AI Explained - Demystifying the Technology behind ChatGPT to shape our Future

Prof. Benjamin F. Grewe

Institute of Neuroinformatics

5th. Sept. 2023, ISPF Lucerne

ETH zürich

institute of neuroinformatics

institute of neuroinformatics

Mission of the Institute

The mission of the Institute is to discover the key principles by which brains work and to implement these in artificial systems that interact intelligently with the real world.

ETH zürich

institute of neuroinformatics

Mission of the Institute

The mission of the Institute is to discover the key principles by which brains work and to implement these in artificial systems that interact intelligently with the real world.

ETH zürich

institute of neuroinformatics

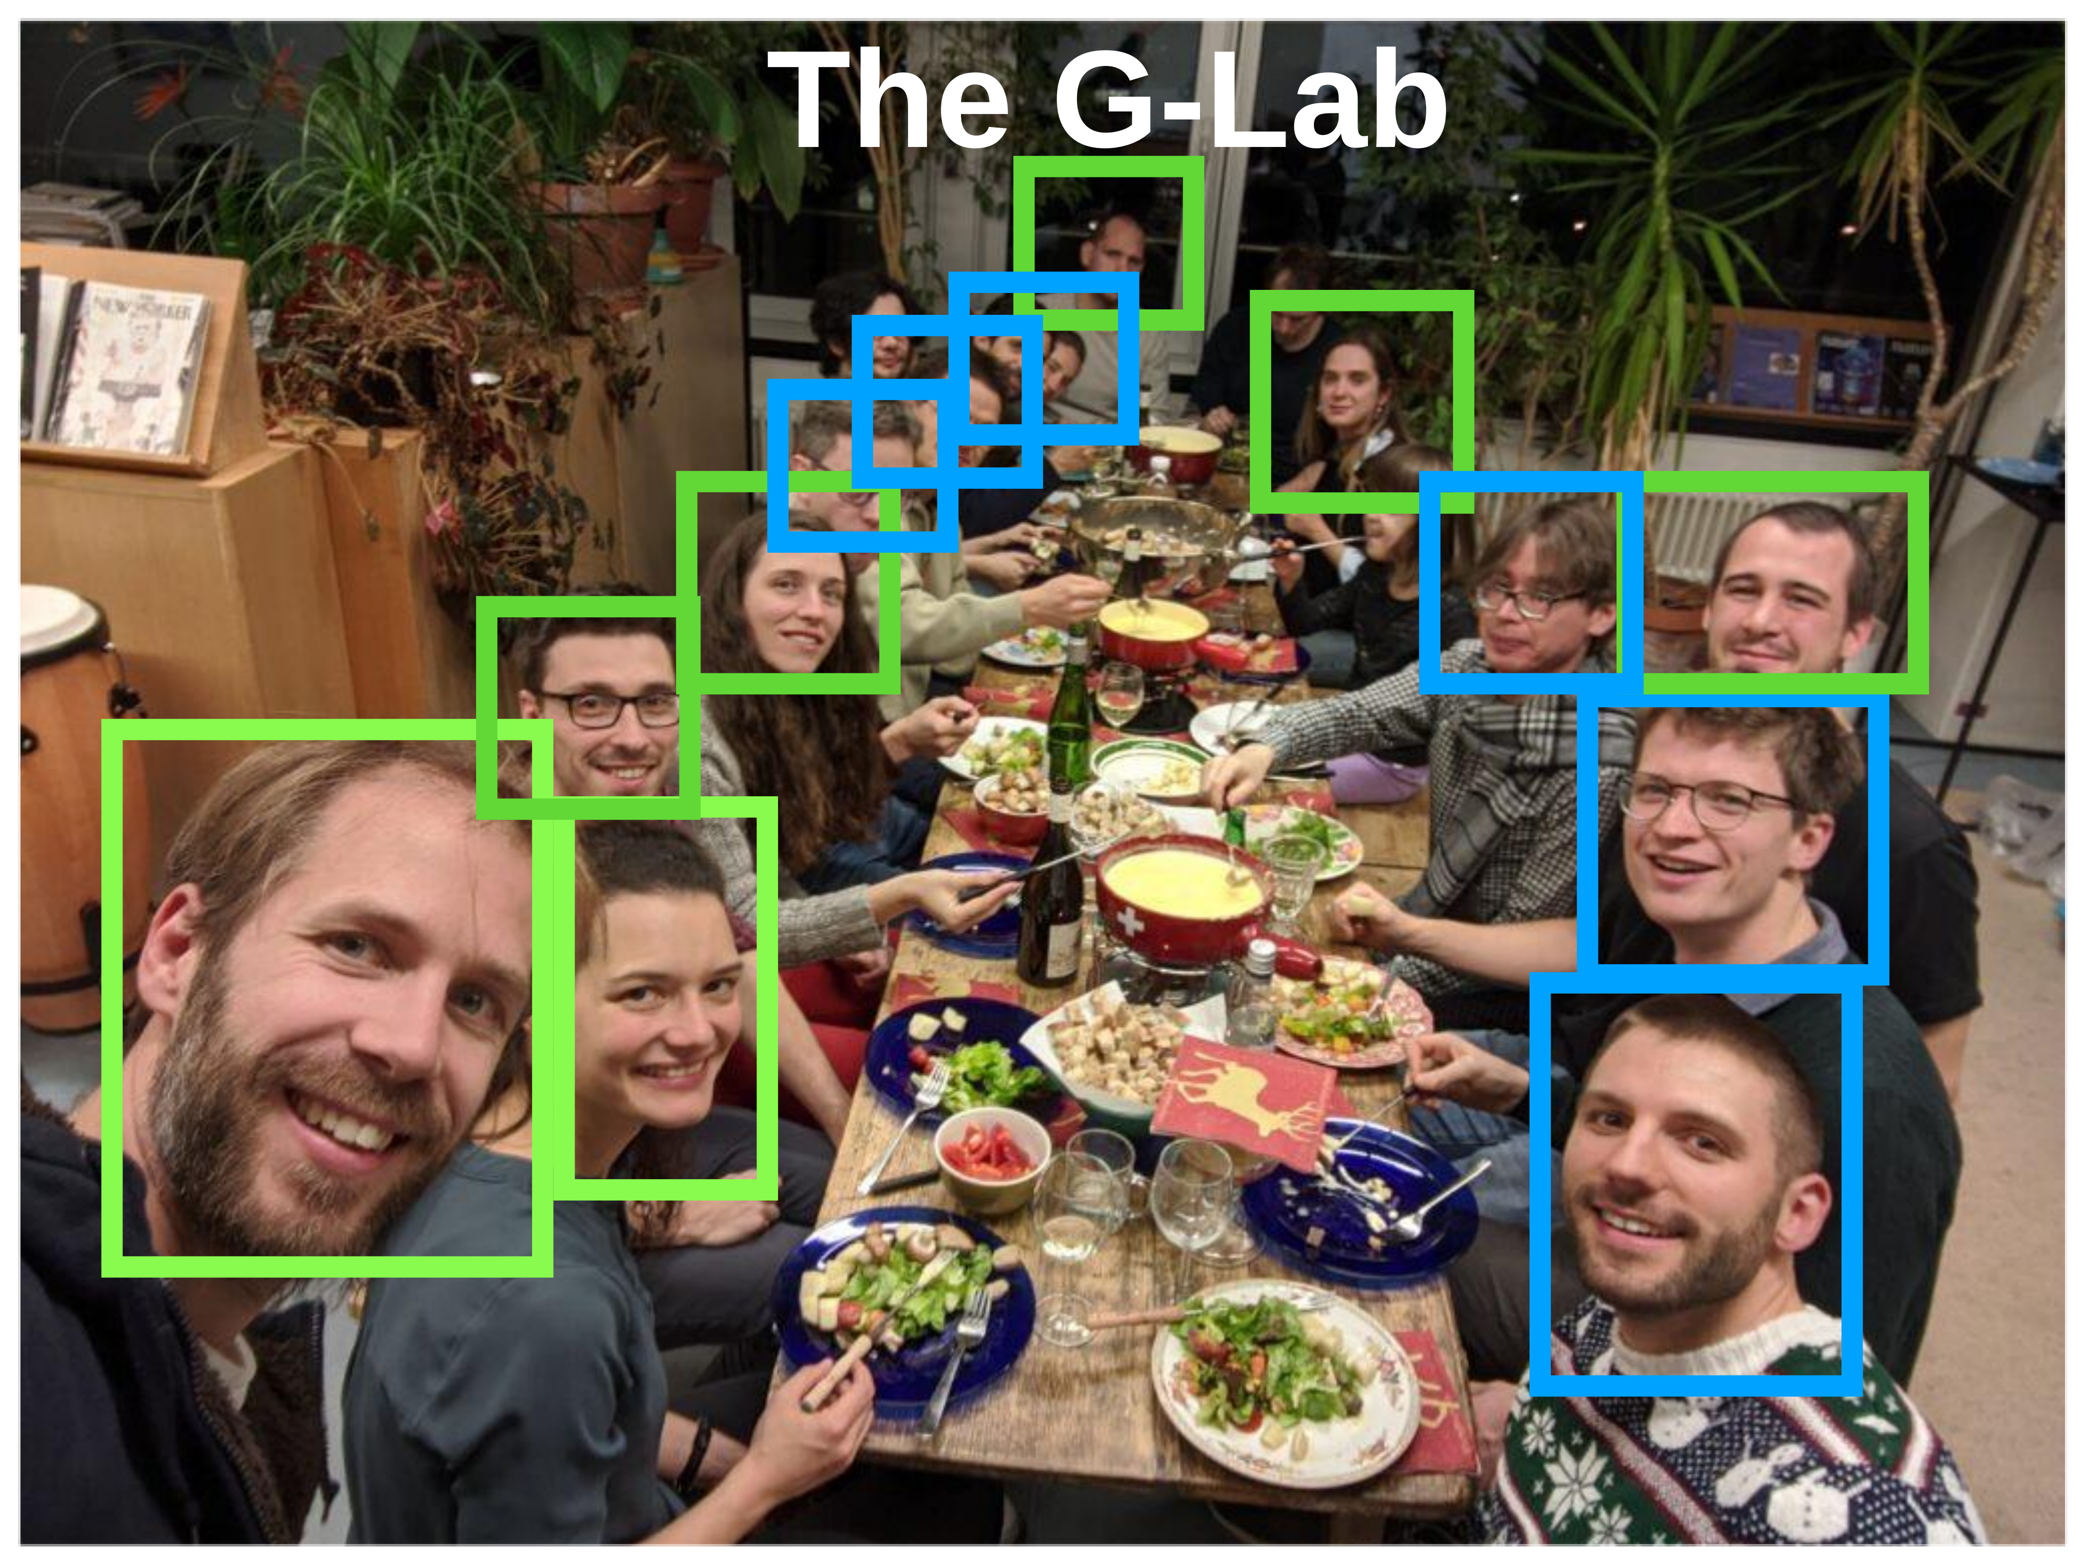

Research Background

- Computational Neuroscience
- AI / Computational Neuroscience

G-Lab: Creating Synergies to Advance Neuroscience and AI

AI

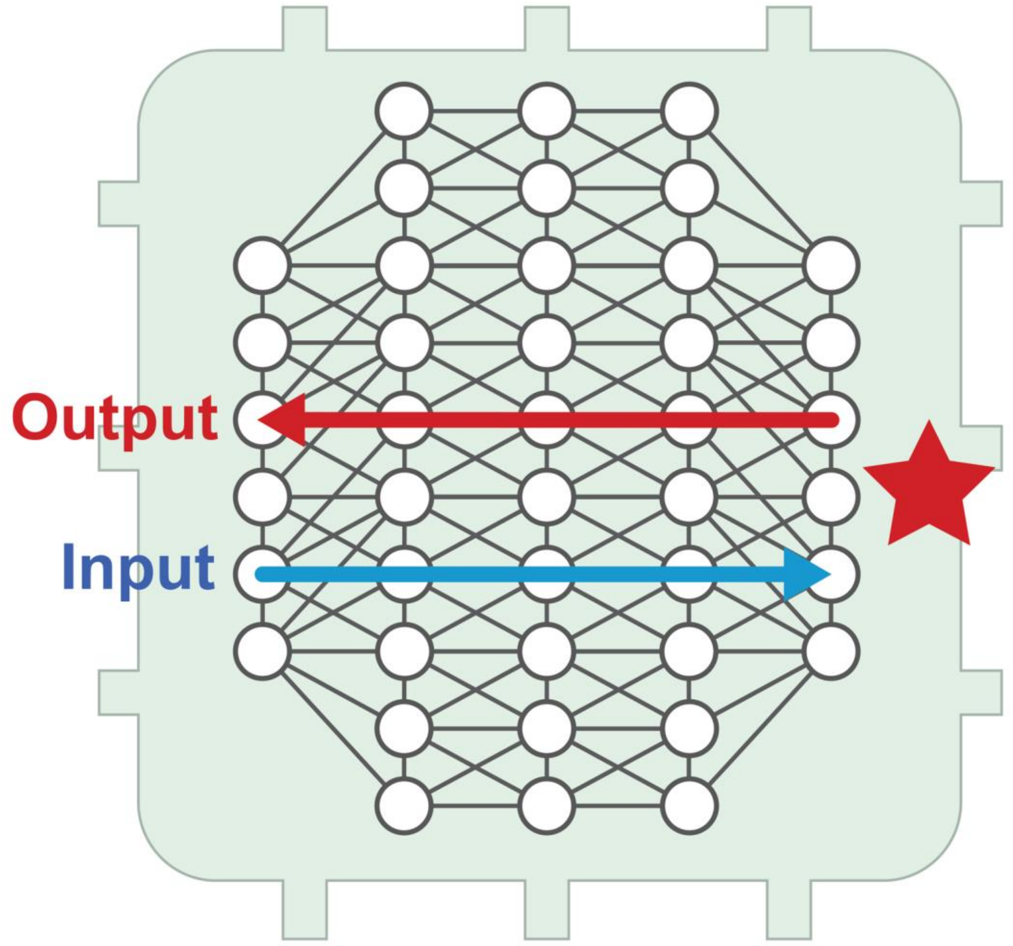

Inspire Hypotheses

Inspire Technologies

Neuroscience

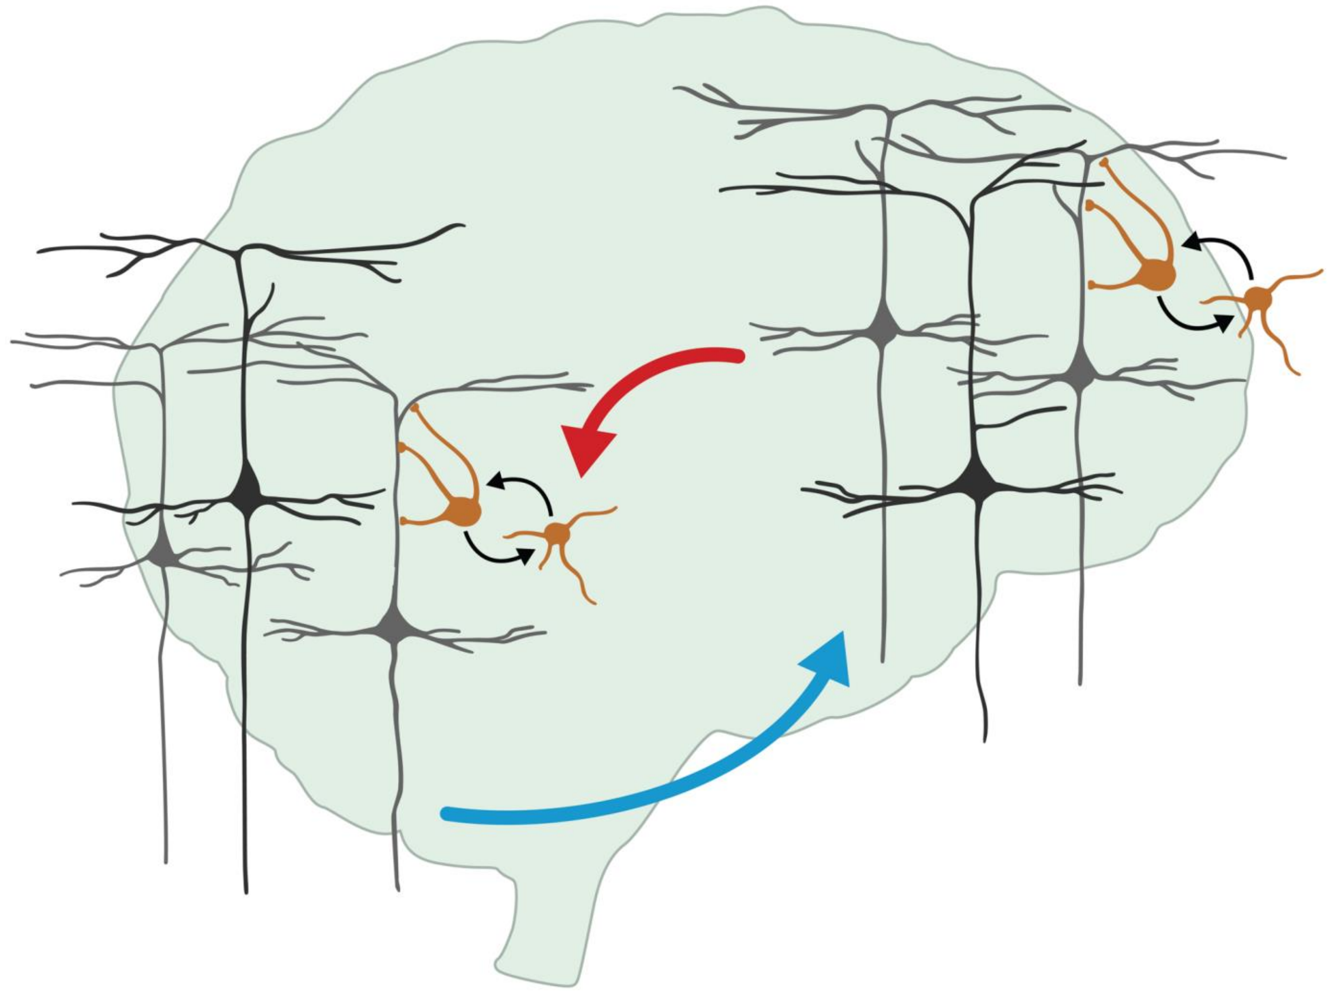

Artificial Intelligence	Natural Intelligence
Deep Networks, Transformers	**Biological Neuronal Networks**
High energy consumption, hundreds of GPUs	**Highly energy efficient (20W),**
Requires huge amounts training data	**Learns extremely efficient**

Content of Today's Talk:

Part I: Learning in Hierarchical (Deep) Cortical Networks

AI → **Neuroscience**

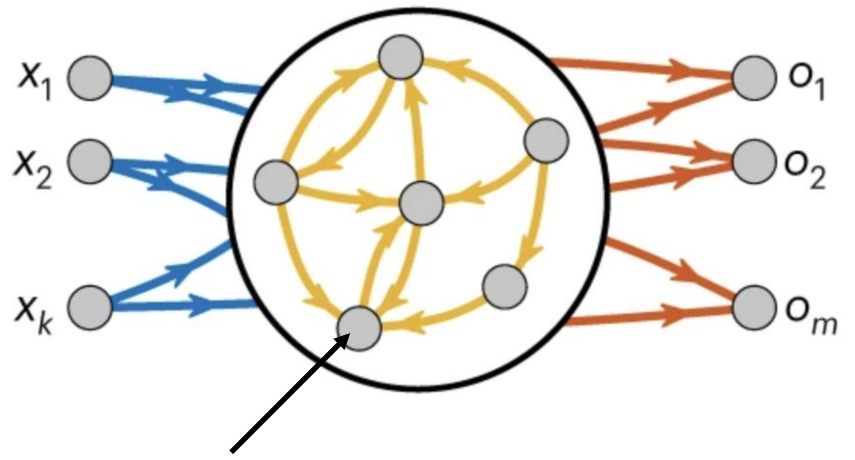

Scientific Question:
How does Credit Assignment
in the Brain work?

ETH zürich

Content of Today's Talk:

institute of neuroinformatics

Part I: Learning in Hierarchical (Deep) Cortical Networks

AI → Neuroscience

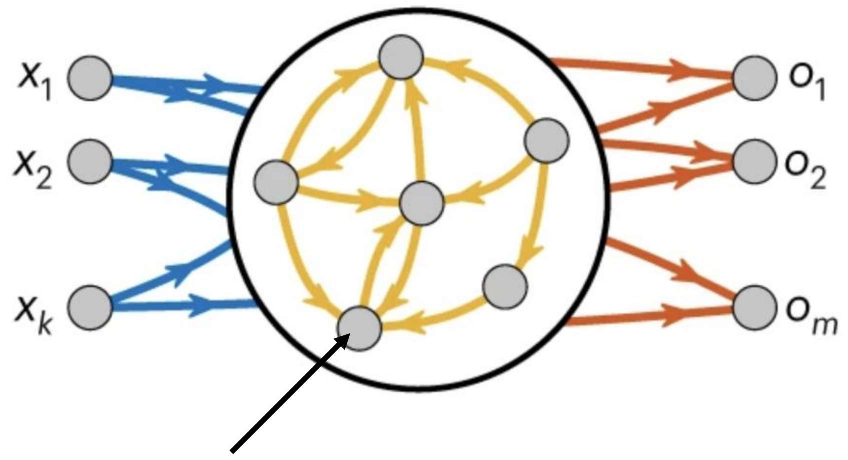

Scientific Question:
How does Credit Assignment
in the Brain work?

