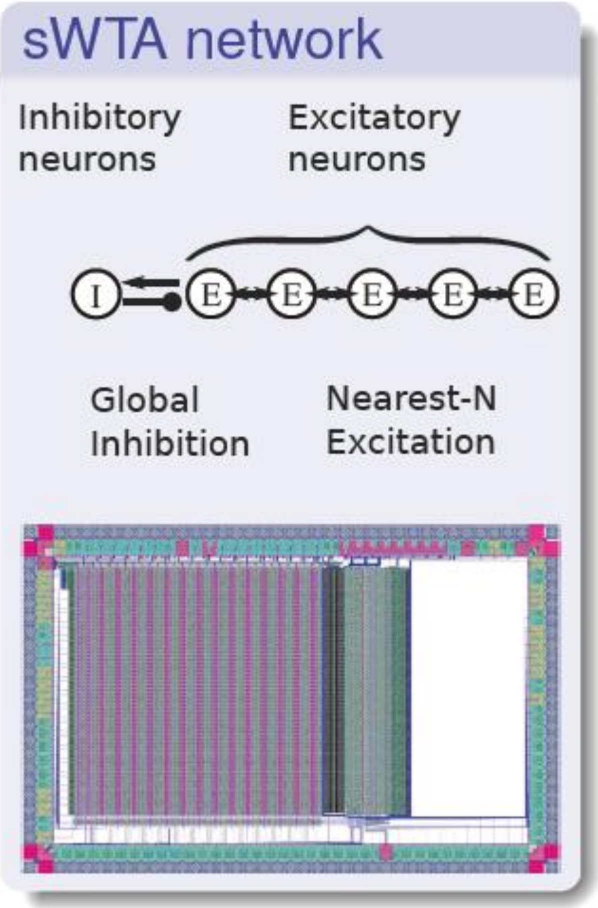

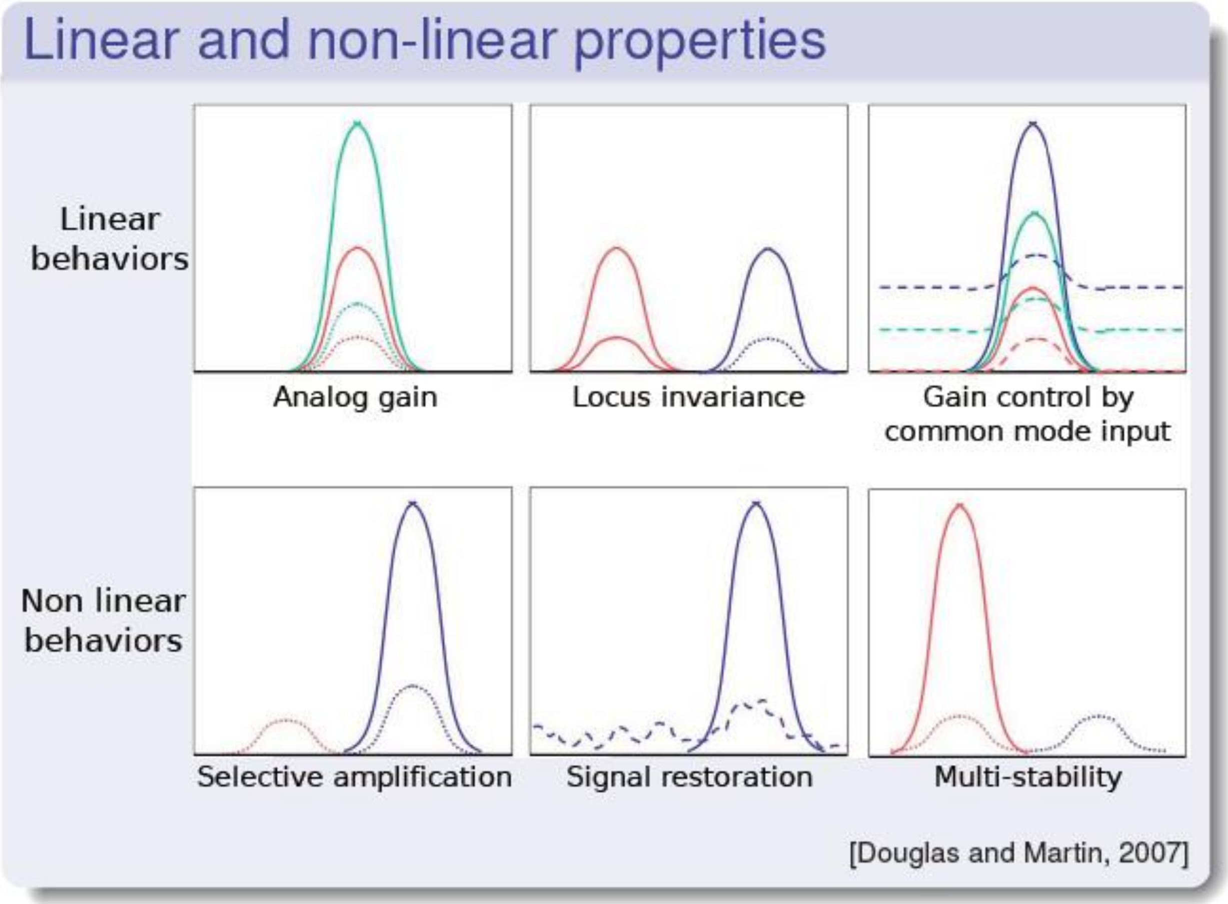

sWTA networks in neuromorphic VLSI

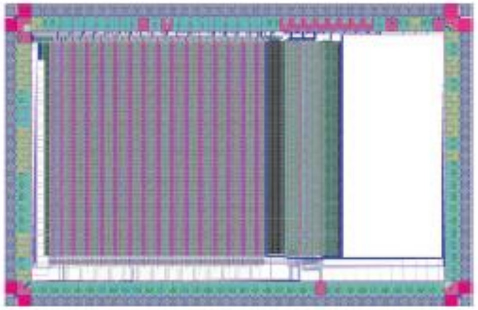

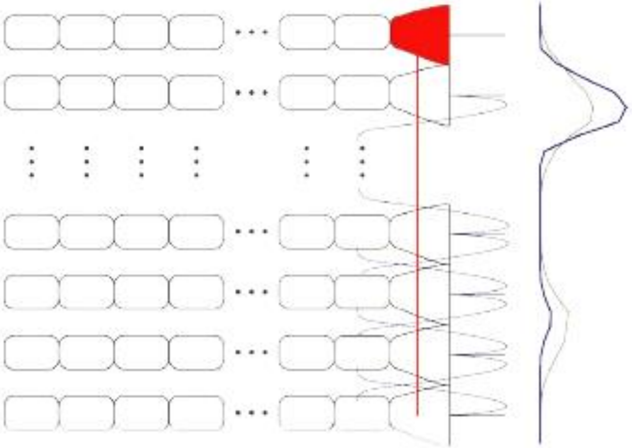

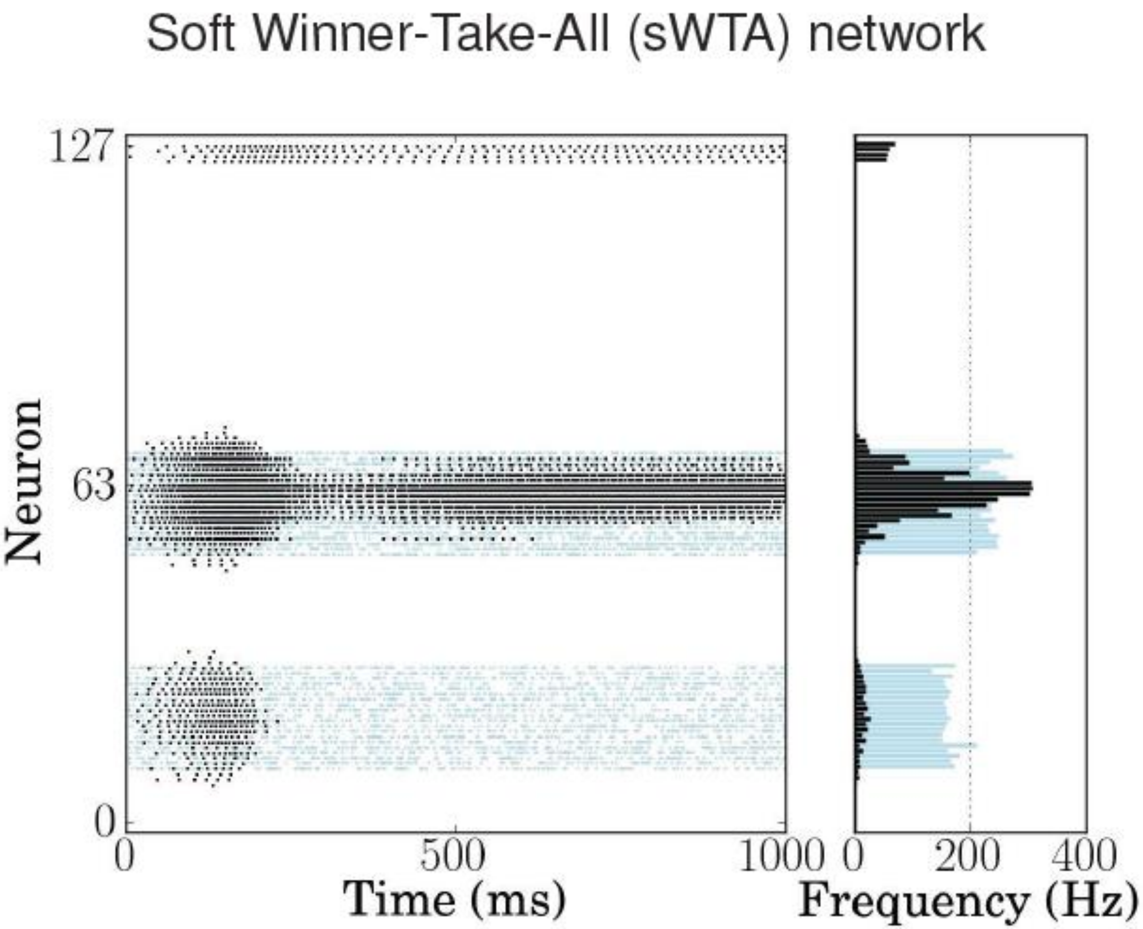

Soft-state-machines

Single WTA

Inhibitory neurons

Excitatory neurons

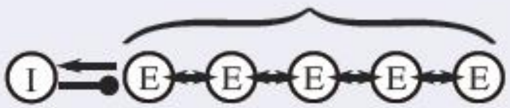

Global Inhibition

Nearest-N Excitation

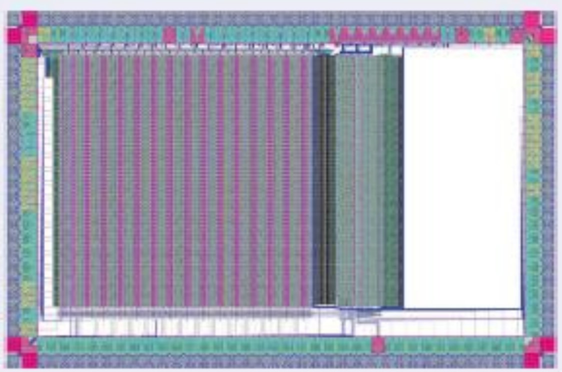

Coupled WTAs

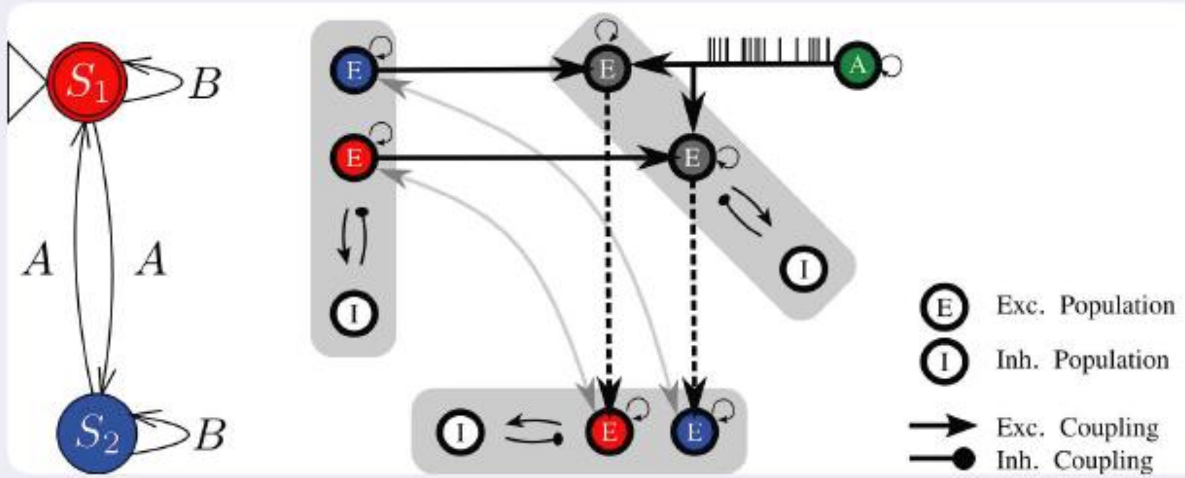

