Life Sciences und
Facility Management

Real-time ML and neuromorphic computing for smart robots

Yulia Sandamirskaya

Institute of Computational Life Sciences

ZHAW

Fast ML for Science 02.09.2025

Why do we build Artificial Intelligent systems?

- Curious to better understand intelligence

Playing Chess Deep RL

Playing Go

Talking LLM

Writing (an Email)

useful

- Create new technology

Making staff

Lifting staff

Harmful, 3D work

Repetitive work

Life-long assistance

Need for robots

Care at home (Spitex):

- Number of clients: 458.578 (2022)
- Number of workers: 61.086 (2022)
- Cost: 3.177.450.000 CHF (2022)

Robots!

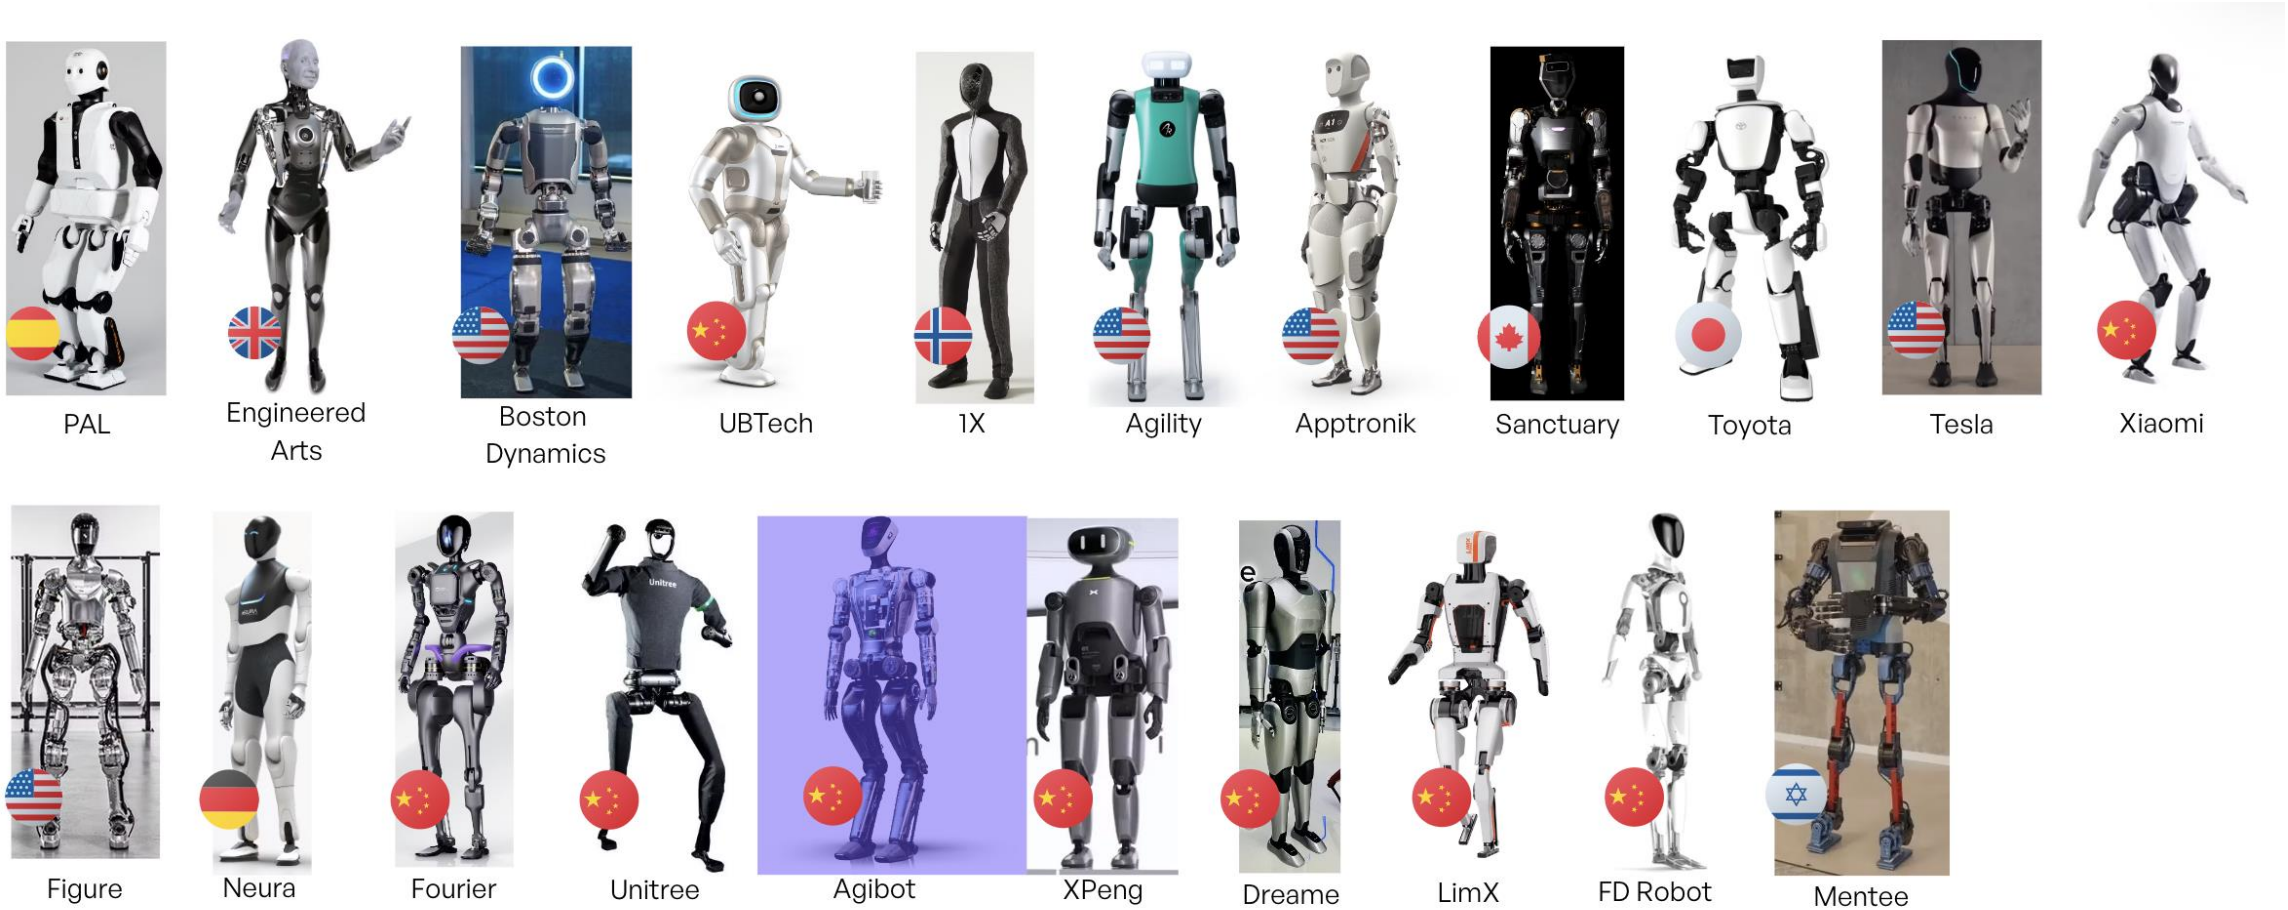

What is a robot?

- An electro-mechanical system

Robots today: a closer look

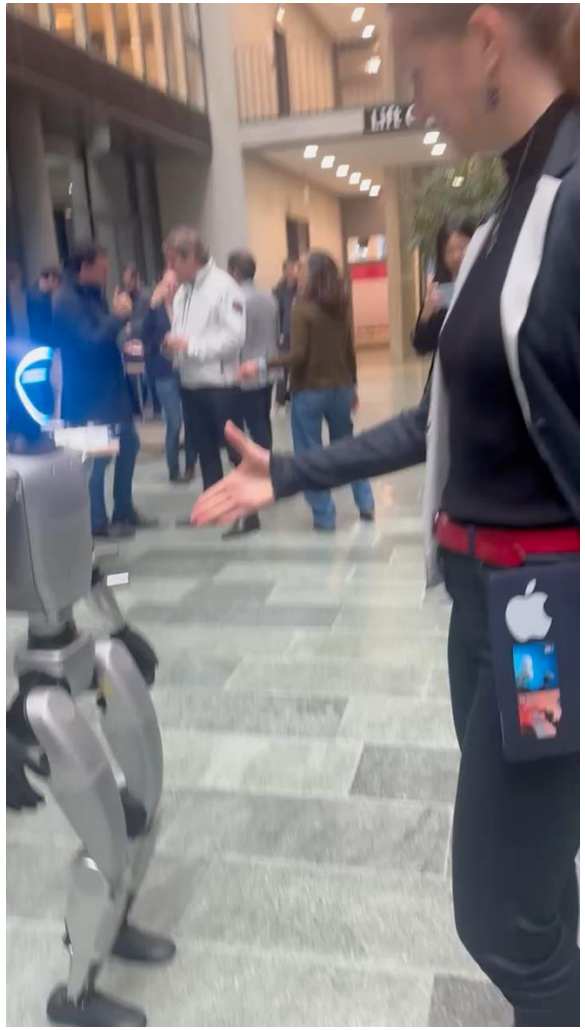

Zurich, November 2024,
UniTree G1 Humanoid

- Why doesn’t it move towards my hand?
- Would it be able to grasp and bring me objects?
- Can it learn new stuff?
- Is it safe?

What does it take to control a robot?

Mobility

- L
- 3D
- SLAM.
- Path planning
- Motion planning

Manipulation

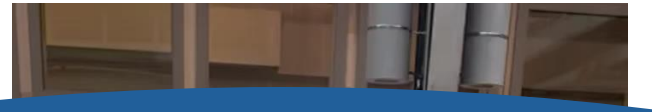

- 6D pose estimation
- Tactile sensing
- Force control

Action planning

+ HRI
+ Learning
+ Adaptation

A lot of (classical) software

Can AI help us to make robots smarter?

Artificial Intelligence today

Image recognition

Autonomous driving

ChatGPT

Automatic Translation
Automatische Übersetzung
Traduction automatique
Traducción automática
Автоматический перевод

Let’s look under the hood!

What is an Artificial Neural Network

Multilayer Perceptron (1950s)

Input

(image)

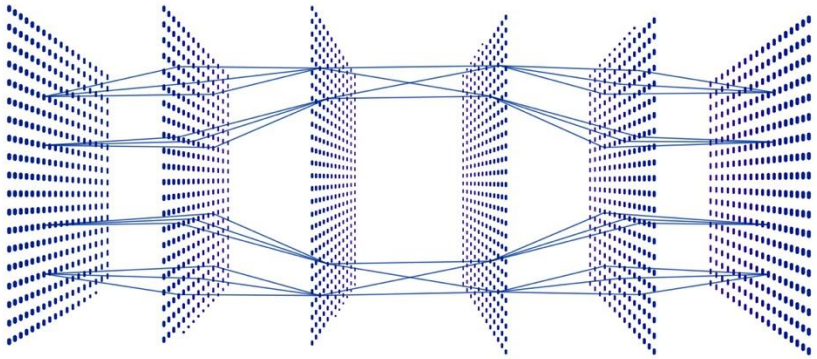

Output

(label, pose, ROI, mask)

A LOT of data:

inputs and correct outputs

(need to cover all corner cases!)

«Neuron»:

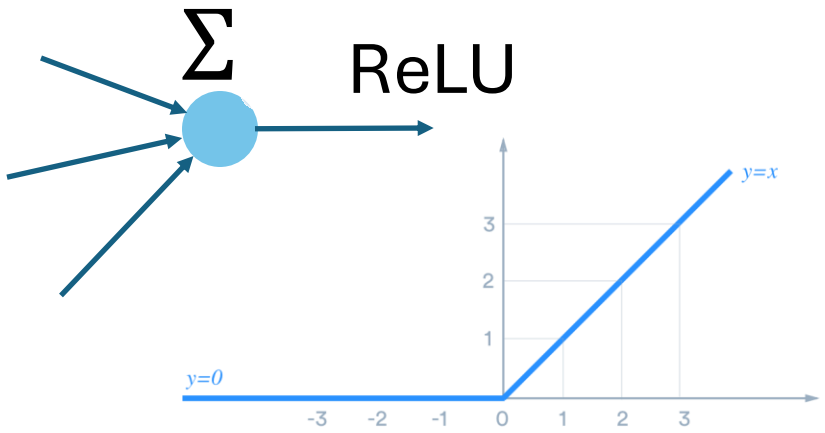

Slow, long training

(Error Back-propagation)

«Bag of tricks»

(norm, batch, training schedule, tuning)

What are the limitation of AI / DL (in robotics)?

- One task at a time
  - Need a much more multi-modal approach (not just images and text)
    - Different sampling rates
    - Different dimensionalities
    - Different reliability
    - Different (changing?) geometrical arrangement
  - Orchestrating different tasks is a challenge
- Data hungry
  - Physical world is fundamentally complex and dynamic
- High latency
  - Control loop might need <10ms latency and confidence
- Reliability: accuracy and robustness

Success came at a high cost

Chat GTP-4 Training:

- 100 Million Dollar
- 1.287 Gigawatt hours
- ??? Human labor

Behind impressive videos

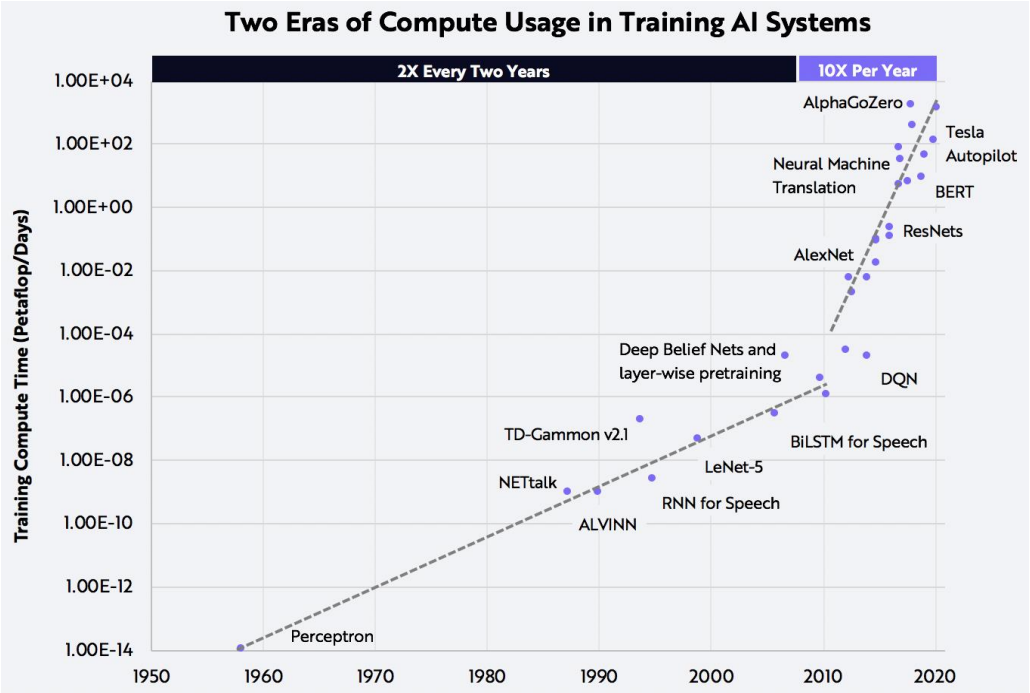

300,000x more power (2014-2020)
vs **8x** by Moore's Law

Transparency
Robustness
Adaptivity

Source: https://techhq.com/2023/03/data-center-energy-usage-chatgpt
Source: ARK Investment Management LLC, "AI and Compute." OpenAI, https://arkinv.st/2ZOH2Rr.

Success of today's AI relies on hardware!

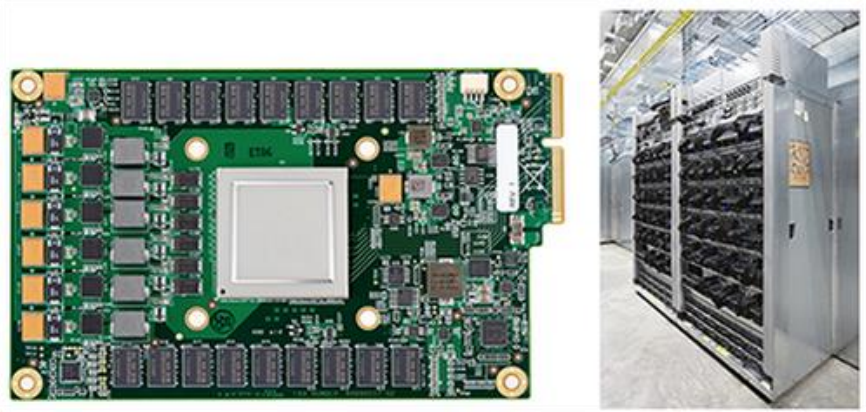

- Optimized for matrix multiplication
- Optimized for batched data processing
- Optimized for tensor processing

On the market today

CPU

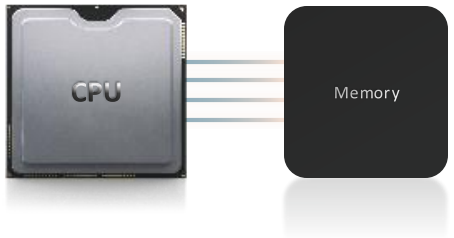

GPU

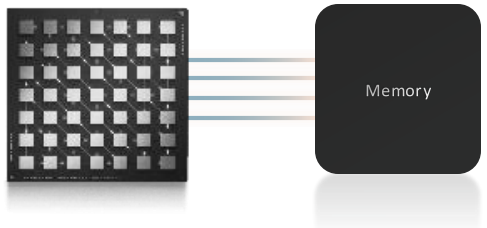

SEQUENTIAL THREADS OF CONTROL	PARALLEL DENSE COMPUTE
SYNCHRONOUS CLOCKING	SYNCHRONOUS CLOCKING
PROGRAMMING BY ENCODING ALGORITHMS	TRAINING USING LABELED DATASETS

```
if X then
  …
else
  …
```

01100
11010
00100

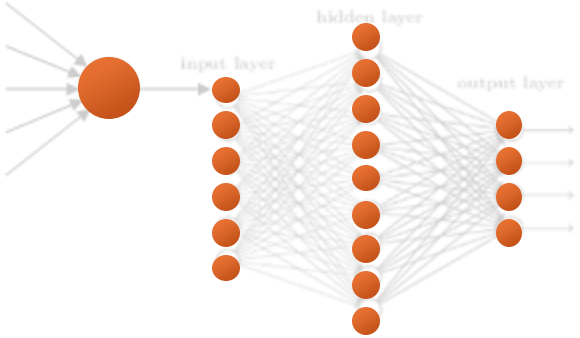

We have examples of neural networks doing better…

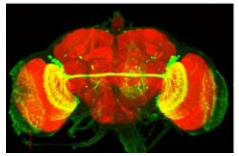

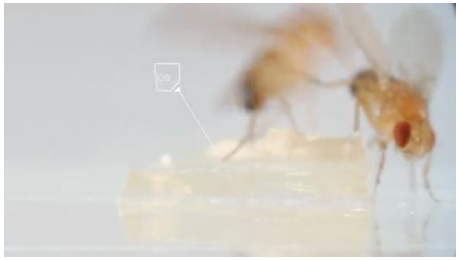

- 0.1 g
- 100K neurons
- 20M connections

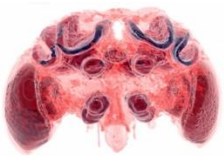

- 1g brain, 1M neurons, 1mW
- Navigates and learns in unknown environments "on the fly"

- 1000g brain, 100 B neurons, 20 W
- Can do "everything we want those robots to do"

Biological brains:

- Adaptive
- Flexible
- Fast
- Precise
- Efficient

- Deal with real-world complexity
- Learn new tasks
- "Cognitive"

How do brains compute?