Part II: Understanding Hierarchical Neuronal Representations in Brain

AI ← Neuroscience

Scientific Question:
What are the neuronal representations of the
sensory input (e.g. image of bus) that allows our brain
to generate goal directed actions?

ETH zürich

Content of Today's Talk:

institute of neuroinformatics

Part I: Learning in Hierarchical (Deep) Cortical Networks

AI → Neuroscience

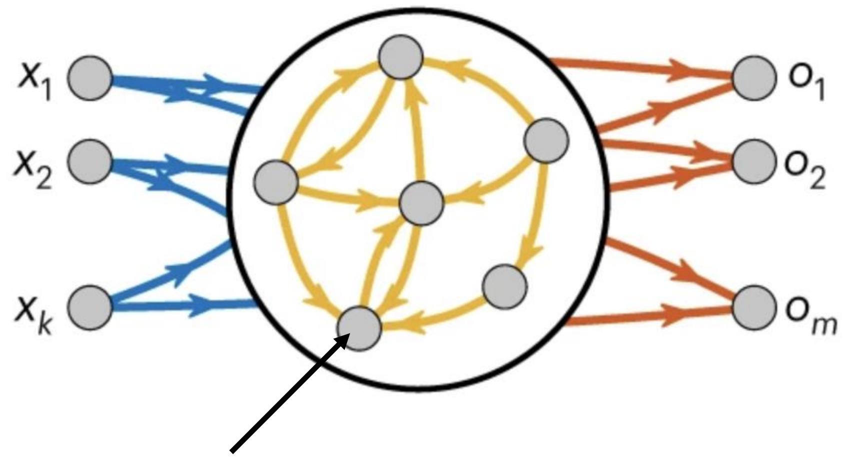

Scientific Question:
How does Credit Assignment in the Brain work?

Content of Today's Talk:

Part I: Learning in Hierarchical (Deep) Cortical Networks

AI → Neuroscience

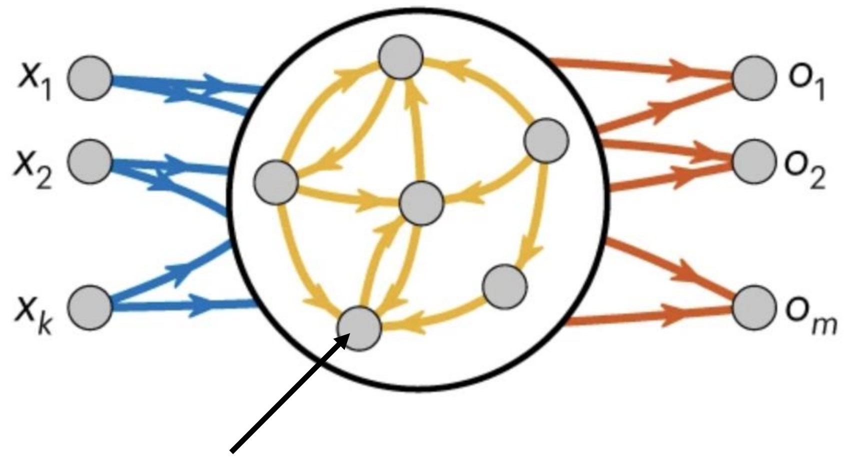

Scientific Question:
How does Credit Assignment
in the Brain work?

Our Approach: Make Deep Network Learning more Biologically Plausible.

Why?

Content of Today's Talk:

Part I: Learning in Hierarchical (Deep) Cortical Networks

AI → **Neuroscience**

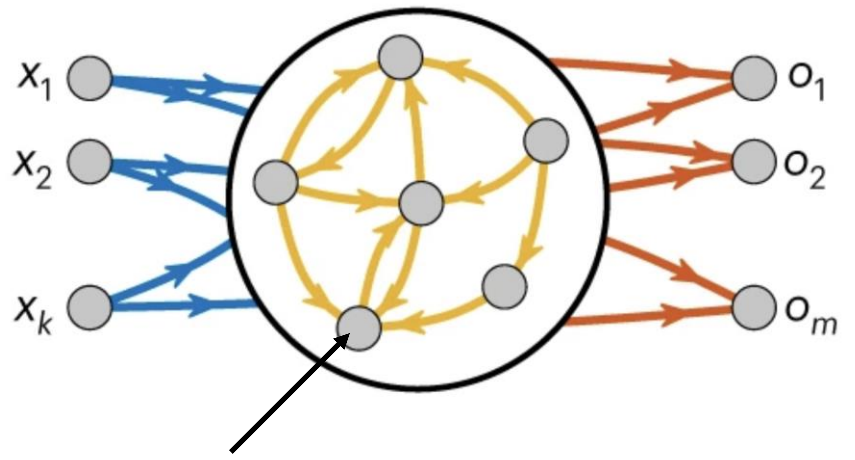

Scientific Question:
How does Credit Assignment
in the Brain work?

Engineering:

- To develop Deep Networks that are more energy efficient.
- To train Deep Networks with less training data.
- To enable Deep Continual Learning.
- To train Deep Networks with dynamic data.

Neuroscience:

- We currently cannot use deep learning algorithms to explain how credit assignment in the brain works.

Background and Motivation - Part I

The Brain's Visual Pathway

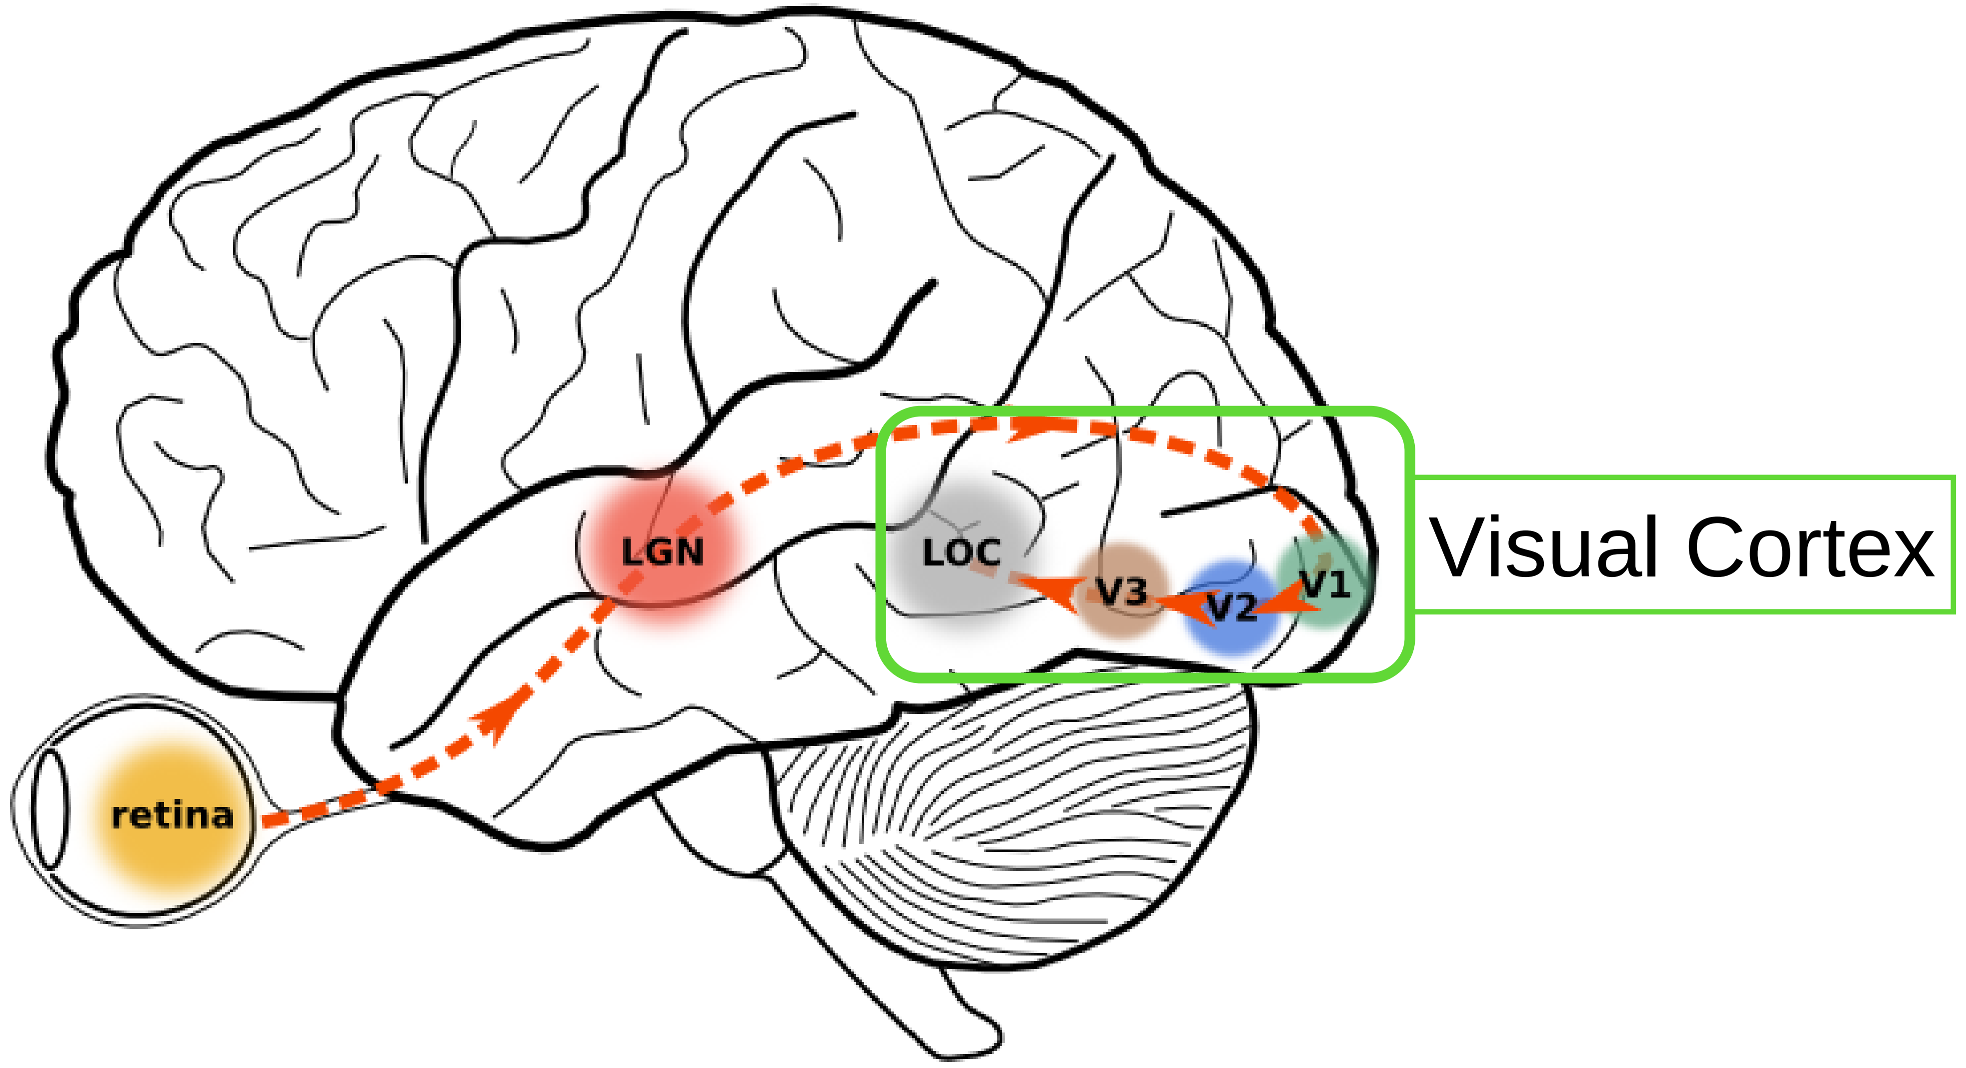

LGN: Lateral Geniculate Nucleus
V1: Primary Visual Cortex
V2. Secondary Visual Cortes
V3: Third Visual Cortex
LOC: Lateral Occipital Cortex