[Rutishauser Douglas, 2009; Rutishauser et al., 2010; Neftci et al., 2010]

Synthesizing Cognition in VLSI

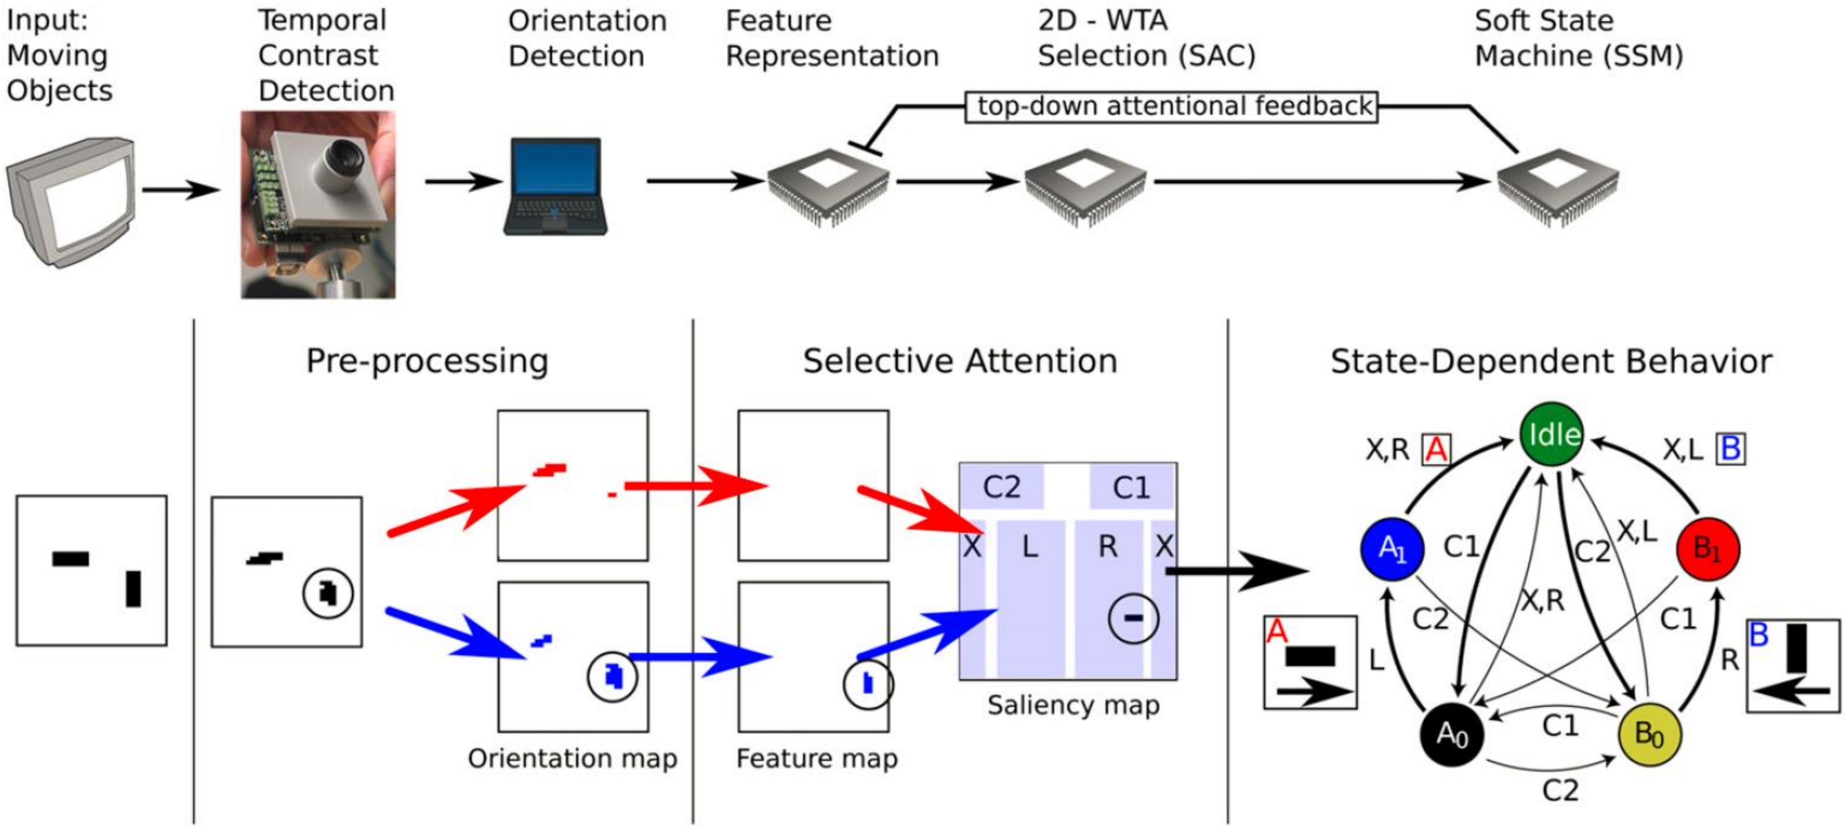

Real-time context-dependent visual processing on multi-chip neuromorphic system, using neuromorphic vision sensors (Neftci et al. PNAS 2013)

Building Blocks…

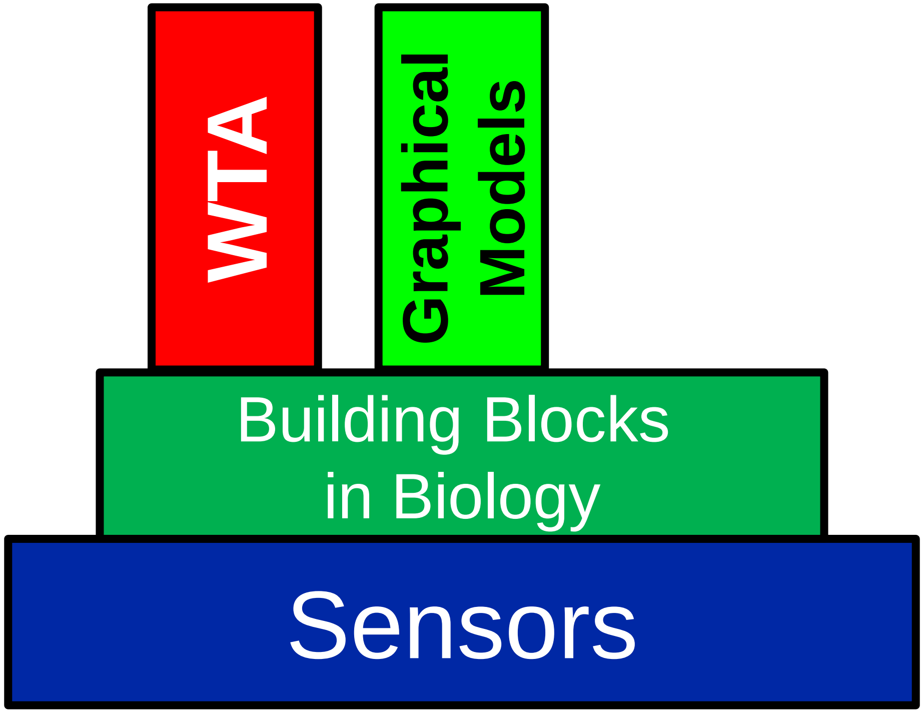

Spike-based learning of Bayesian models

(with Bernhard Nessler, Wolfgang Maass; TU Graz)

Graphical model (Bayesian network)