1. Six cortical layers form “predictive” circuits.
   - Consistence check (spatial, temporal)
   - Continual learning and adaptation

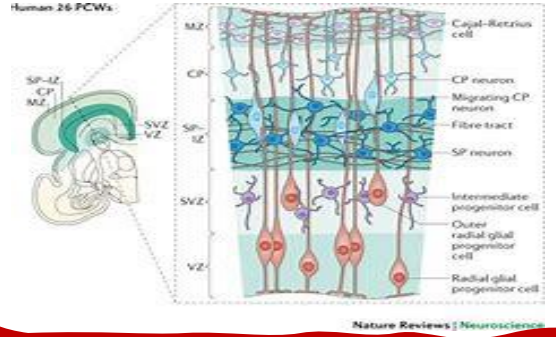

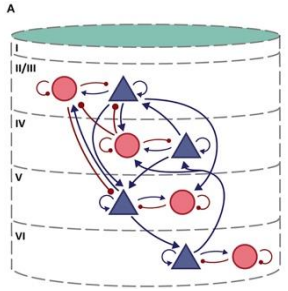

2. Thalamo-cortical and Cortical-BG loops
   - Adaptive resonance, active sensing
   - Cognitive control of behaviors

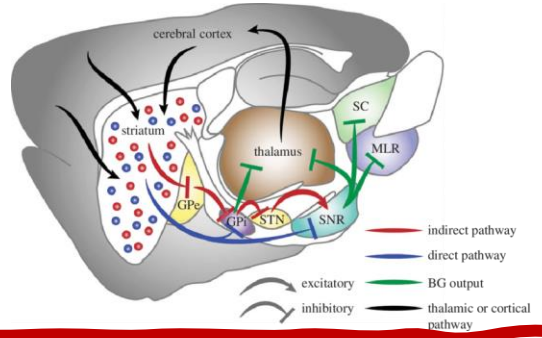

3. Cerebellum
   - Adaptive learnable mappings
   - Muscle coordination, movement fine-tuning

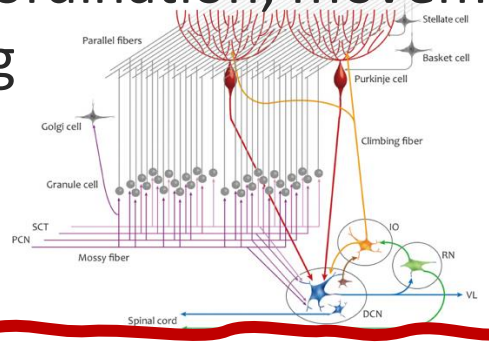

4. Hippocampus, ethorhinal cortext
   - Place-cells, grid cells
   - Odometry, SLAM, episodic memory

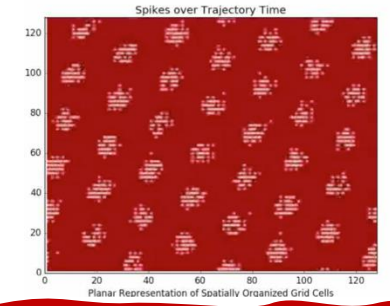

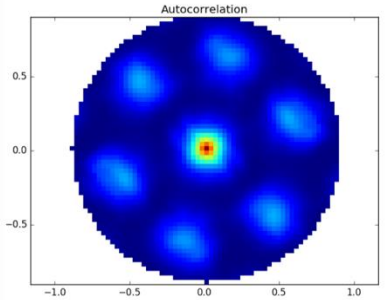

- Many different “algorithms”
- Continual learning and adaptation

Why CPU and GPU are not a good match for brain-inspired algorithms?

- Parallel, but not “images” or matrices or (full-rank) tensors
  - Sparse neural networks (graphs)
  - Recurrent neural networks (loops!)

- Different transfer functions (not just ReLU)
  - Statefull neural networks (local memory!)
  - Plasticity (more local memory!)

- Need asynchronous processing
  - Computation is not clocked, rather event-driven
  - Dirrent sampling rates

Neuromorphic Computing Hardware

CPU

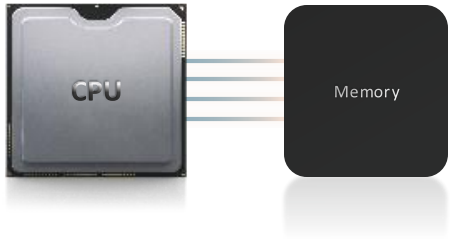

GPU

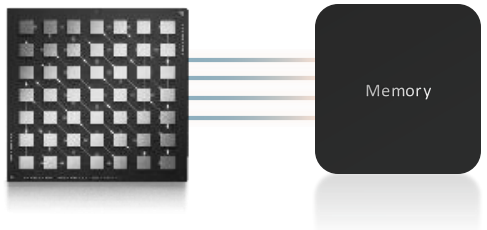

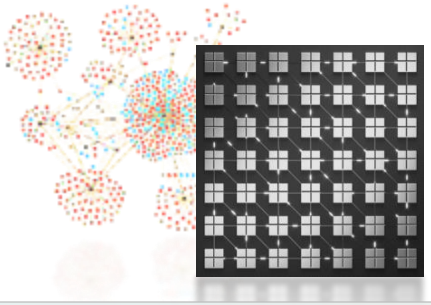

NCH

SEQUENTIAL THREADS OF CONTROL	PARALLEL DENSE COMPUTE	PARALLEL SPARSE COMPUTE
SYNCHRONOUS CLOCKING	SYNCHRONOUS CLOCKING	ASYNCHRONOUS EVENT-BASED SPIKES
PROGRAMMING BY ENCODING ALGORITHMS	OFFLINE TRAINING USING LABELED DATASETS	LEARN ON THE FLY THROUGH LOCAL LEARNING RULES

```
if X then
  ...
else
  ...
```

01100
11010
00100

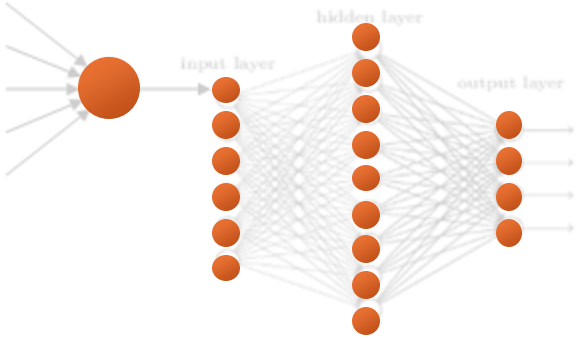

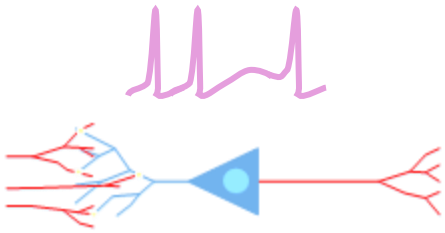

What is a neuromorphic chip?

- Mixed-signal (analog-digital) circuits

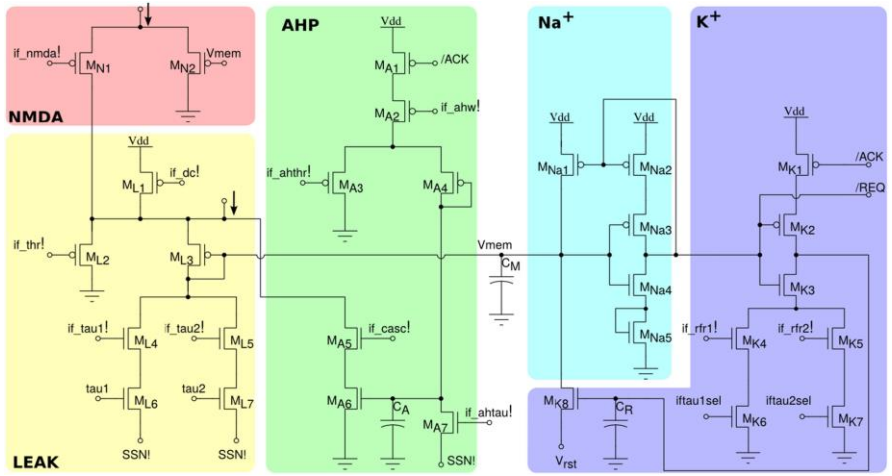

- memristors
- optical
- bio
- nanowires
- nanodots

- Digital

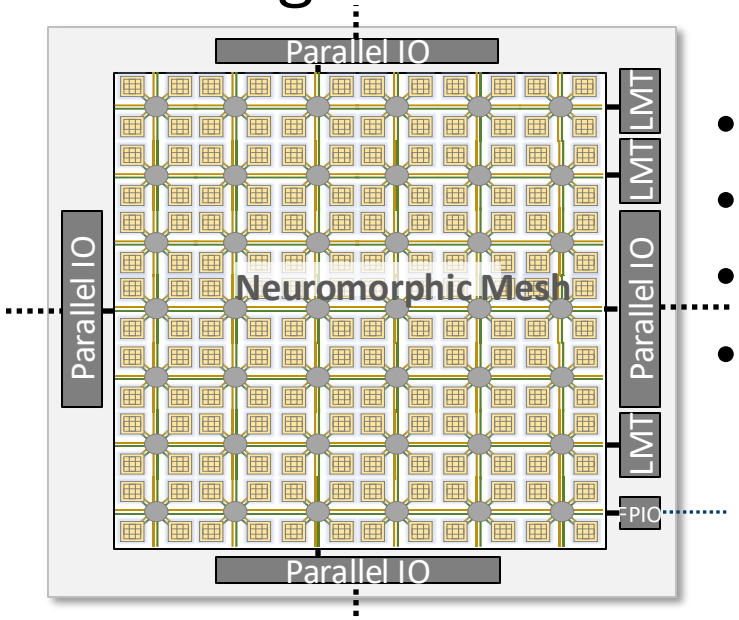

- asynchronous
- fix-point
- smart routing
- optimized ALU

- VLSI chips (alt: in-memory computing)

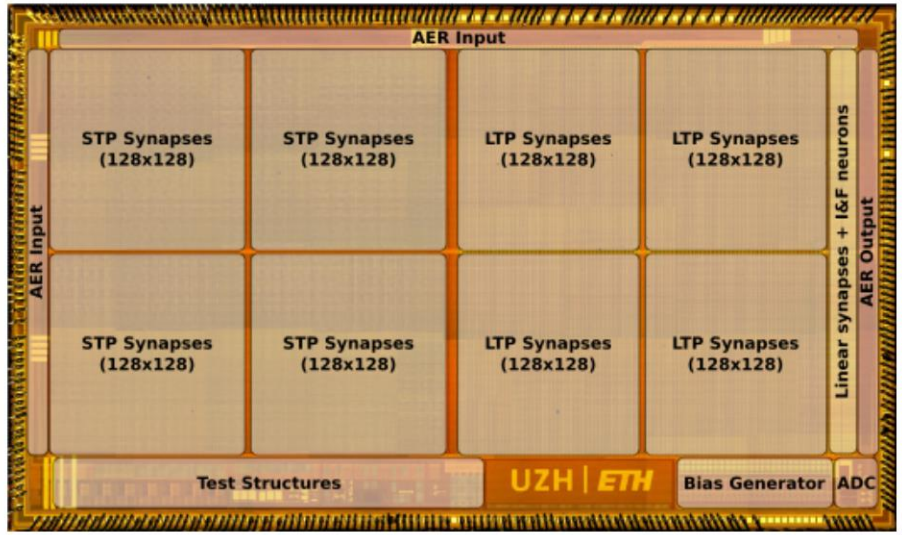

- Digital chips (+I.O, co-processors,...

Recurrent networks with novel bio-inspired properties give the best gains

Reference architecture

- CPU (Intel Core/Xeon)
- GPU (Nvidia)
- Movidius (NCS)
- TrueNorth

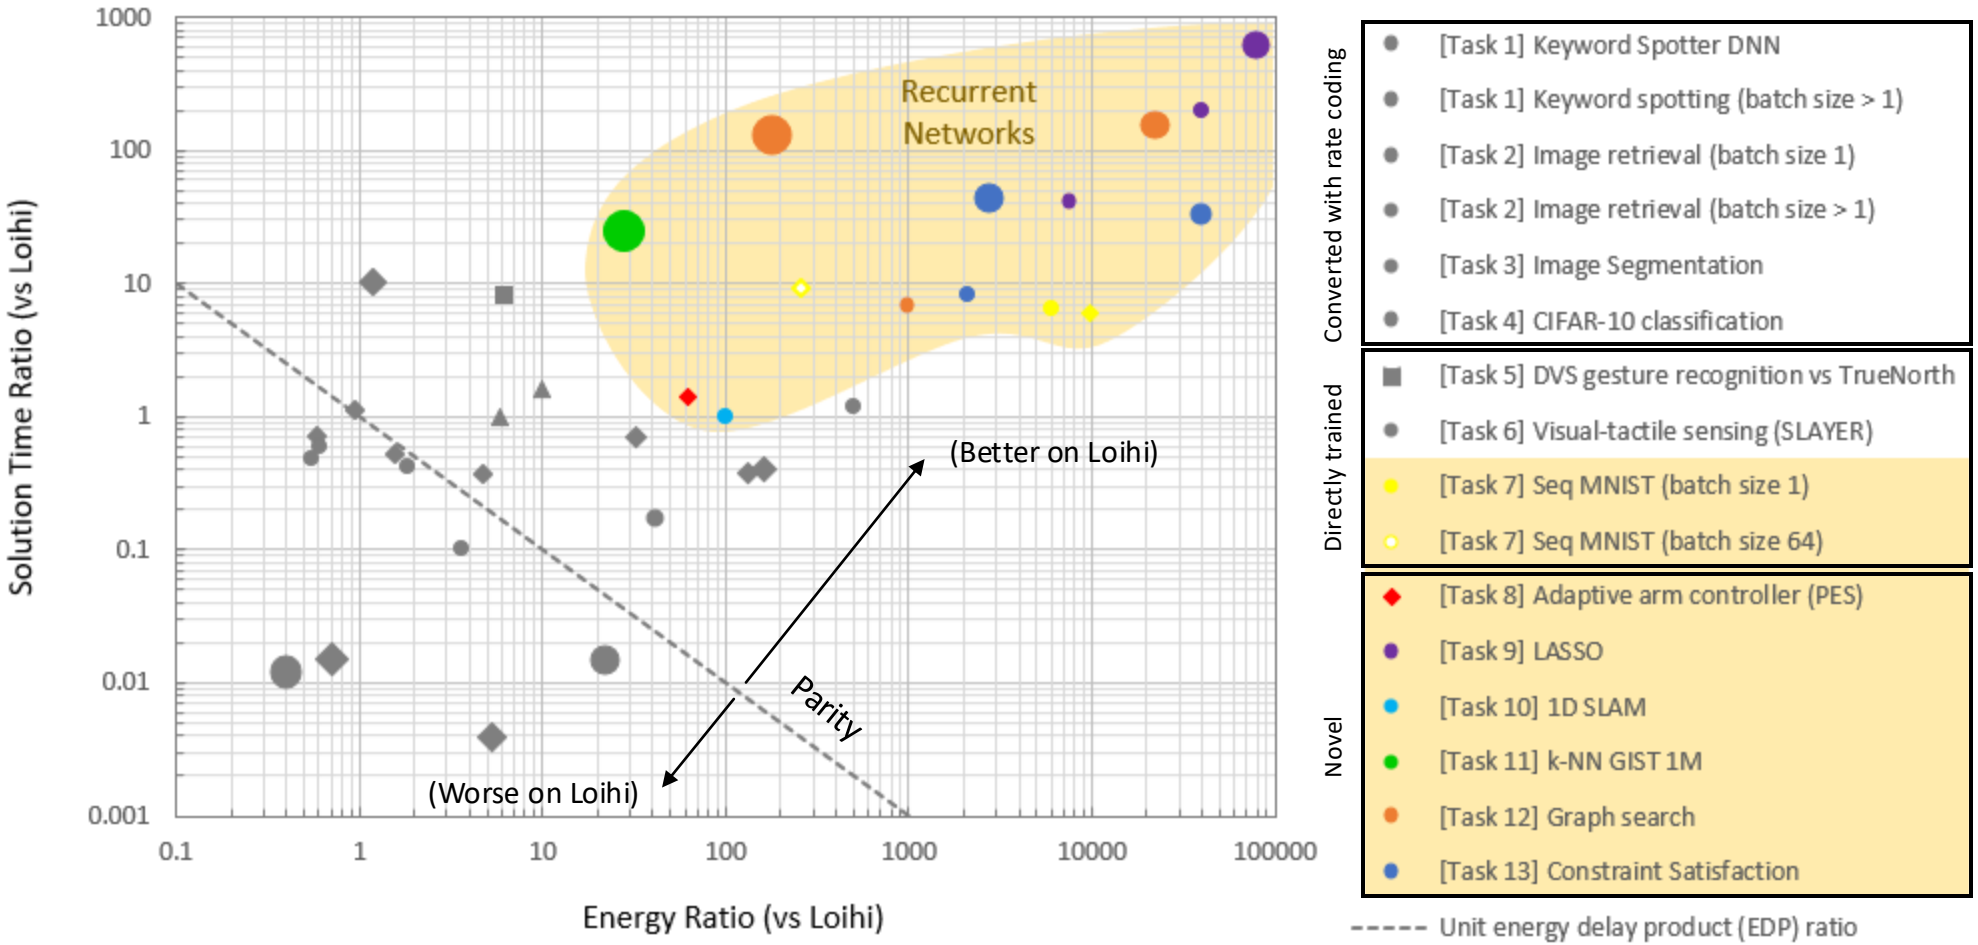

Davies et al, Proc. IEEE 2021. Results may vary.

Neuromorphic chips on the (research) market

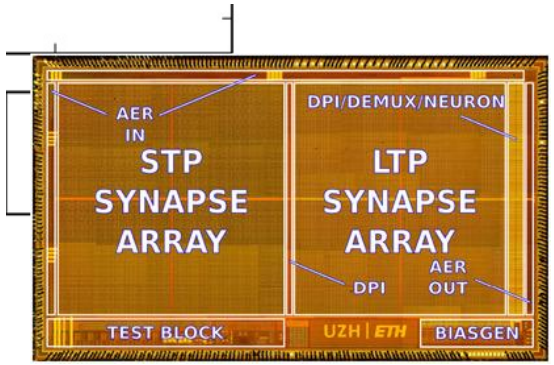

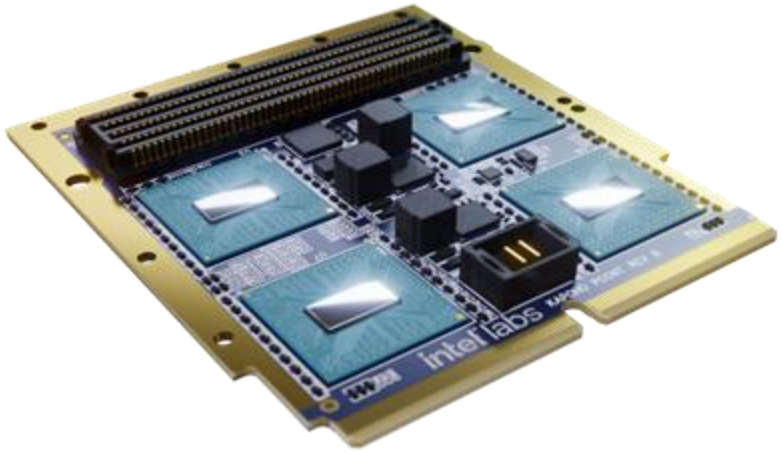

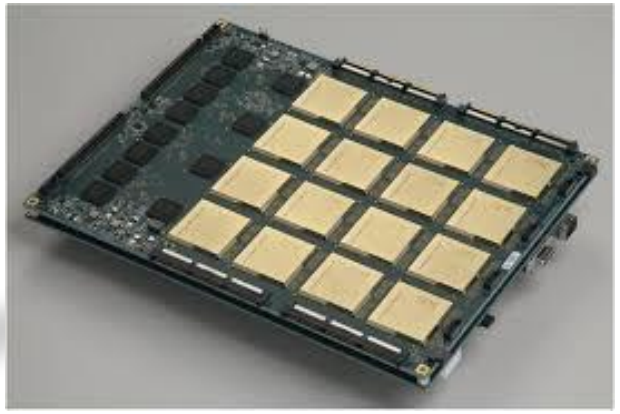

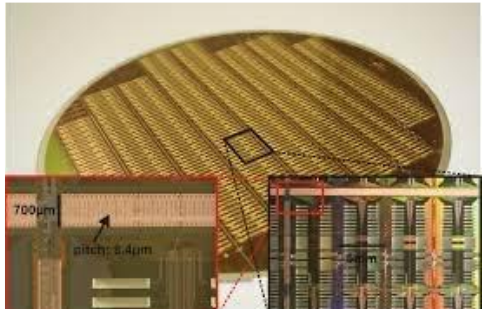

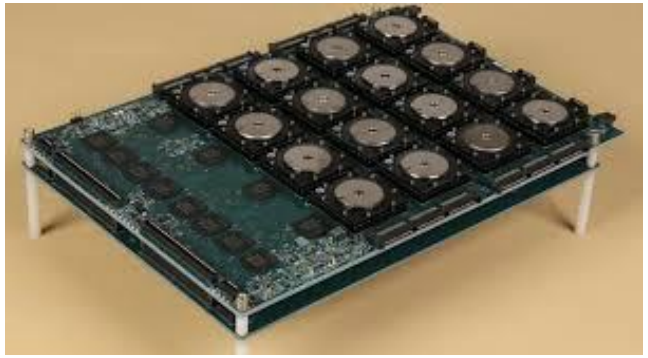

Hardware-algorithm co-design!