Image credit Jonas Kubilinas

ETH zürich

Background and Motivation - Part I

institute of neuroinformatics

The Brain's Visual Pathway

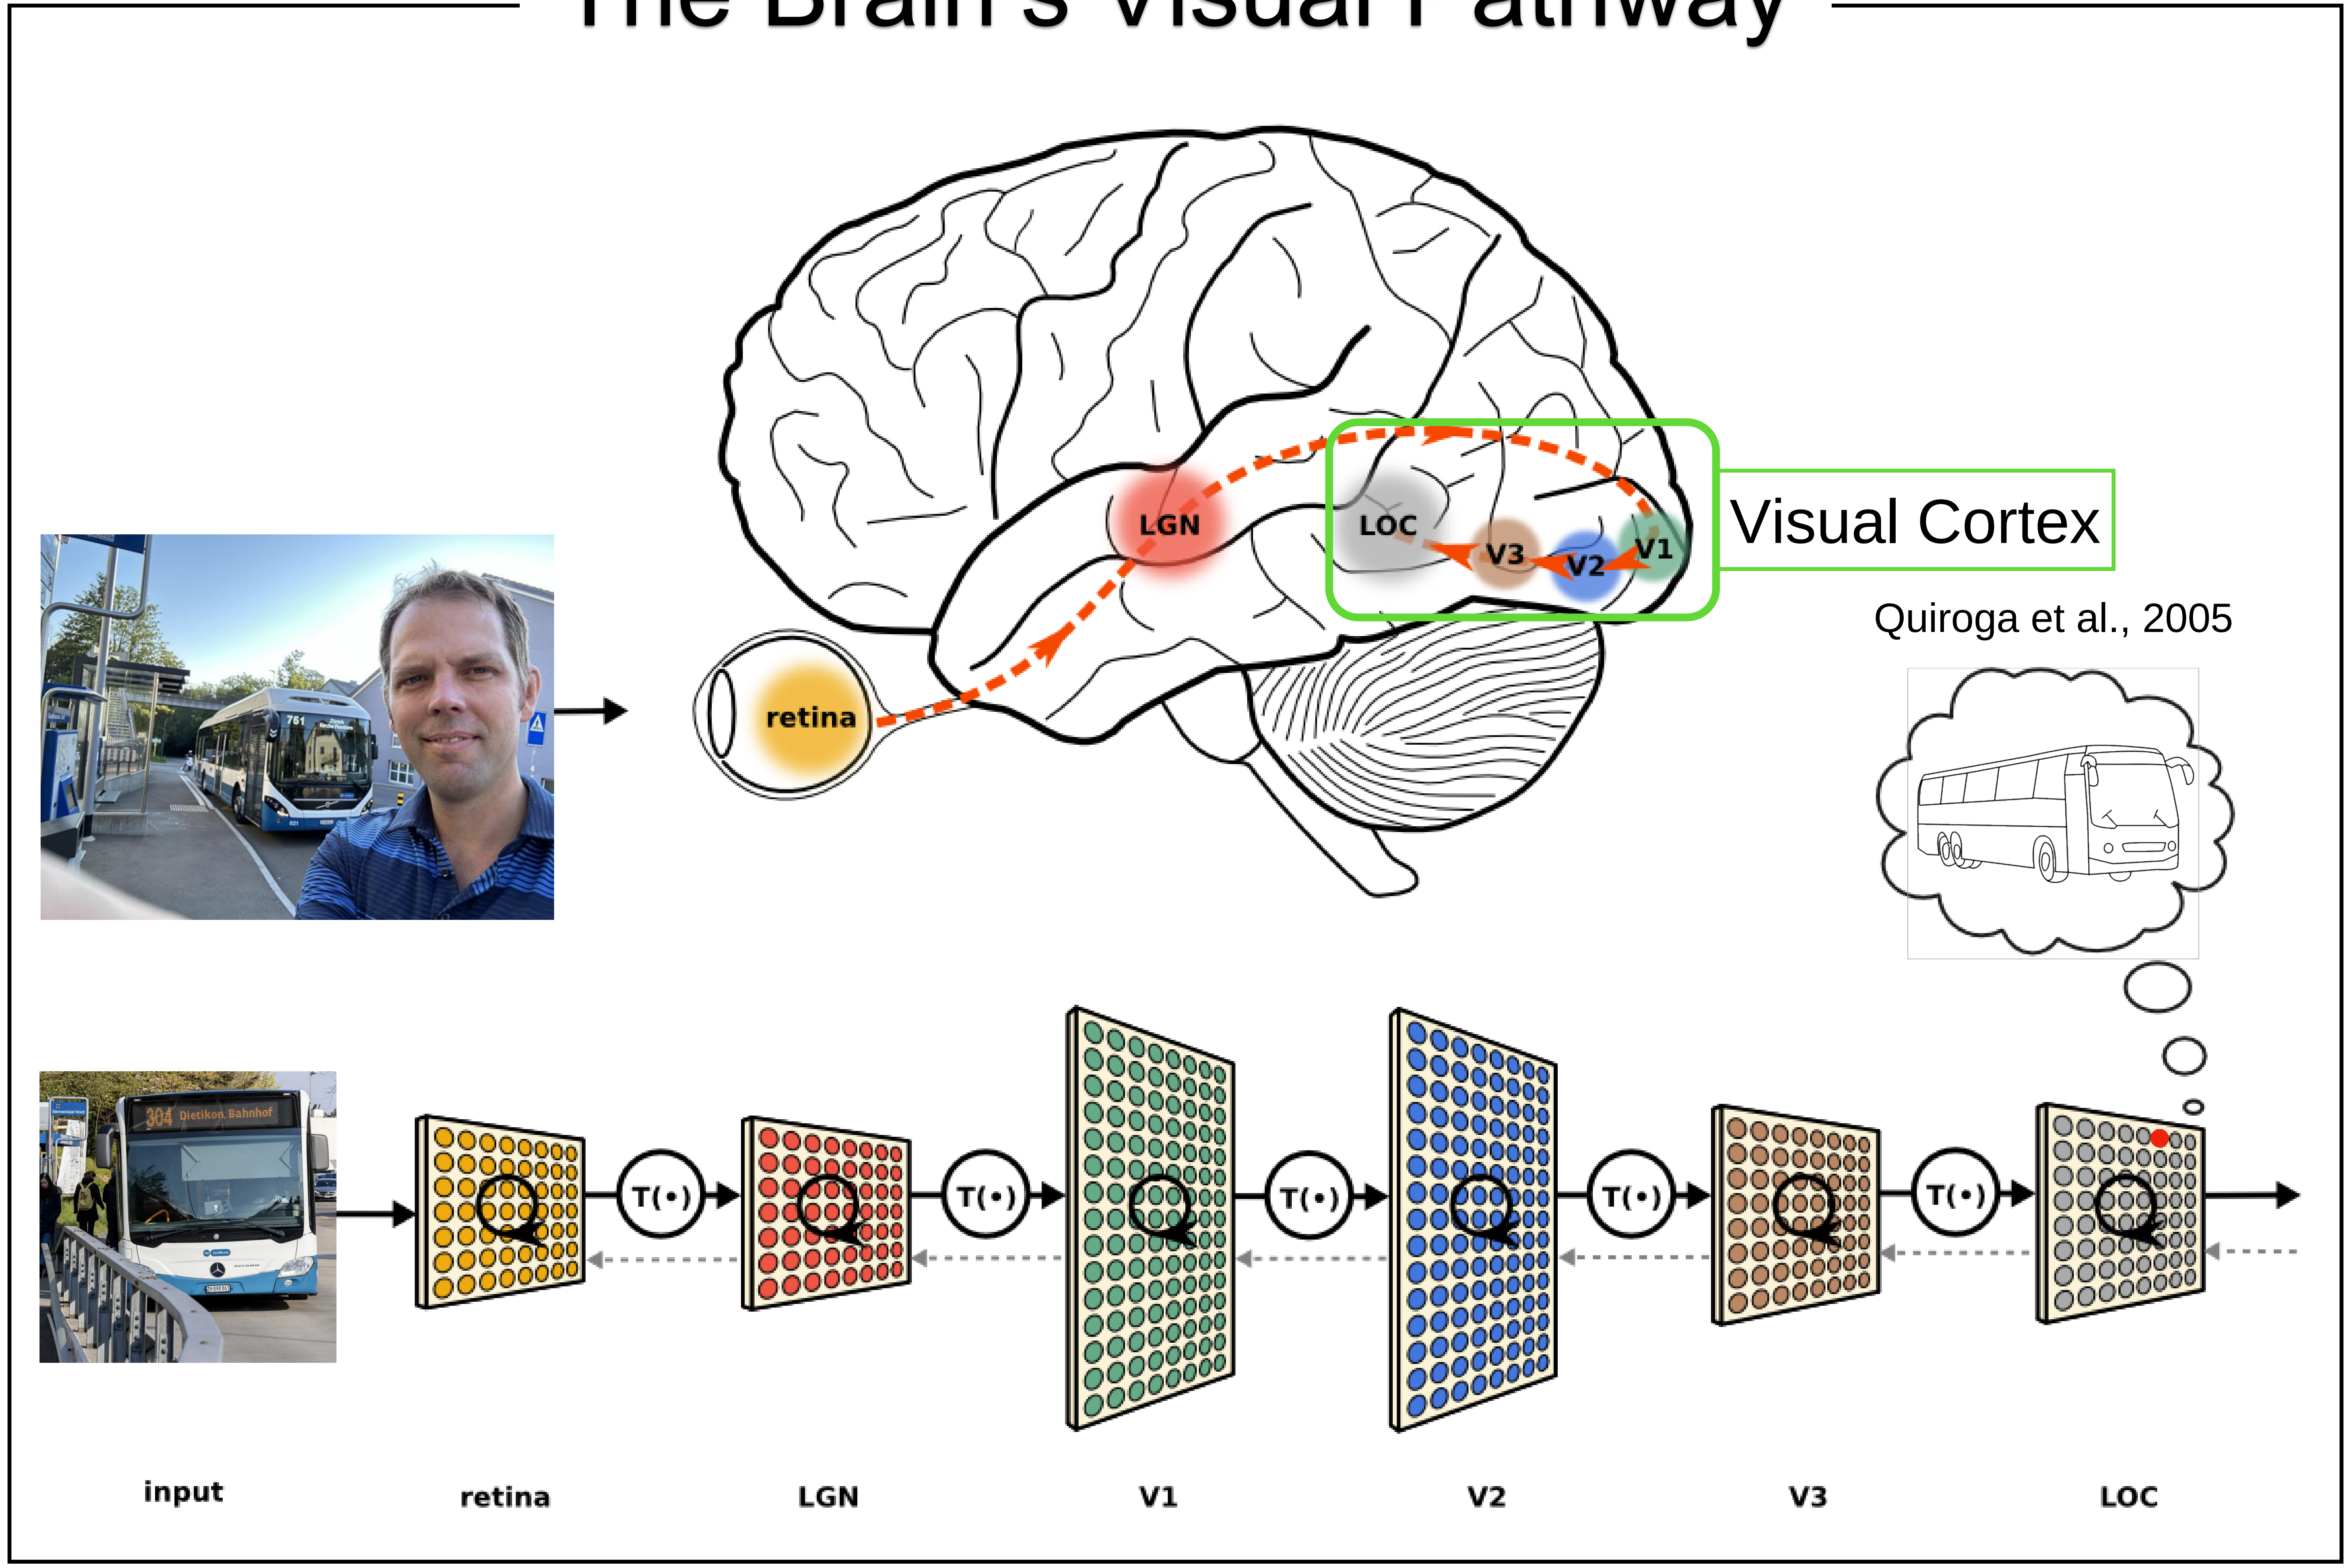

LGN: Lateral Geniculate Nucleus
V1: Primary Visual Cortex
V2. Secondary Visual Cortes
V3: Third Visual Cortex
LOC: Lateral Occipital Cortex