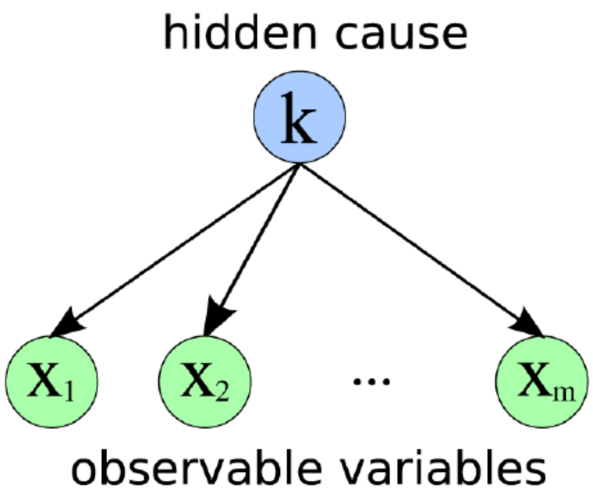

Winner-take-all architecture

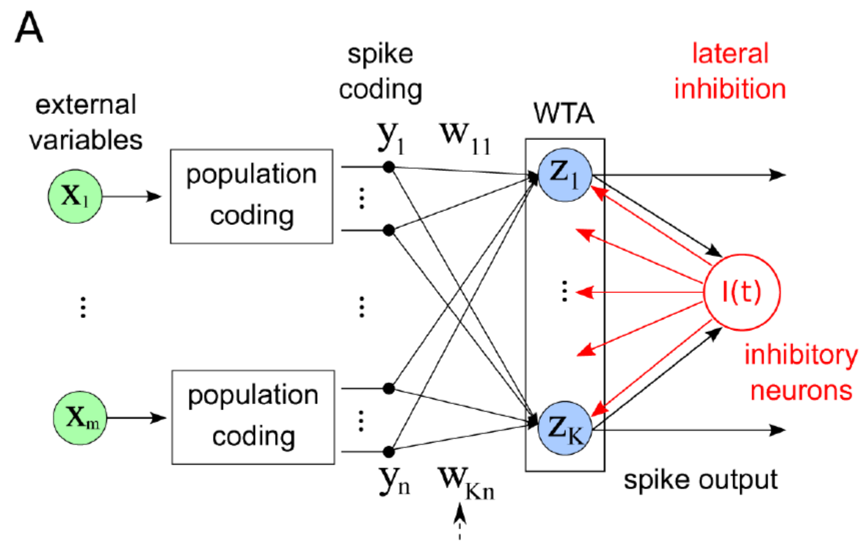

$$p(z_k \text{ fires at time } t|\mathbf{y}) = \frac{e^{u_k(t)}}{\sum_{l=1}^{K} e^{u_l(t)}}$$

$$u_k(t) = \sum_{i=1}^{n} w_{ki}\tilde{y}_i(t) + w_{k0}$$

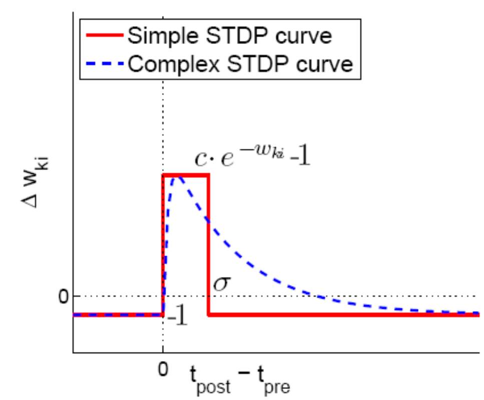

Weight-dependent STDP learning

Nessler et al. 2009: STDP enables spiking neurons to detect hidden causes of their inputs. NIPS 2009

Nessler et al. 2013: Bayesian Computation emerges in generic cortical microcircuits through STDP. PLoS CB

Spike-based EM learning

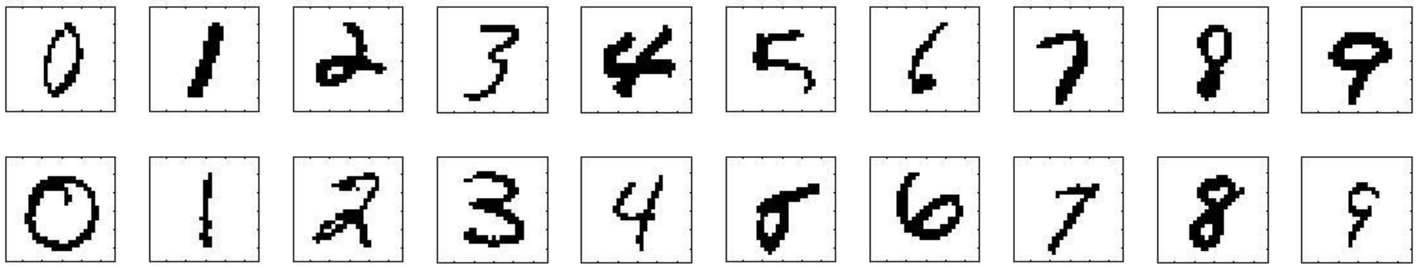

20 random samples from the 70 000 samples in the MNIST dataset.

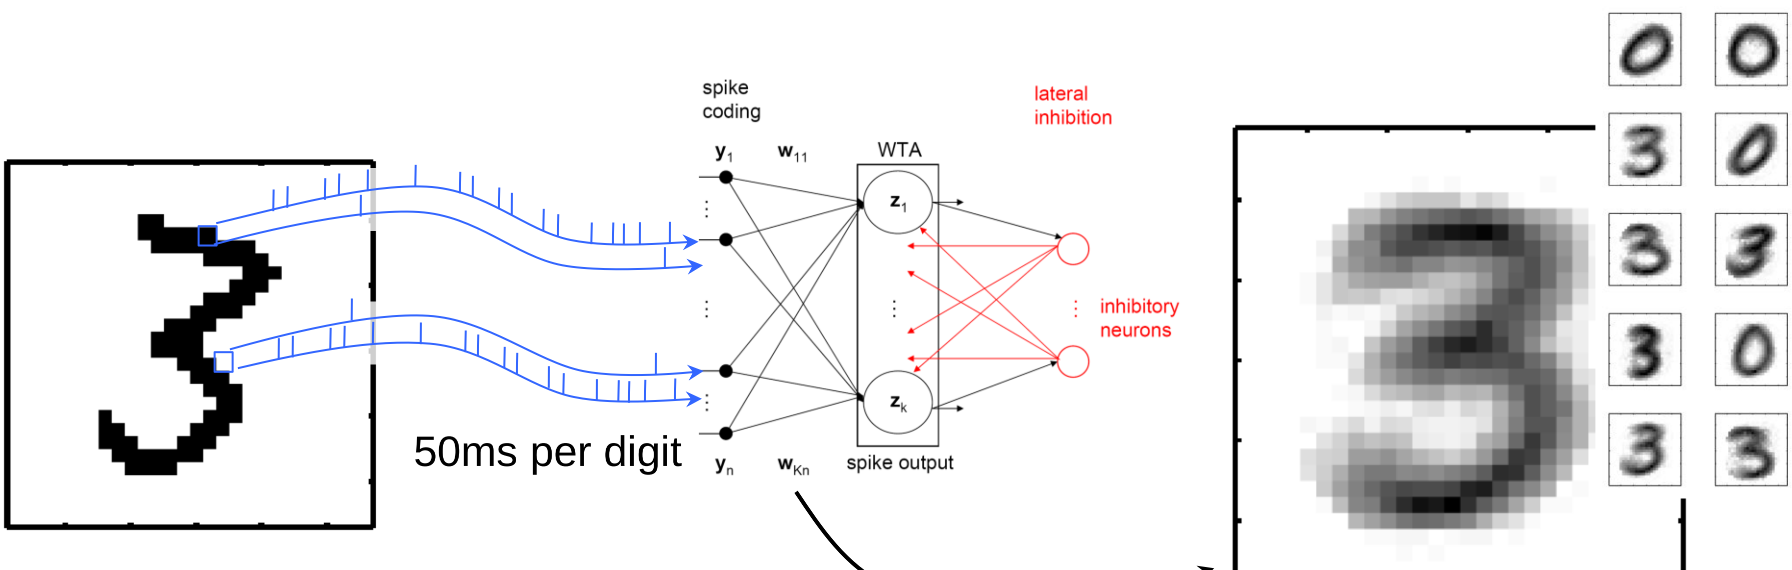

weights of one neuron

Learning of generative models with STDP

We can rigorously prove that this STDP curve in this circuit approximates the **Expectation-Maximization (EM)** algorithm

- Most general and most widely used tool for unsupervised machine learning (clustering, HMM learning, …)
- **Spike-based Expectation Maximization (SEM)**

Weights converge to conditional log-probabilities:

log p(presyn. neuron has fired just before time t / postsyn. neuron fires at time t)

$$w_{ki}^* := \log p_{\mathbf{w}}^*(y_i = 1 | z_k = 1) \text{ and } w_{k0}^* := \log p_{\mathbf{w}}^*(z_k = 1)$$

A spike-based view of Bayesian computation

- Synapses learn generative models of their inputs
- Output spike is **probabilistic sample from posterior distribution**
- **Building block** for learning and inference

Nessler et al. 2013: Bayesian Computation emerges in generic cortical microcircuits through STDP. PLoS CB