Programming neuromorphic hardware…

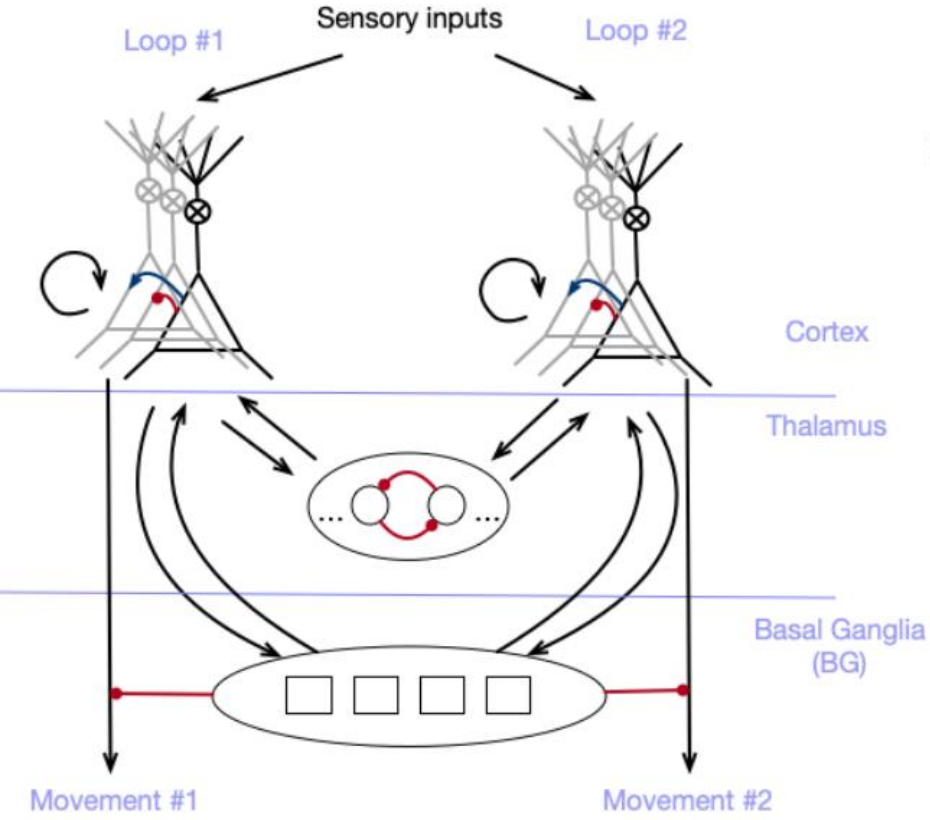

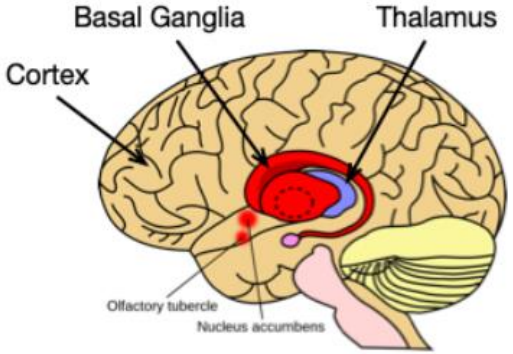

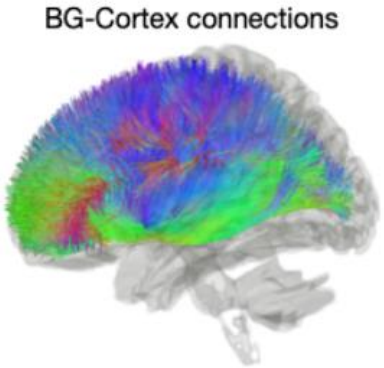

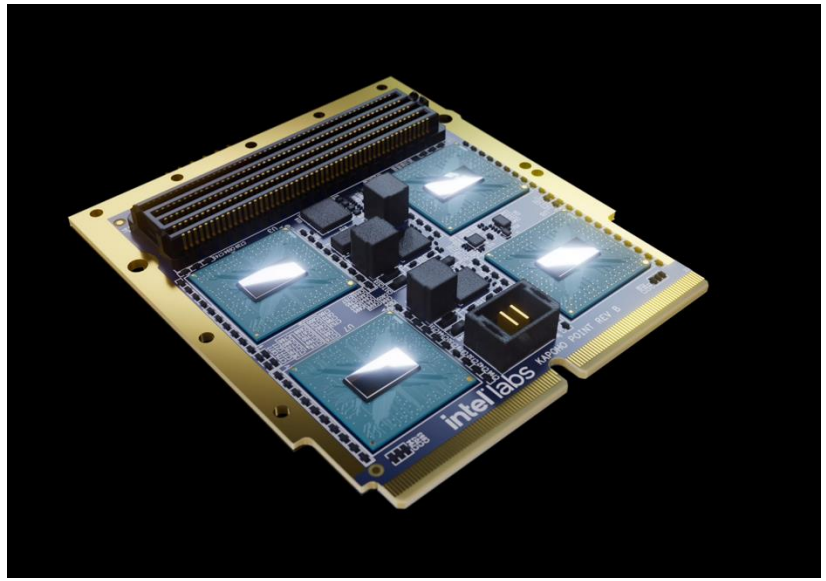

…to control (assistive) robots

Not all neuroscience insights are made equal…

Brain structures and dynamics

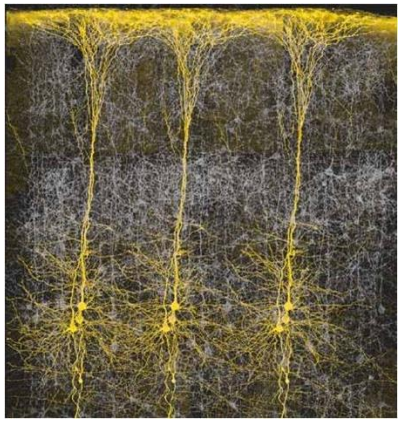

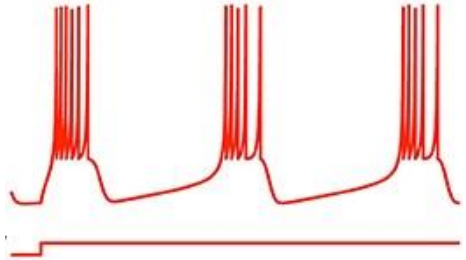

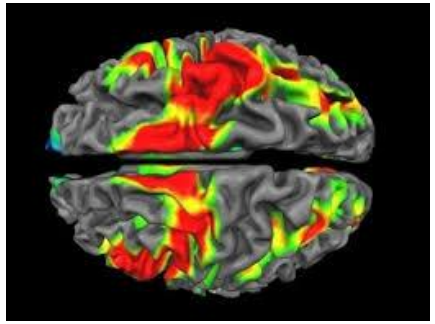

Behavior

Neural Principles of behavior and cognition

1. Groups of neurons form (attractor) dynamics

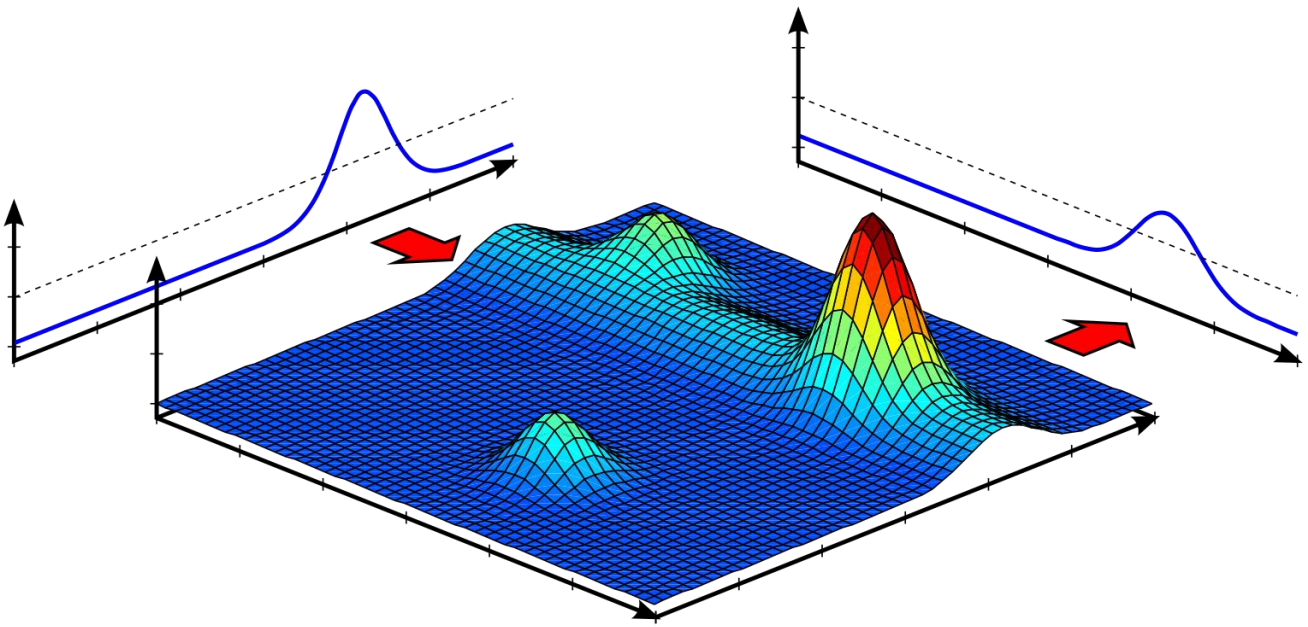

https://dynamicfieldtheory.org/

Behavioral studies

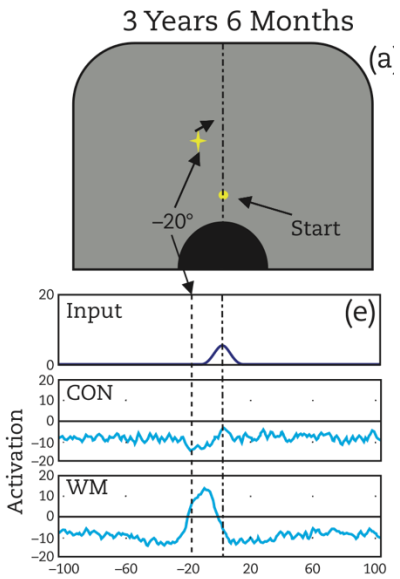

Neural studies

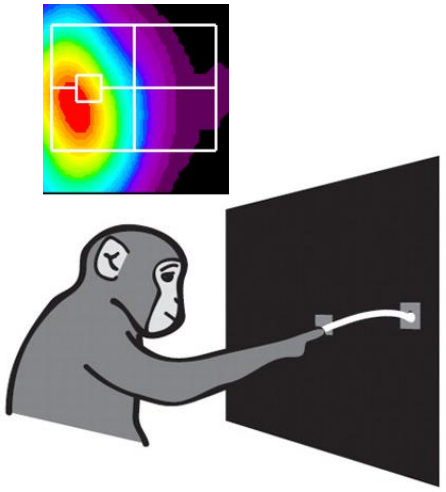

Model for **Working Memory**

- Basic element of
  - perception
  - intentional action

Neural Principles of behavior and cognition

2. Intentional structure of neural states (Neural state machines)

- Attractors need to be actively destabilized
- Anticipating the sensory consequences of actions

6 cortical layers

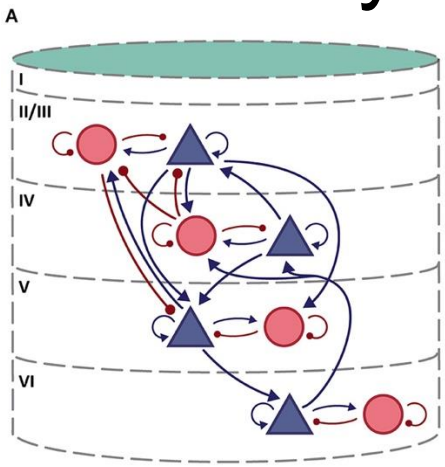

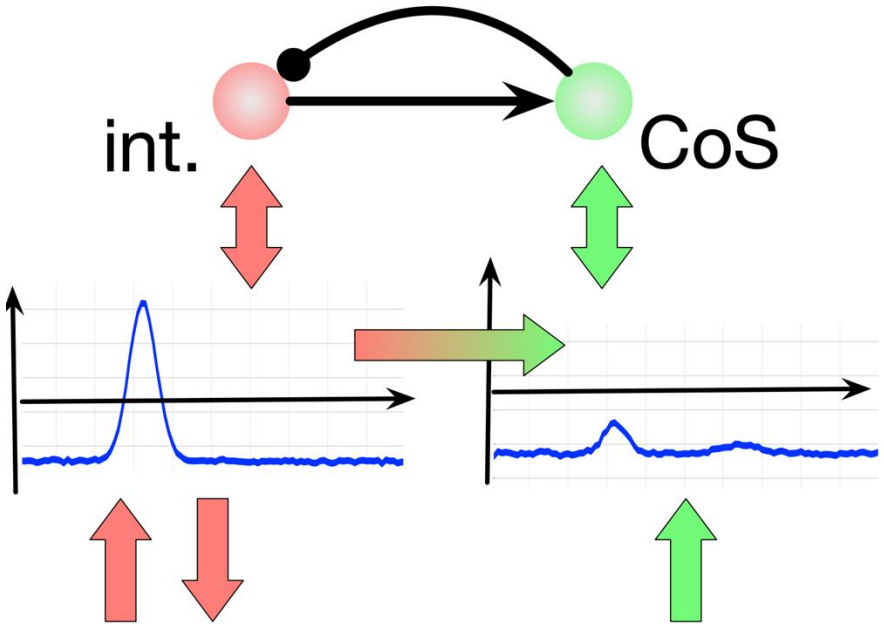

Structure of an "elementary behavior"