Image credit Jonas Kubilinas

Background and Motivation - Part I

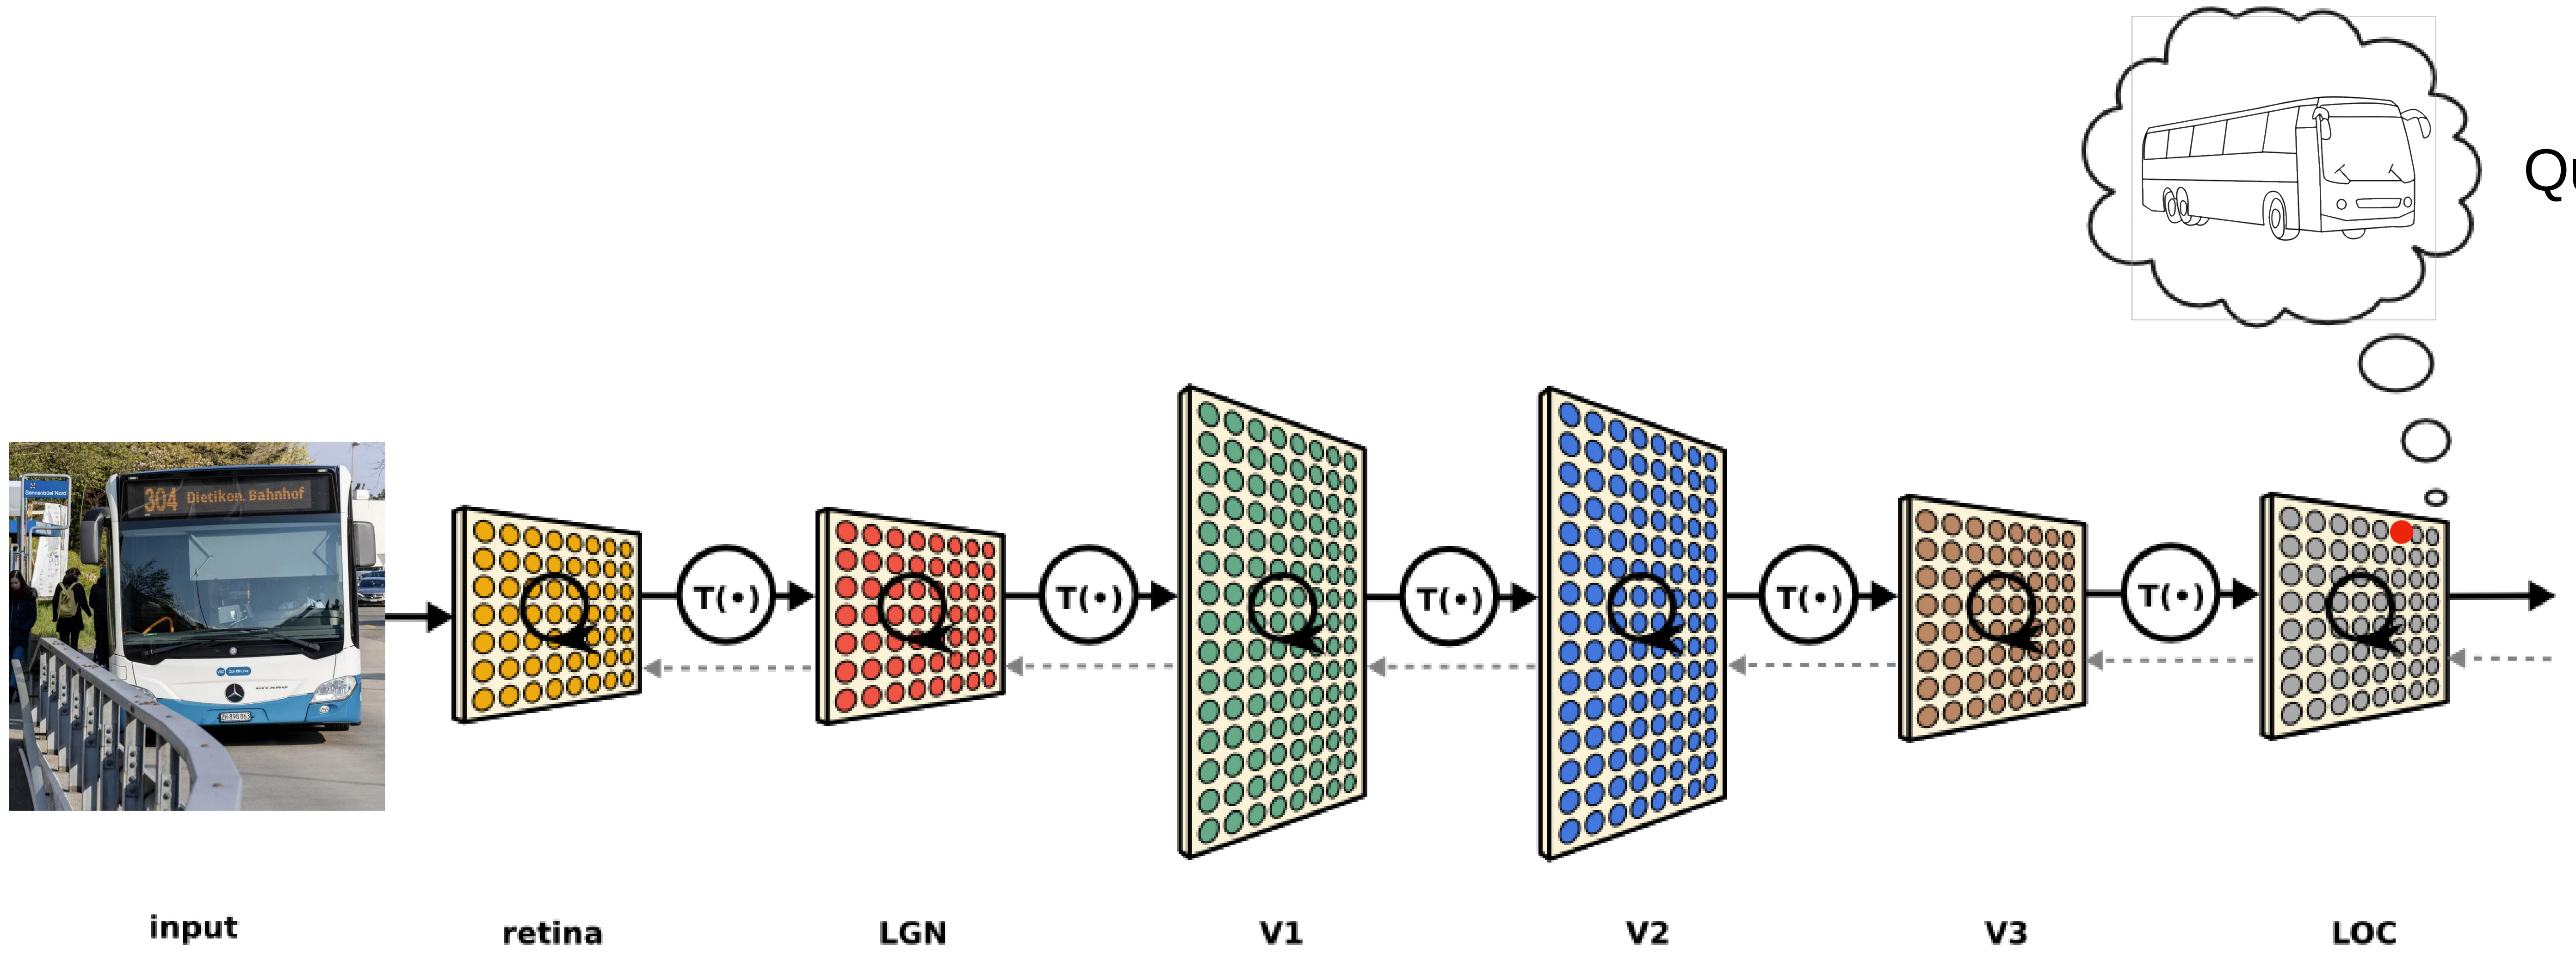

Quiroga et al., 2005

Brain Visual Pathway

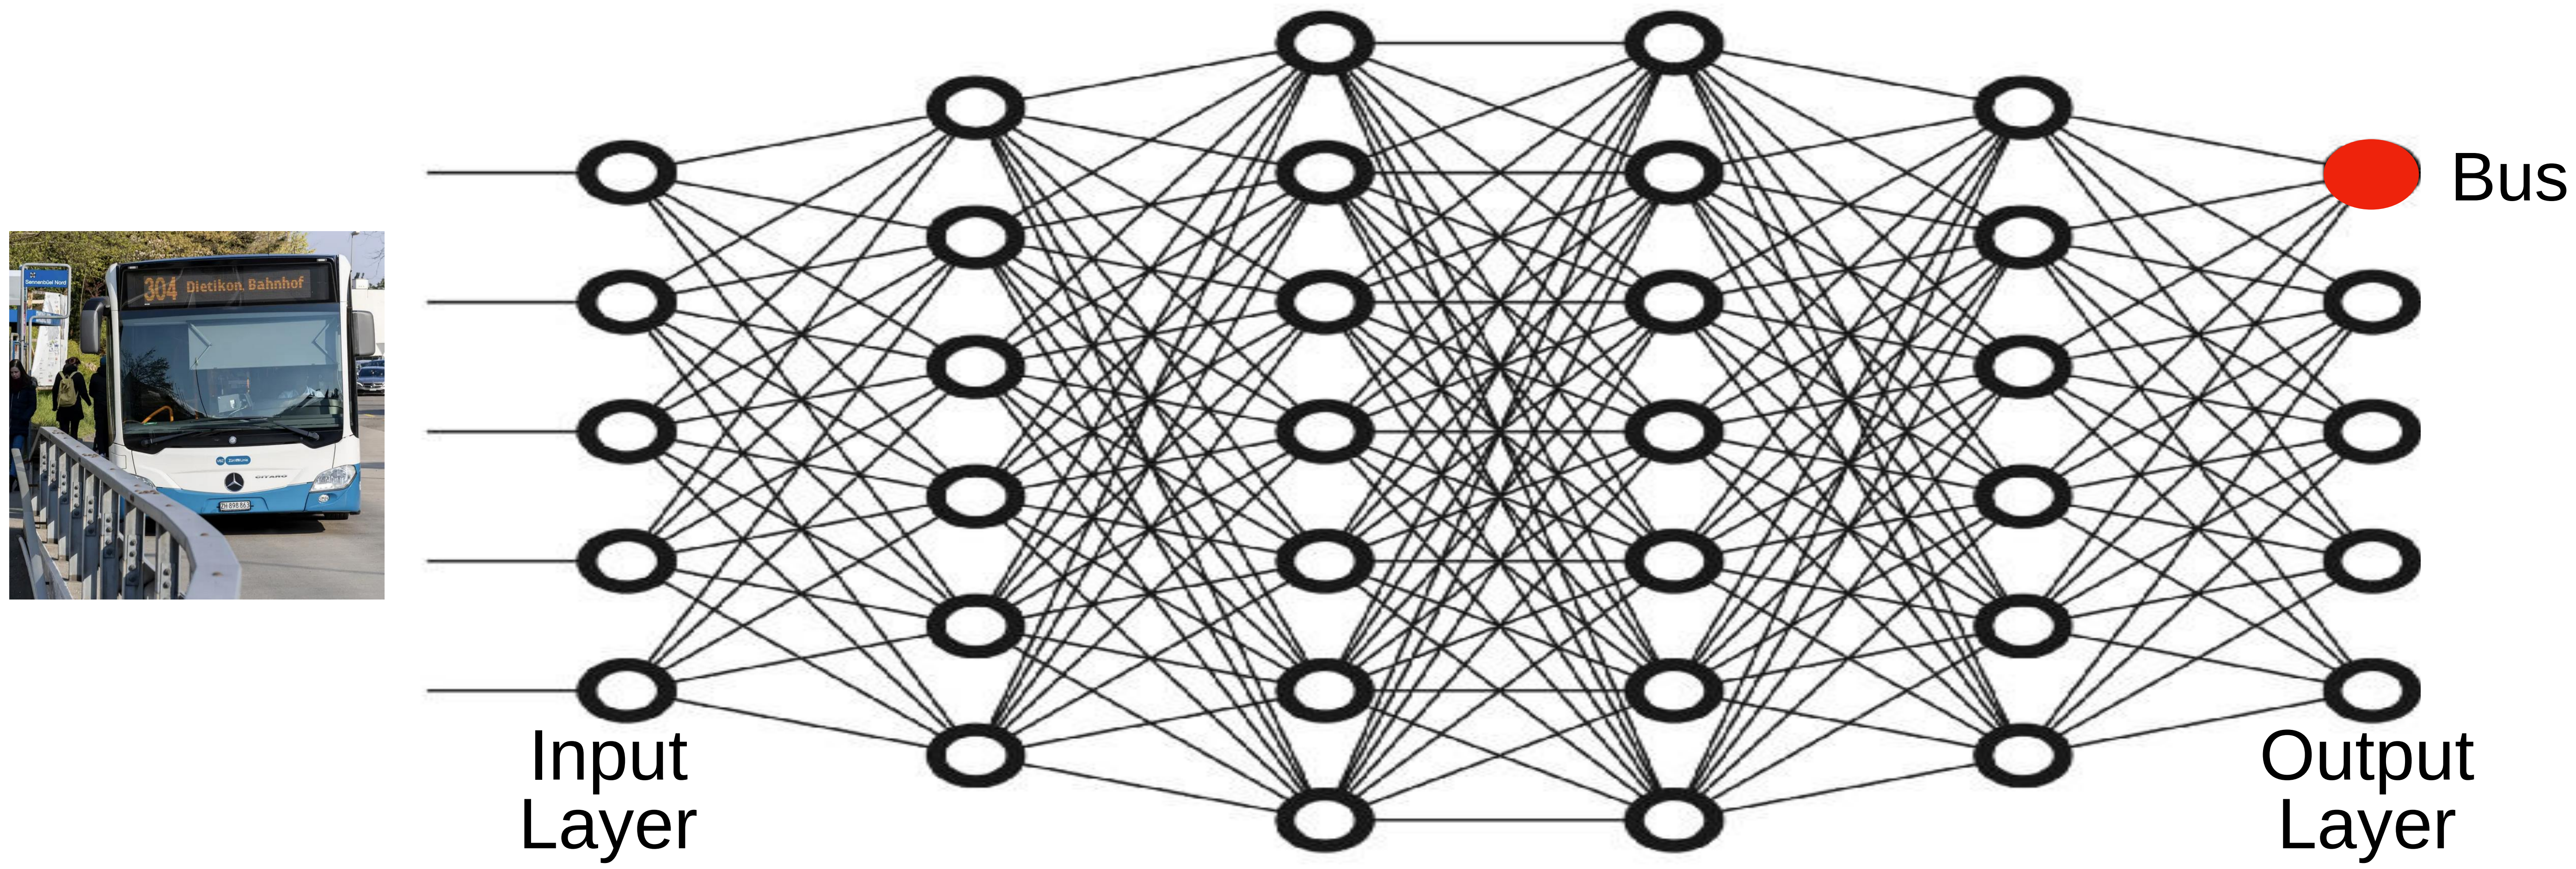

Deep Neuronal Network

Background and Motivation - Part I

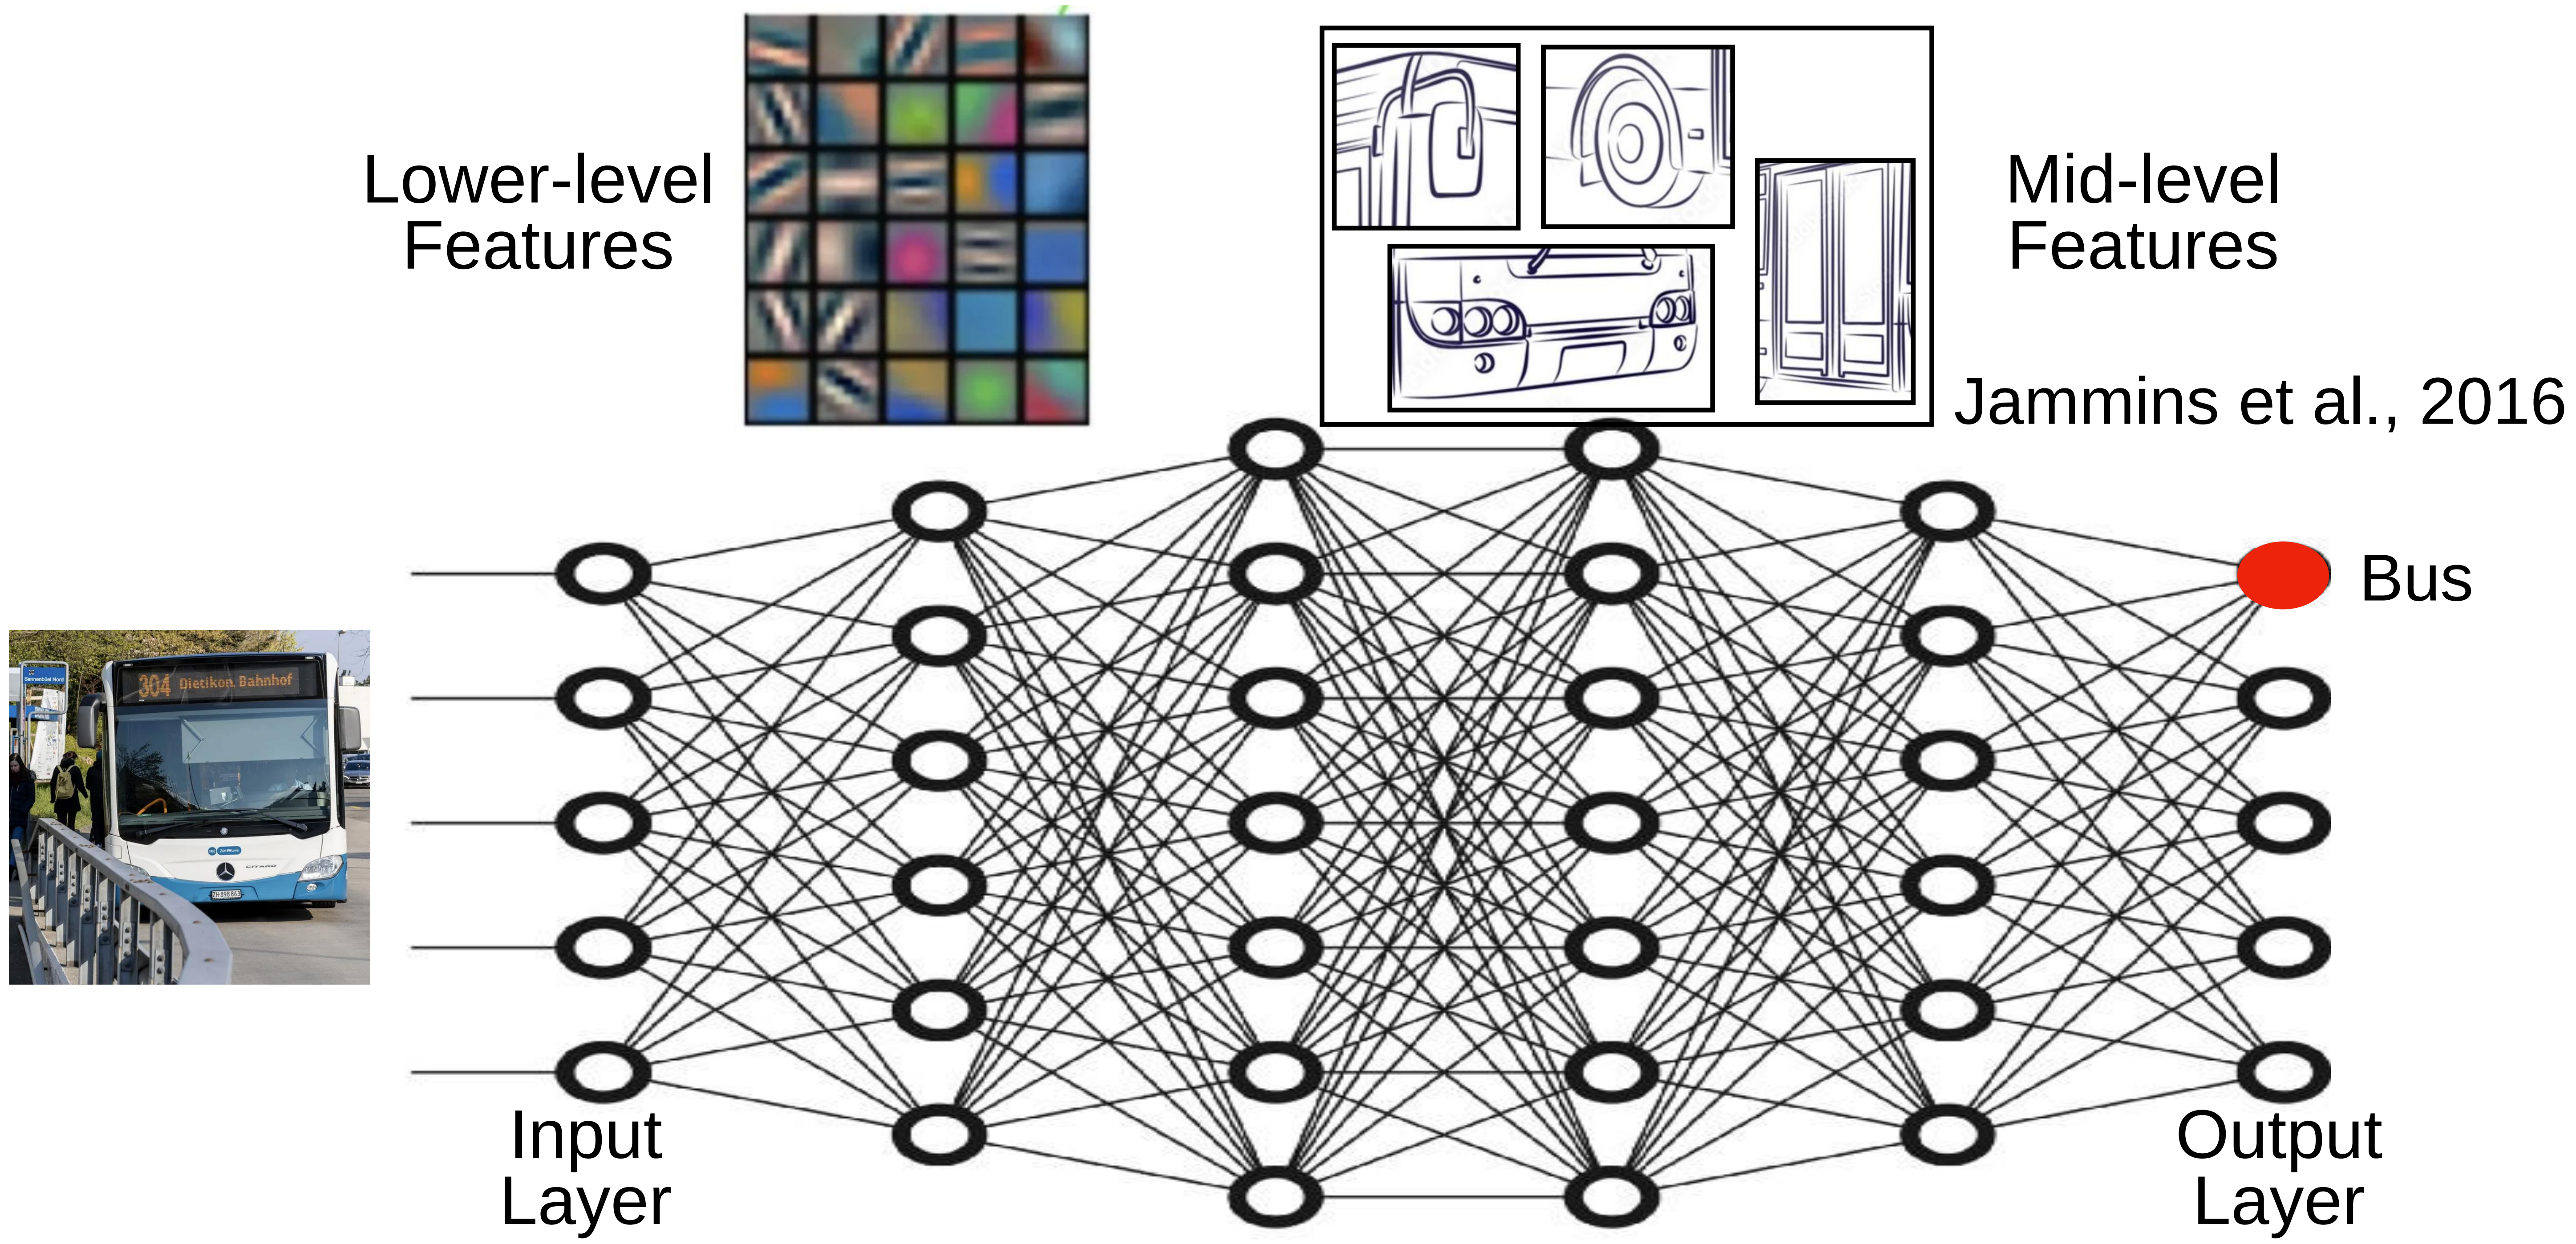

Deep Neuronal Network

Background and Motivation - Part I
The Error Backpropagation Method

1. <u>The Forward Pass</u>

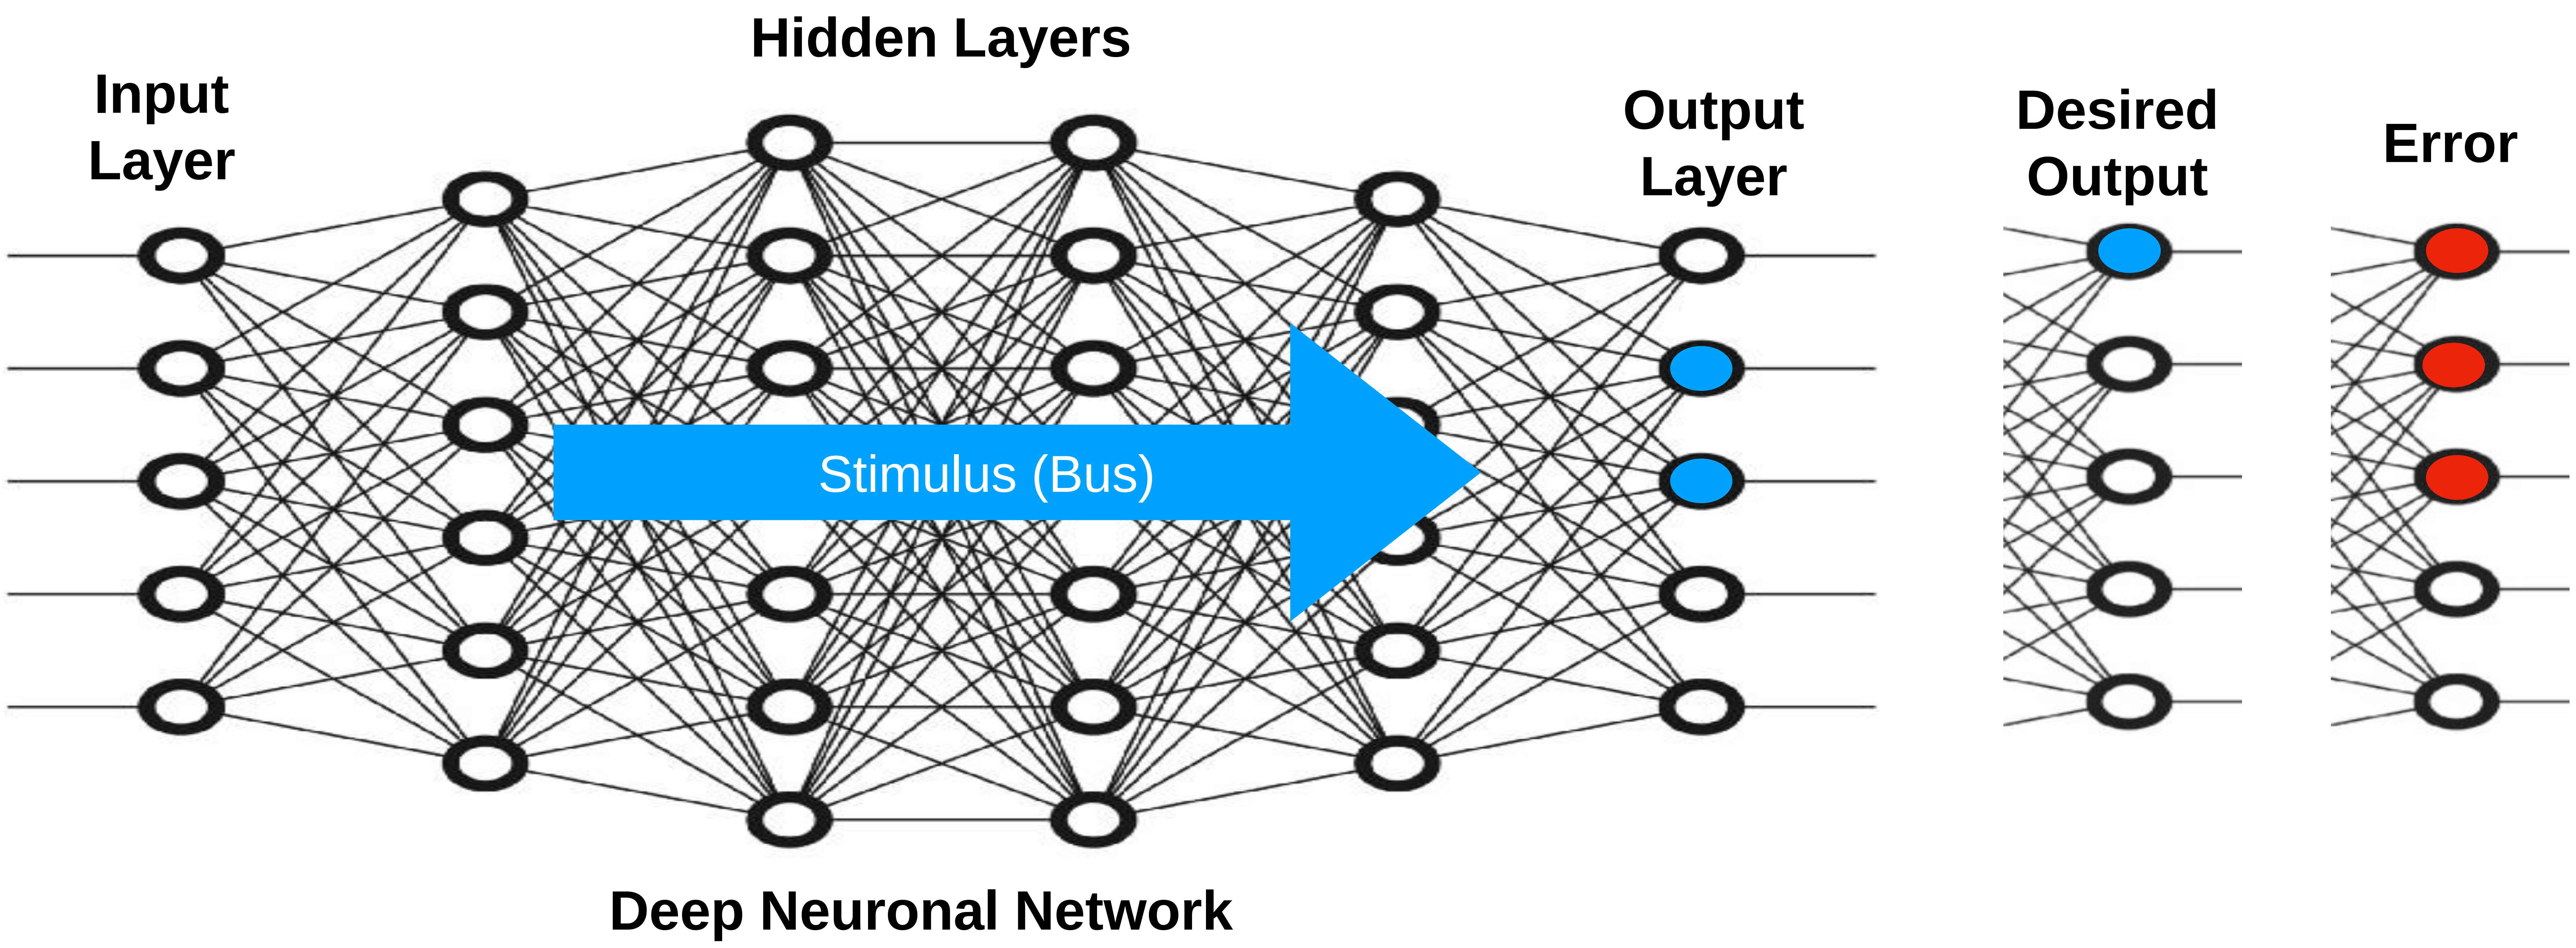

ETH zürich

Background and Motivation - Part I
The Error Backpropagation Method

institute of neuroinformatics

1. The Forward Pass
2. <u>The Backward (Error) Pass</u>