Learning of long spatio-temporal patterns

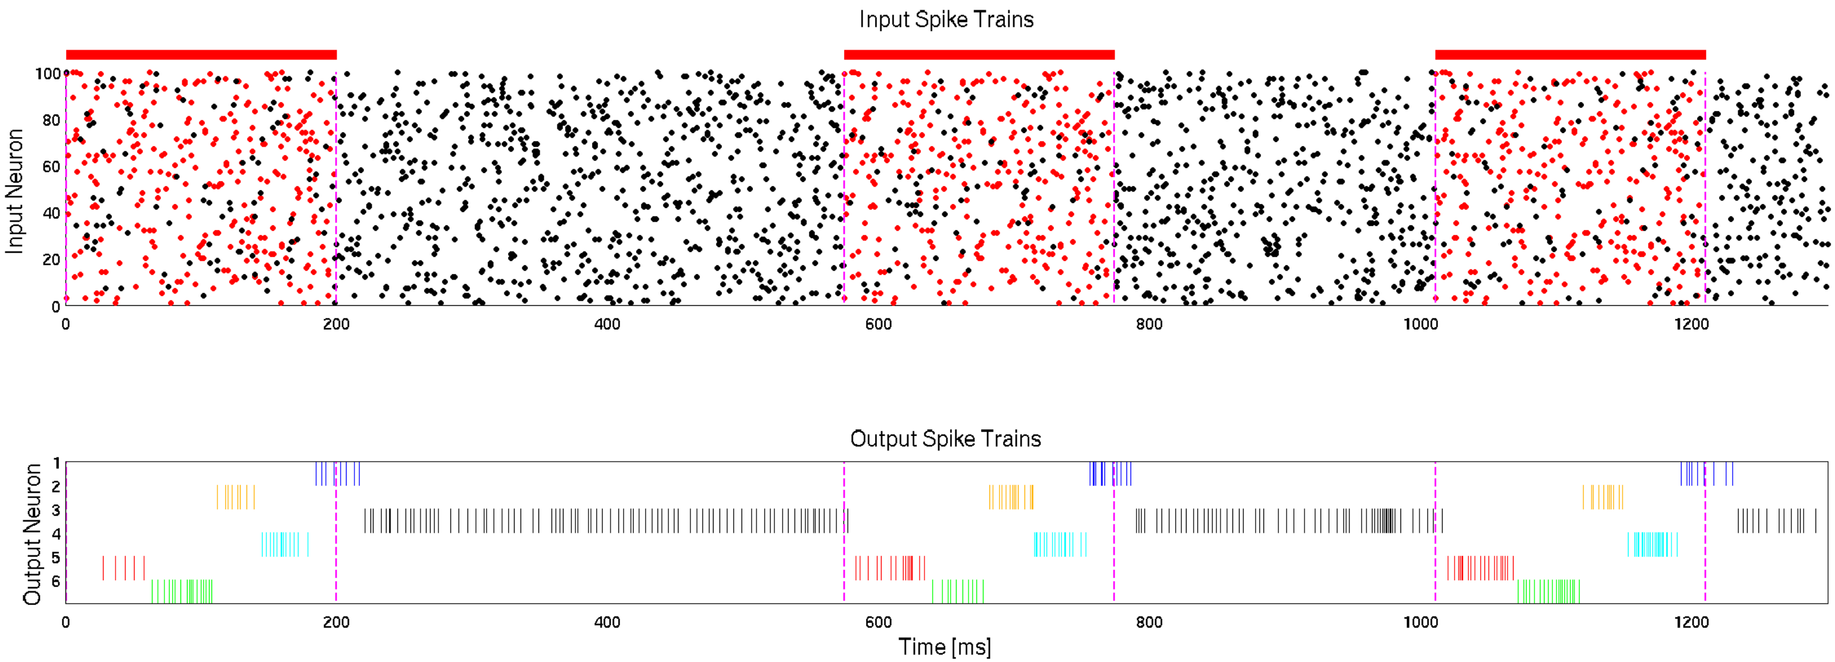

Output neurons learn to fire in characteristic sequence

- A state-machine or HMM-like approach can learn to recognize such sequences [Corneil et al., Cosyne 2014]

Building Blocks…

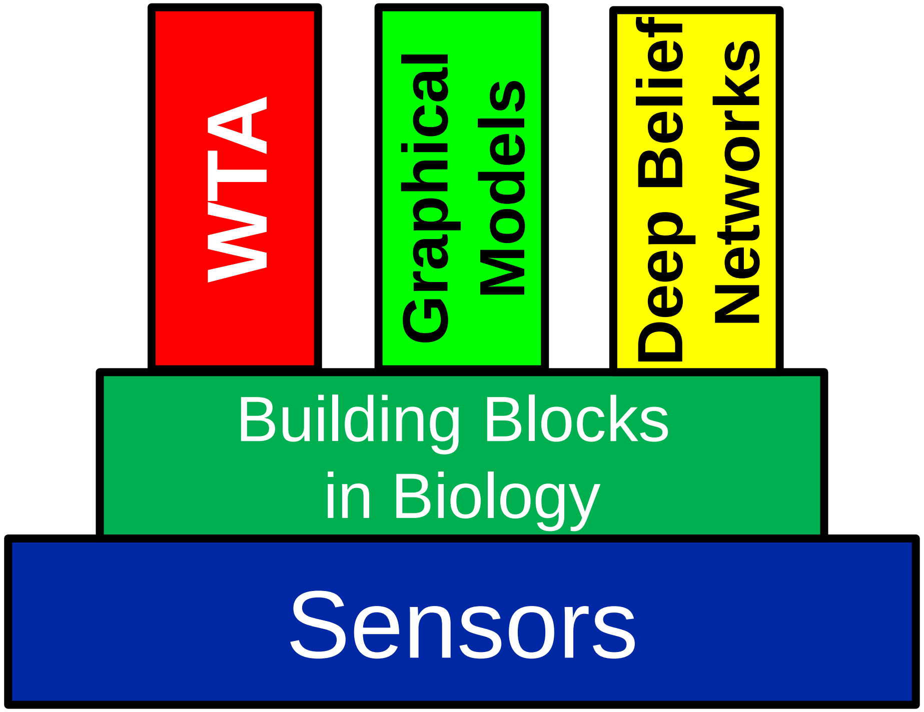

Building Blocks…

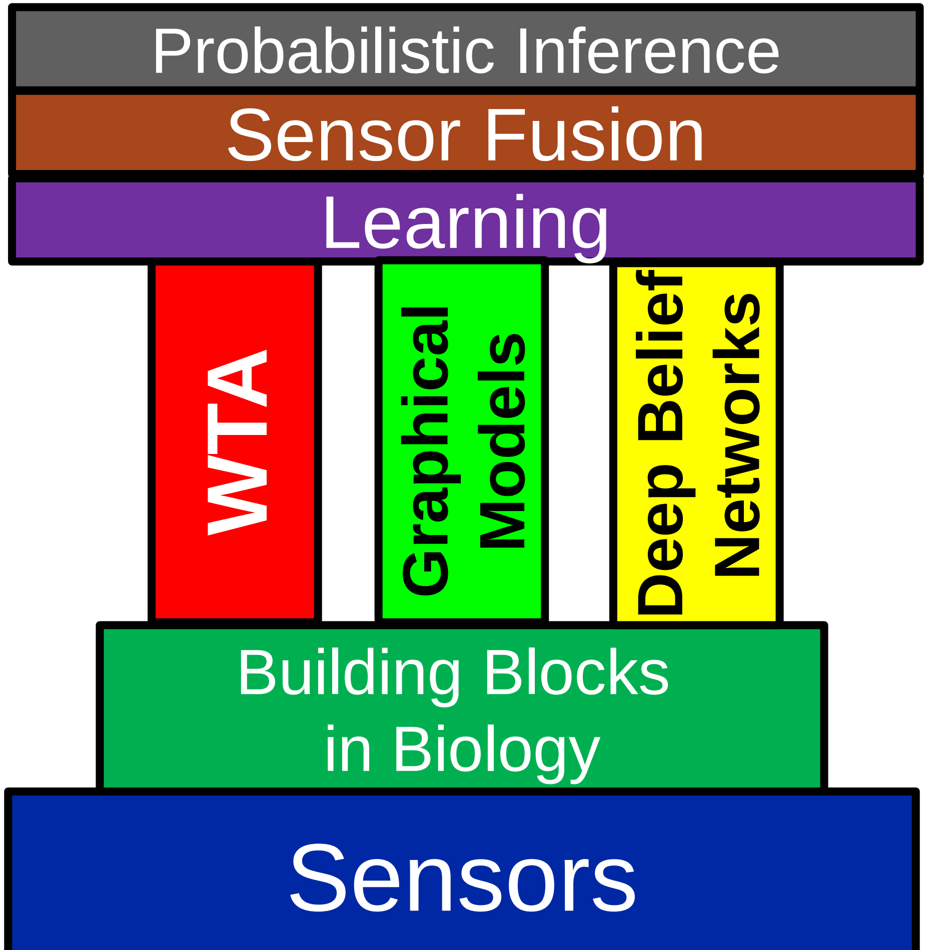

Principles of Deep Learning

Unlabeled Data

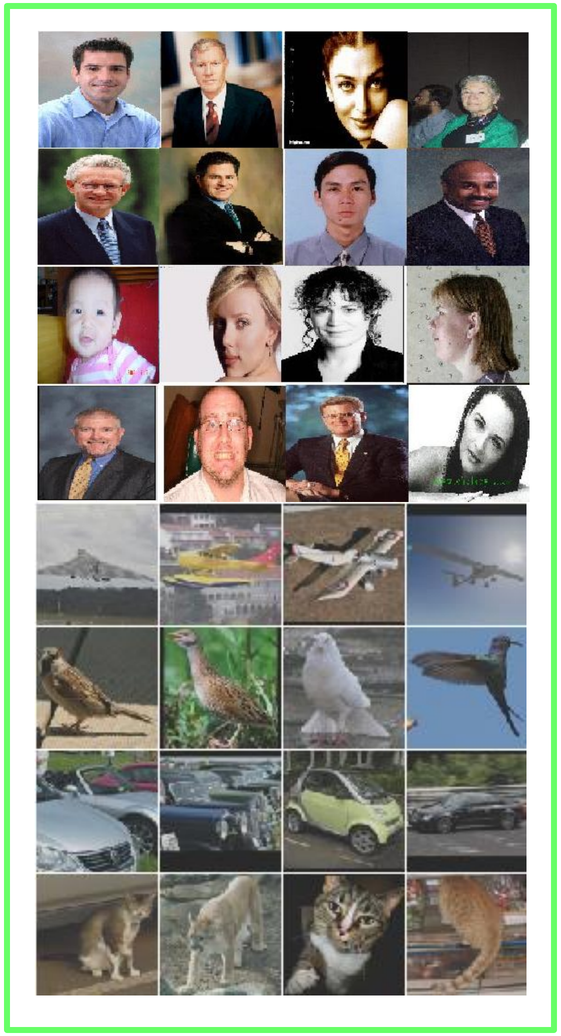

1. Use Joint Dataset to learn a hierarchy of task-independent features
    - Restricted Boltzmann Machines (RBMs)
    - Deep Belief Network (DBN)