Sandamirskaya, Schoner, 2010; Richter et al 2012; Billing et al 2015

Neural Principles of behavior and cognition

3. Reference frames. Transformations.

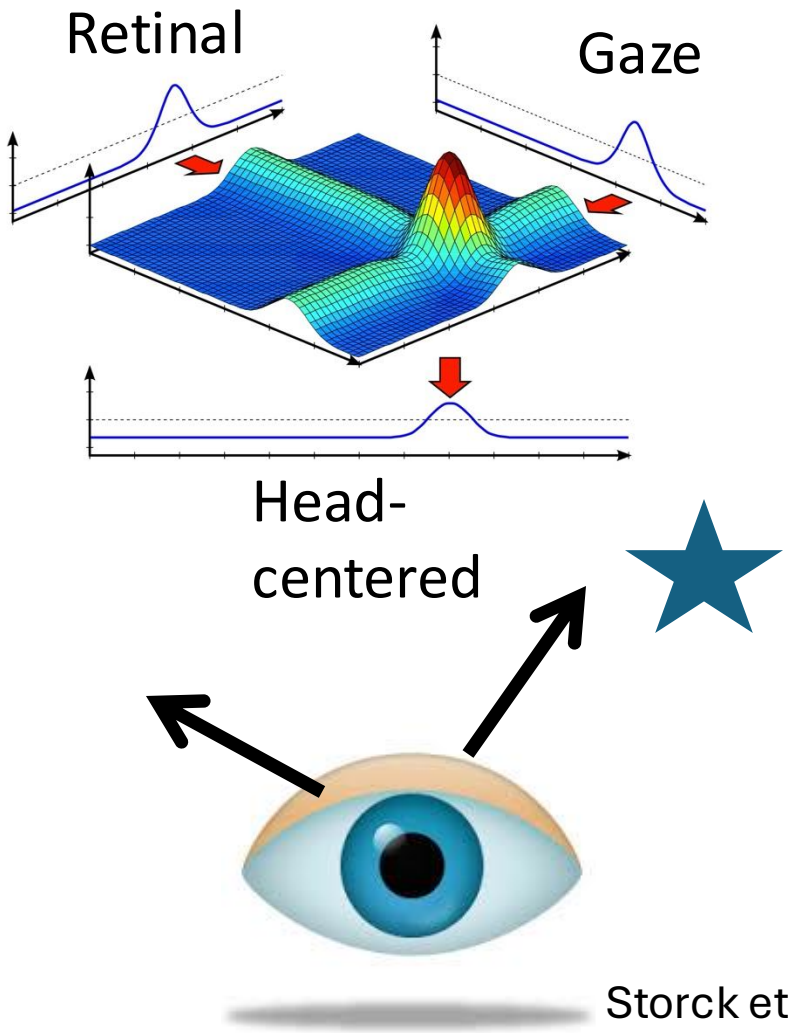

Storck et al 2014

Spatial relations

Lipinski et al 2011

Learning relations

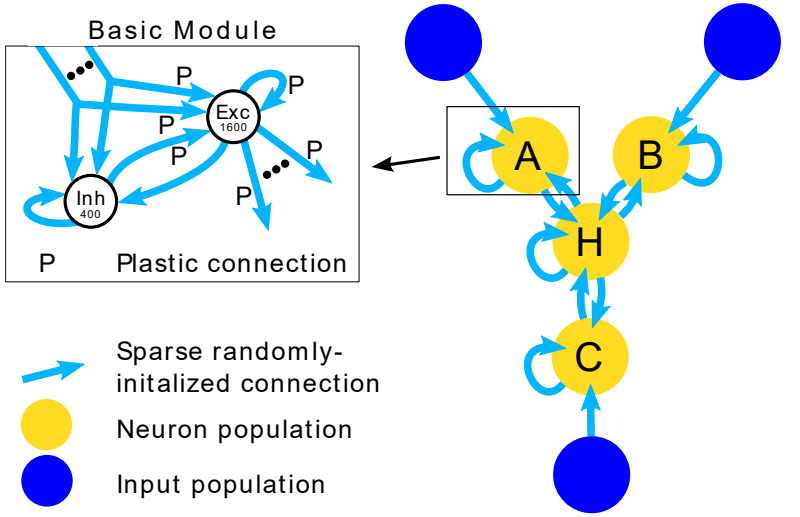

Cook & Diehl, 2016

Neural Principles of behavior and cognition

4. Autonomous learning

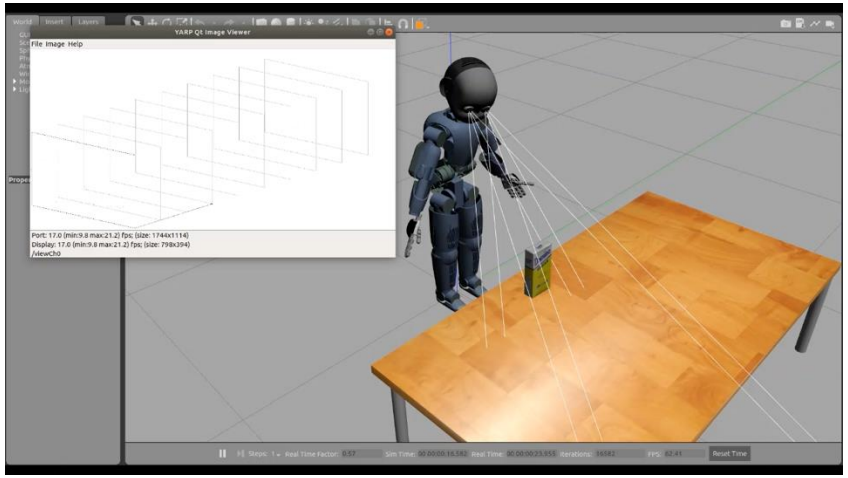

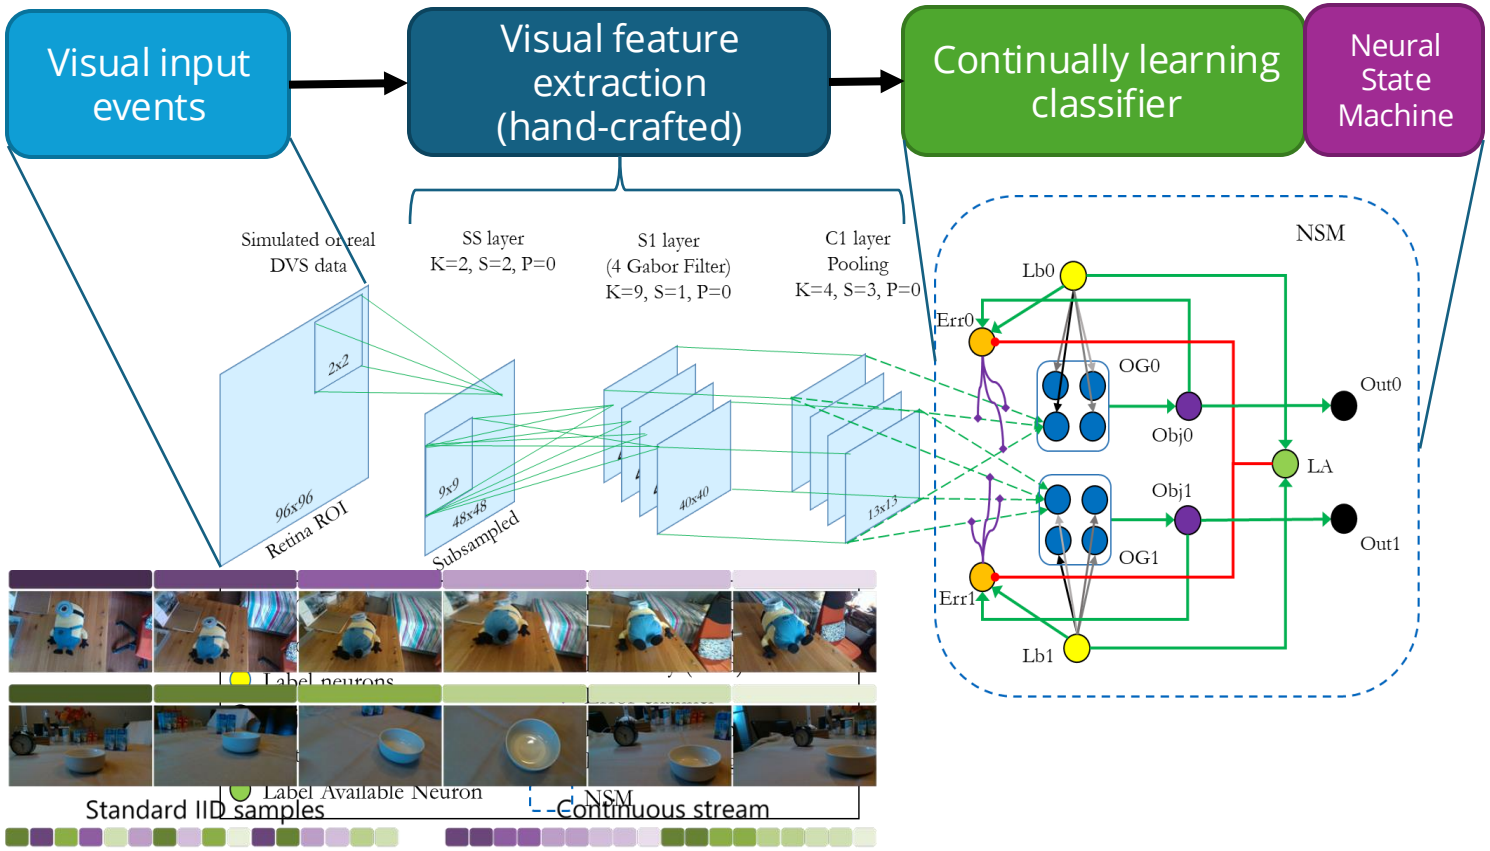

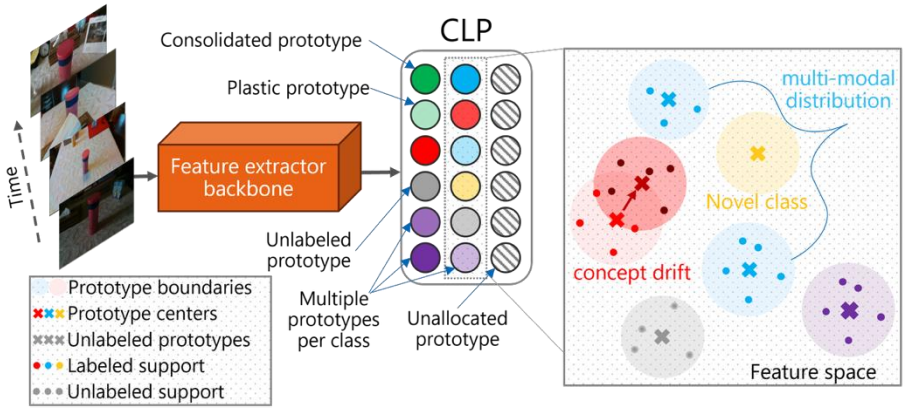

Continually learning prototypes (CLP)

Hajizada et al 2022, ICONS (best paper award)

It all needs to come together

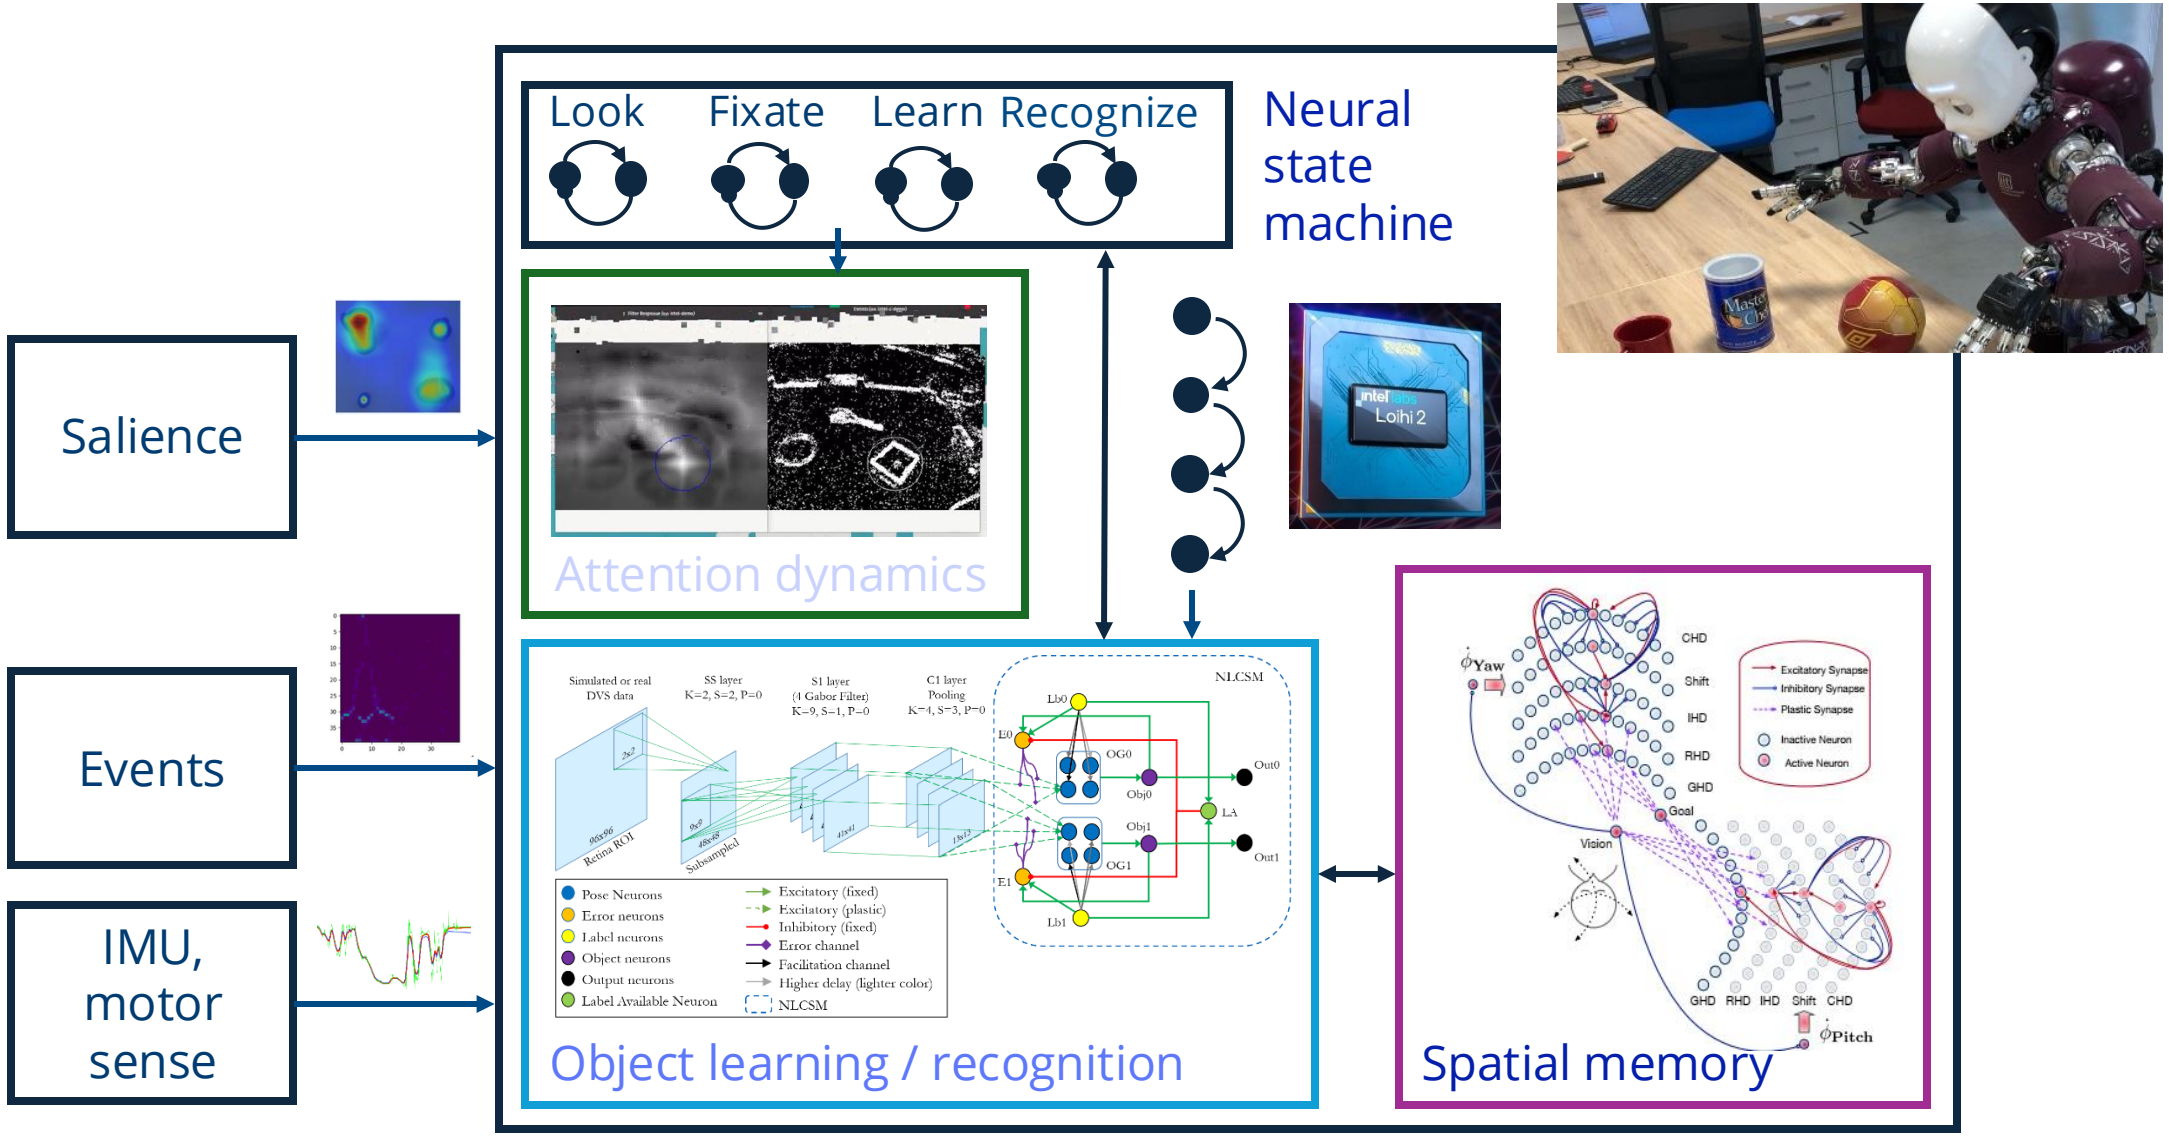

New computing framework!

Neural dynamics (attractors)

- Stateful neurons
- Recurrent networks
  - «Small world»
    - Some sparce, some dense
- Plastic connections
- Structural plasticity

Intentional structure

- Consistency check (in time and space)
- Hierarchy or event-driven processing
- Autonomous learning

Reference frames

- Transformations
- Explicit design of constraints

Hardware-algorithm co-design

- Efficient computing
  - Power
- Fast computing
  - Latency matters
- Efficient learning
  - Data efficiency
  - Adaptivity

The goal

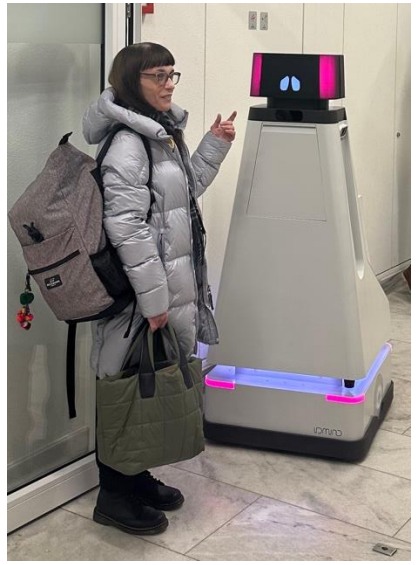

Mobile robots

Mobile manipulation

Enabling embodied AI for assitive robotics

User studies

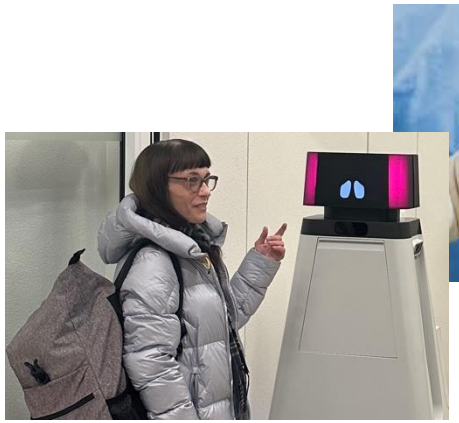

&
Neuromorphic technology

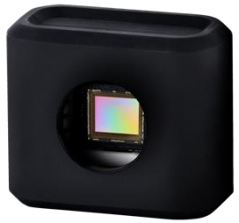

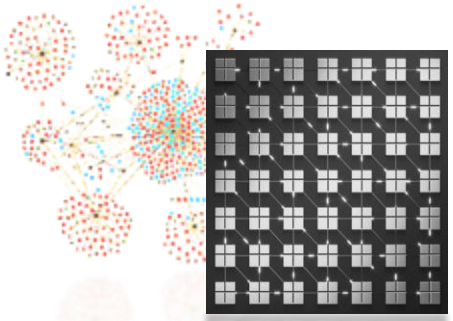

RobotHub

DIZH

Zürcher Hochschule
für Angewandte Wissenschaften

Connect!