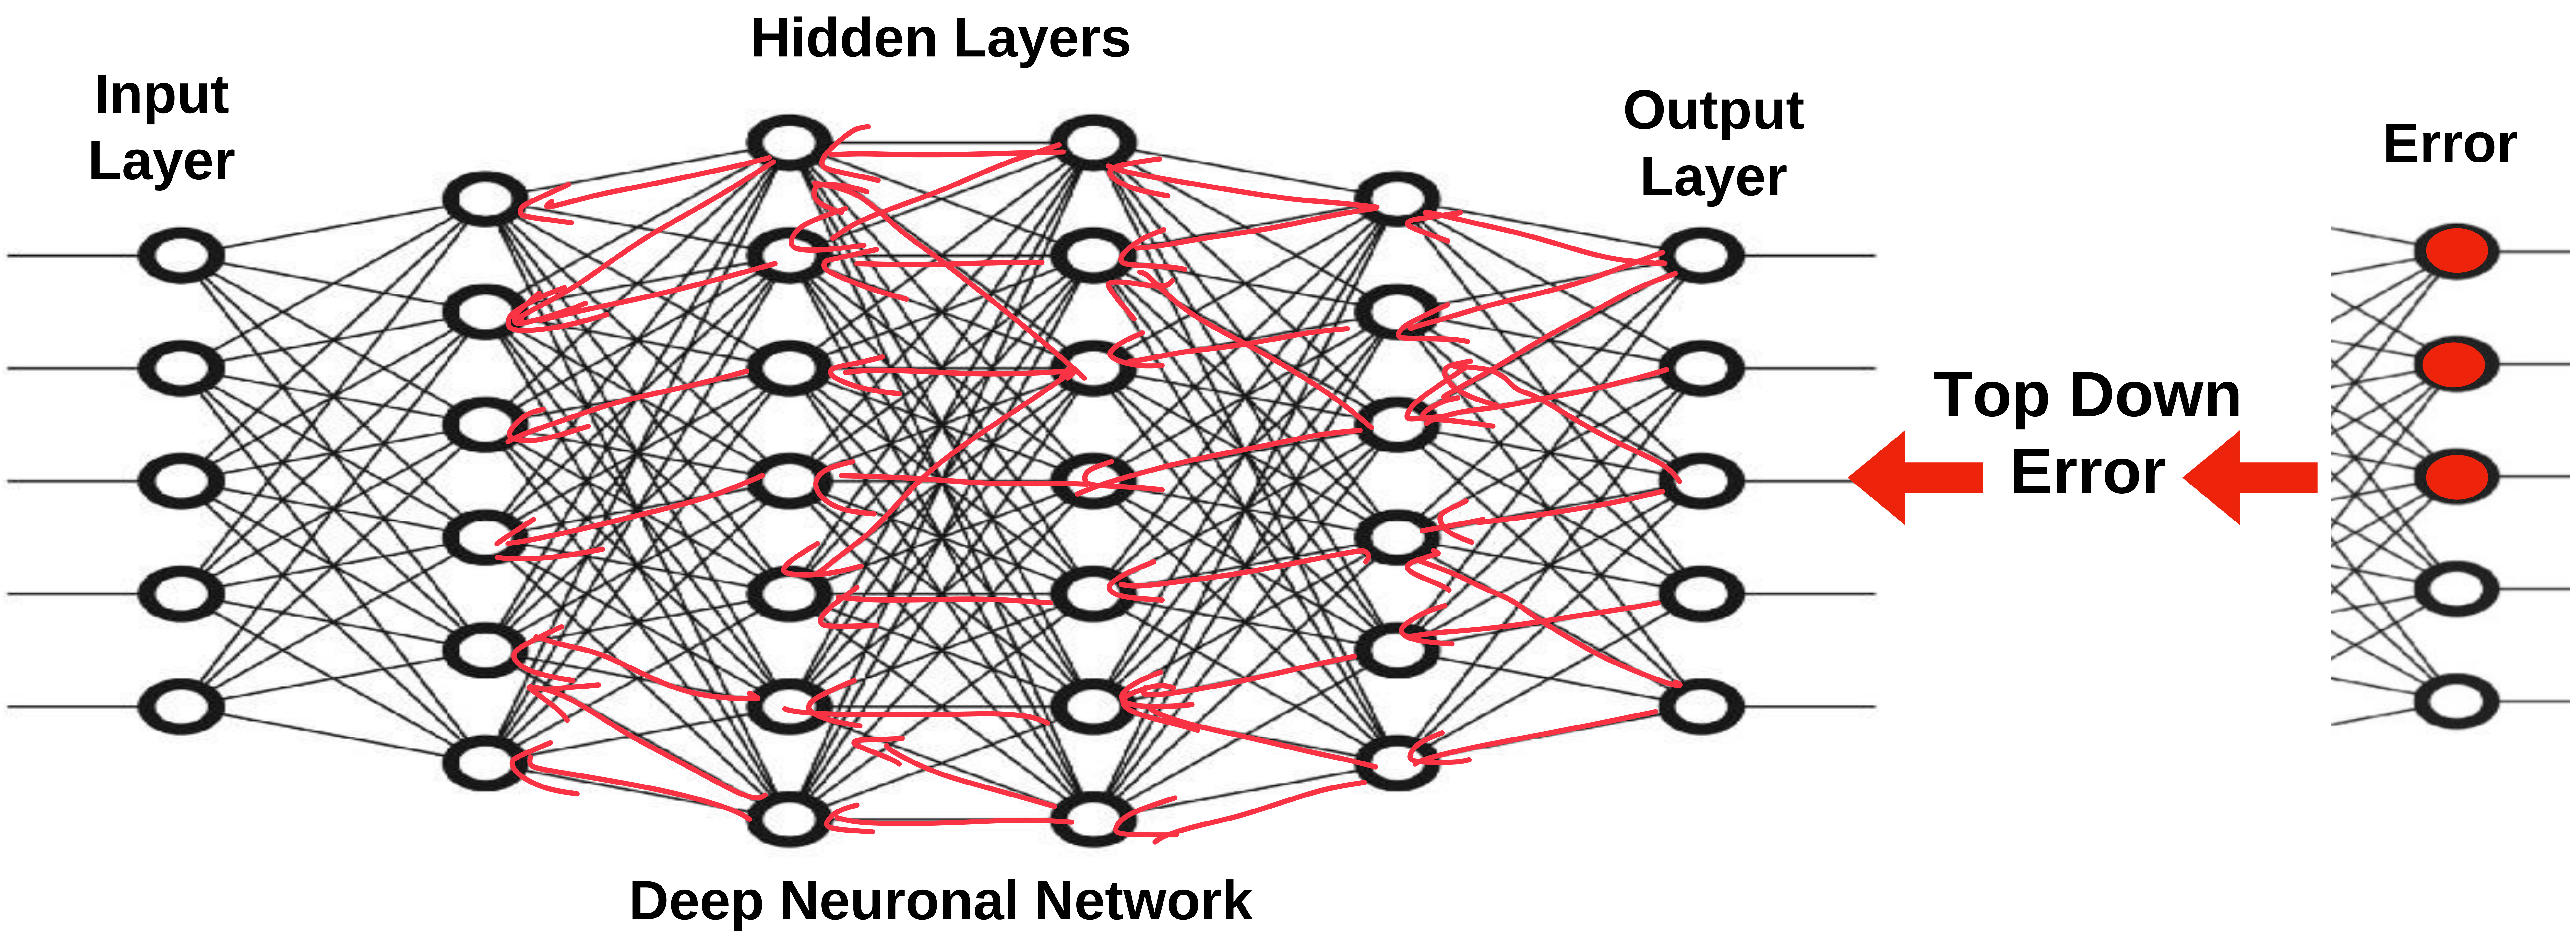

Background and Motivation - Part I
The Error Backpropagation Method

1. The Forward Pass
2. <u>The Backward (Error) Pass</u>

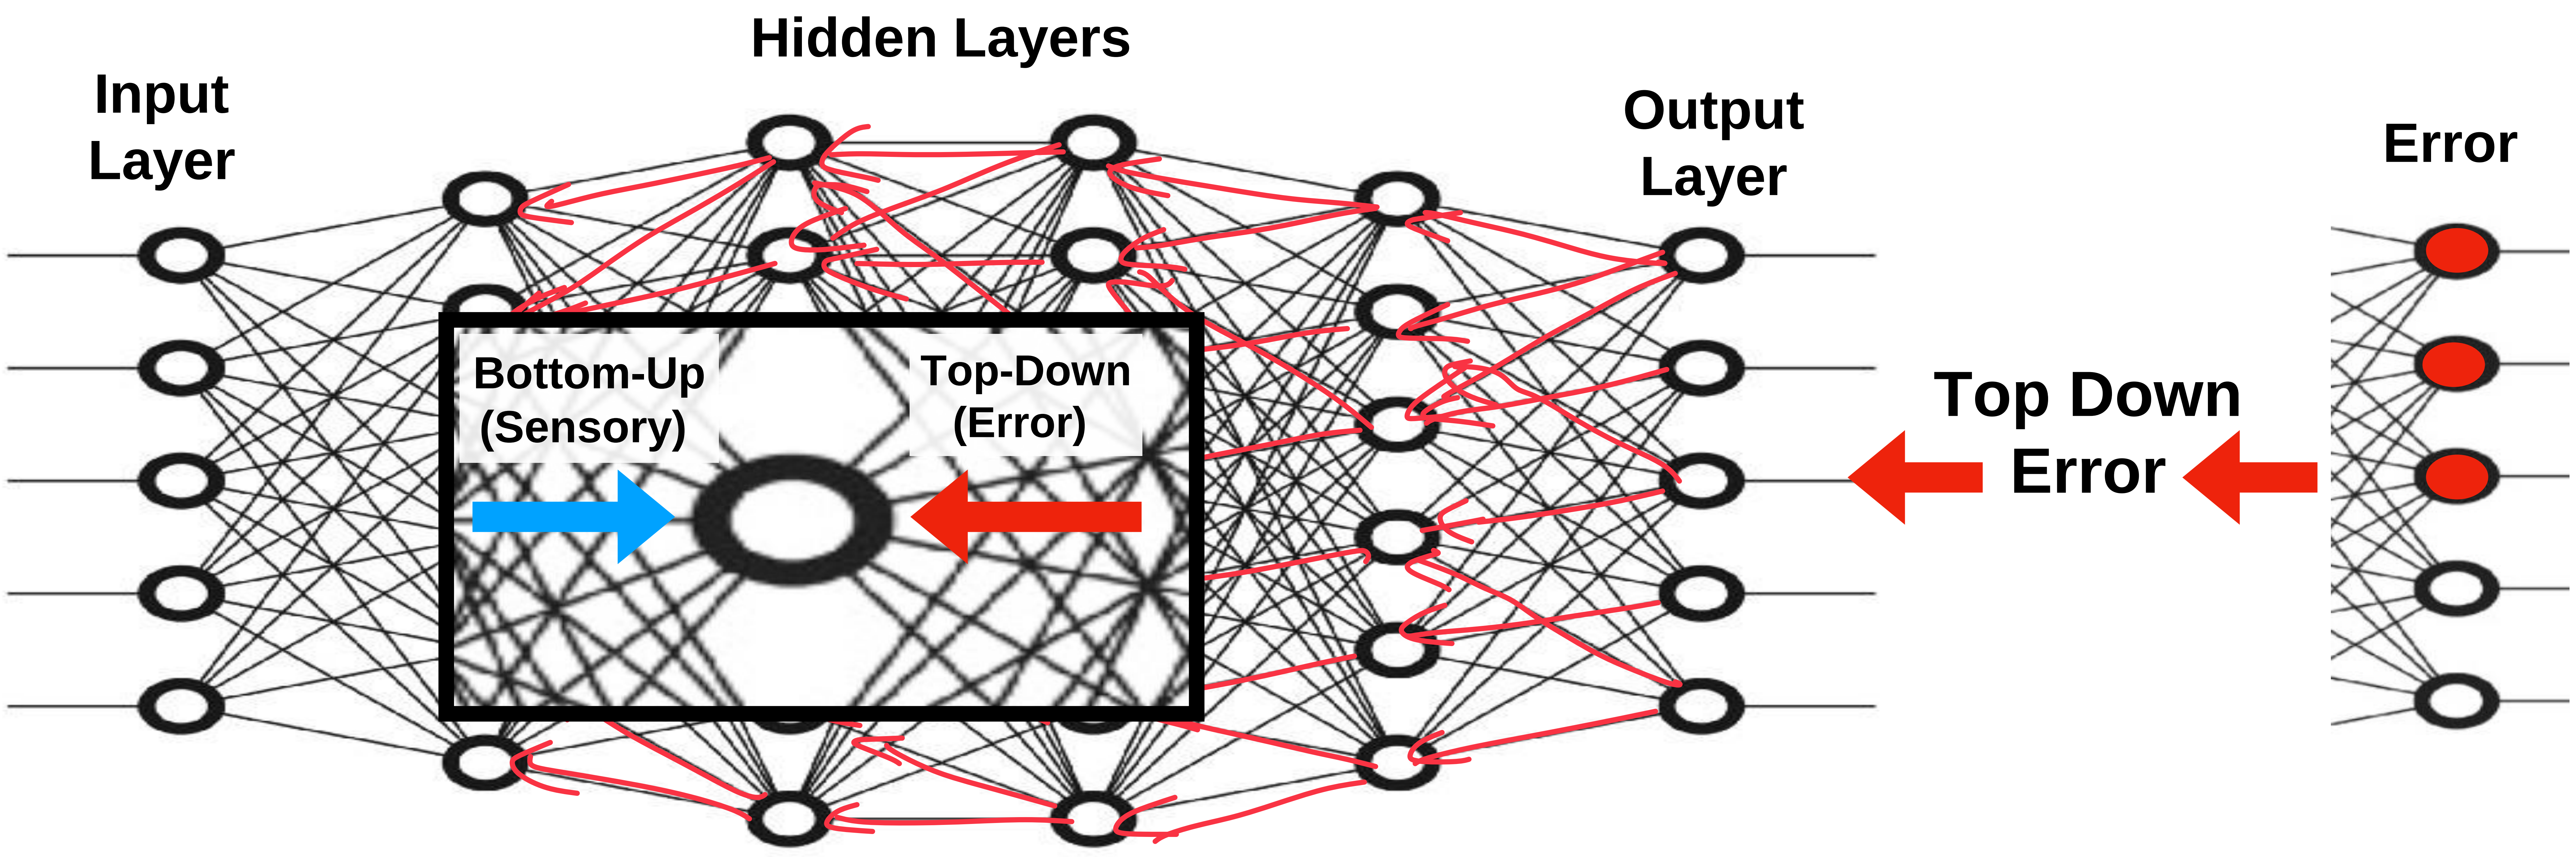

Deep Neuronal Network

Background and Motivation - Part I
The Error Backpropagation Method

Why does Backpropagation not explain how credit assignment in biology works?

- Requires separate forward and backward phases.
- Sends sensory information forward, but error backward.
- Sends errors backward though the SAME weights.
- Is based on discrete stepwise computations.
- Neuron update/plasticity not biologically observed.