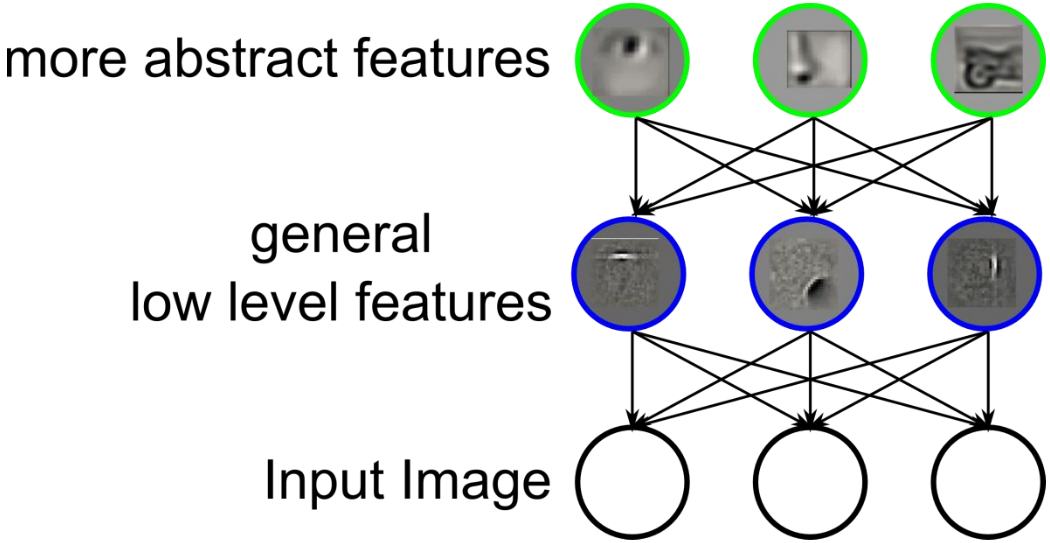

Principles of Deep Learning

Human Faces

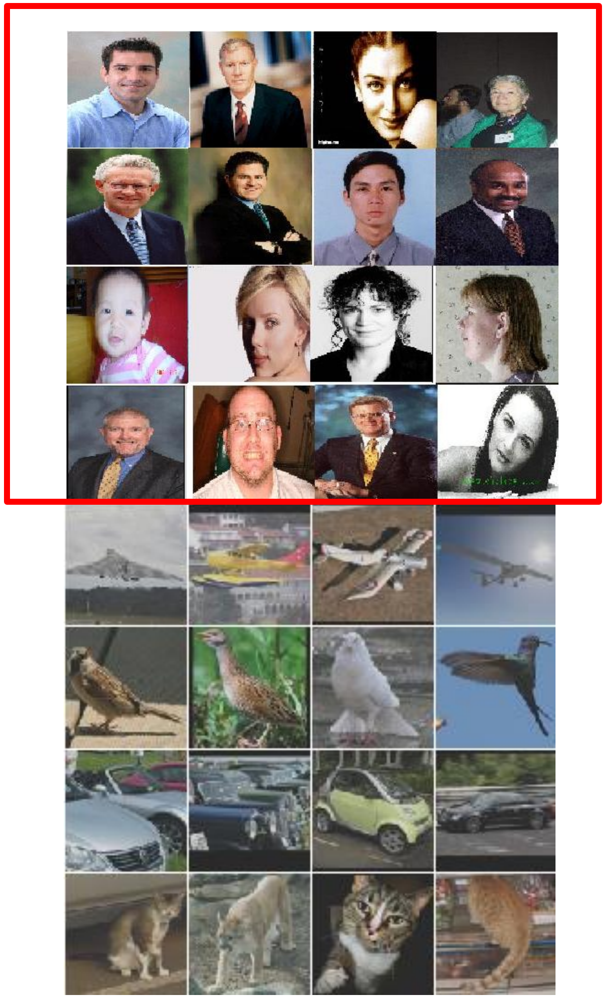

1. Use Joint Dataset to learn a hierarchy of task-independent features
2. Optimize for specific task

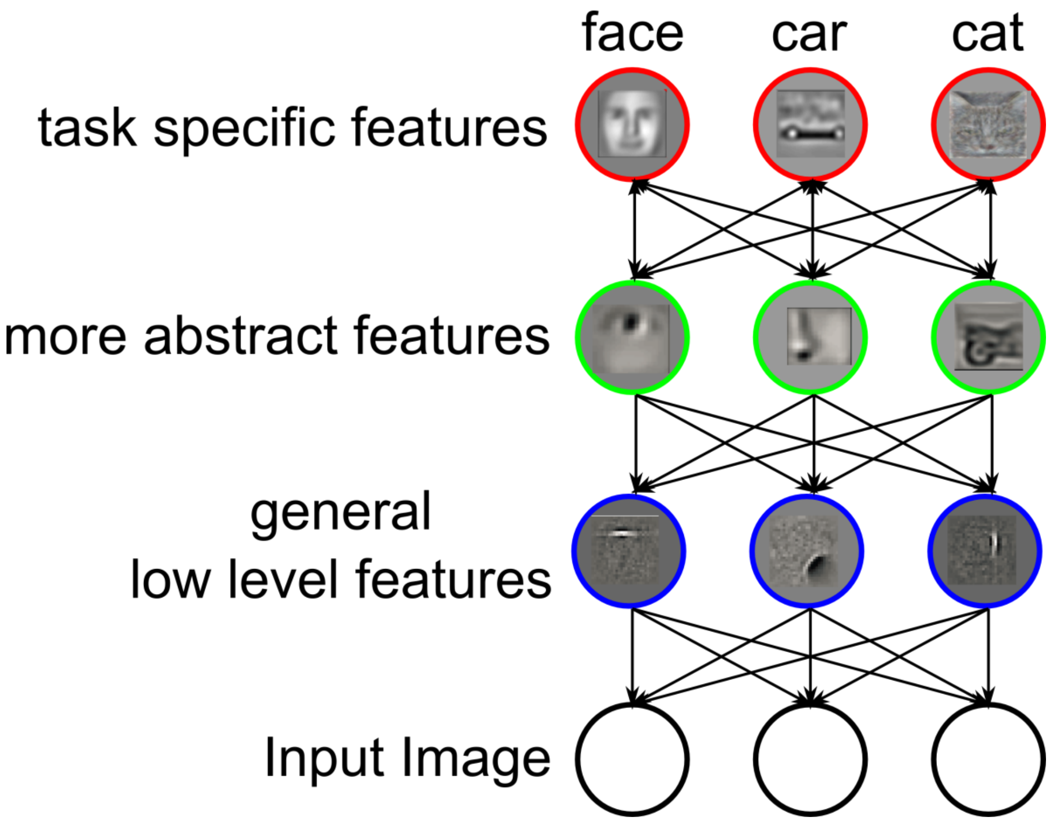

Unsolved Problems of Deep Learning

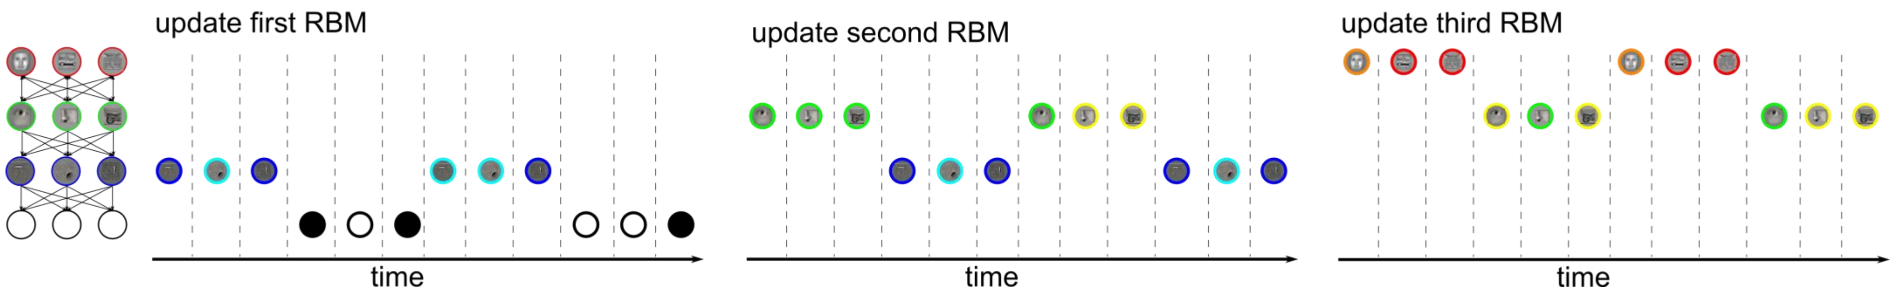

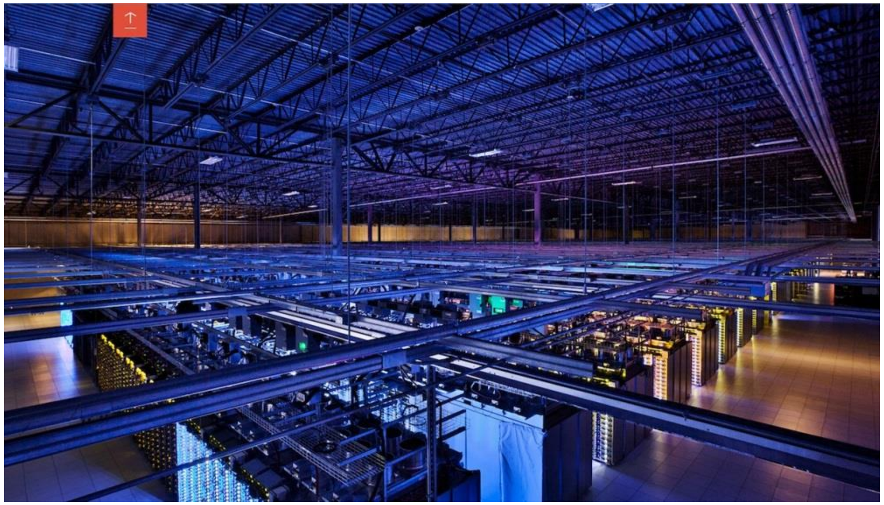

Google Data Center (2013)