Yulia Sandamirskaya

We have examples of neural networks doing better…

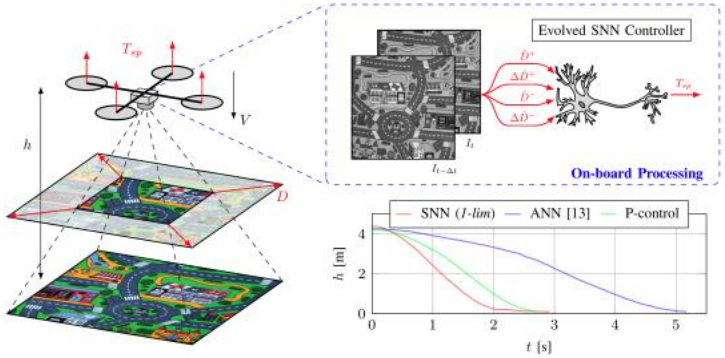

De Croon, 2021

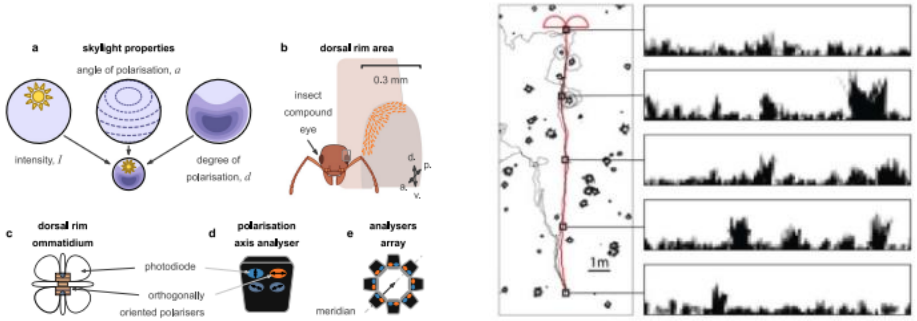

Webb, 2023

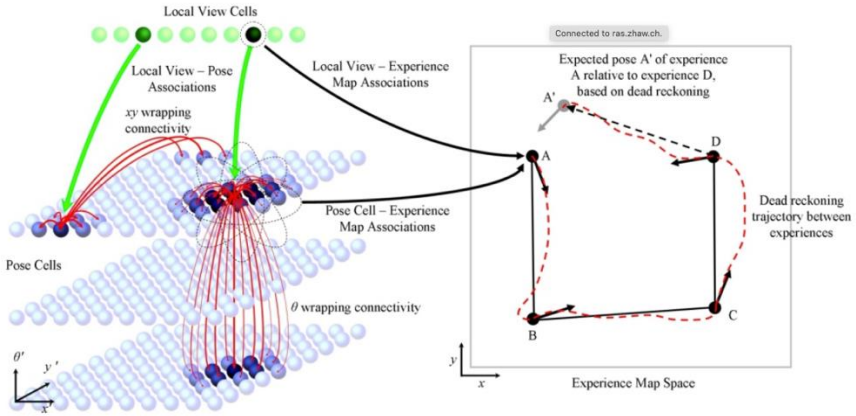

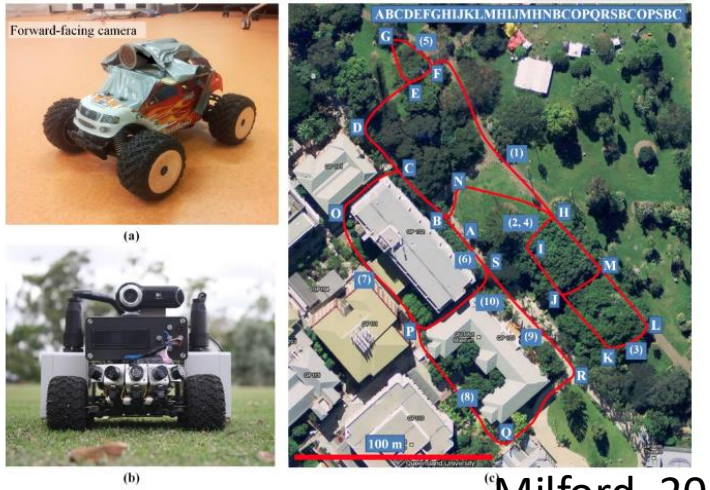

Milford, 2014

Our workplan

Model building

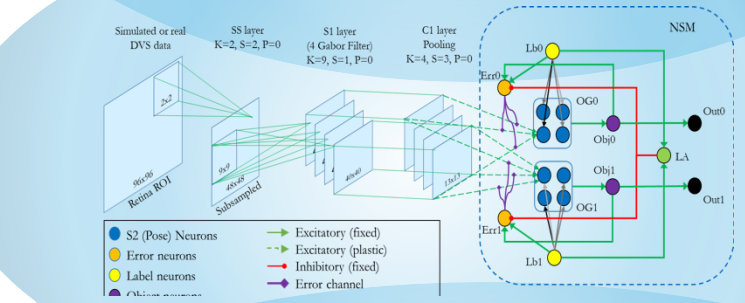

Cognition

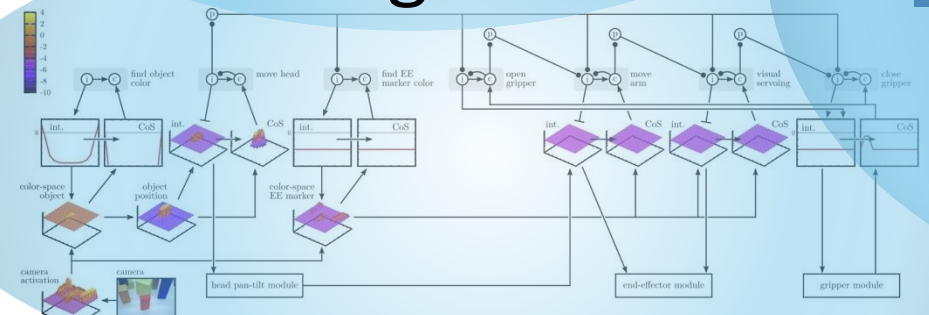

Sensing

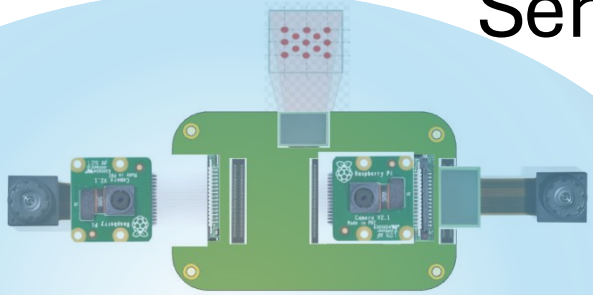

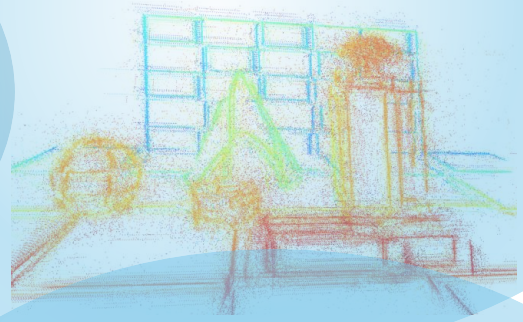

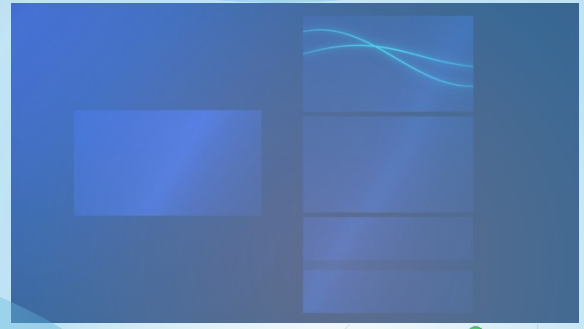

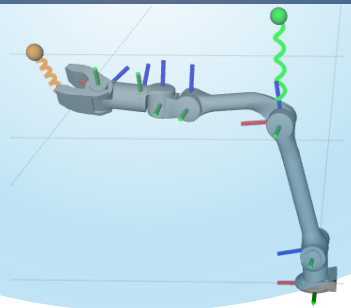

Control

What does it take to make a robot smart?

How can we make humanoids as smart[1] as humans?

Where do robots come from?

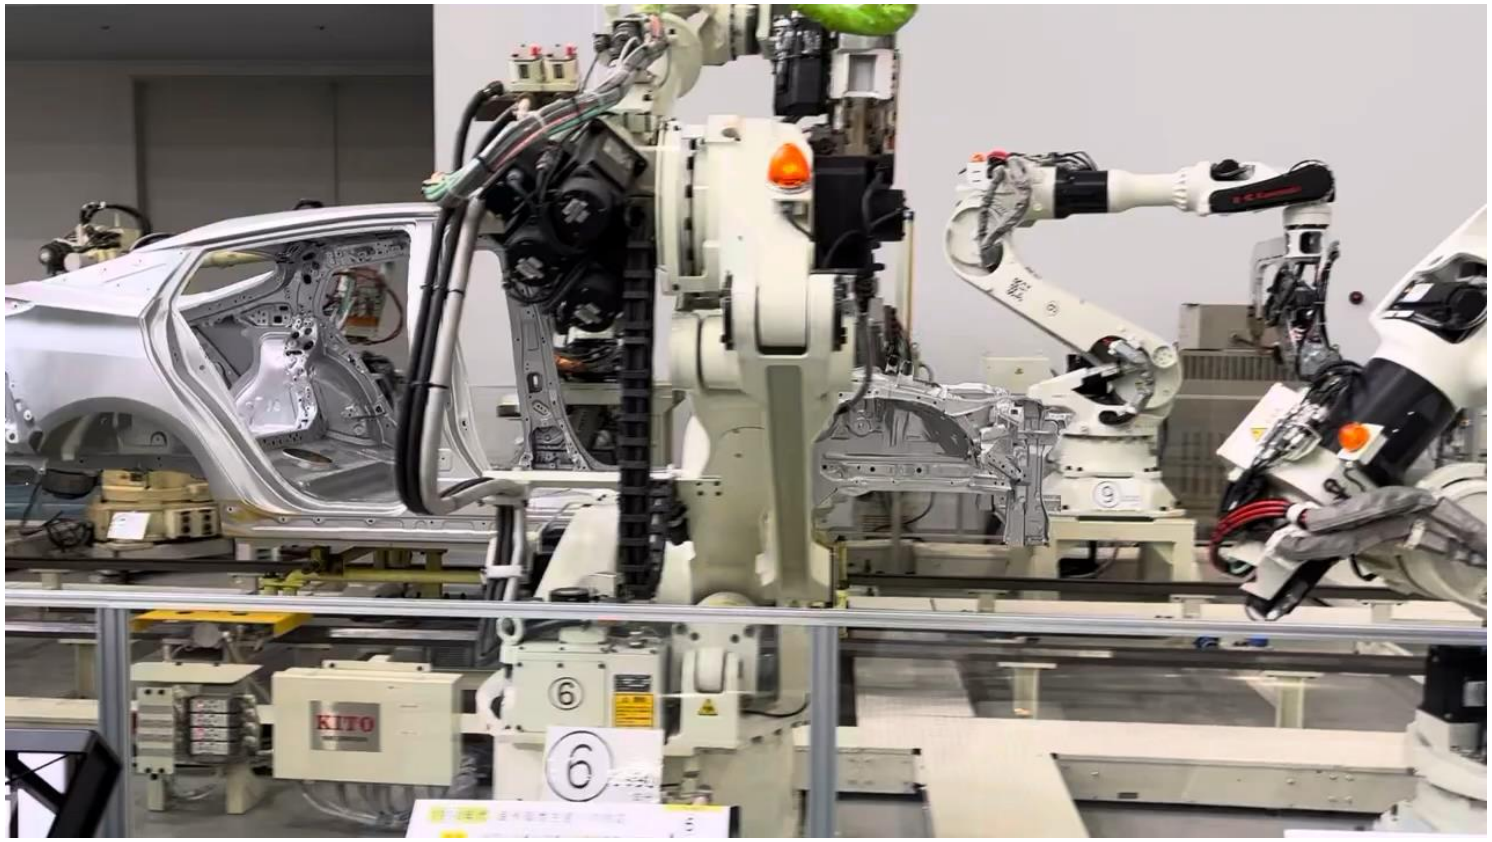

(Autonomotive) Industry

[1]Agile, flexible, able to learn quickly. Not: Wordy, forgetful, hallucinating.

So what is the problem?

Robot control loop:

- <5ms: command update

Vision:

- 30ms: get a frame
- 100ms: detect and track the object

Motion planning:

- 100-300ms

Why humanoids?

Human embodiment is so versatile

"Reminds" of perception, thinking, and motion control, required to solve complex tasks

So, what is the problem?

"Where to move?" for manipulation

Requires "Discrete" perception:

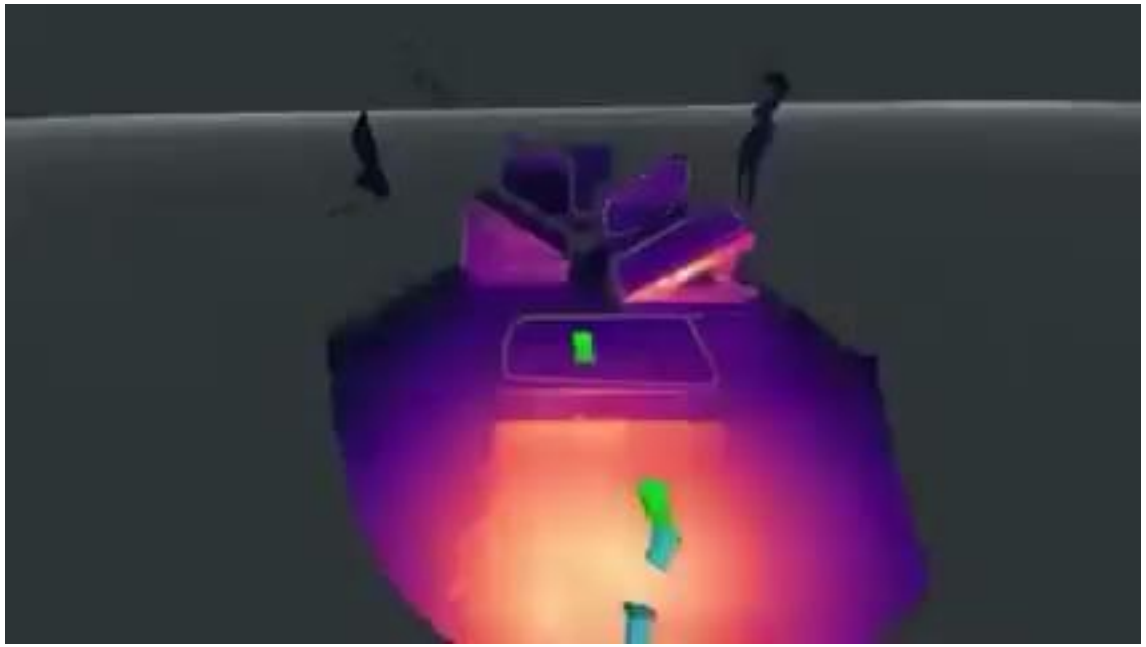

- Need to detect, "recognize" objects
- Need to locate them in 3D
- And estimate their pose
- Many different objects
- Select the best / right target

We need "AI"

What can AI do today?

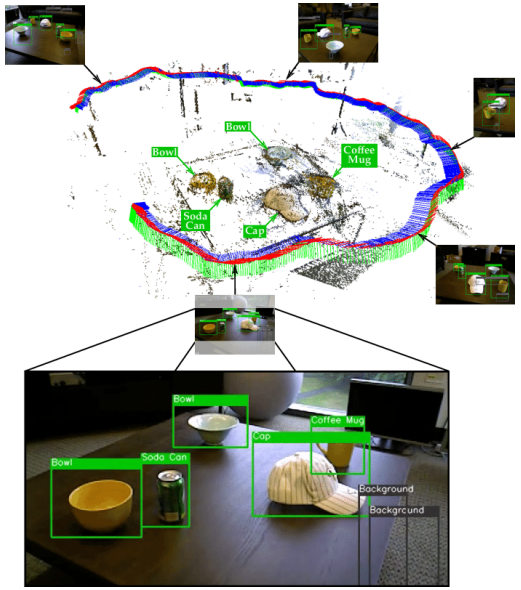

Assistive robots require “Discrete” Sensing

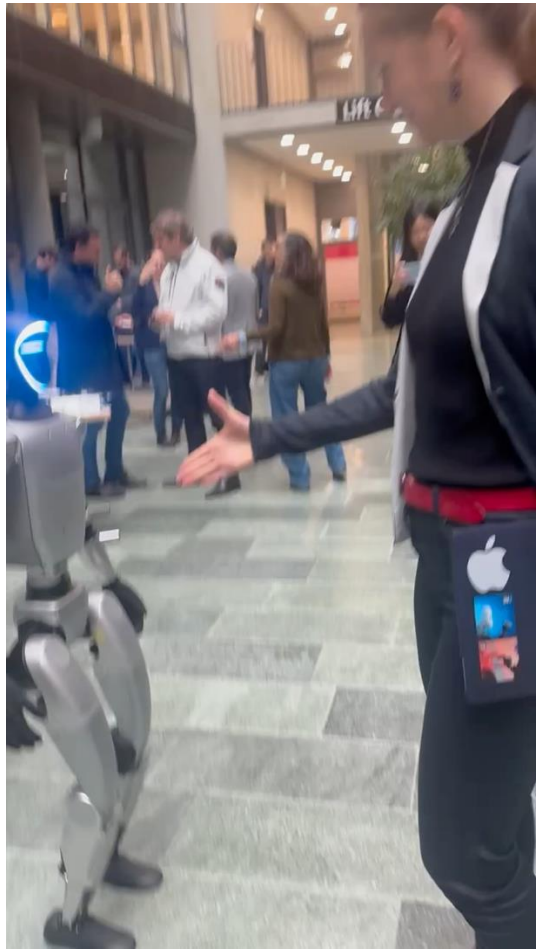

November 2024

- Is this an object? Is this an object I know? Is this an object I was looking for?
- Can I estimate its size/weight, position, pose?
- Do I know how this object can be manipulated?
- Is this a new object I should learn?

→ “Intelligence”!

Artificial Intelligence today

The trend: end-to-end

Image → 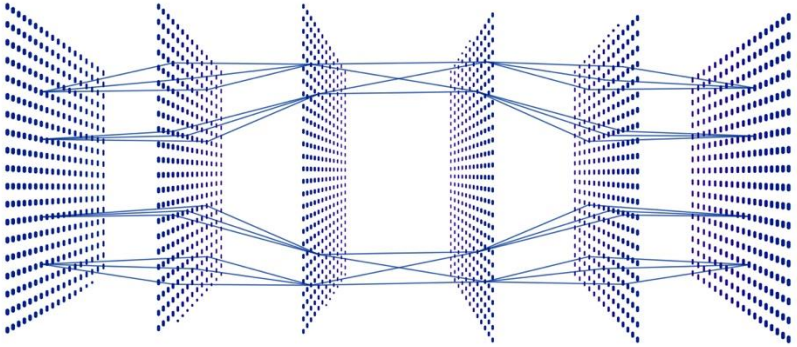 → Behavior

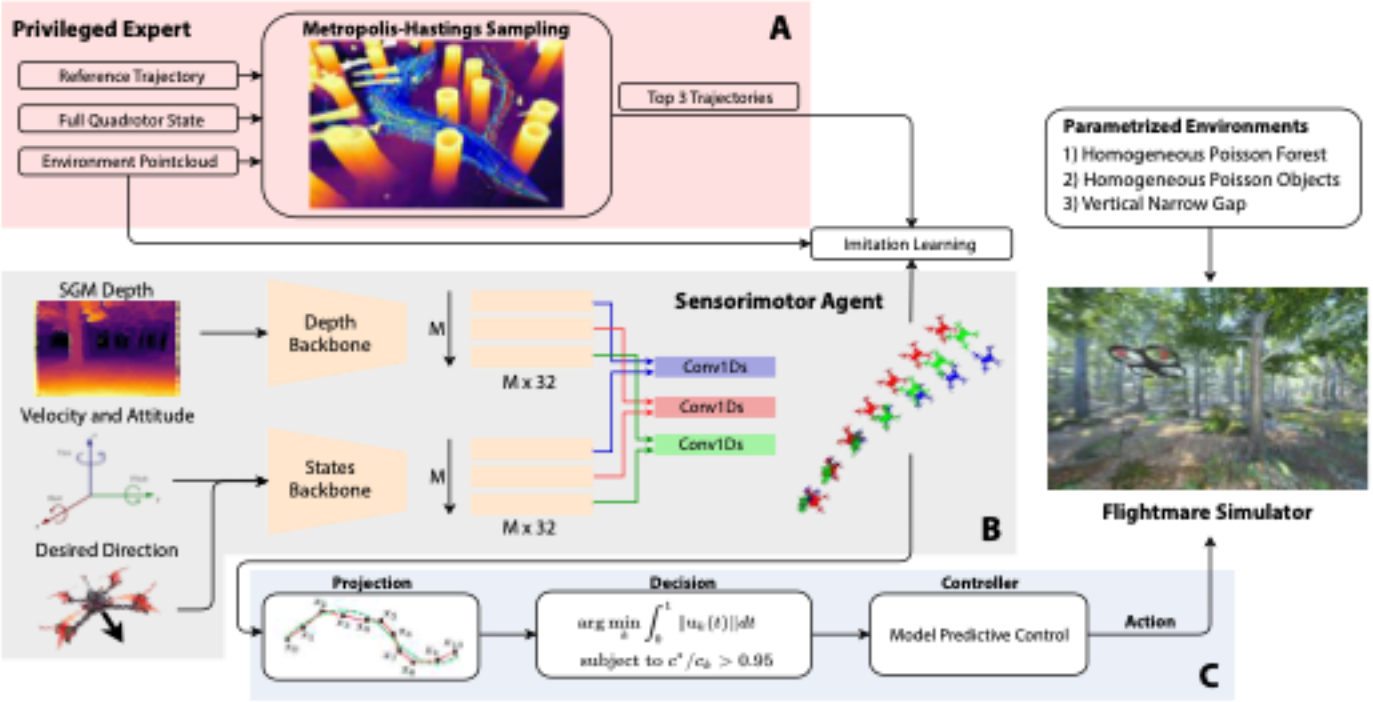

Successful example:

- DNN to control a drone

Loquercio, A., Kaufmann, E., Ranftl, R., Müller, M., Koltun, V., & Scaramuzza, D. (2021). Learning high-speed flight in the wild. *Science Robotics*, *6*(59), eabg5810.

Artificial Intelligence today

The trend: end-to-end

Image → 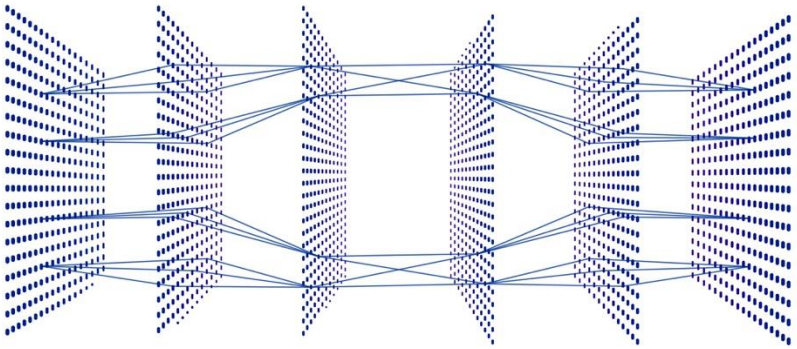 → Behavior

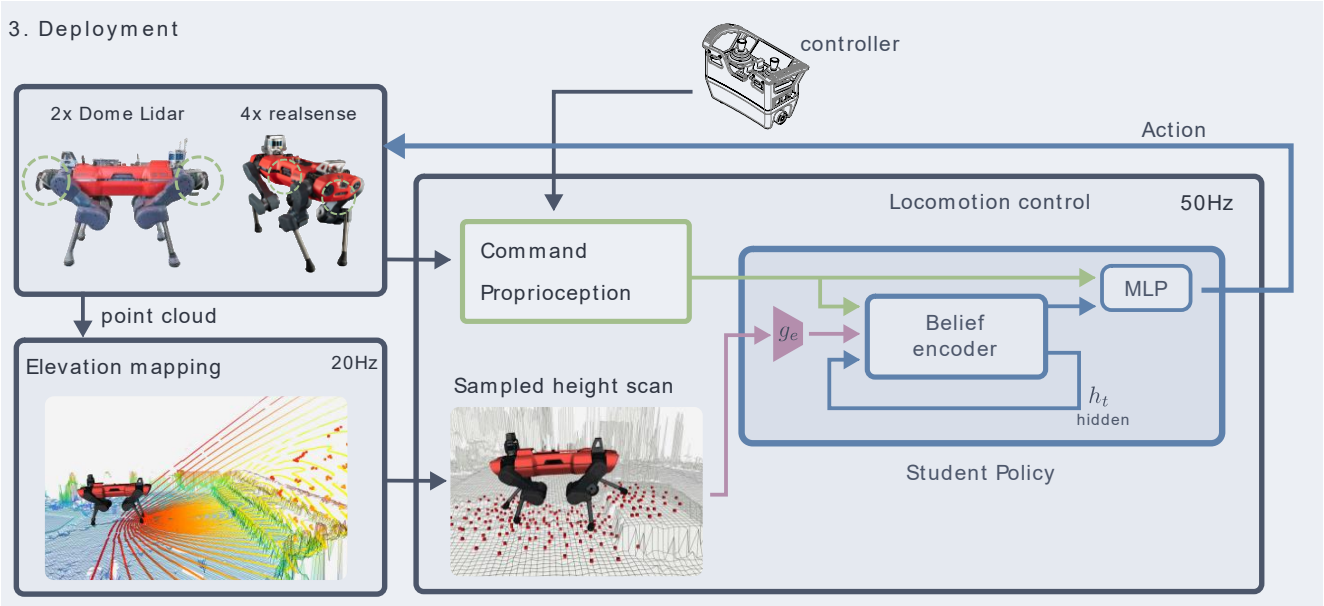

Successful example:

- Skills on a quadruped

Miki, T., Lee, J., Hwangbo, J., Wellhausen, L., Koltun, V., & Hutter, M. (2022). Learning robust perceptive locomotion for quadrupedal robots in the wild. *Science robotics*, *7*(62), eabk2822.