Background and Motivation - Part I
Seeking Inspiration from Biology

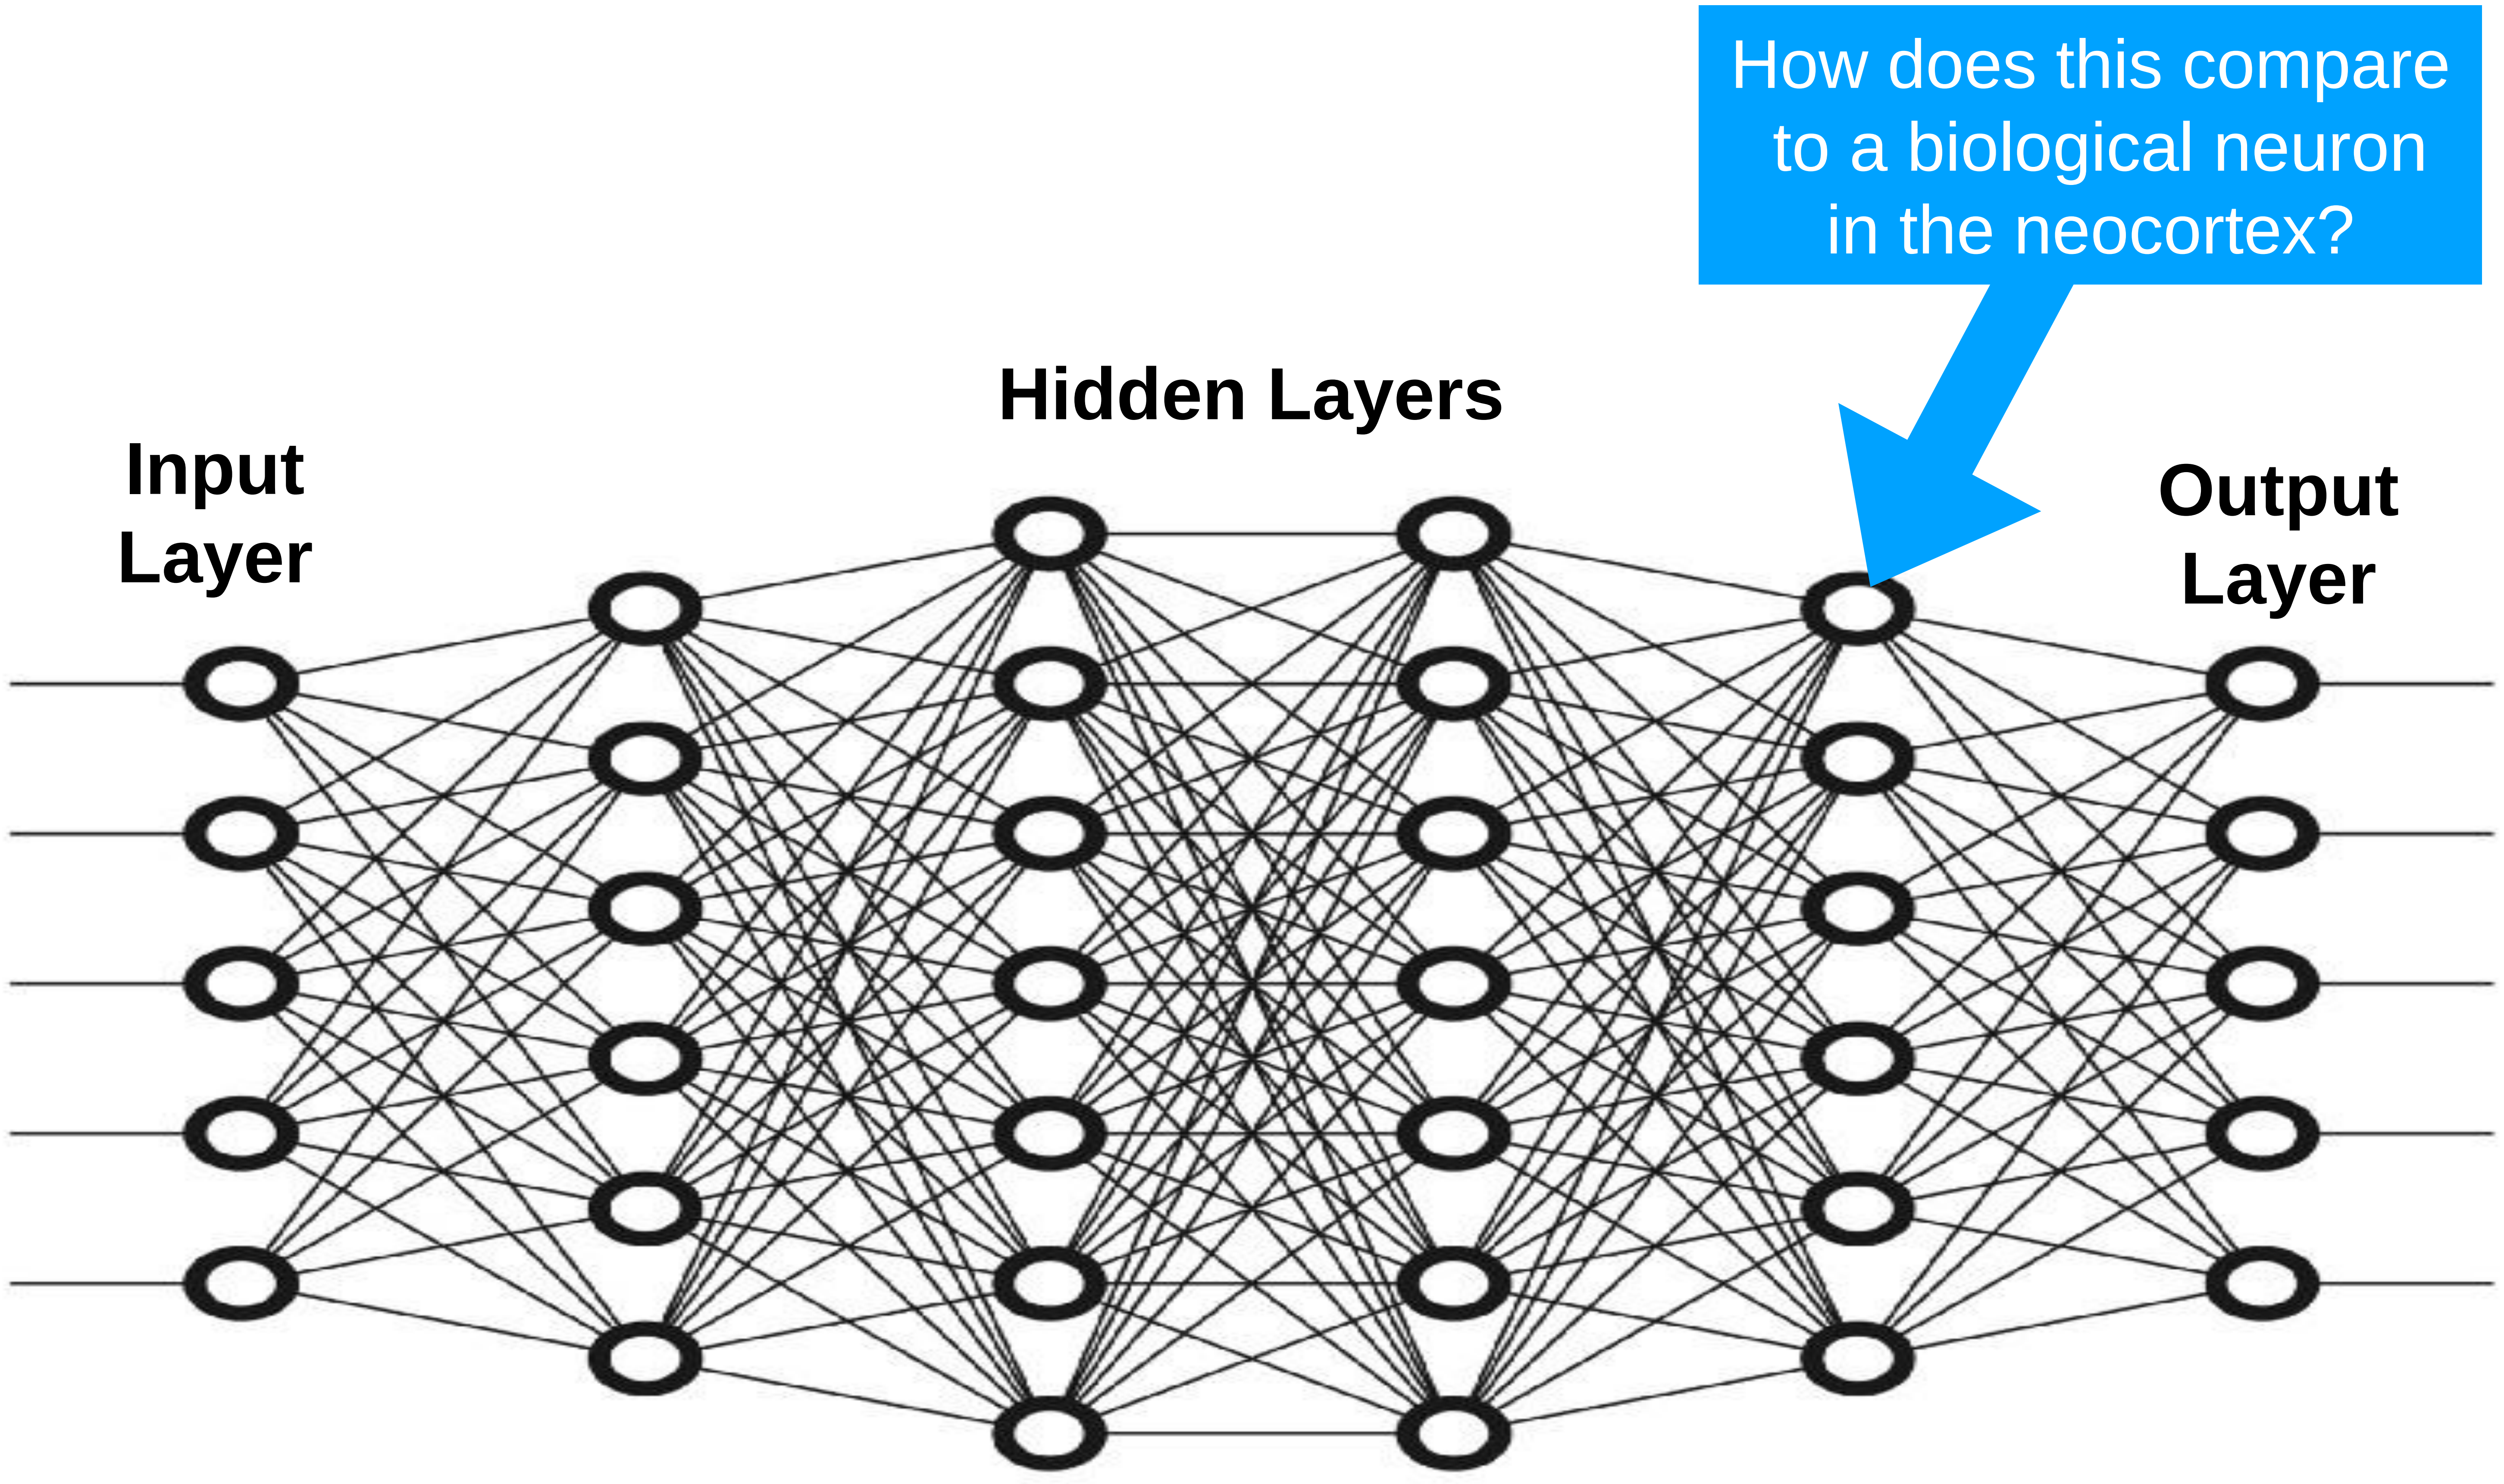

Background and Motivation - Part I
Seeking Inspiration from Biology

Neocortex

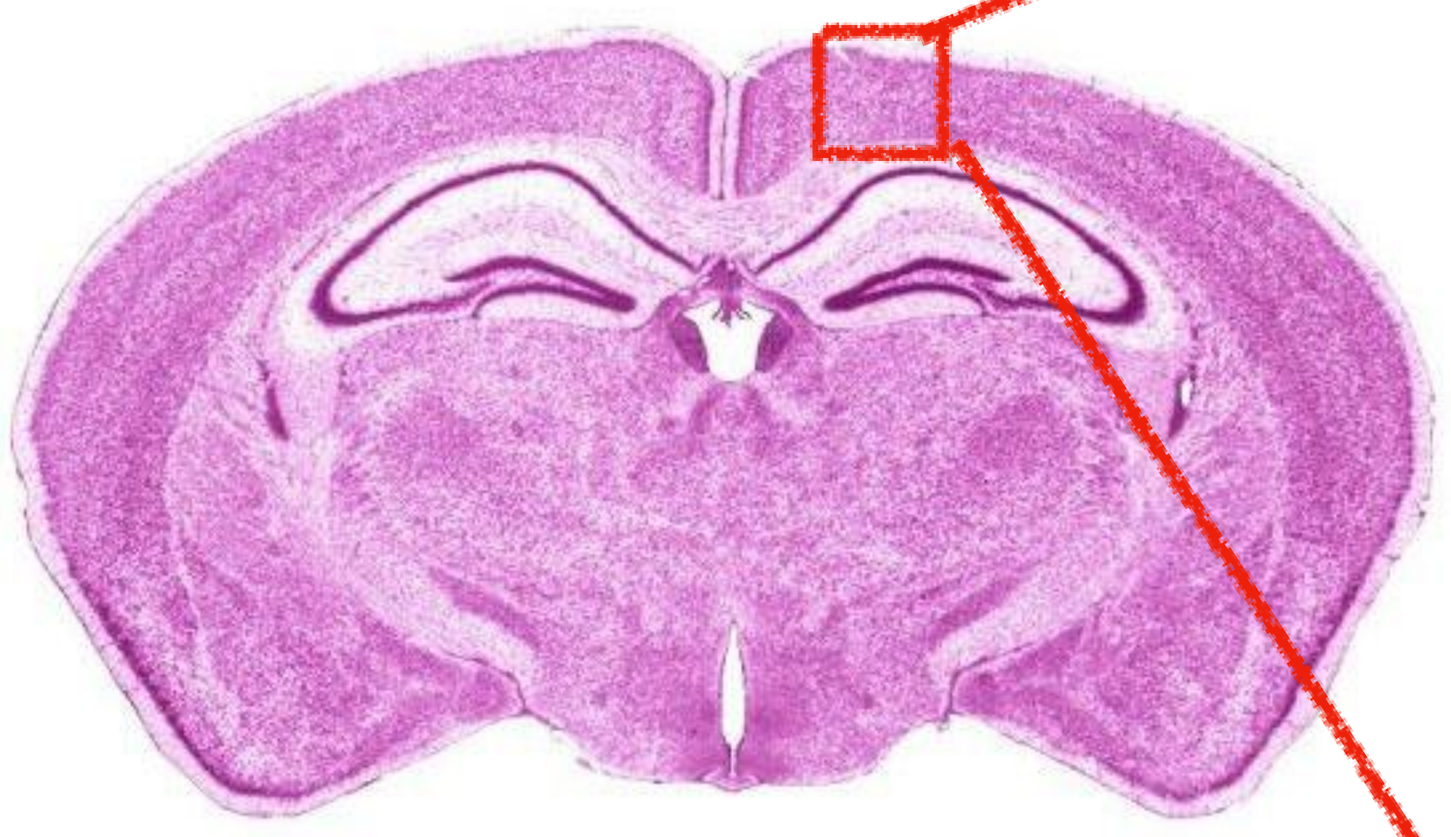

Mouse Brain Slice

Neocortex

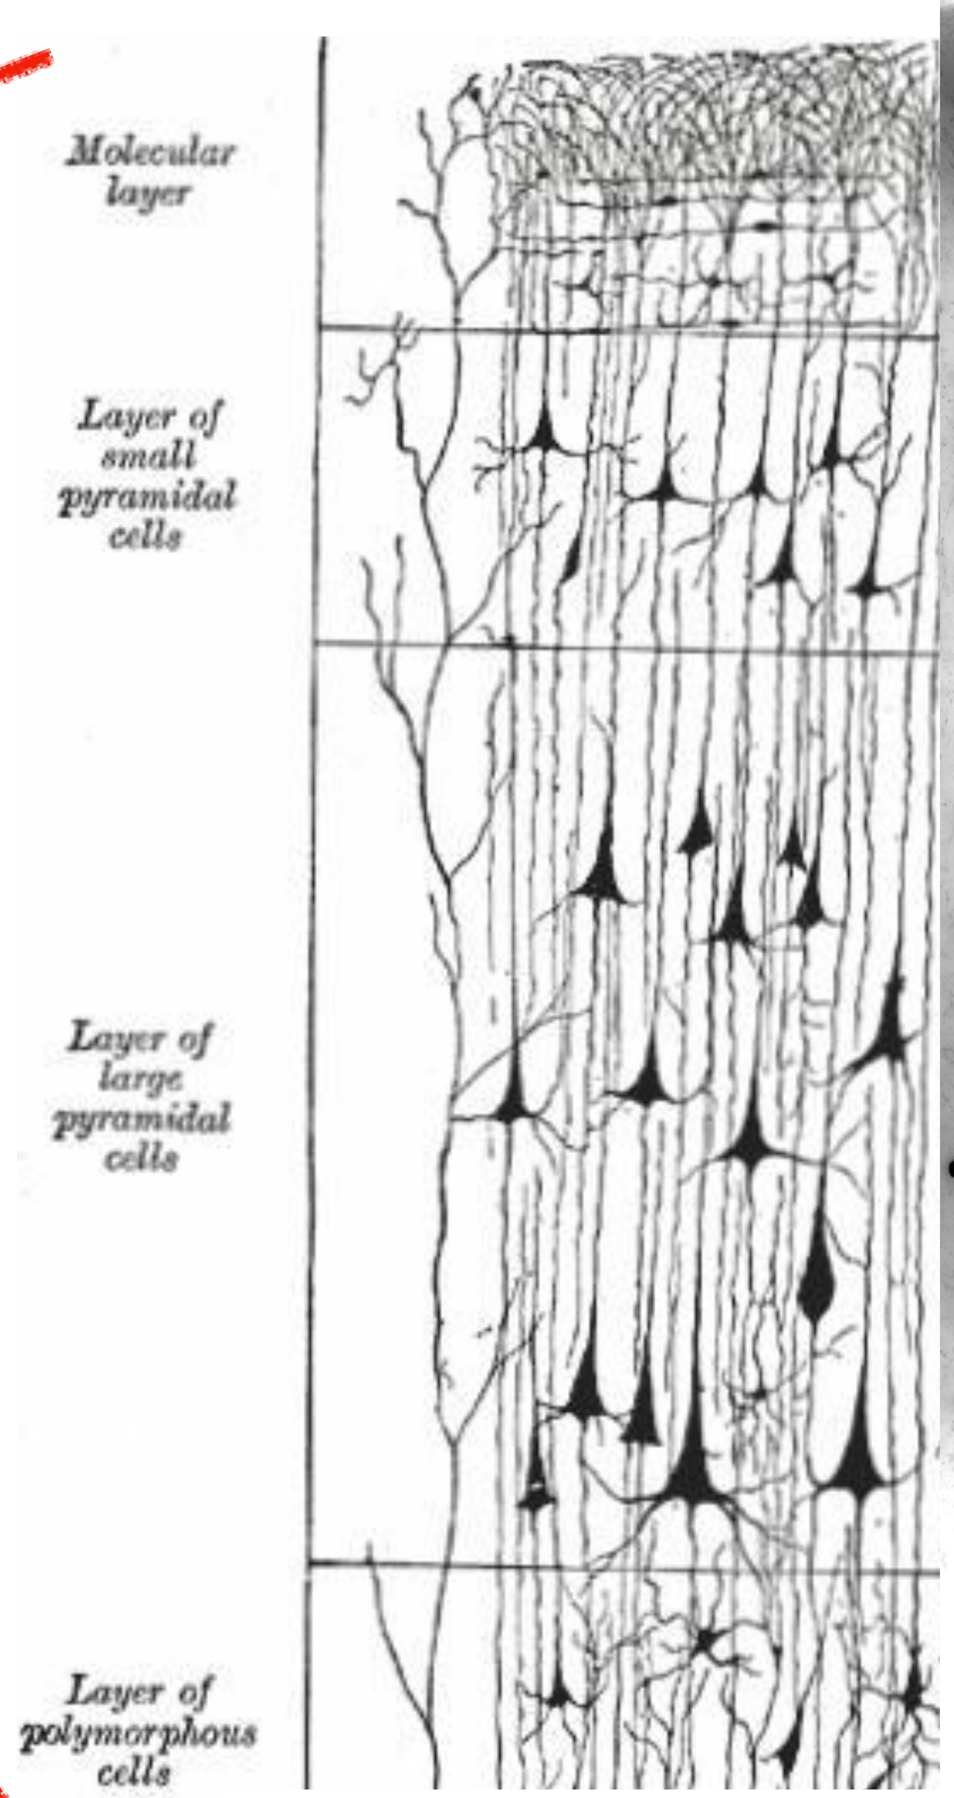

Roman y Cajal, 1906

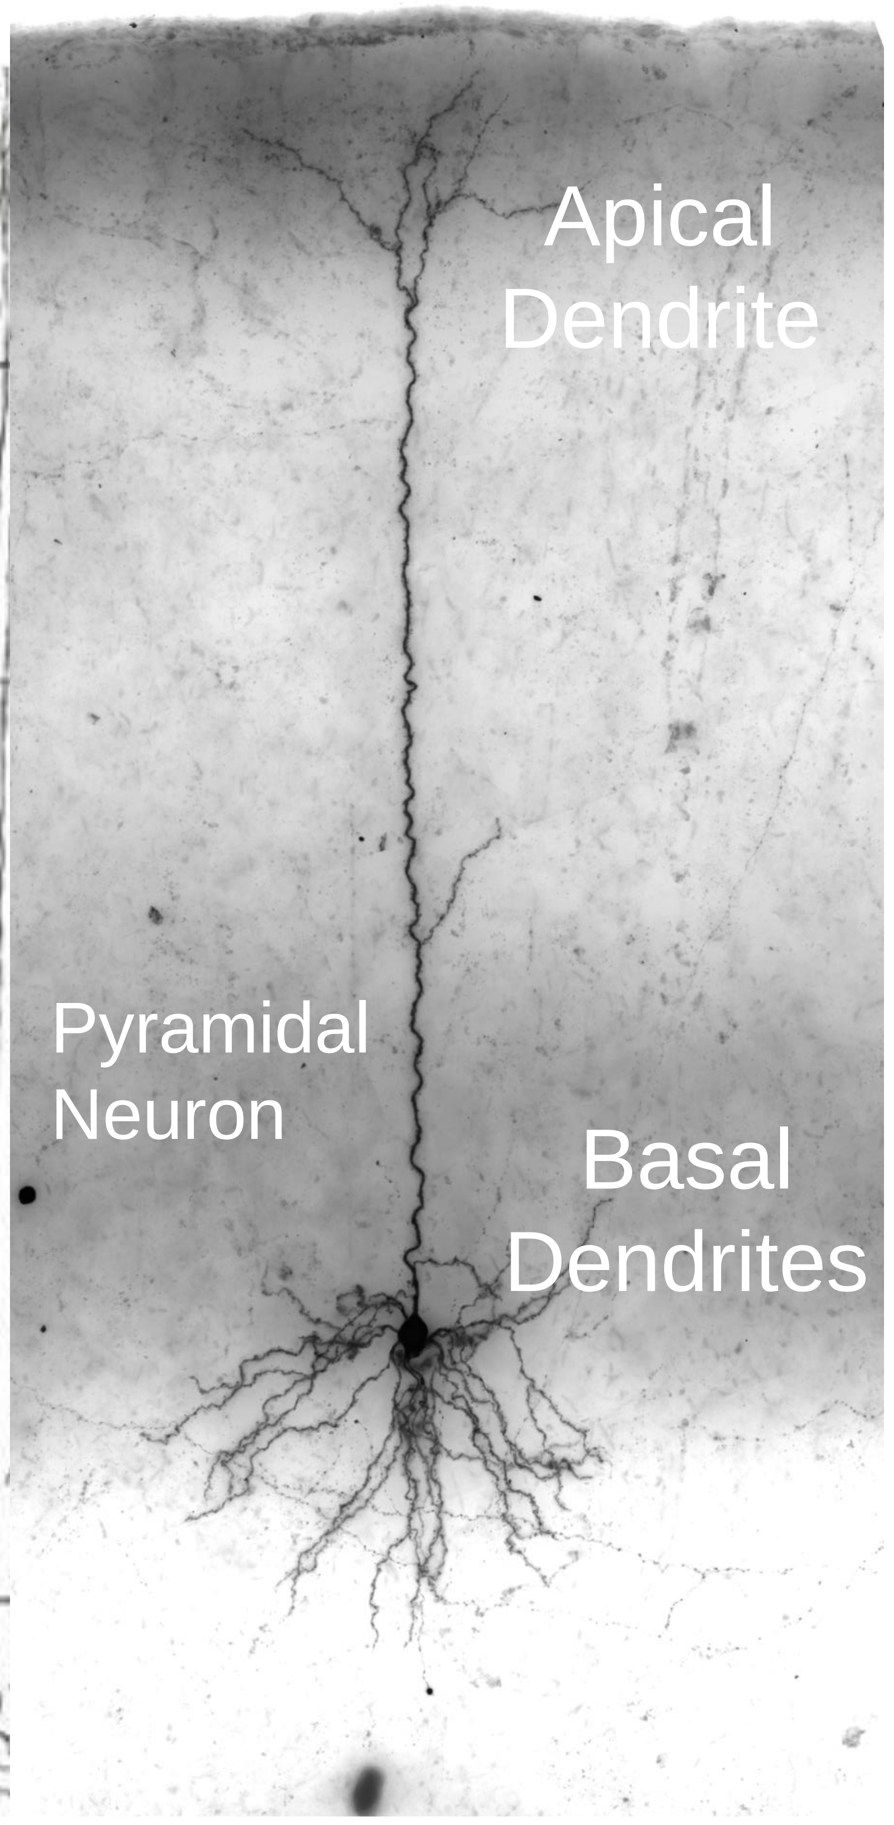

Grewe et al. 2010

ETH zürich

Background and Motivation - Part I
Seeking Inspiration from Biology

institute of neuroinformatics

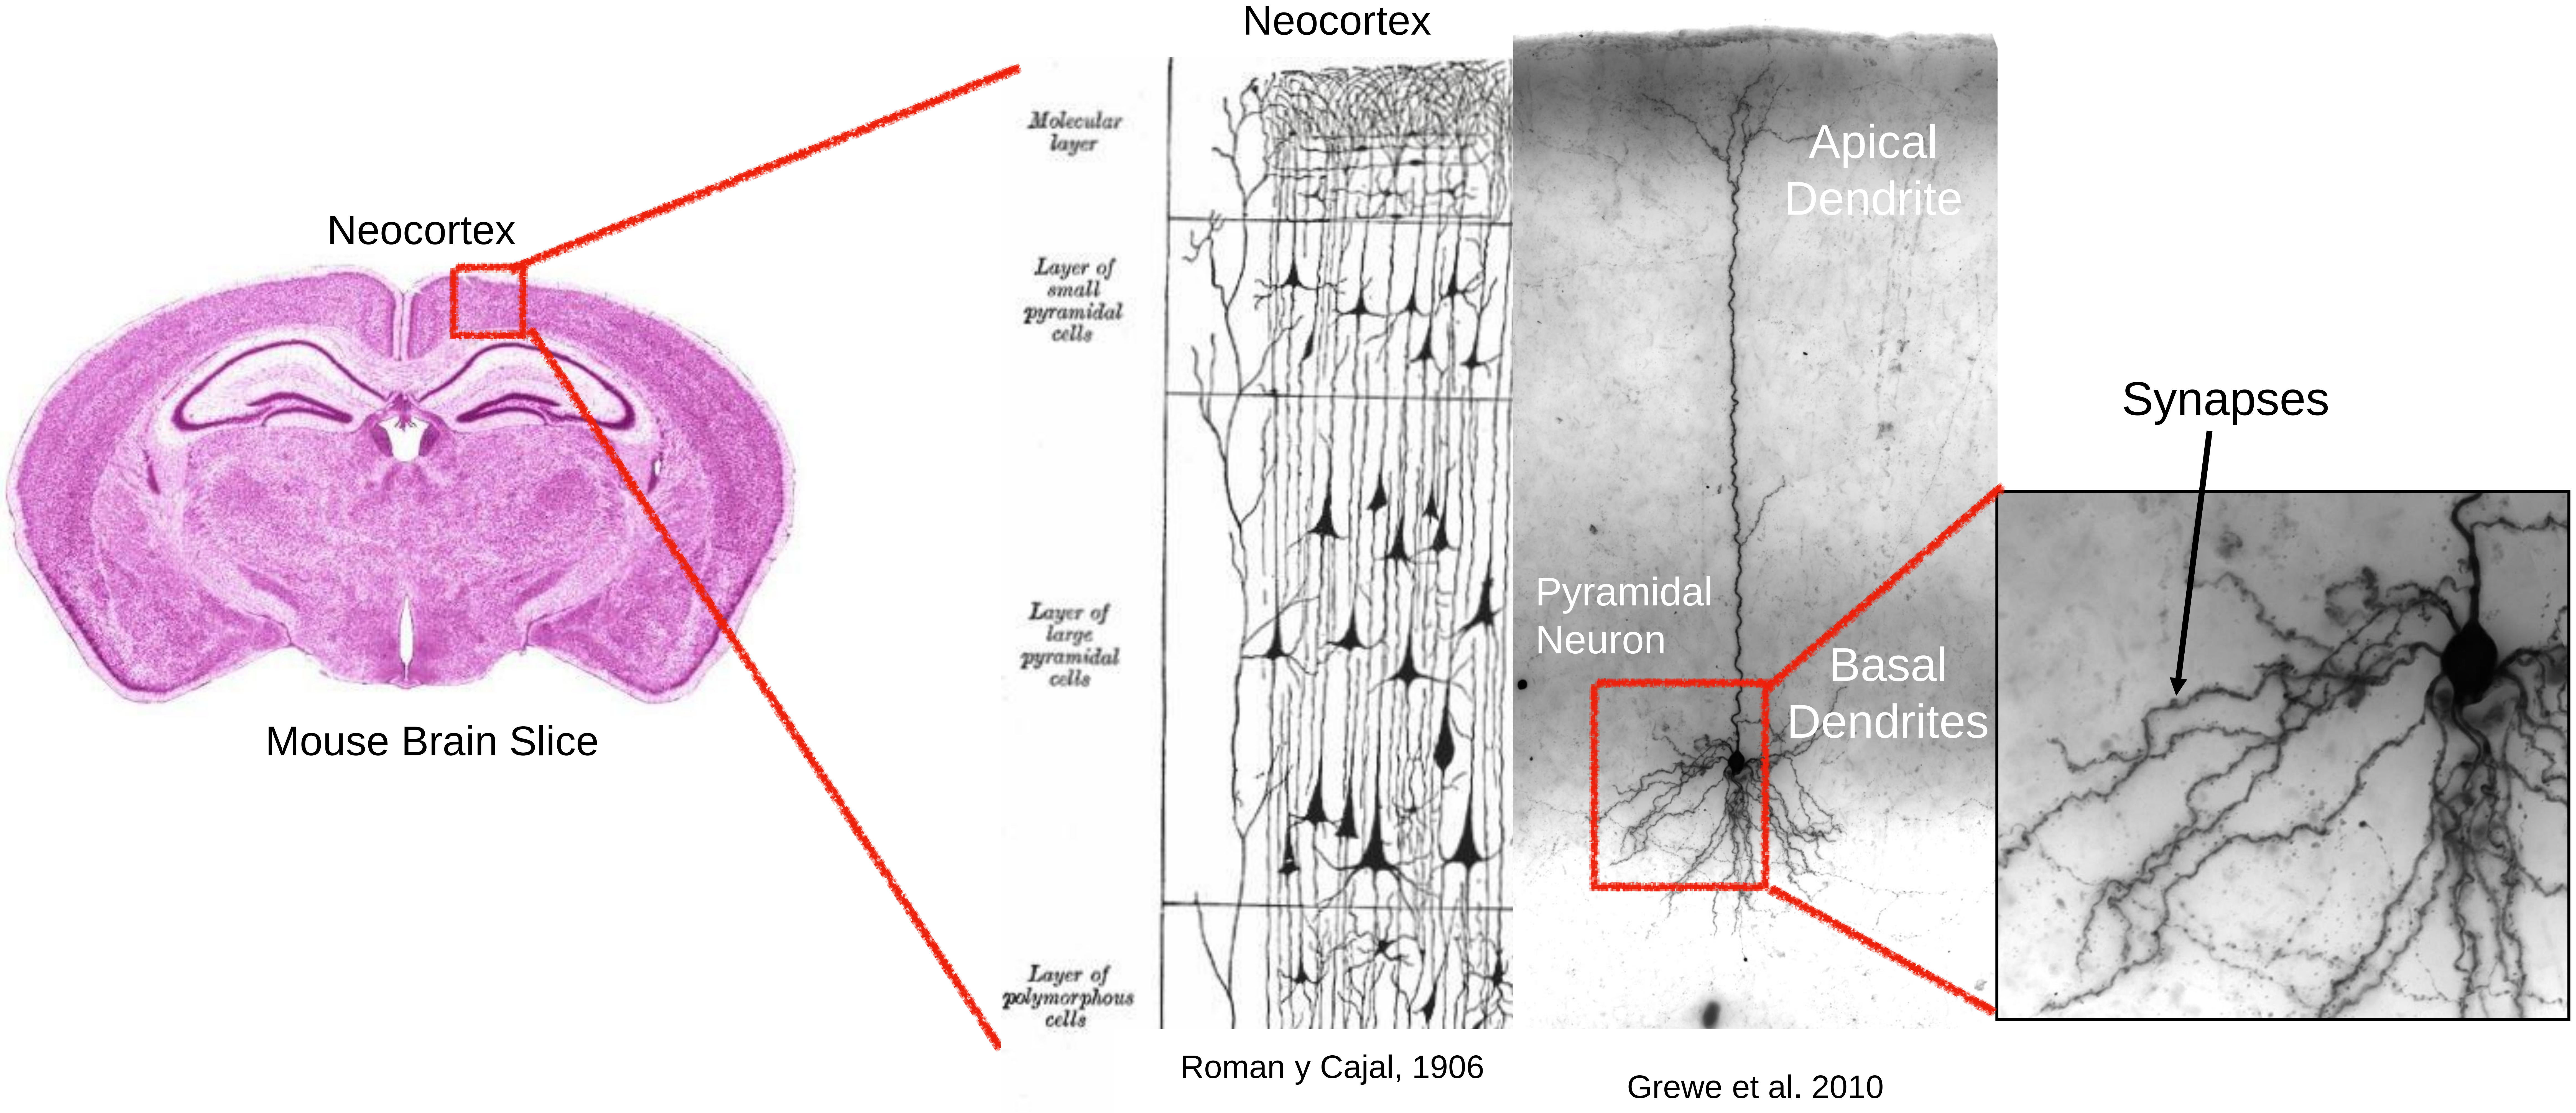

Background and Motivation - Part I
Seeking Inspiration from Biology

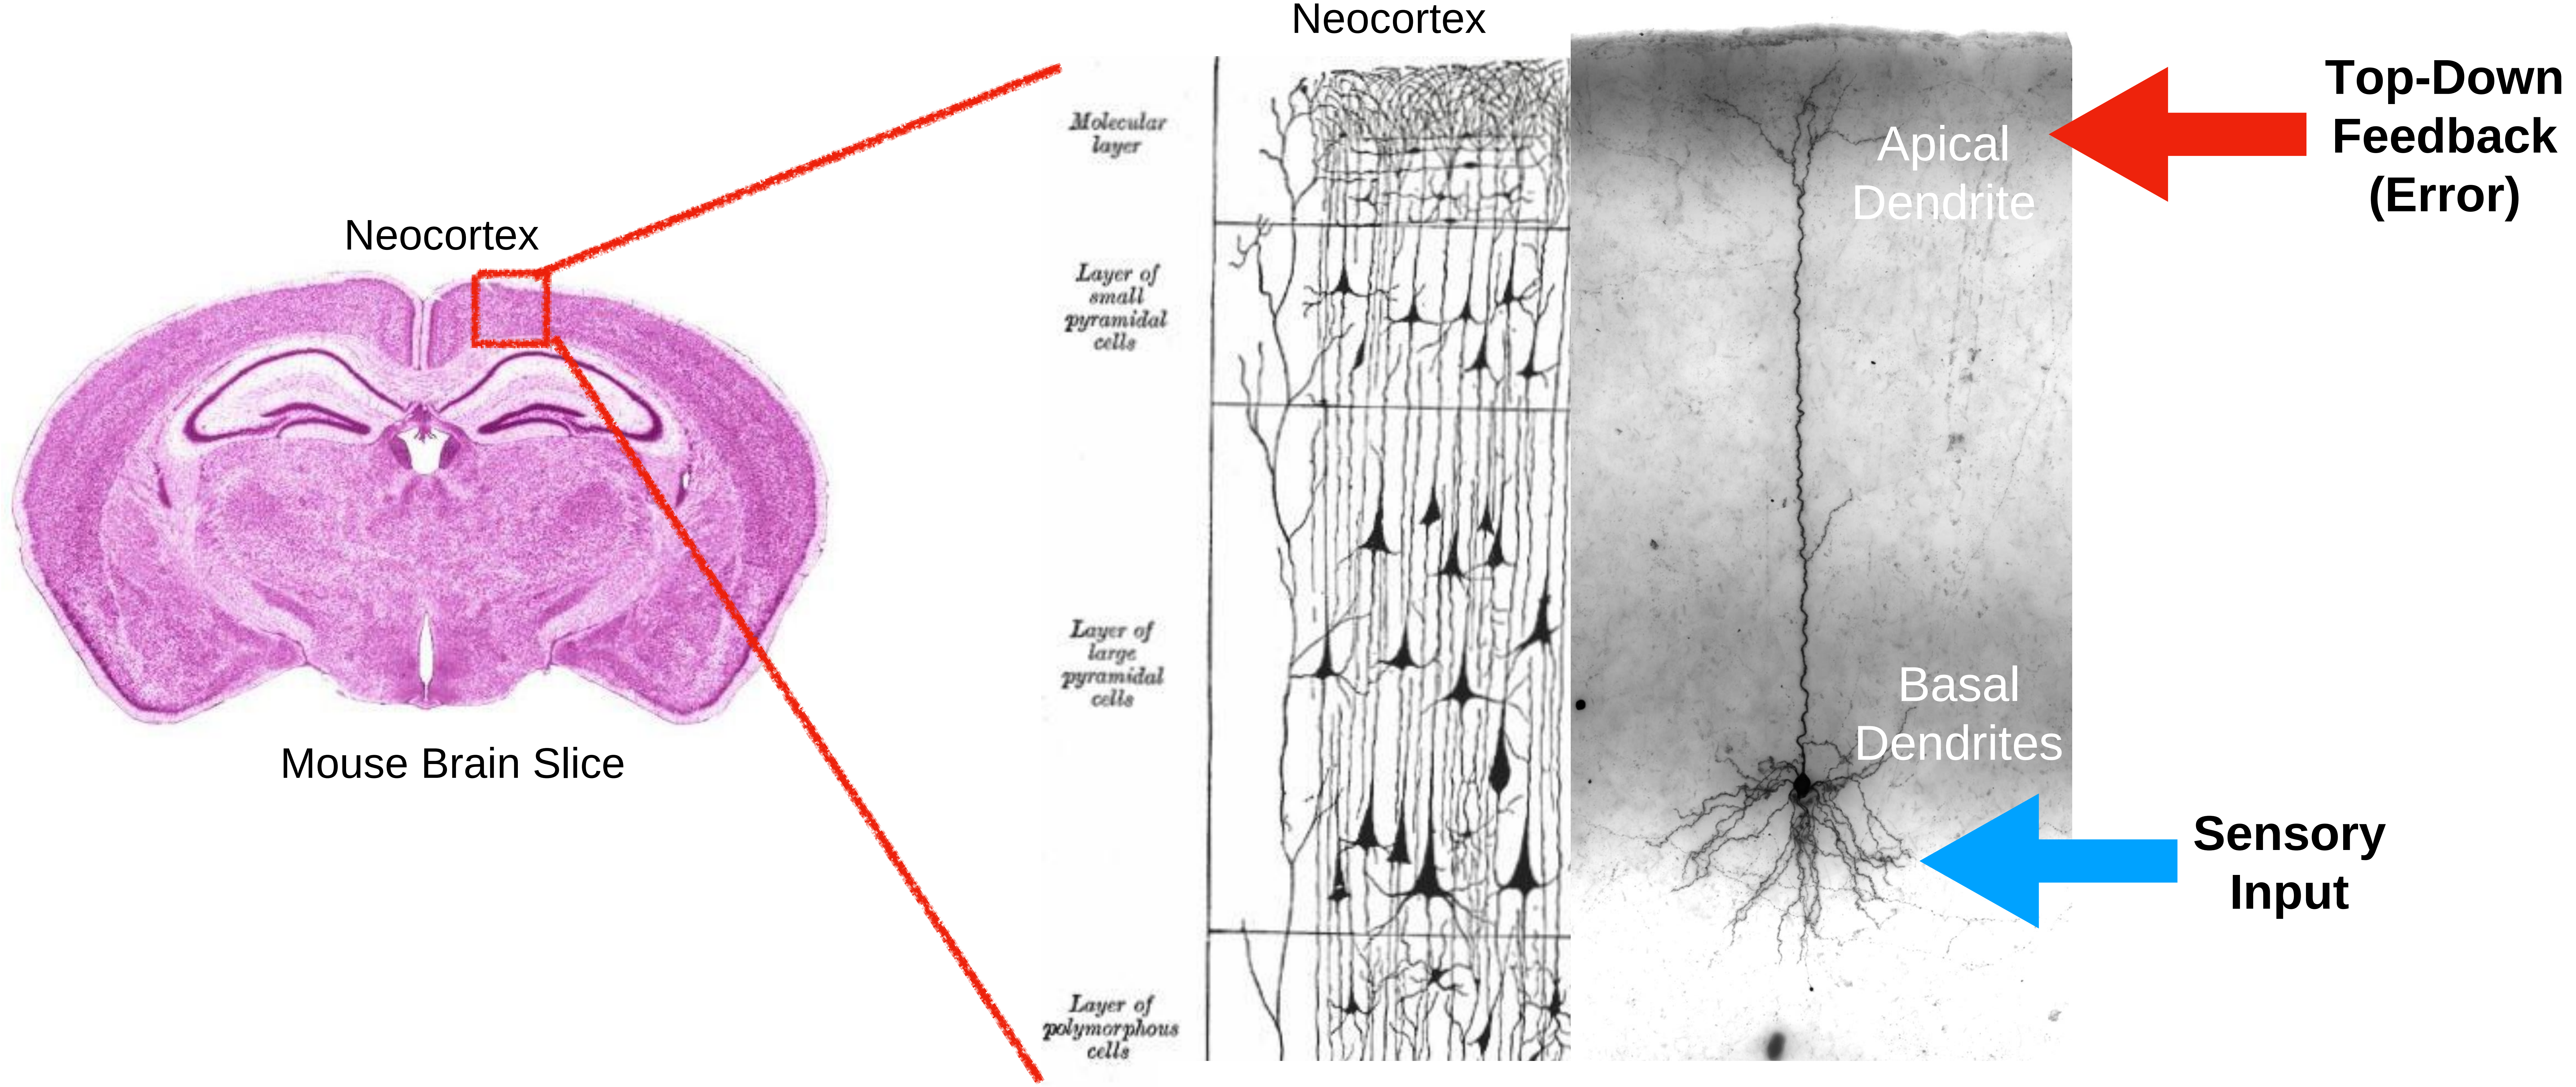

Roman y Cajal, 1906

Grewe et al. 2010

ETH zürich

Deep Learning Through Feedback Control

Matilde T. Farinha