Unsolved Problems of Deep Learning

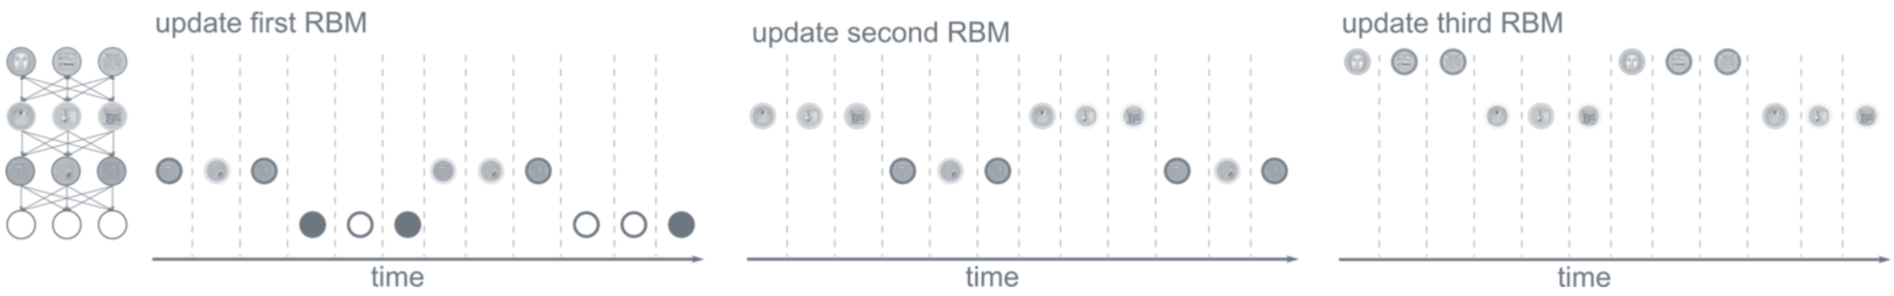

Proposed Solution: Event-based Deep Belief Networks

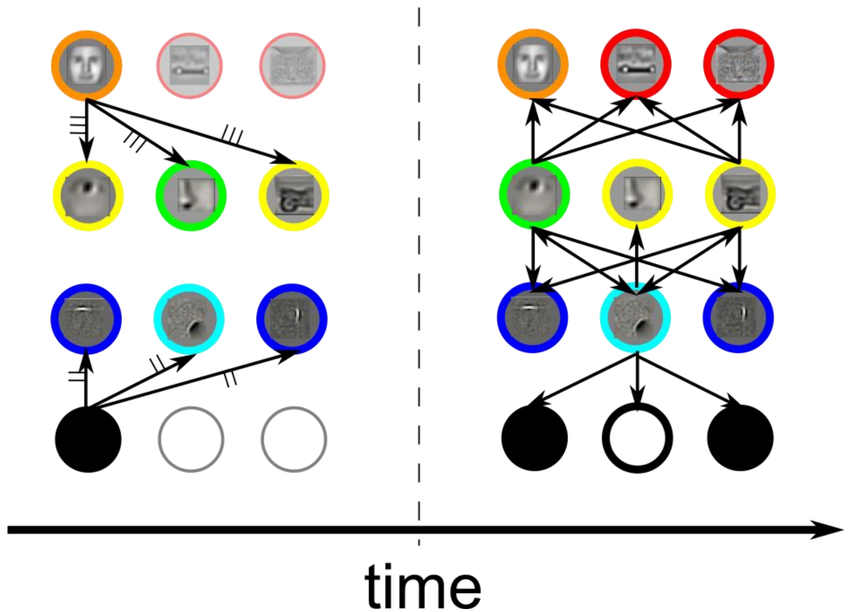

- Massively parallel
- Asynchronous
- Sparse updates
- Online learning
- Scalable

Analogies and Advantages of Brain-like computation

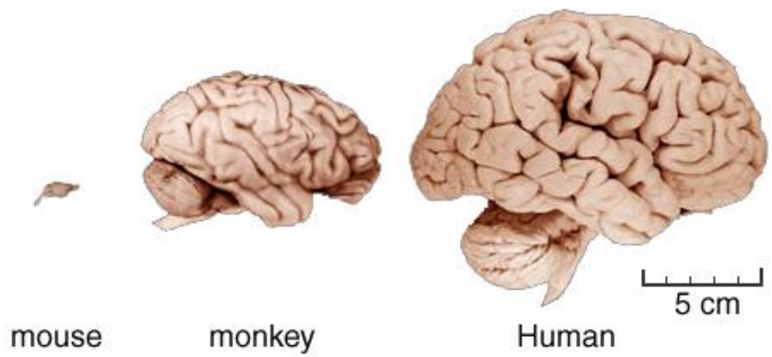

Scaling up without slowing down

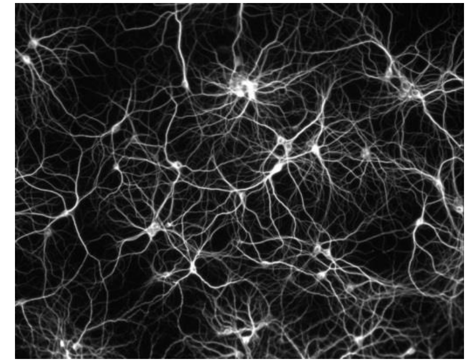

Massively parallel computation
of independent units

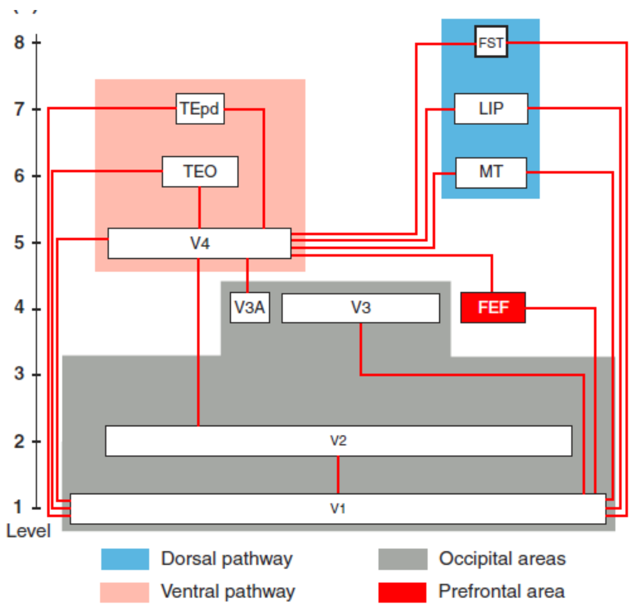

Hierarchical organization

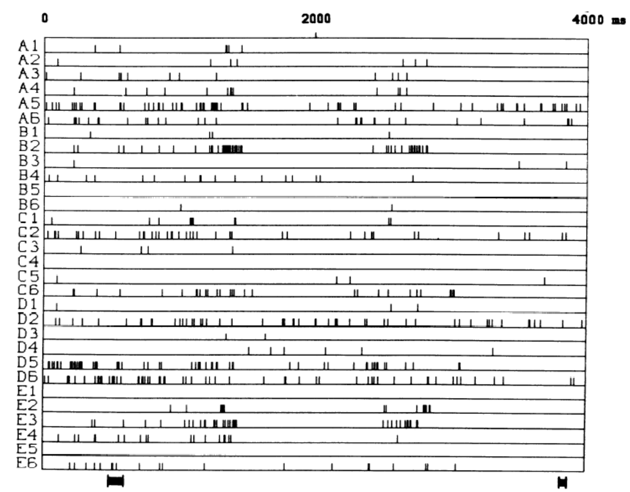

Asynchronous,
sparse distributed
event codes

Figures from: Markov and Kennedy, 2013; Paul DeKoninck Lab; Krüger and Aiple, 1988

Event-based Deep Belief Networks

DVS 128 Vision Sensor

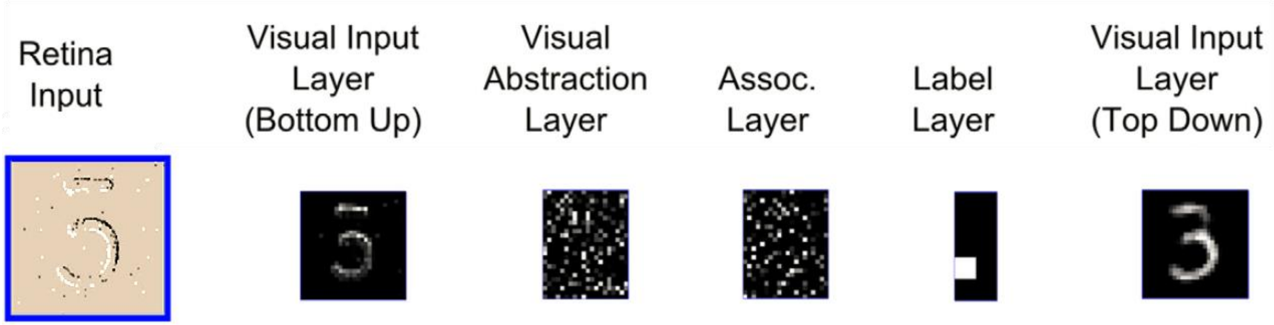

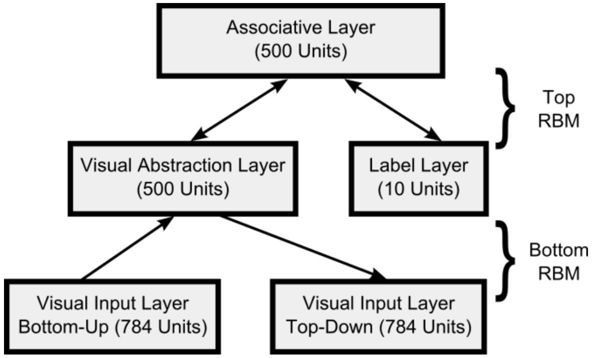

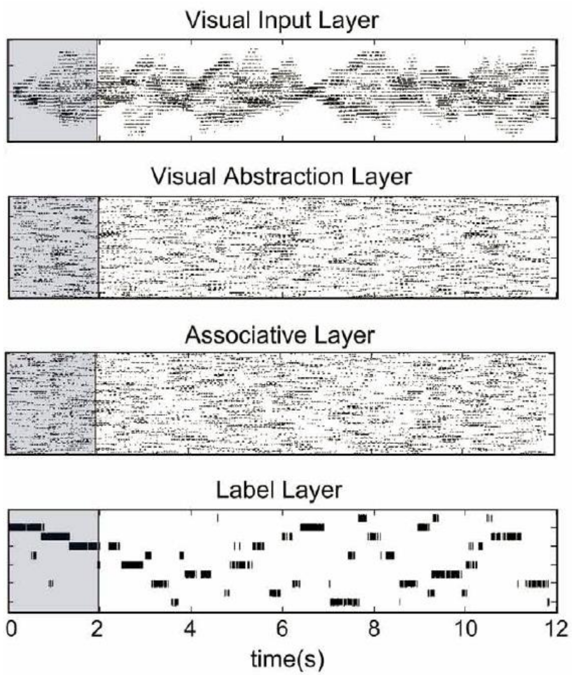

Software simulation (jAER):

- 5.8ms latency
- 94.1% accuracy

O'Connor, Neil, Liu, Delbruck, Pfeiffer: Frontiers in Neuromorphic Engineering (2013)