Is this the only way?

300,000x more power (6 years) vs **8x** by Moore's Law

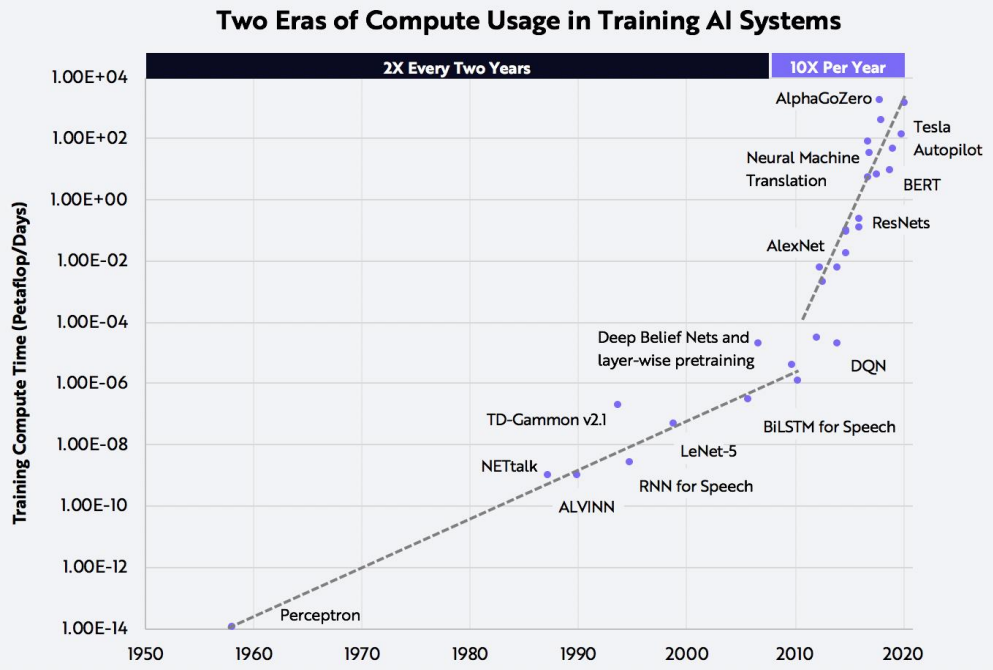

Chat GTP-4 Training:

100 Million Dollar

1.287 Gigawatt hours

??? Human labor

Transparency

Robustness

Adaptivity

Source: https://techhq.com/2023/03/data-center-energy-usage-chatgpt

Source: ARK Investment Management LLC, "AI and Compute." OpenAI, https://arkinv.st/2ZOH2Rr.

Neuromorphic Computing Hardware

CPU

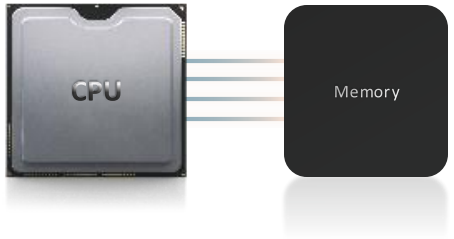

GPU

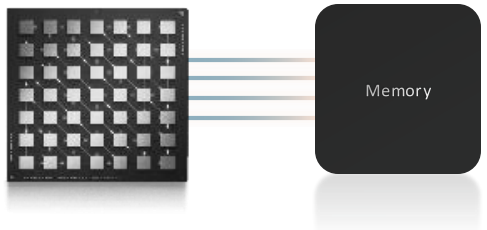

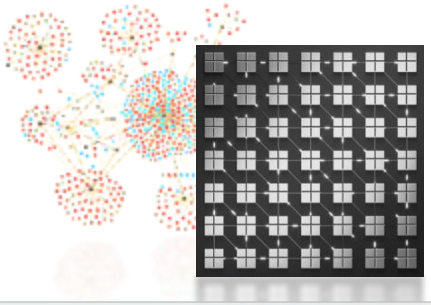

NCH

SEQUENTIAL THREADS OF CONTROL	PARALLEL DENSE COMPUTE	PARALLEL SPARSE COMPUTE
SYNCHRONOUS CLOCKING	SYNCHRONOUS CLOCKING	ASYNCHRONOUS EVENT-BASED SPIKES
PROGRAMMING BY ENCODING ALGORITHMS	OFFLINE TRAINING USING LABELED DATASETS	LEARN ON THE FLY THROUGH LOCAL LEARNING RULES

```
if X then
  ...
else
  ...
```

01100
11010
00100

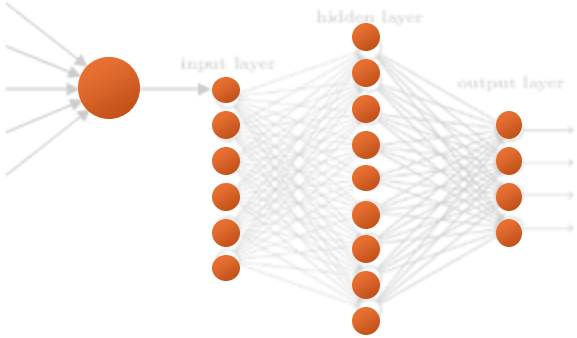

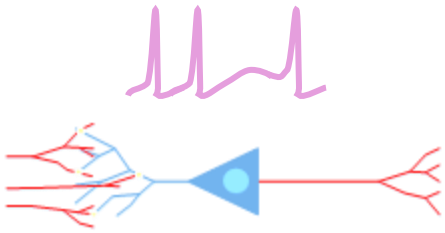

Neuromorphic Computing Hardware

CPU

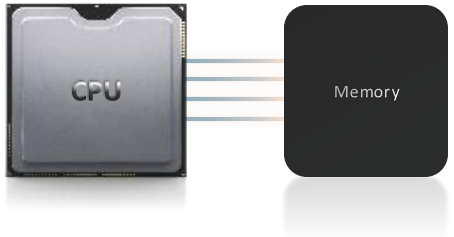

GPU

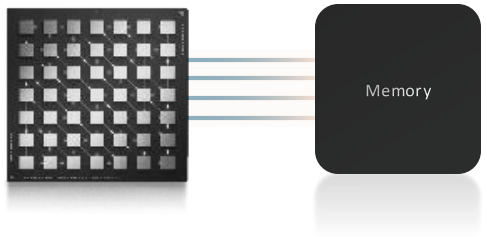

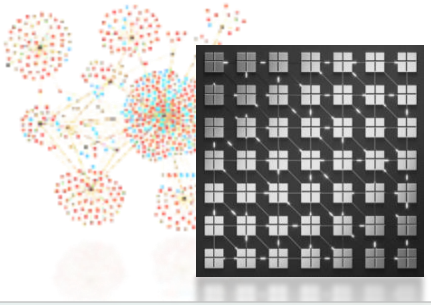

NCH

CPU	GPU	NCH
SEQUENTIAL THREADS OF CONTROL	PARALLEL DENSE COMPUTE	PARALLEL SPARSE COMPUTE
SYNCHRONOUS CLOCKING	SYNCHRONOUS CLOCKING	ASYNCHRONOUS EVENT-BASED SPIKES
PROGRAMMING BY ENCODING ALGORITHMS	OFFLINE TRAINING USING LABELED DATASETS	LEARN ON THE FLY THROUGH LOCAL LEARNING RULES

```
if X then
  ...
else
  ...
```

01100
11010
00100

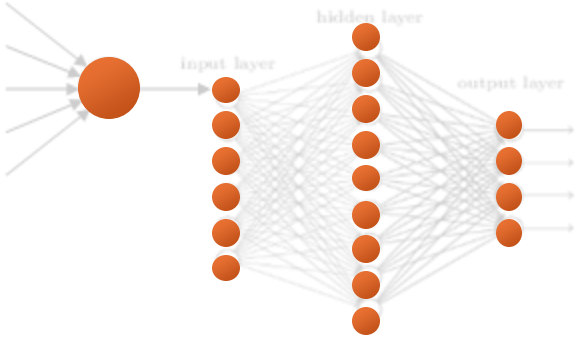

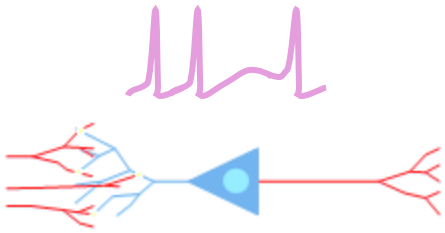

Technology: The body

Mechanics and motion control

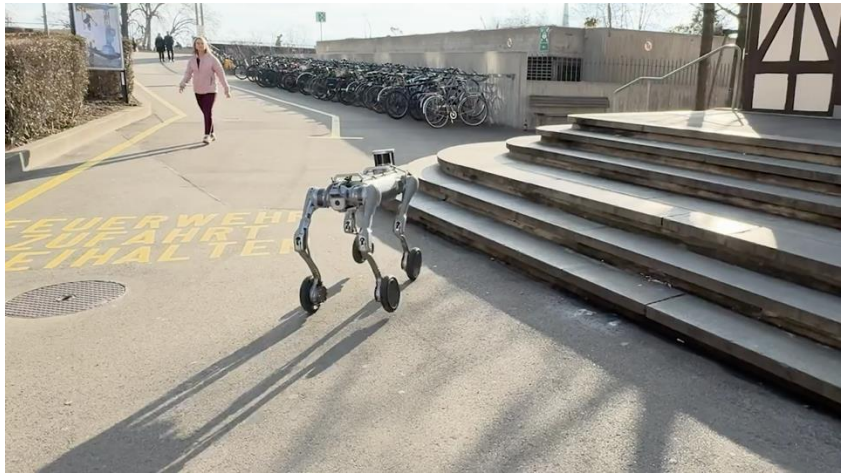

Sensing: “continuous” perception

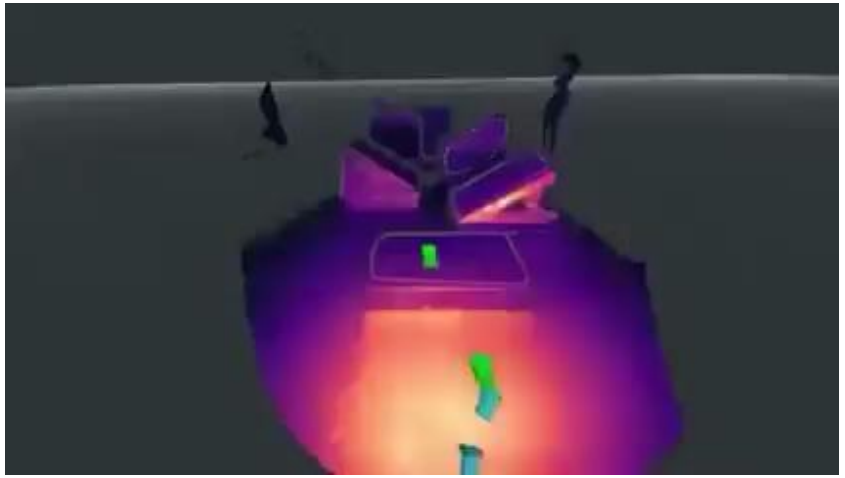

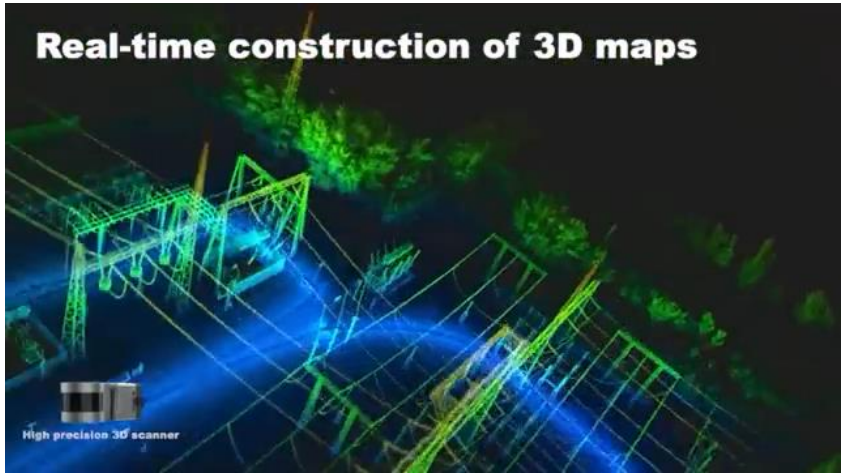

Technology for assistive robotics: The brain and classical robotics

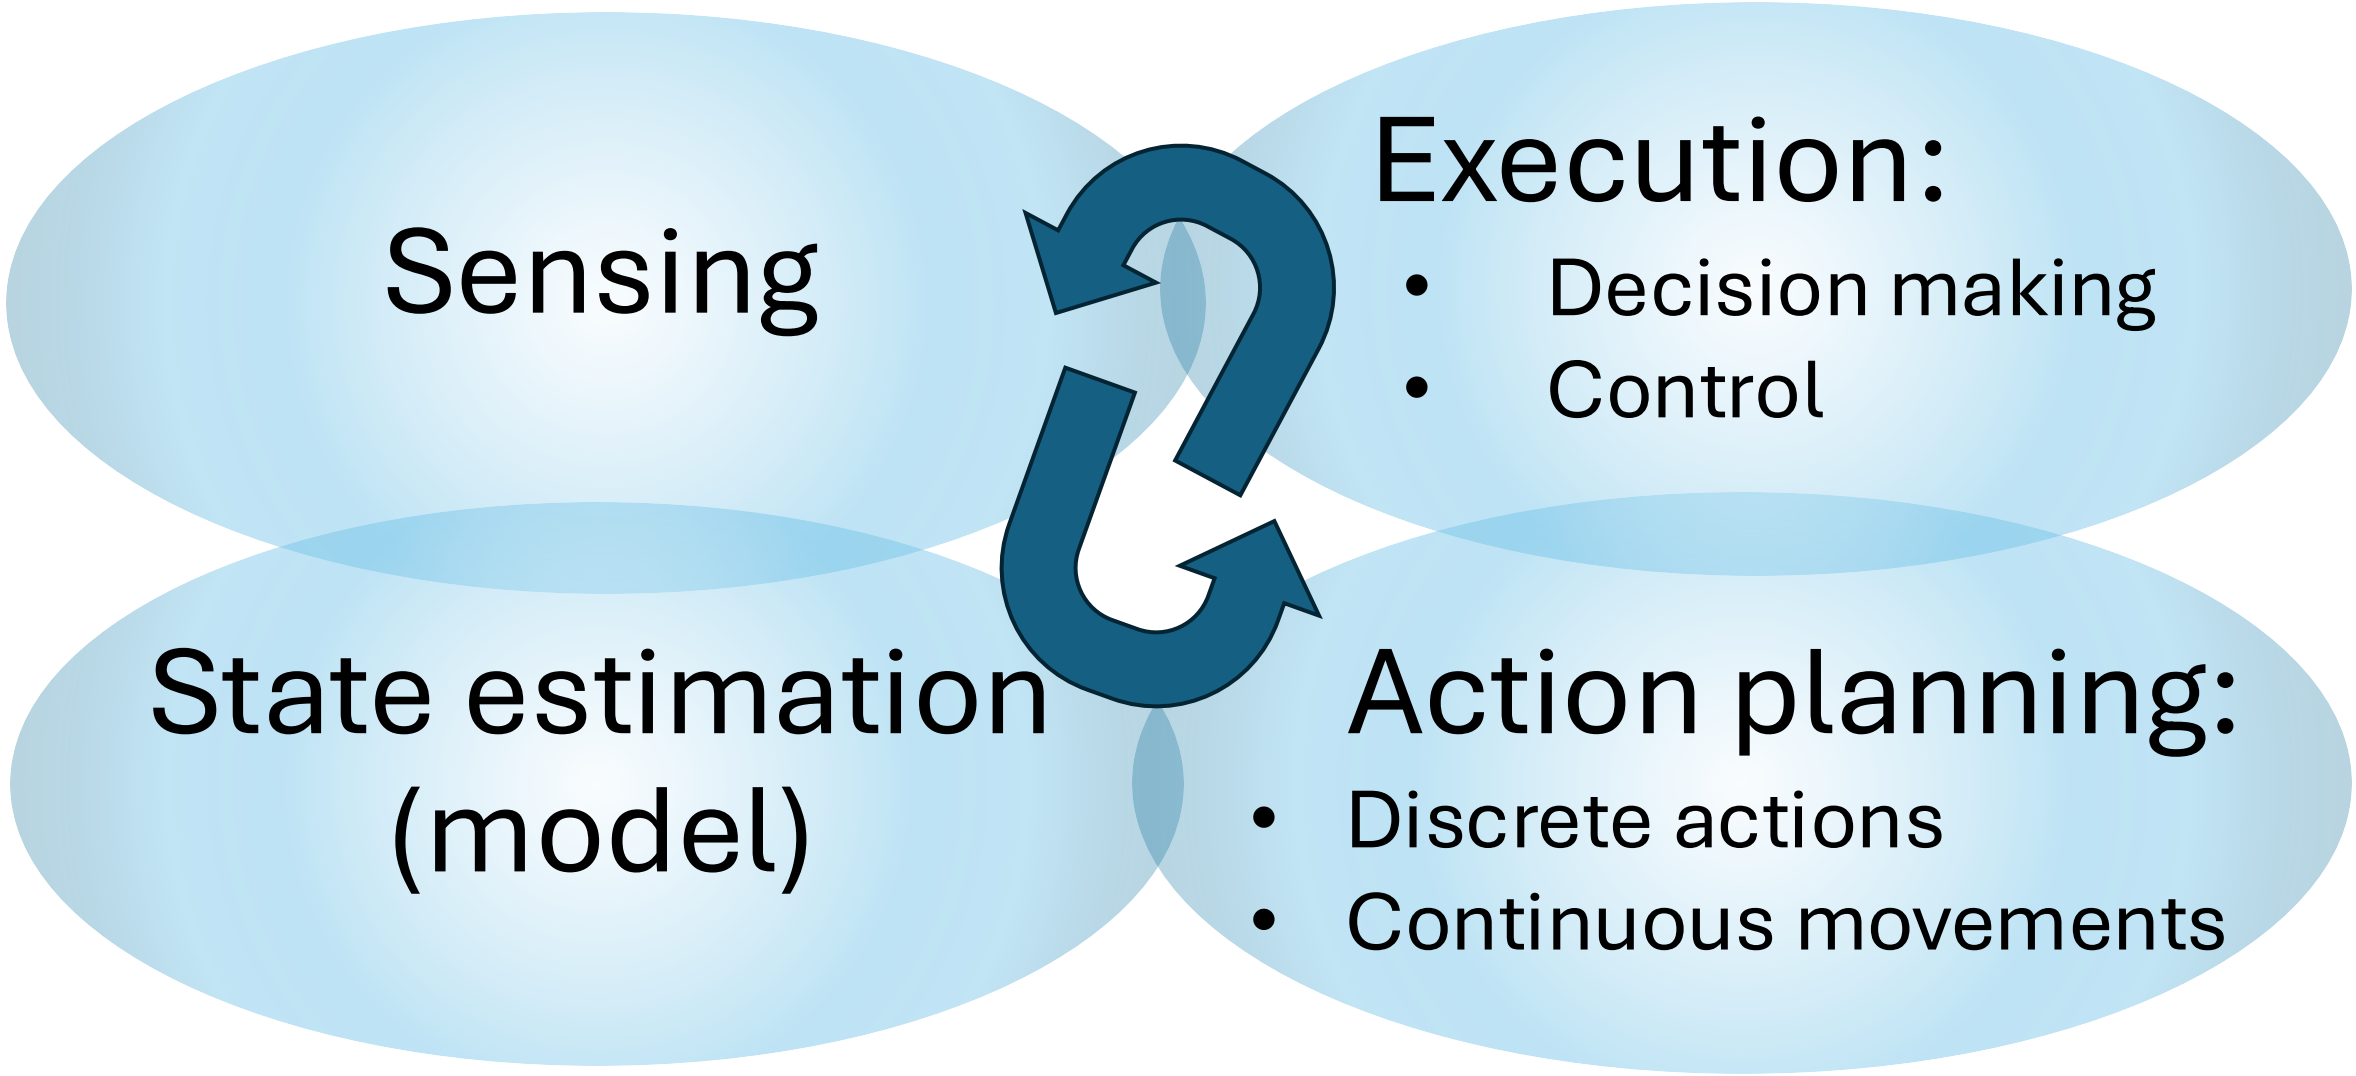

How do brains compute?