Alexander Meulemans

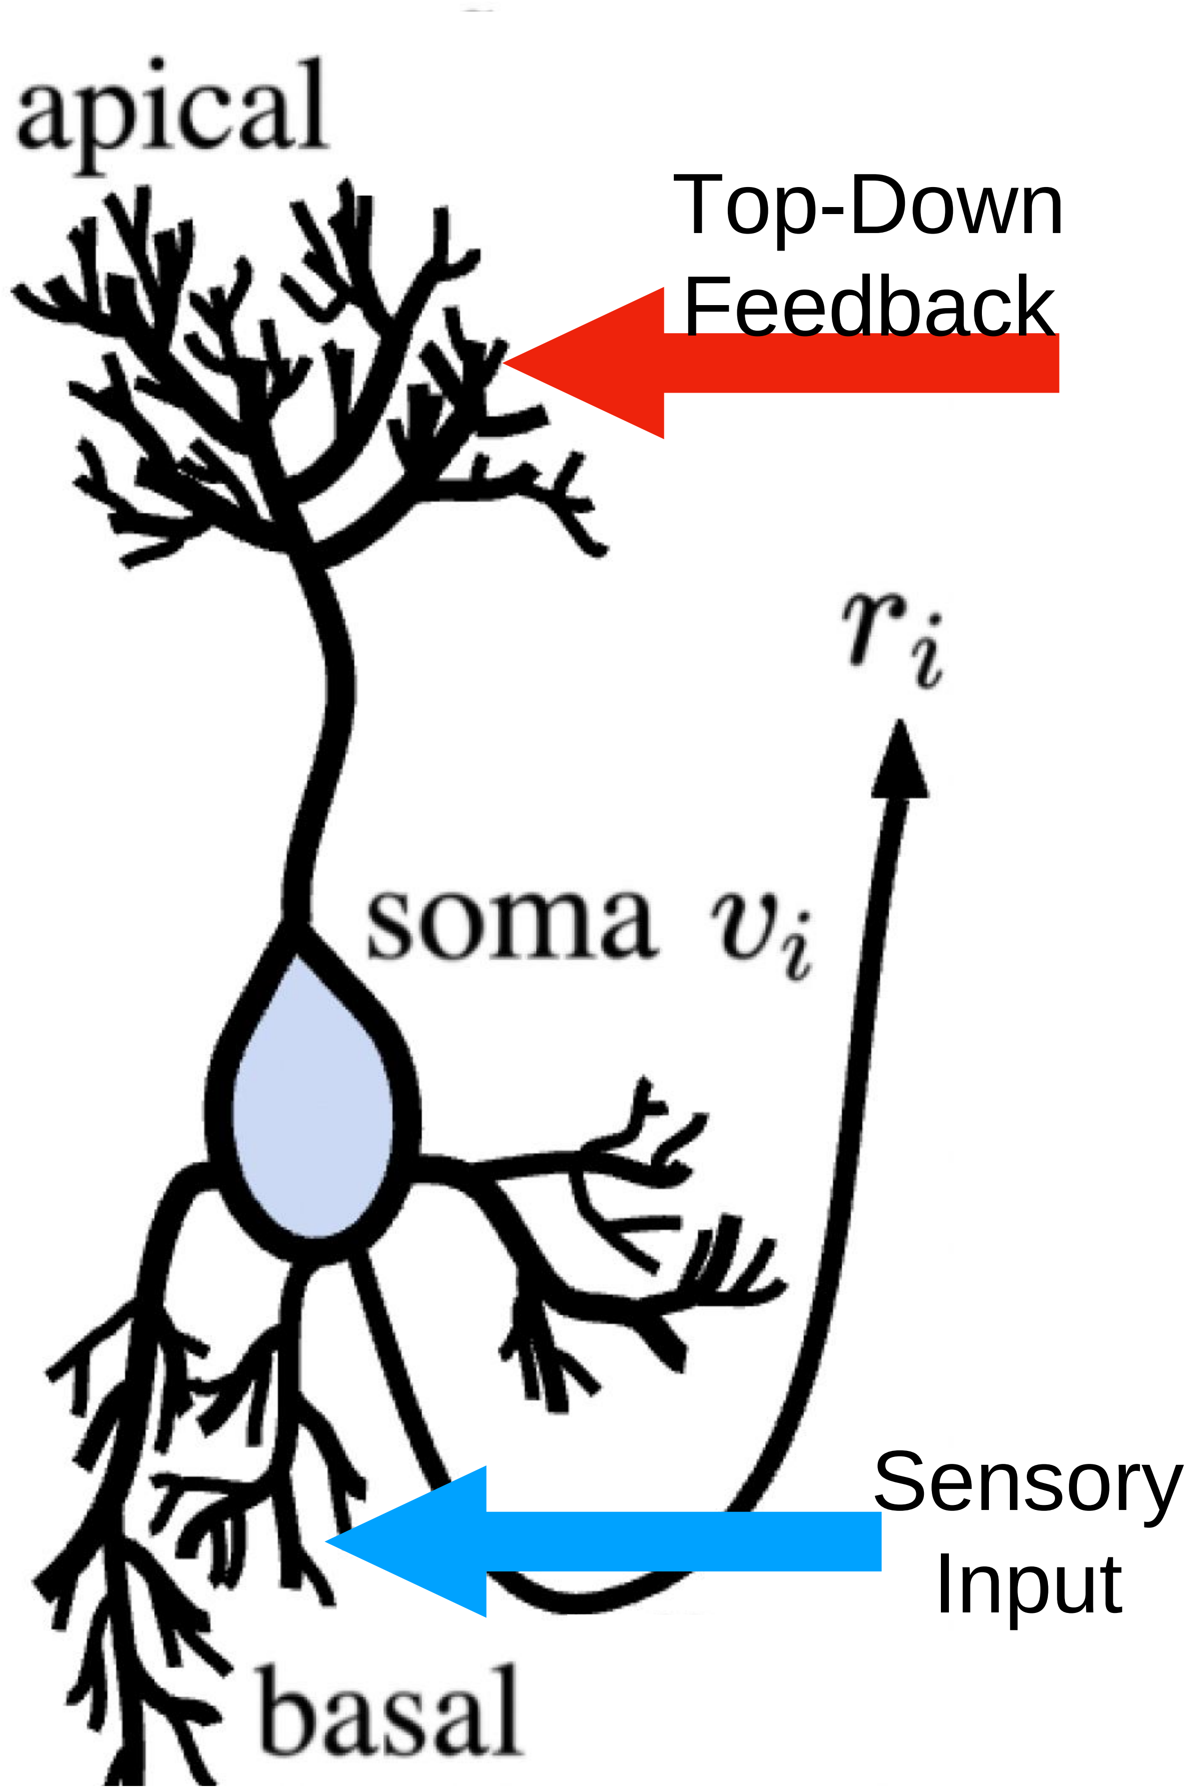

v_i membrane potential
r_i neuron firing rate

Meulemans et al. 2020, 2021

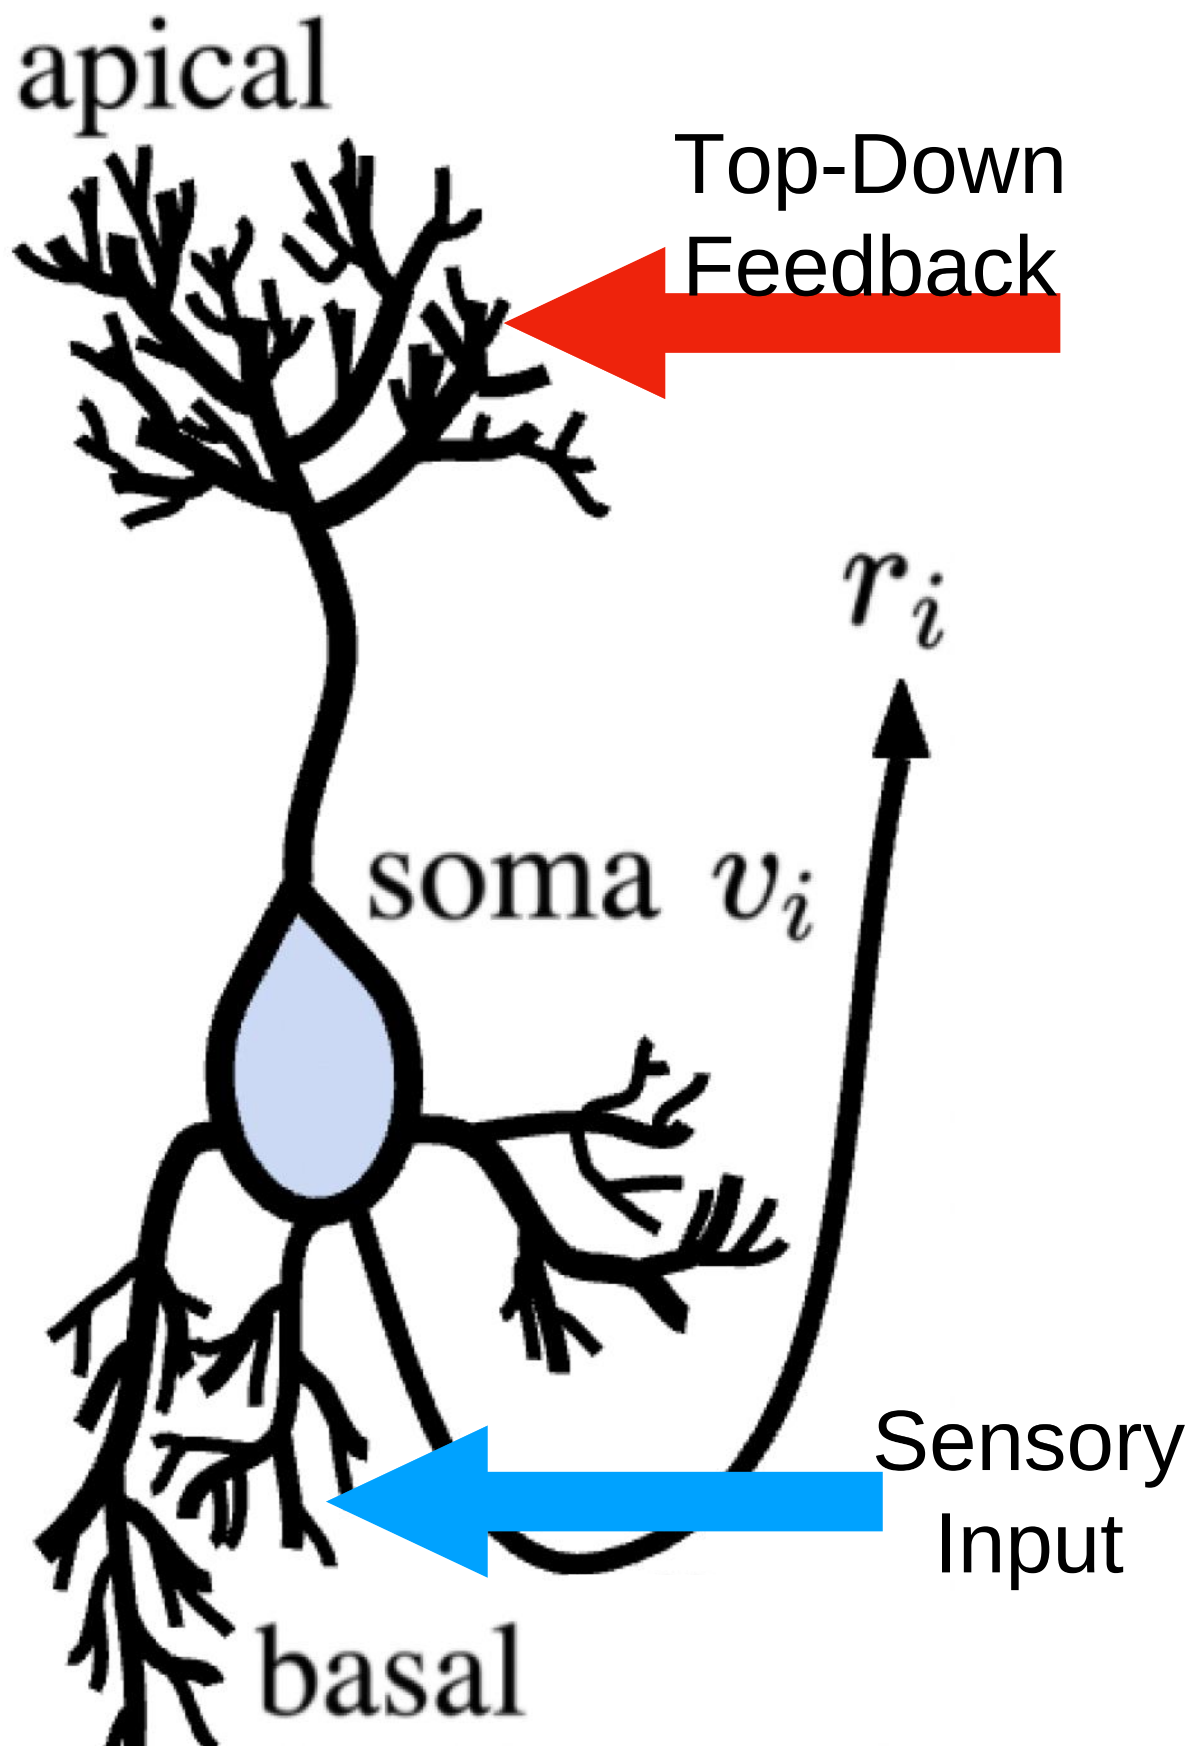

Neuron Output / Firing Rate

$$\mathbf{r}_i = \phi(\mathbf{v}_i)$$

Non. Linearity

v_i membrane potential
r_i neuron firing rate

Meulemans et al. 2020, 2021

Deep Learning Through Feedback Control

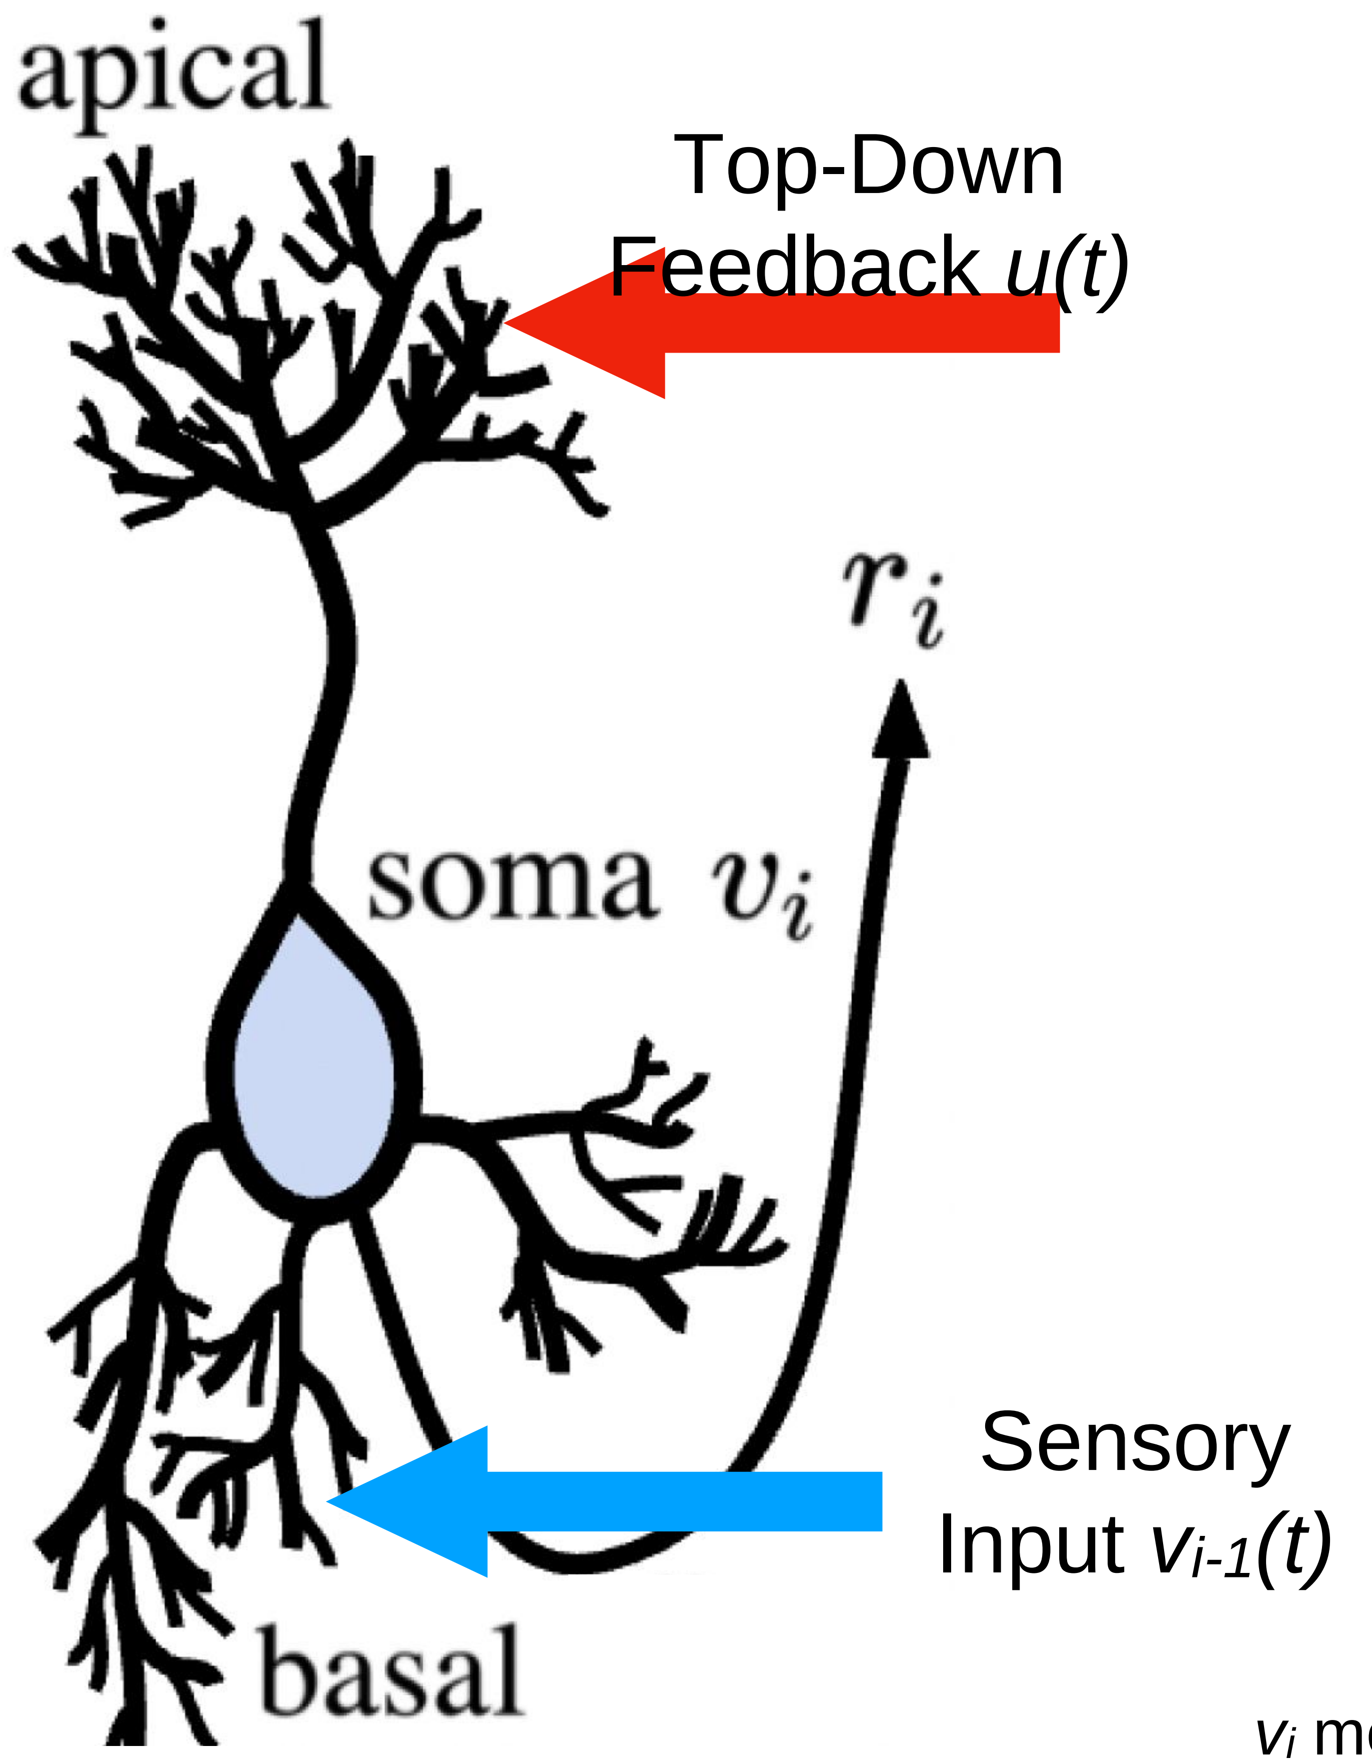

Neuron Output / Firing Rate

$$\mathbf{r}_i = \phi(\mathbf{v}_i)$$

Non. Linearity

Membrane Potential Dynamics

$$\tau_v \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{v}_i(t) = -\mathbf{v}_i(t) + \underbrace{\boxed{W_i\phi\big(\mathbf{v}_{i-1}(t)\big)}}_{\text{Sensory}} + \underbrace{\boxed{Q_i\mathbf{u}(t)}}_{\text{Feedback}}$$