Training Spiking Deep Belief Networks

$$\mu_Q = \tau_m \cdot \mathbf{w} \cdot \lambda_{\mathbf{in}} \qquad \sigma_Q^2 = \frac{\tau_m}{2} \cdot \mathbf{w}^2 \cdot \lambda_{\mathbf{in}} \tag{1}$$

$$\Upsilon = V_{\text{rest}} + \mu_Q \qquad \Gamma = \sigma_Q \tag{2}$$

$$\lambda_{out} = \Phi(\Upsilon, \Gamma) = \left(t_{\text{ref}} + \frac{\tau_m}{\Gamma}\sqrt{\frac{\pi}{2}} \int_{V_{\text{reset}}+k\gamma\Gamma}^{V_{\text{th}}+k\gamma\Gamma} du \cdot \exp\left[\frac{(u-\Upsilon)^2}{2\Gamma^2}\right] \cdot \left[1 + \text{erf}\left(\frac{u-\Upsilon}{\Gamma\sqrt{2}}\right)\right]\right)^{-1} \tag{3}$$

Siegert model: λ_in/out … Poisson input/output rate

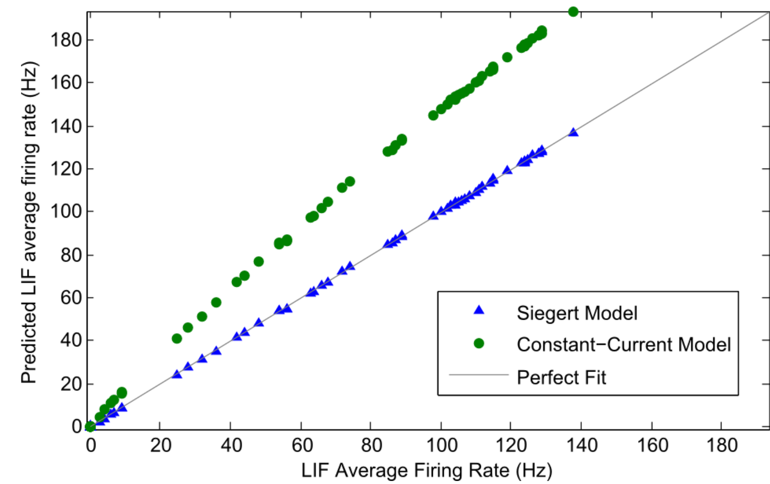

- Offline training of RBMs with Contrastive Divergence
- Use linear-threshold units instead of binary units, replace by LIF
- Approximate LIF firing rate with **Siegert function**
- Usual RBM training, replacing sigmoid transfer function with Siegert
- Transfer trained weights to equivalent spiking DBN

Siegert (1951); Jug et al. (2012)

Multi-sensory Association in Real-time

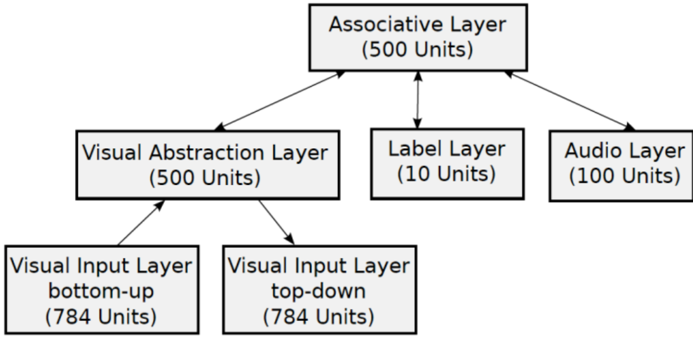

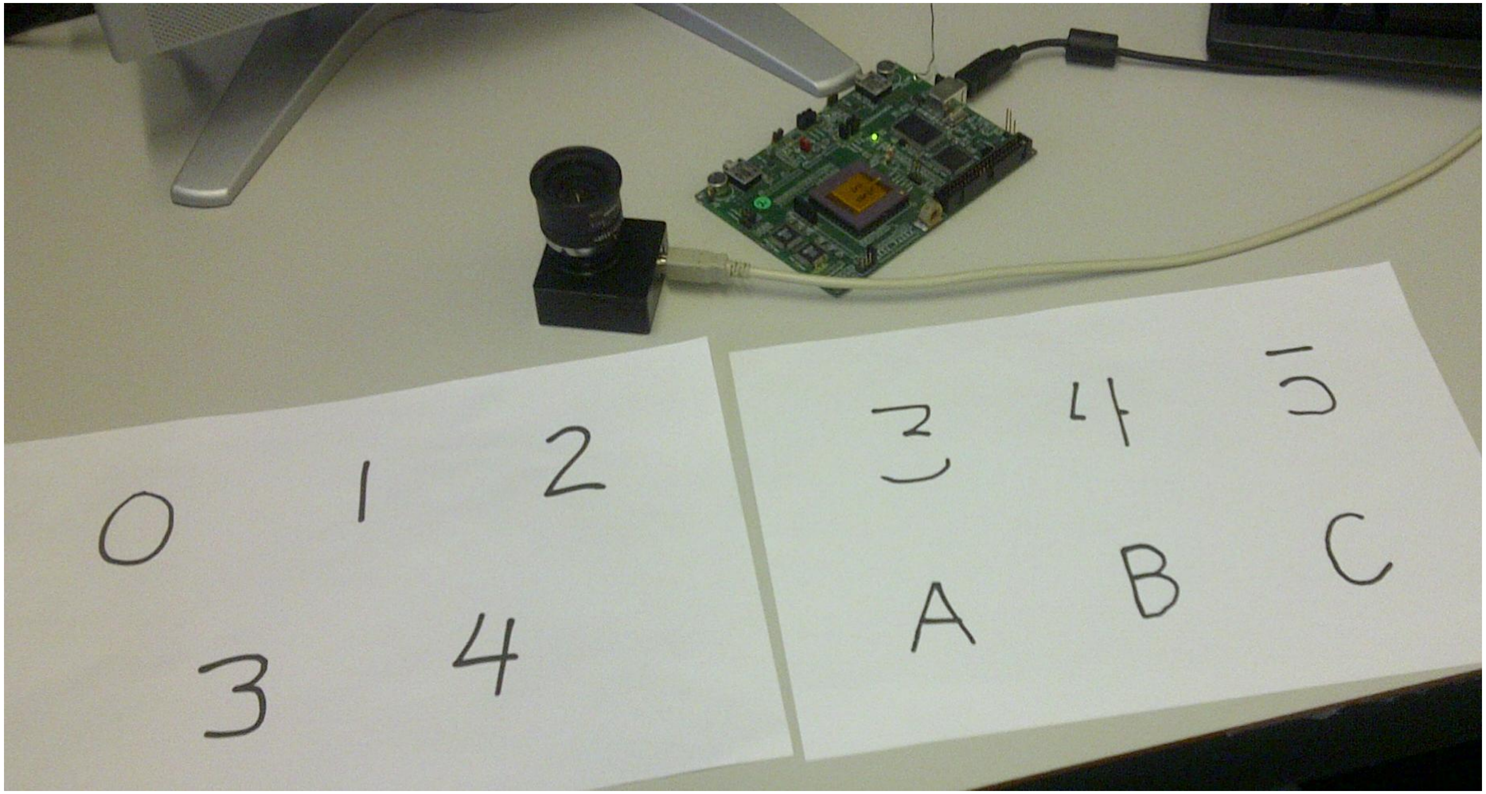

Real-time Classification and Sensor Fusion

Link to event-based RBM / DBN / sensor-fusion videos:

https://sites.google.com/site/thebrainbells/home/event-driven-rbms

Visual Recognition with Distractors

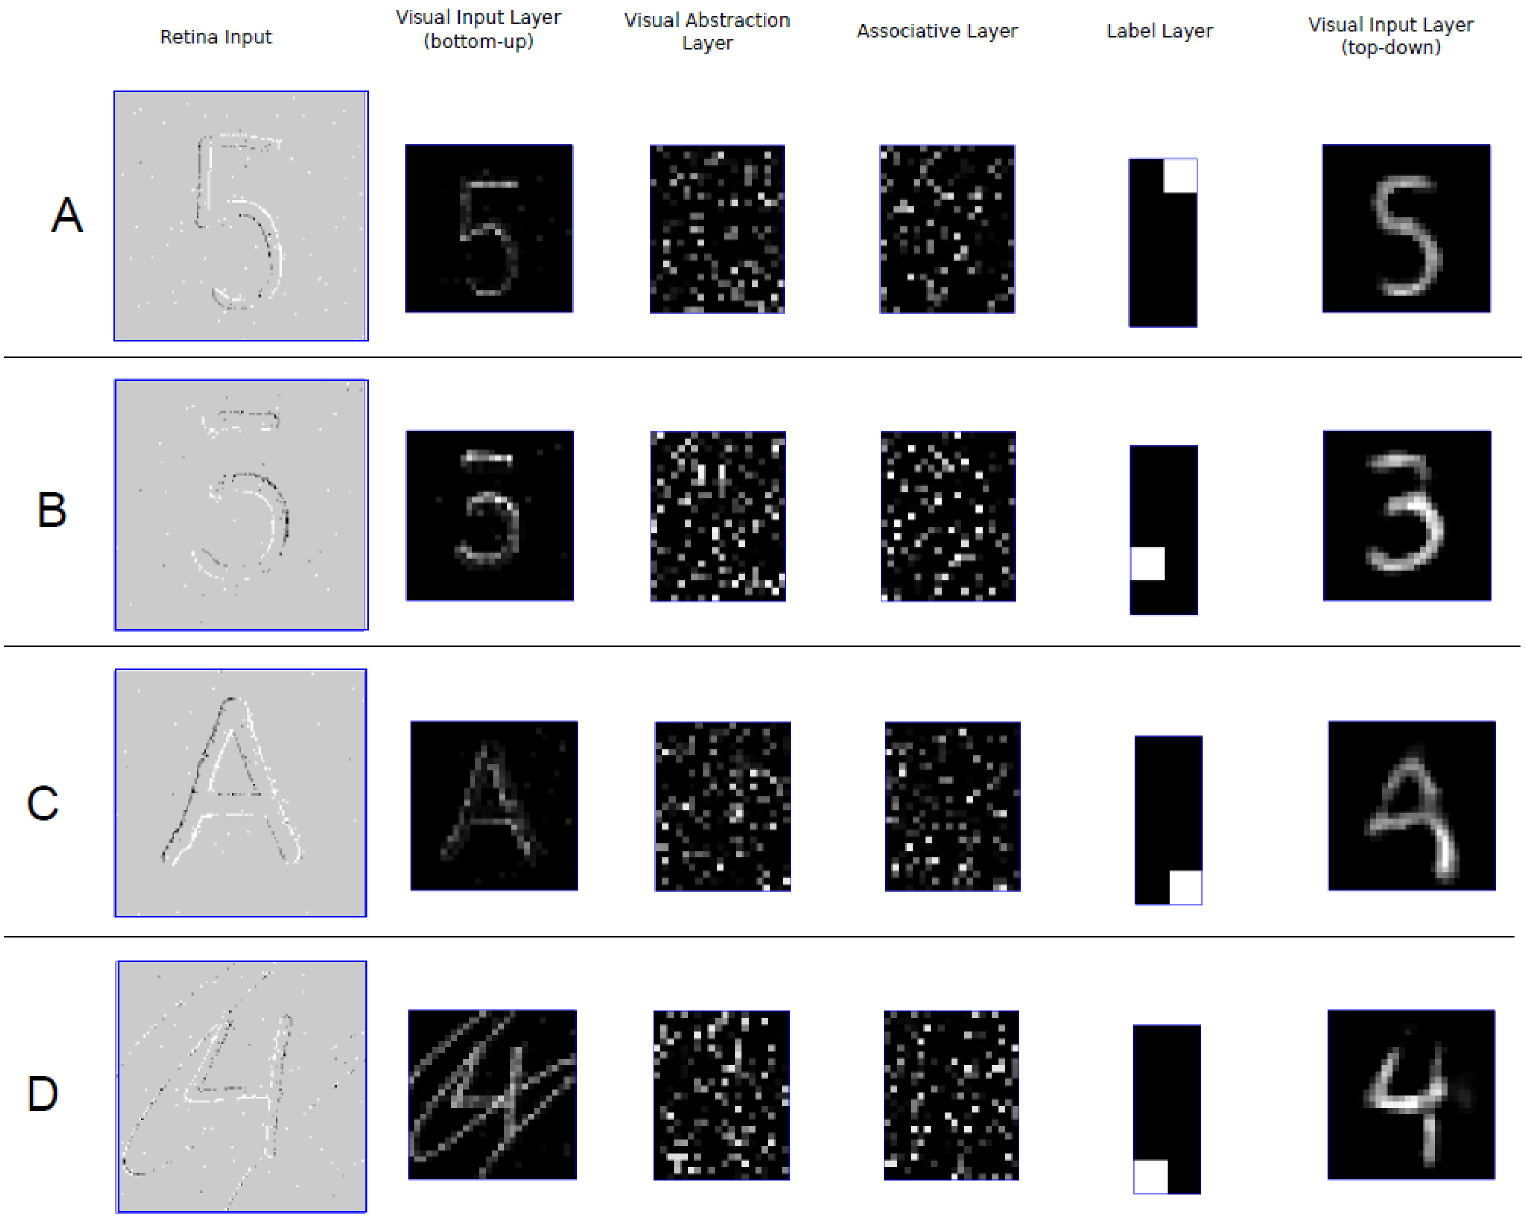

The final block

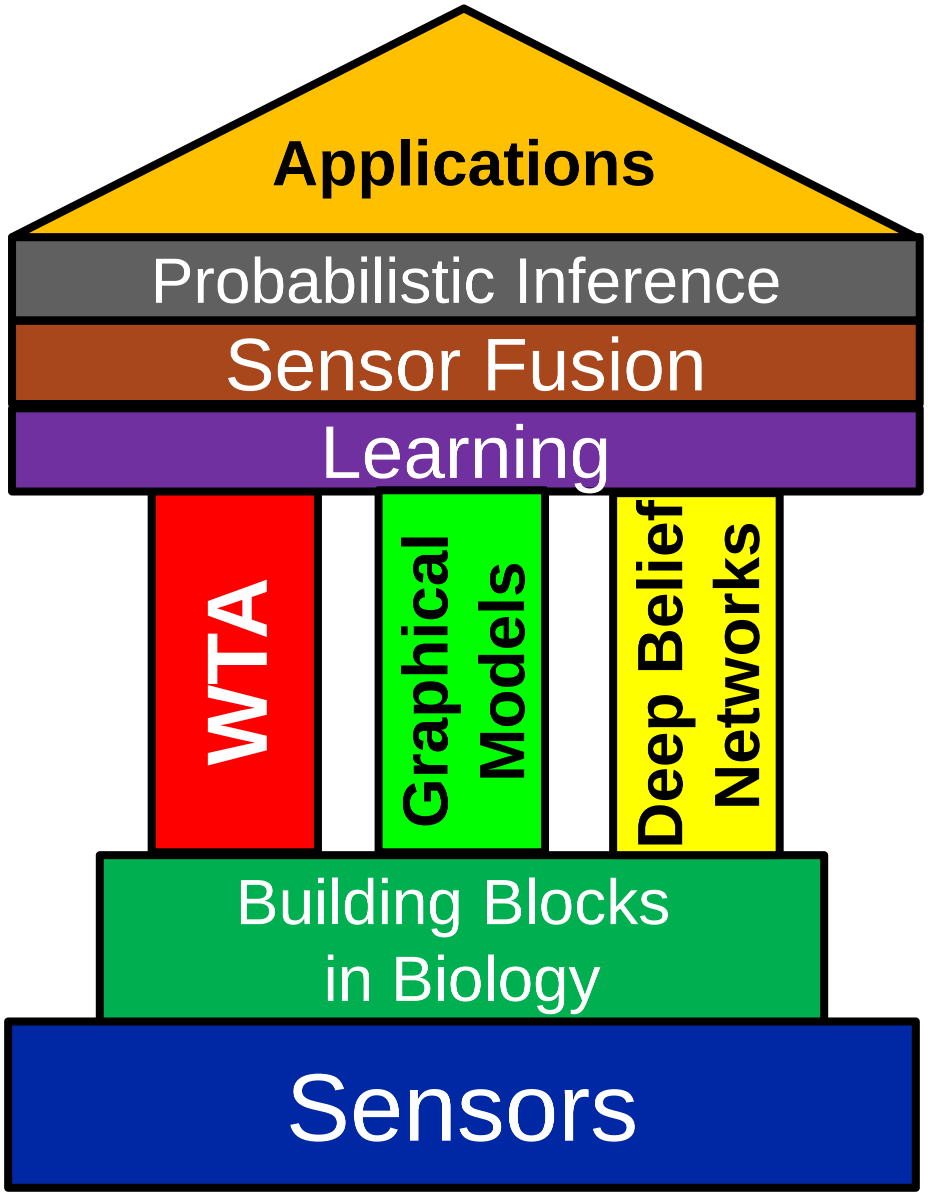

Applications: What to build?

Specialized sensory processing
(e.g. gestures, robotics, fusion, …)

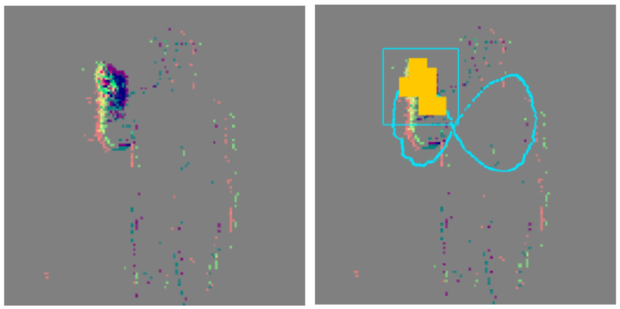

Event-based Machine Learning
(e.g. DBN, EM, …)

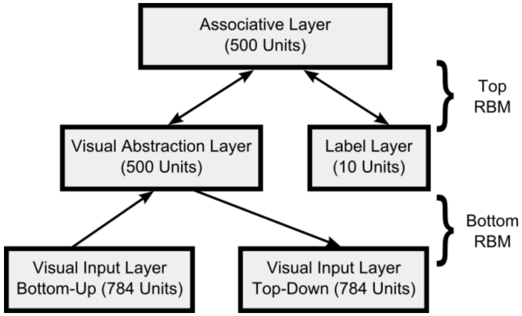

Models of biology
(e.g. WTA, cortical
hierarchies, …)

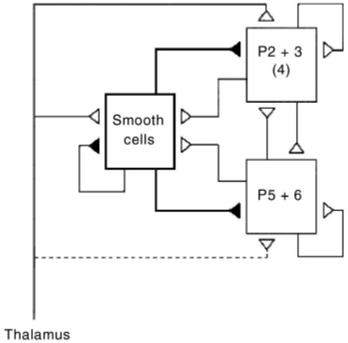

Spatio-temporal processing
(e.g. state-machines, HMM, …)

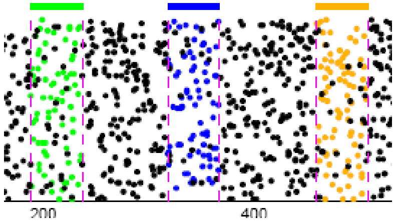

Configuring hardware
(e.g. mismatch, scaling, …)

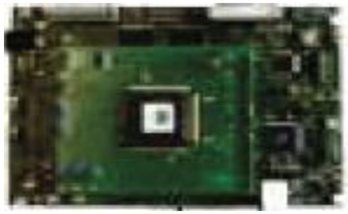

Summary

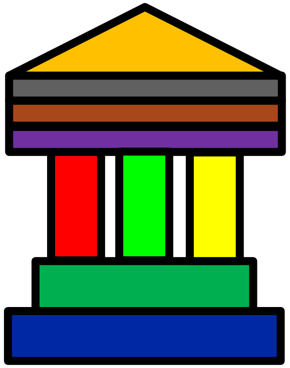

- Building blocks of spiking components for specialized and general purpose applications
  - Sensors as first stage of processing
  - Synthesizing state-machines
  - Learning and probabilistic inference

- Links between machine learning methods and biological plasticity paradigms like STDP

- Deep architectures are more efficient in event-based systems, and can be used for complex classification and sensory fusion tasks

- Suitable for hardware implementation
  - Open issues: realiability, configuration, online adaptation, scaling

Capo Caccia Cognitive Neuromorphic Engineering Workshop

Alghero, Sardinia (Italy)

28 April – 10 May 2014

capocaccia.ethz.ch

Acknowledgements

UZH and ETHZ

Rodney Douglas

Tobi Delbruck

Shih-Chii Liu

Giacomo Indiveri

Danny Neil

Peter O'Connor

Dane Corneil

Emre Neftci

SAIT

Jun Haeng Lee

Hyunsurk Ryu

TU Graz

Wolfgang Maass

Bernhard Nessler

Caltech

Ueli Rutishauser

Funding

University of Zurich[UZH]

SECO Project

TU Graz

Convergent Science Network of biomimetic and biohybrid systems

SAIT