1. Cortical Layers

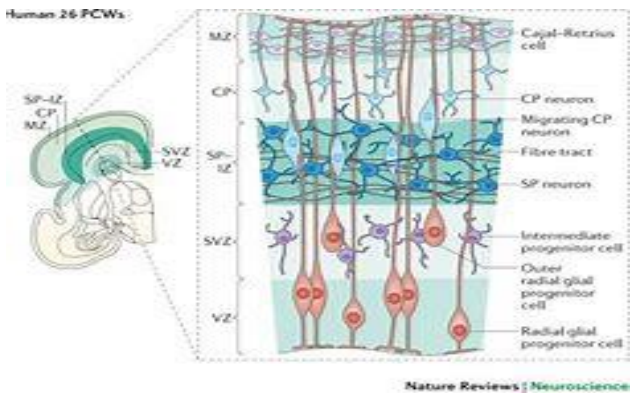

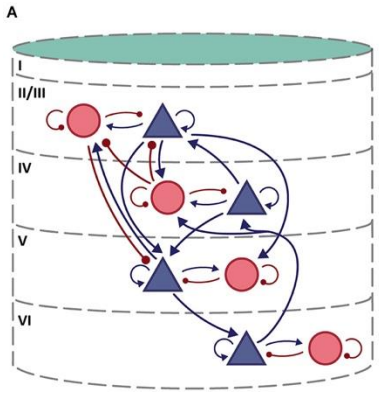

2. Thalamo-cortical and Cortical-BG loops

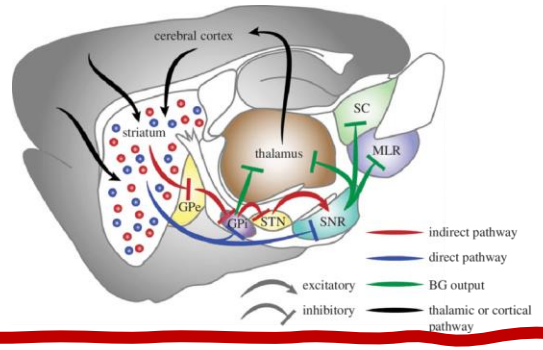

3. Cerebellum

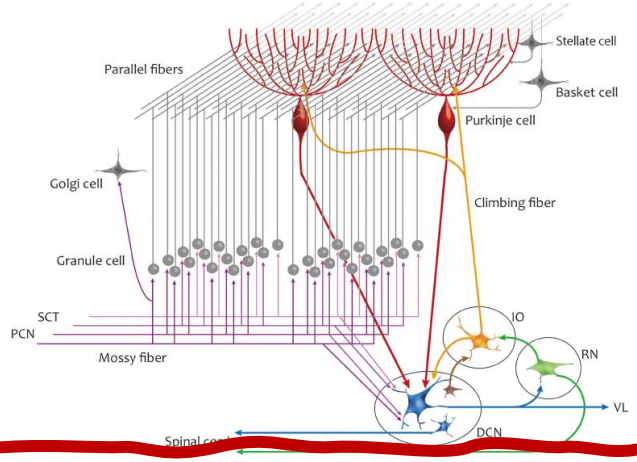

4. Hippocampus

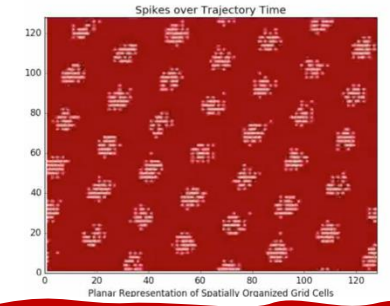

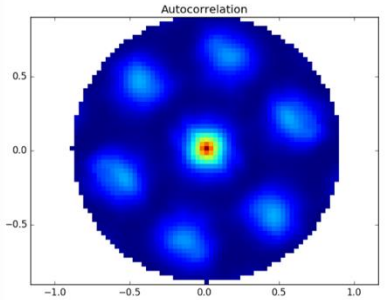

- Many different "algorithms"
- Continual learning and adaptation

Technology: inspired by the (modern understanding of the) brain

Structure of the cortex

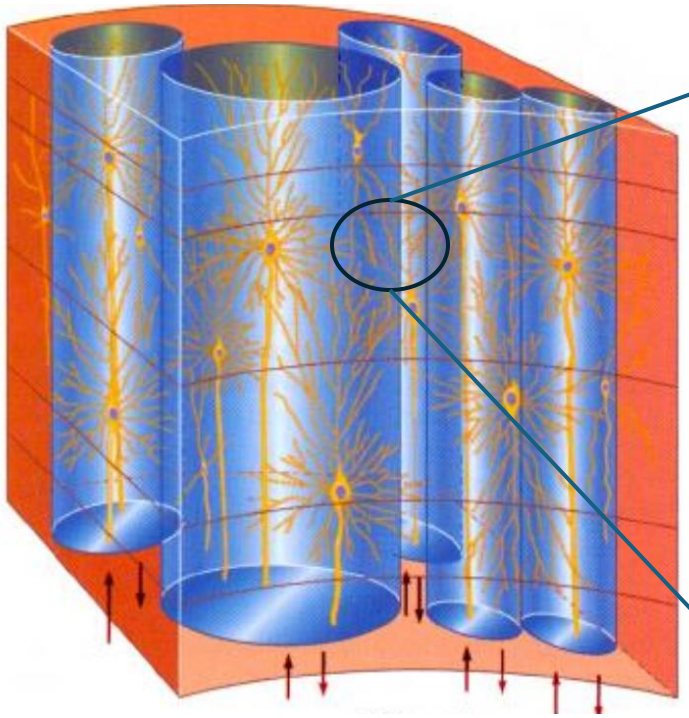

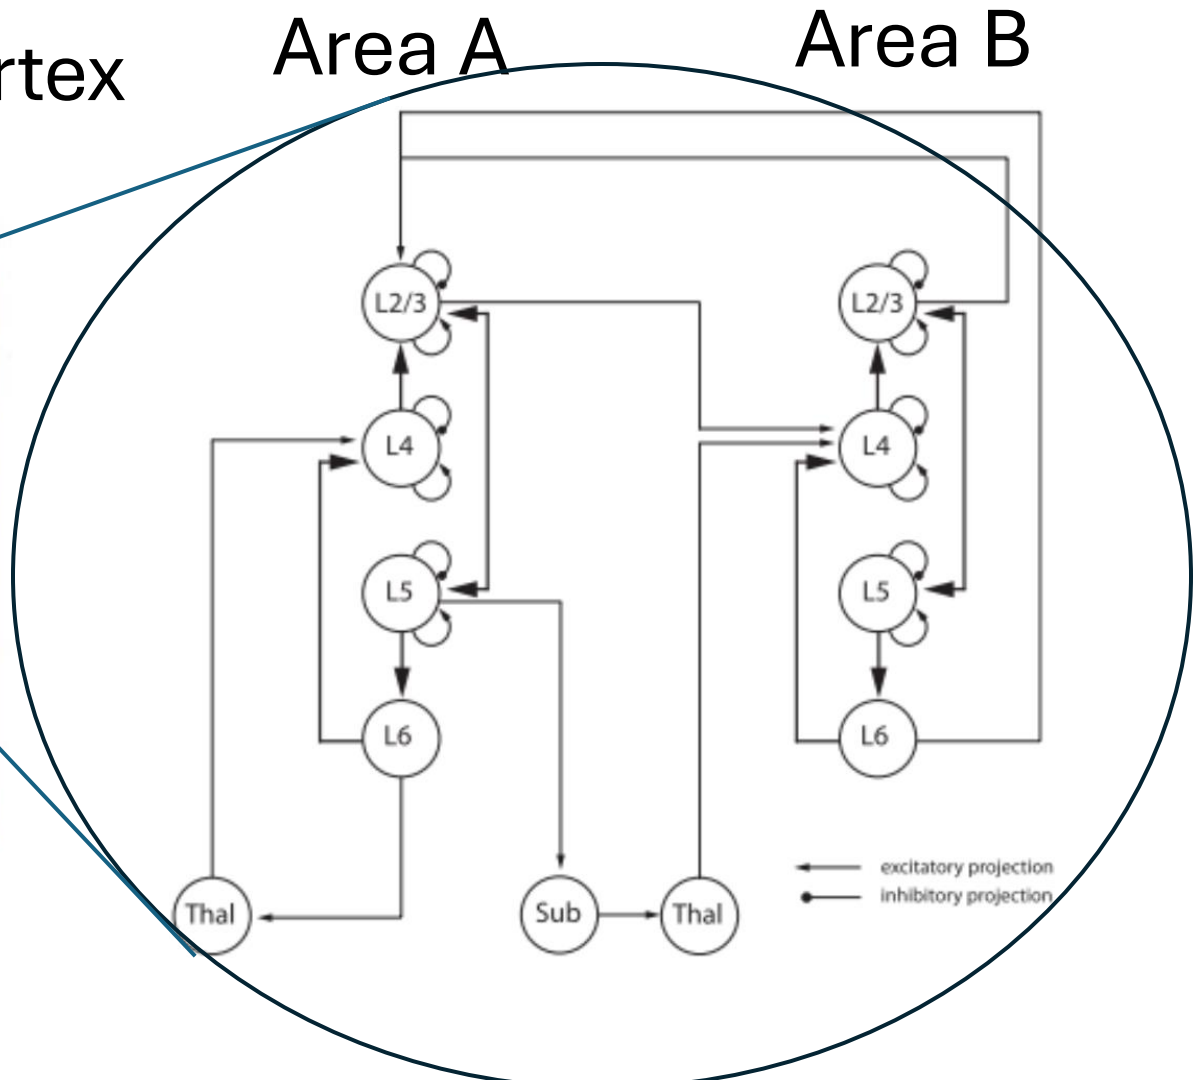

Technology: inspired by the (modern understanding of the) brain

Structure of the cortex

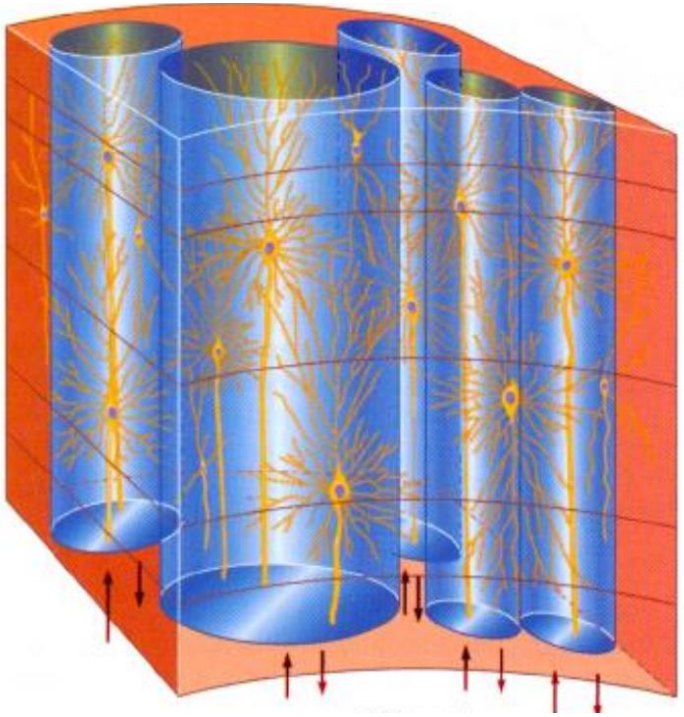

Visual cortex

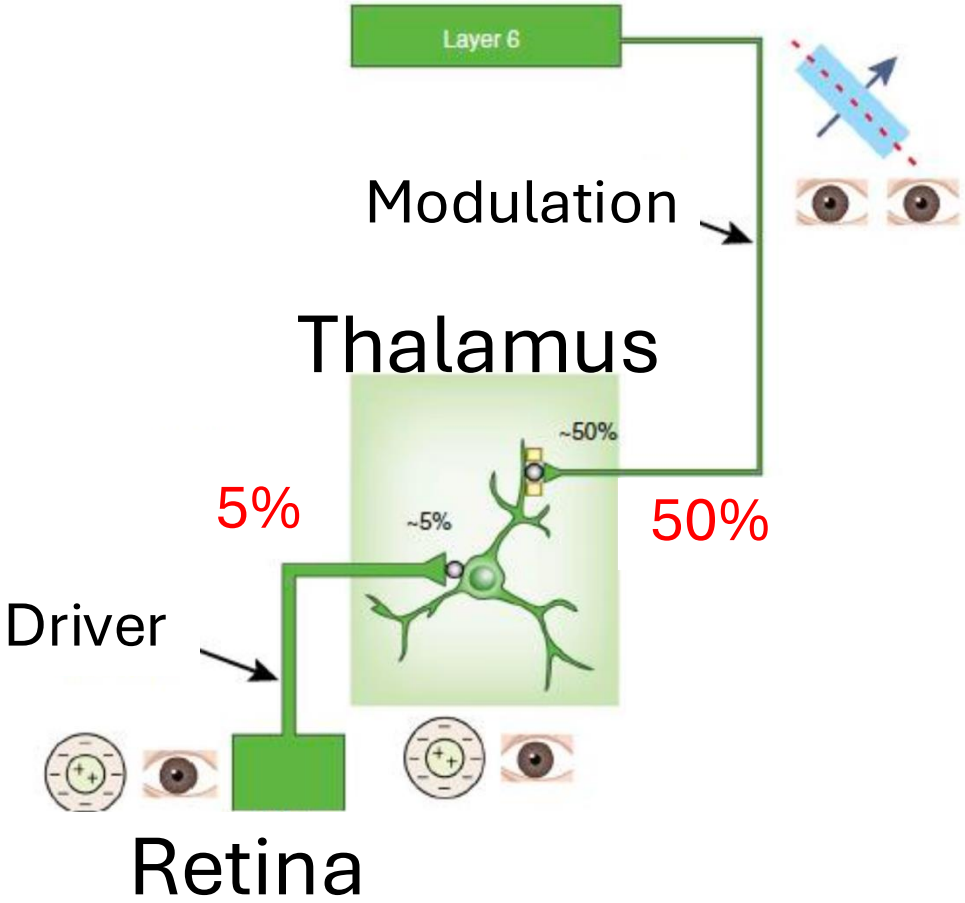

All a bit like MPC
(Model-predictive control)

Predict trajectory based on the model

Choose the best trajectory (optimization)

Sensing and state estimation

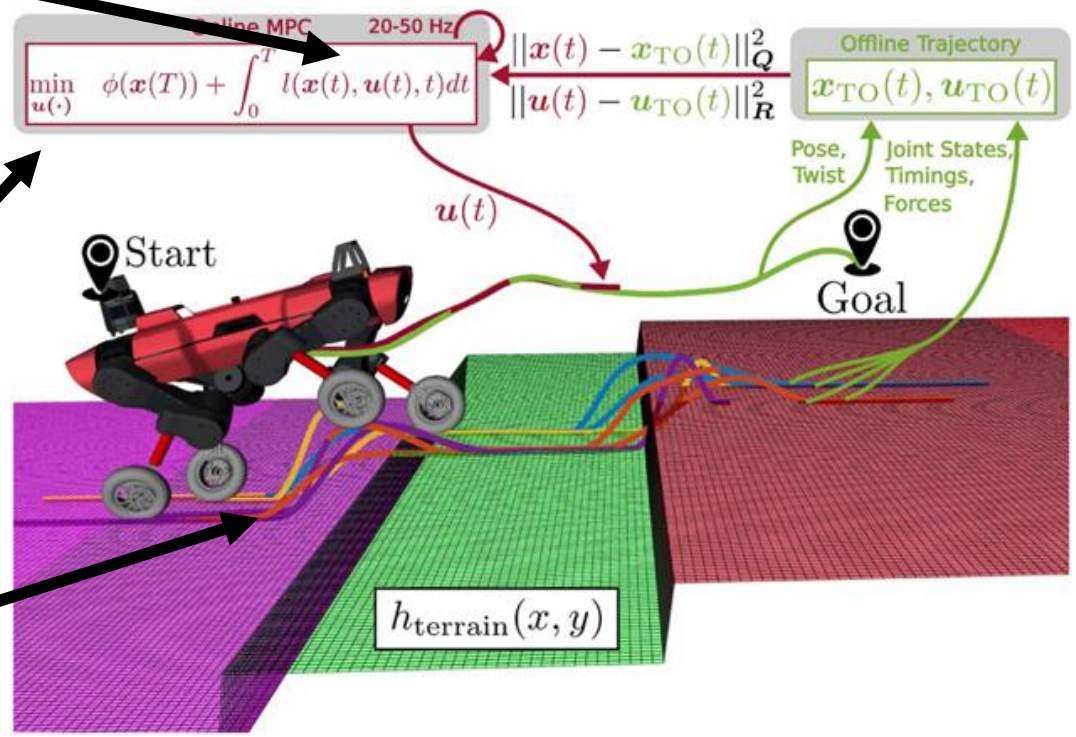

All a bit like MPC
(Model-predictive control)

Predict trajectory based on the model

Choose the best trajectory (optimization)

Sensing and state estimation

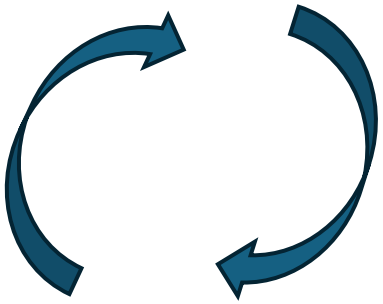

We need many such loops (modules)

A bit like Brook's subsumption architecture

Rodney Brooks: brought the first assistive robot on the market (iRobot; Roomba)

A bit like Brook's subsumption architecture

Modularity

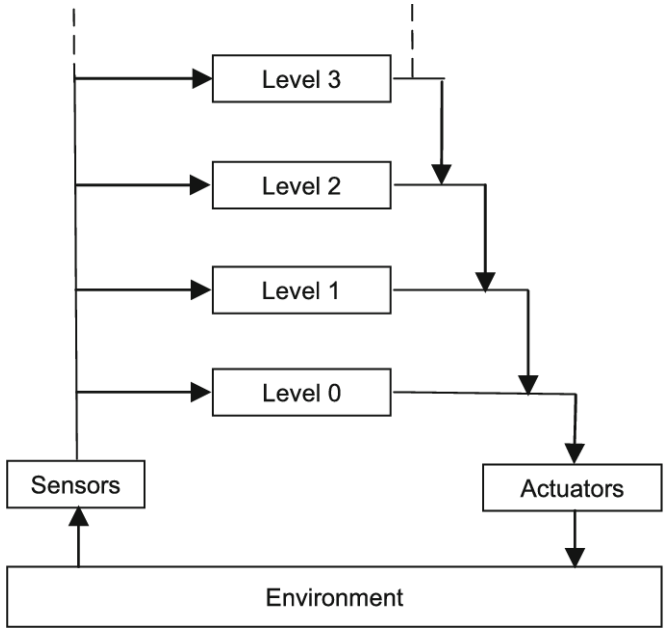

Multiple behavioral modules

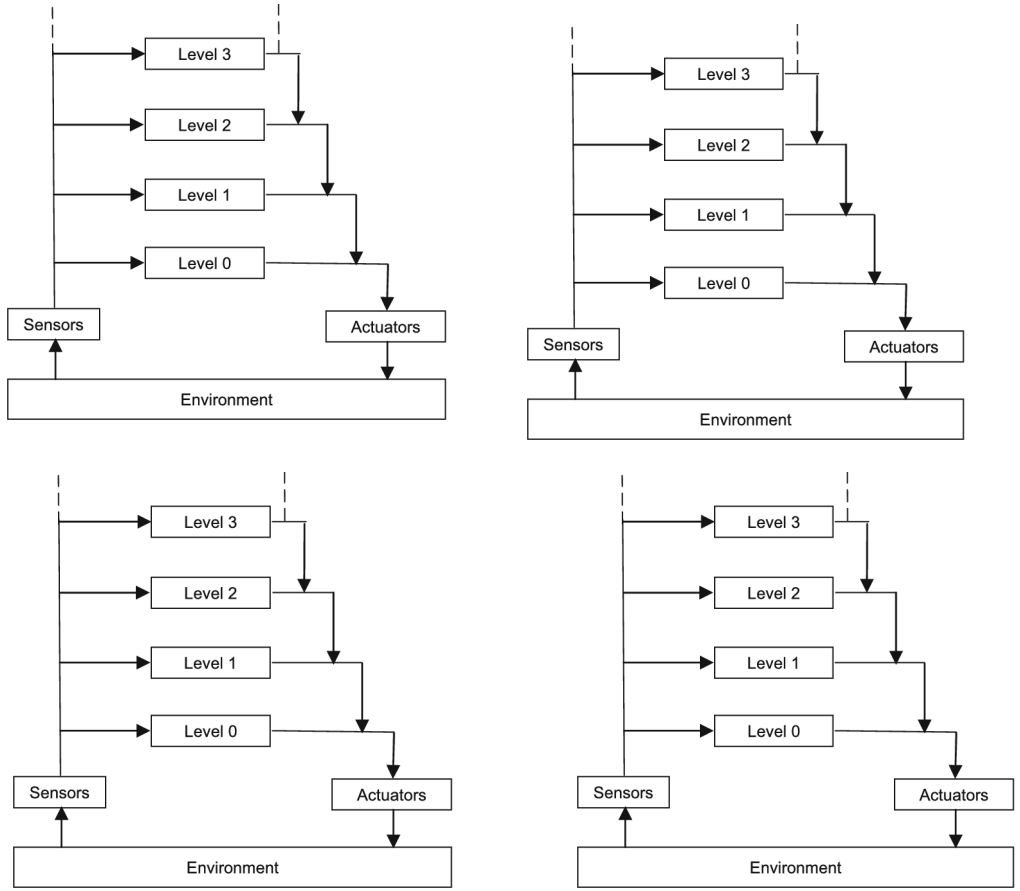

Success of today's AI relies on hardware!

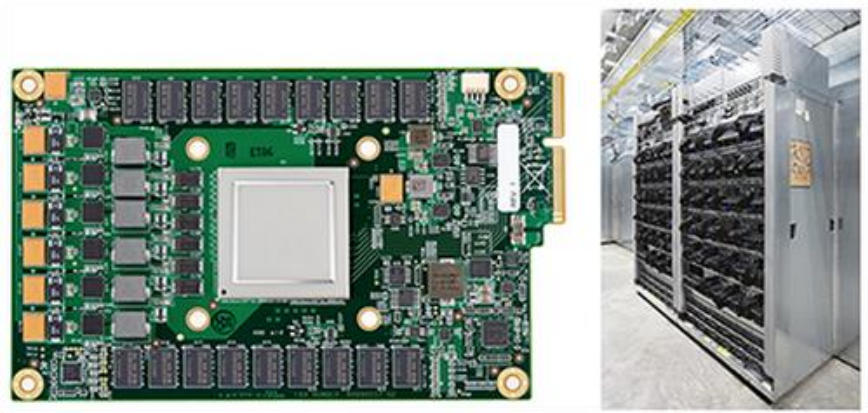

- Optimized for matrix multiplication
- Optimized for batched data processing
- Optimization of number representation

Neuromorphic chips

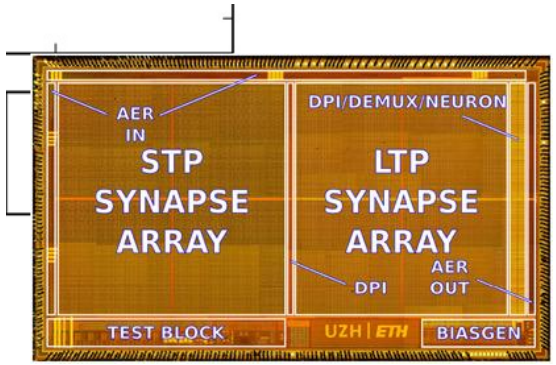

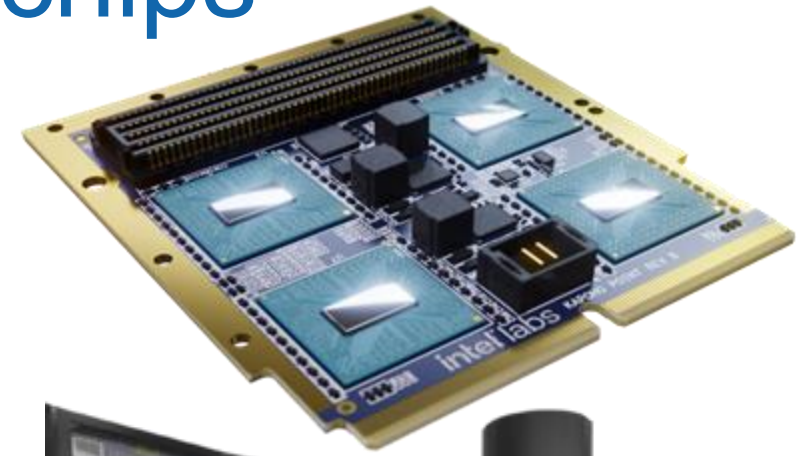

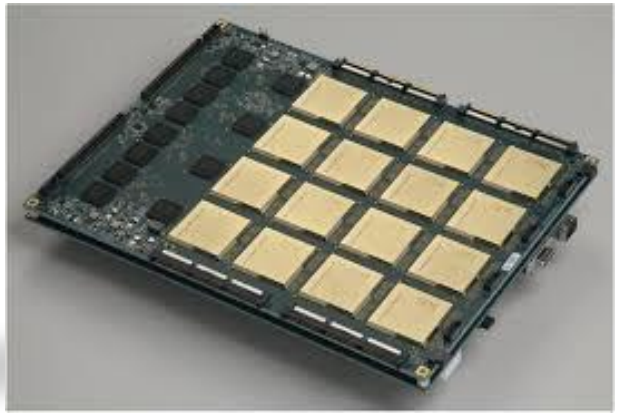

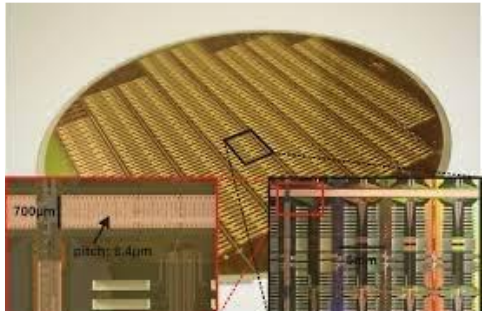

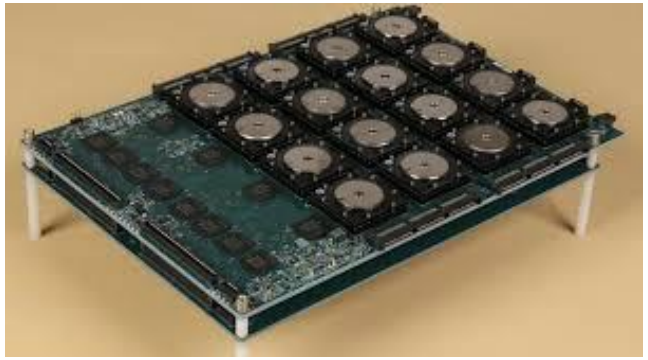

Hardware and algorithm need to match:

- Granular parallelism (local memory)
- Asynchronous processing
- Heterogeous platform

Success of today's AI is rooted in a match with the computing hardware (GPU)

Recurrent networks with novel bio-inspired properties give the best gains

Reference architecture

- CPU (Intel Core/Xeon)
- GPU (Nvidia)
- Movidius (NCS)
- TrueNorth

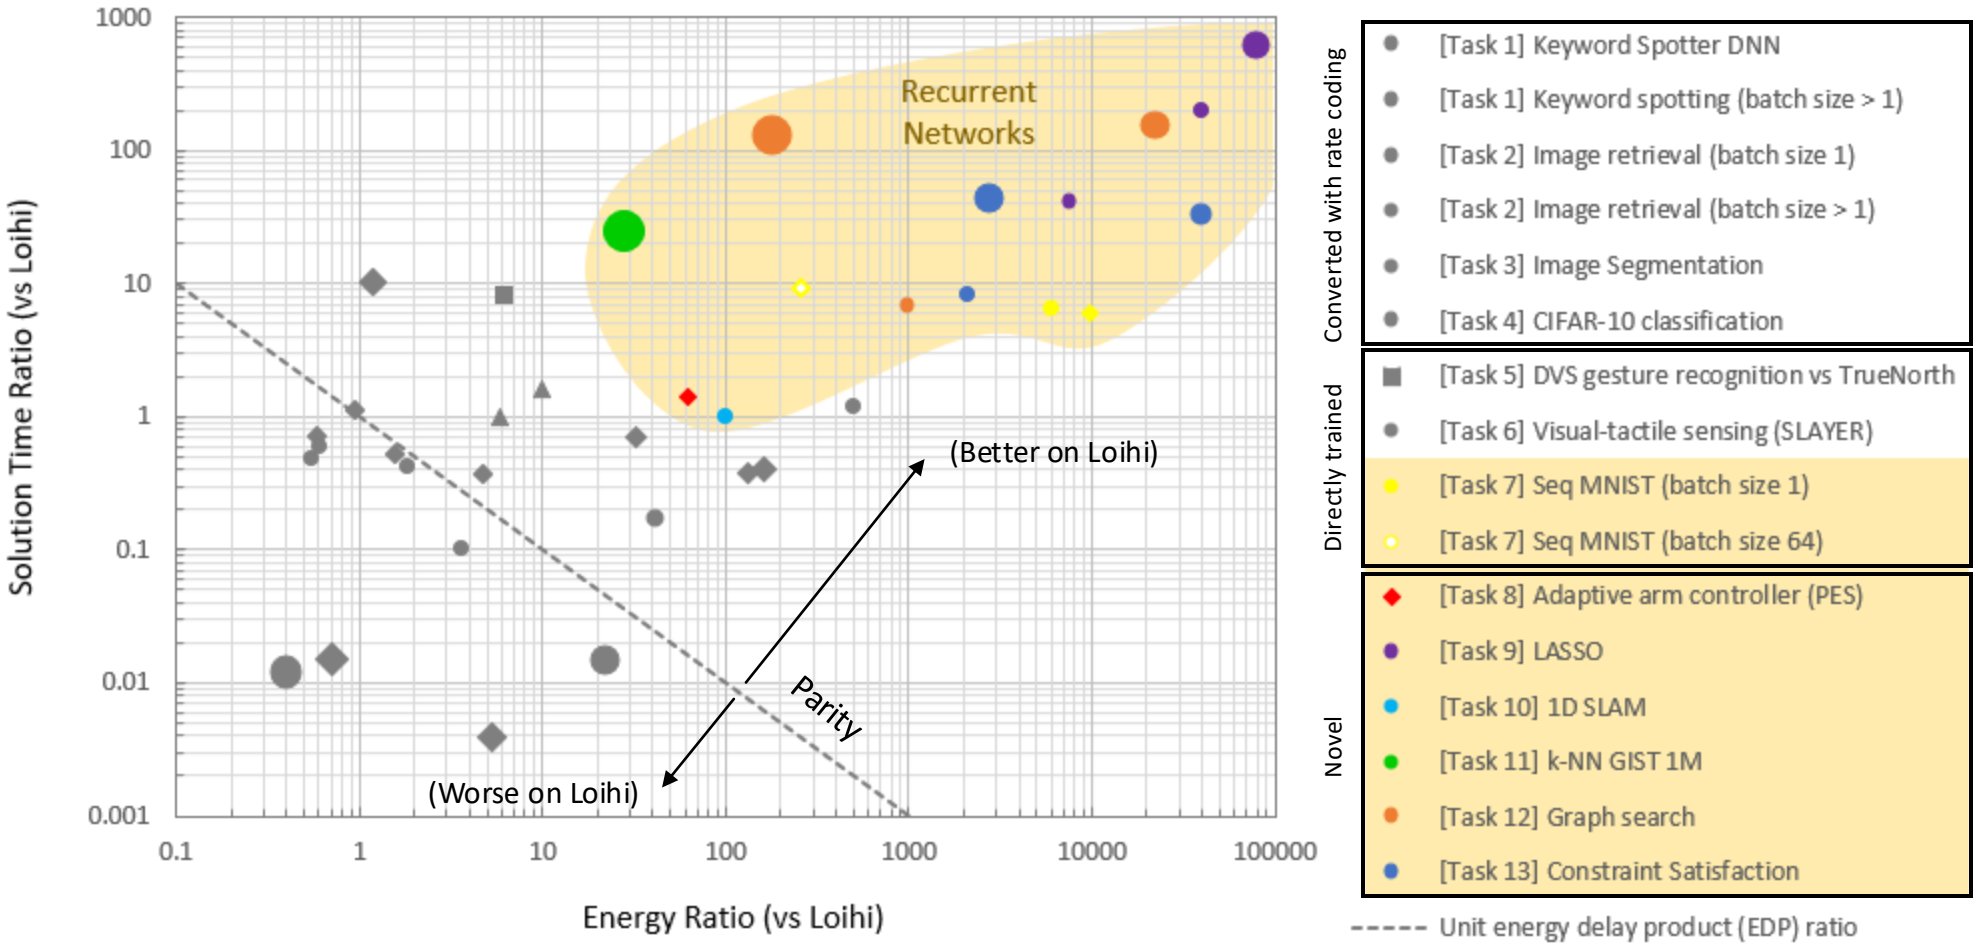

Davies et al, Proc. IEEE 2021. Results may vary.

New theory of embodied AI

- Control loops for different behaviors
  - Sensing needs to be fast for control
- Adaptation
  - Self-calibration
  - Learning the model
  - One-shot (fast) learning of new things
- Scaling up: modular approach
  - Coordinate frames between modalities
  - Behavior organization: RL, sequence learning

Success of today's AI relies on hardware!

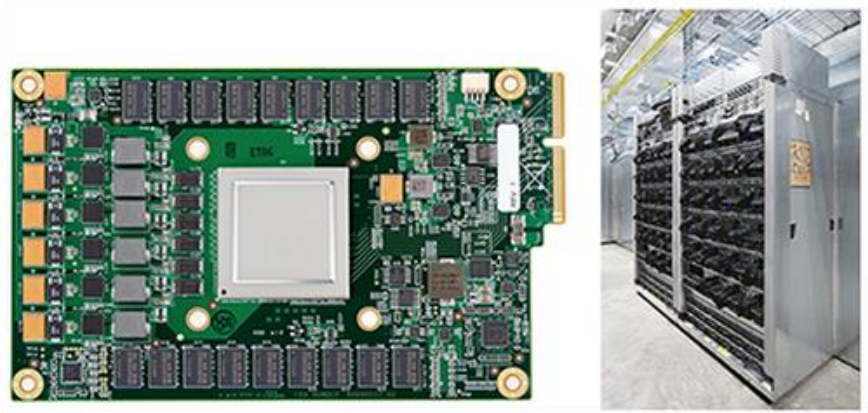

- Optimized for matrix multiplication
- Optimized for batched data processing
- Optimization of number representation

First steps: vision integrated in motion planning and control

Our robots

Dynamixel WindowX

NEURA Robotics, MIPA

UniTree Z1

Camera: RGB+Depth+Event

NEURA Robotics, MAIRA

In January

Exemplar projects on Neuromorphic chips

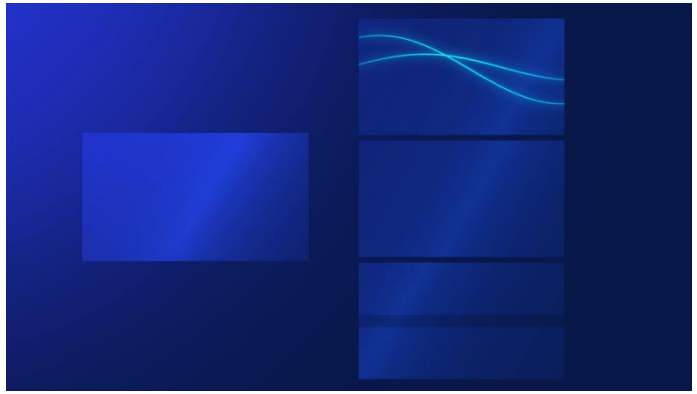

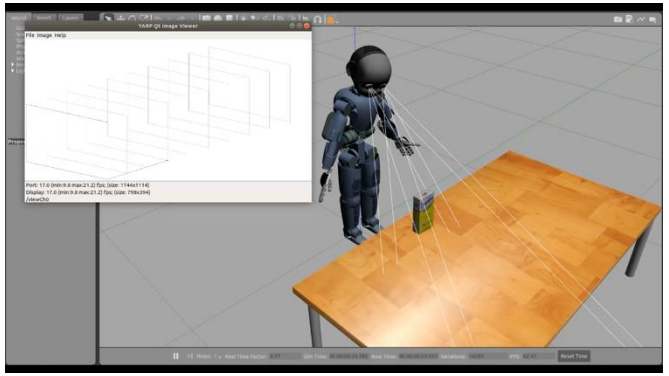

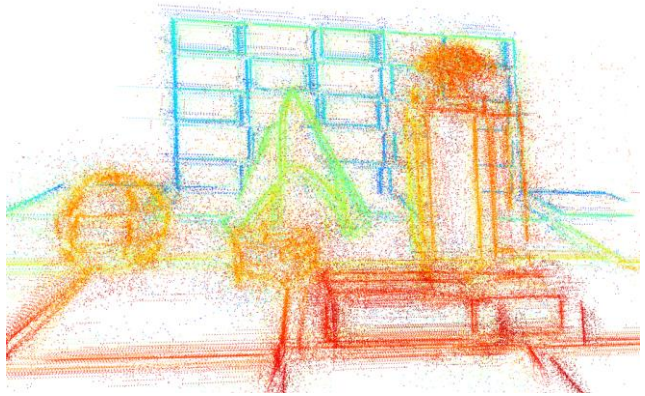

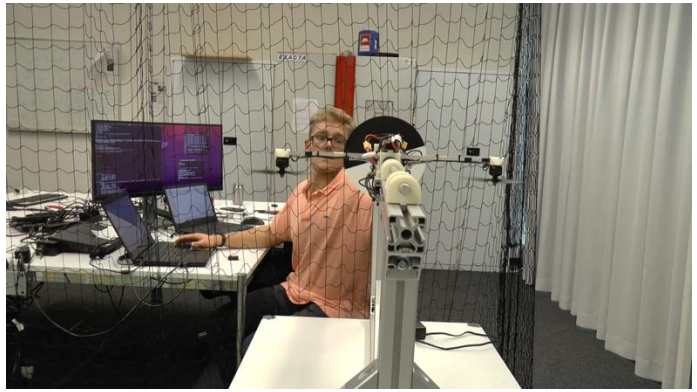

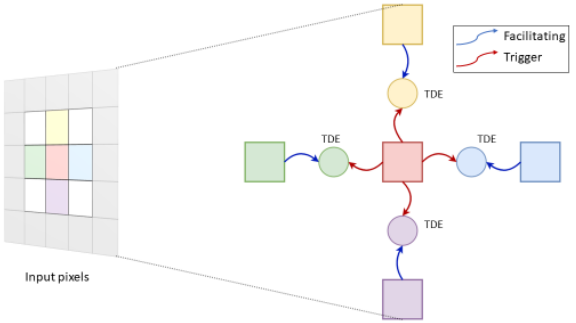

Real-time vision for control (20kHz, 200mks)

Vitale et al, ICRA 2020

Continual object learning (few shot, autonomous learning)

Hajizada et al., 2022, 2024

Depth from motion (300mks)

Chiavazza et al, 2023