v_i membrane potential
r_i neuron firing rate

u(t) Feedback signal
Q_i Feedback weights

Meulemans et al. 2020, 2021

Deep Learning Through Feedback Control

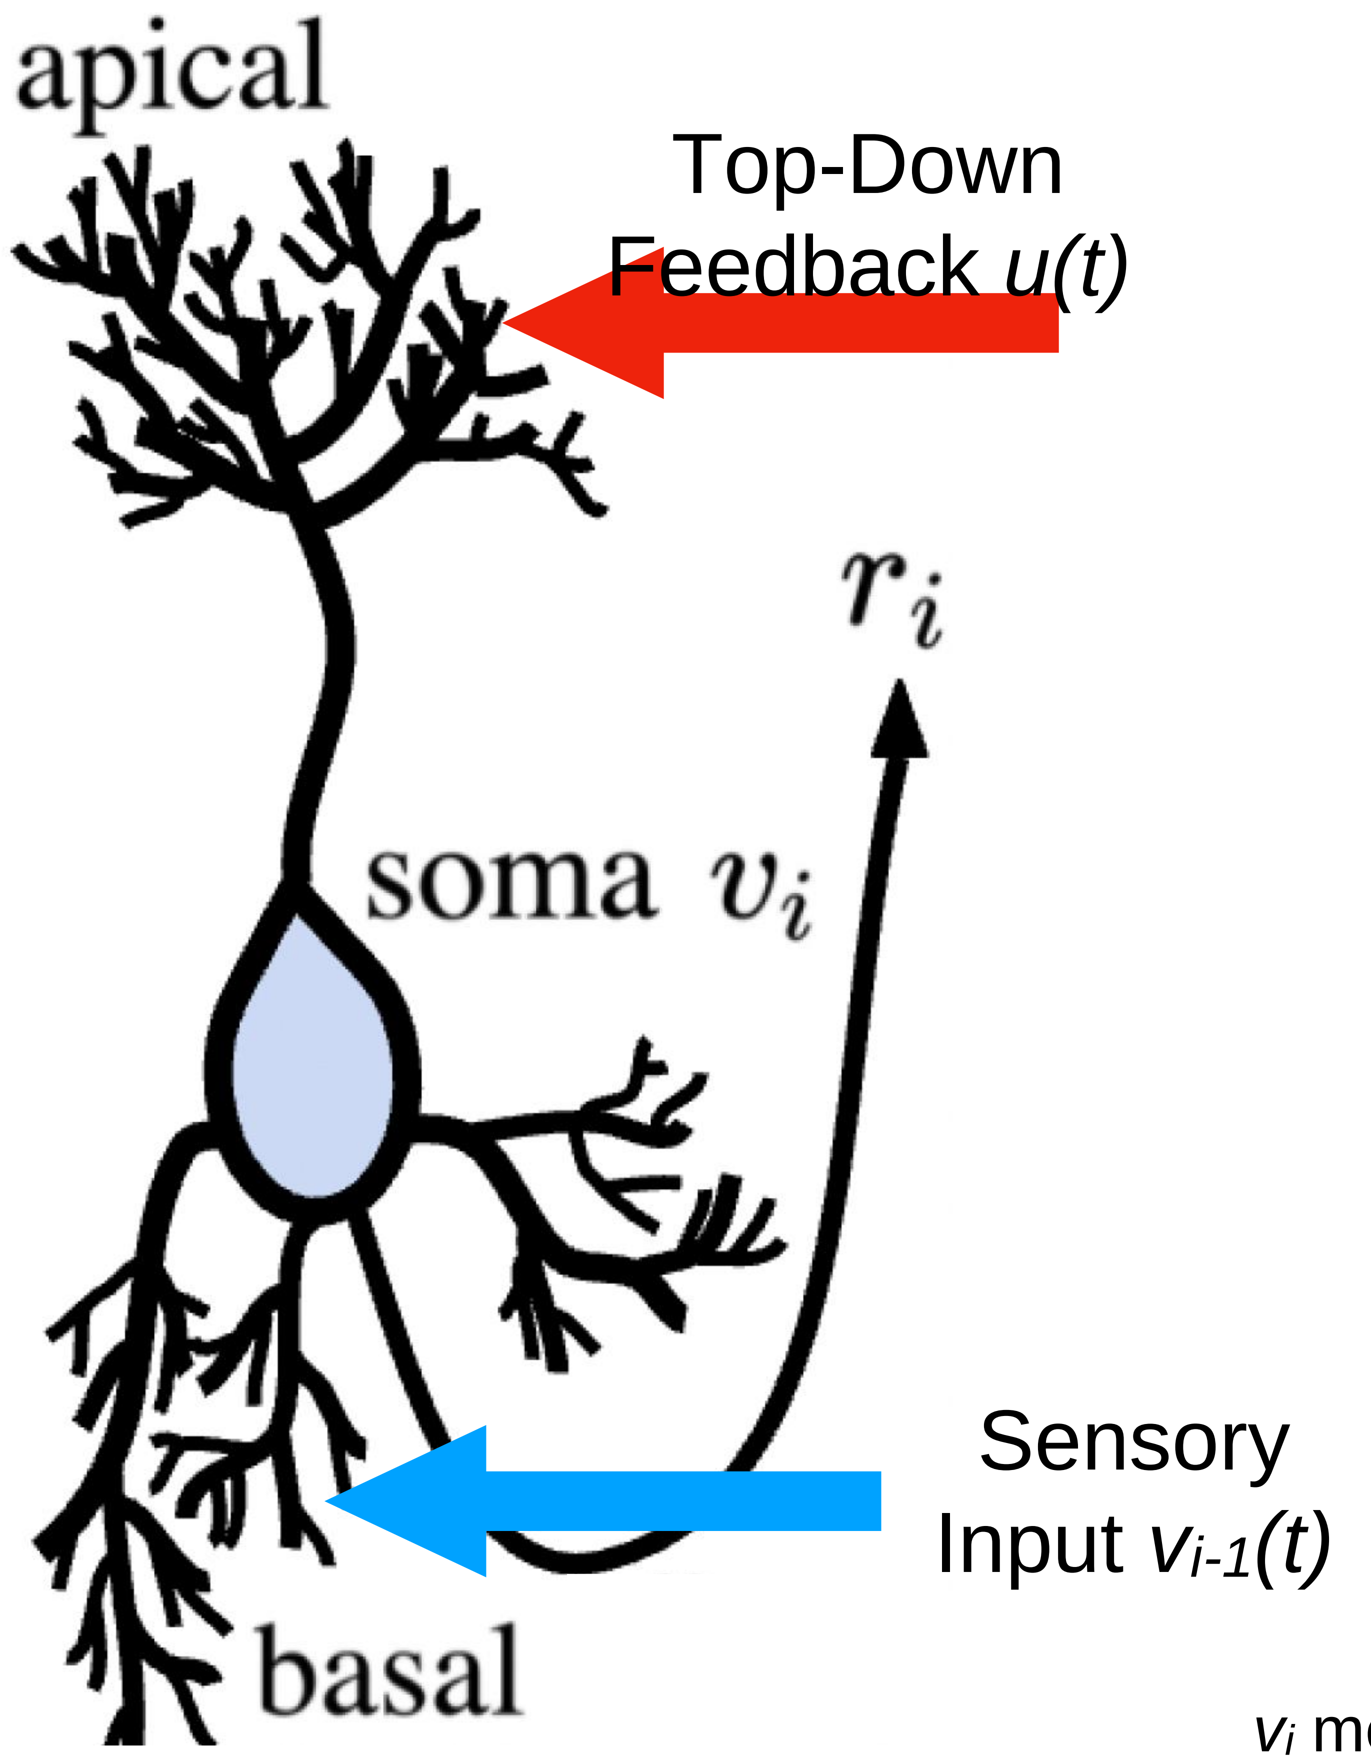

Neuron Output / Firing Rate

$$\mathbf{r}_i = \underbrace{\phi}_{\text{Non. Linearity}}(\mathbf{v}_i)$$

Membrane Potential Dynamics

$$\tau_v \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{v}_i(t) = -\mathbf{v}_i(t) + \underbrace{\boxed{W_i\phi\big(\mathbf{v}_{i-1}(t)\big)}}_{\text{Sensory}} + \underbrace{\boxed{Q_i\mathbf{u}(t)}}_{\text{Feedback}}$$

How can we train this single neuron to detect a specific sensory input pattern?

v_i membrane potential
r_i neuron firing rate

$u(t)$ Feedback signal
Q_i Feedback weights

Meulemans et al. 2020, 2021

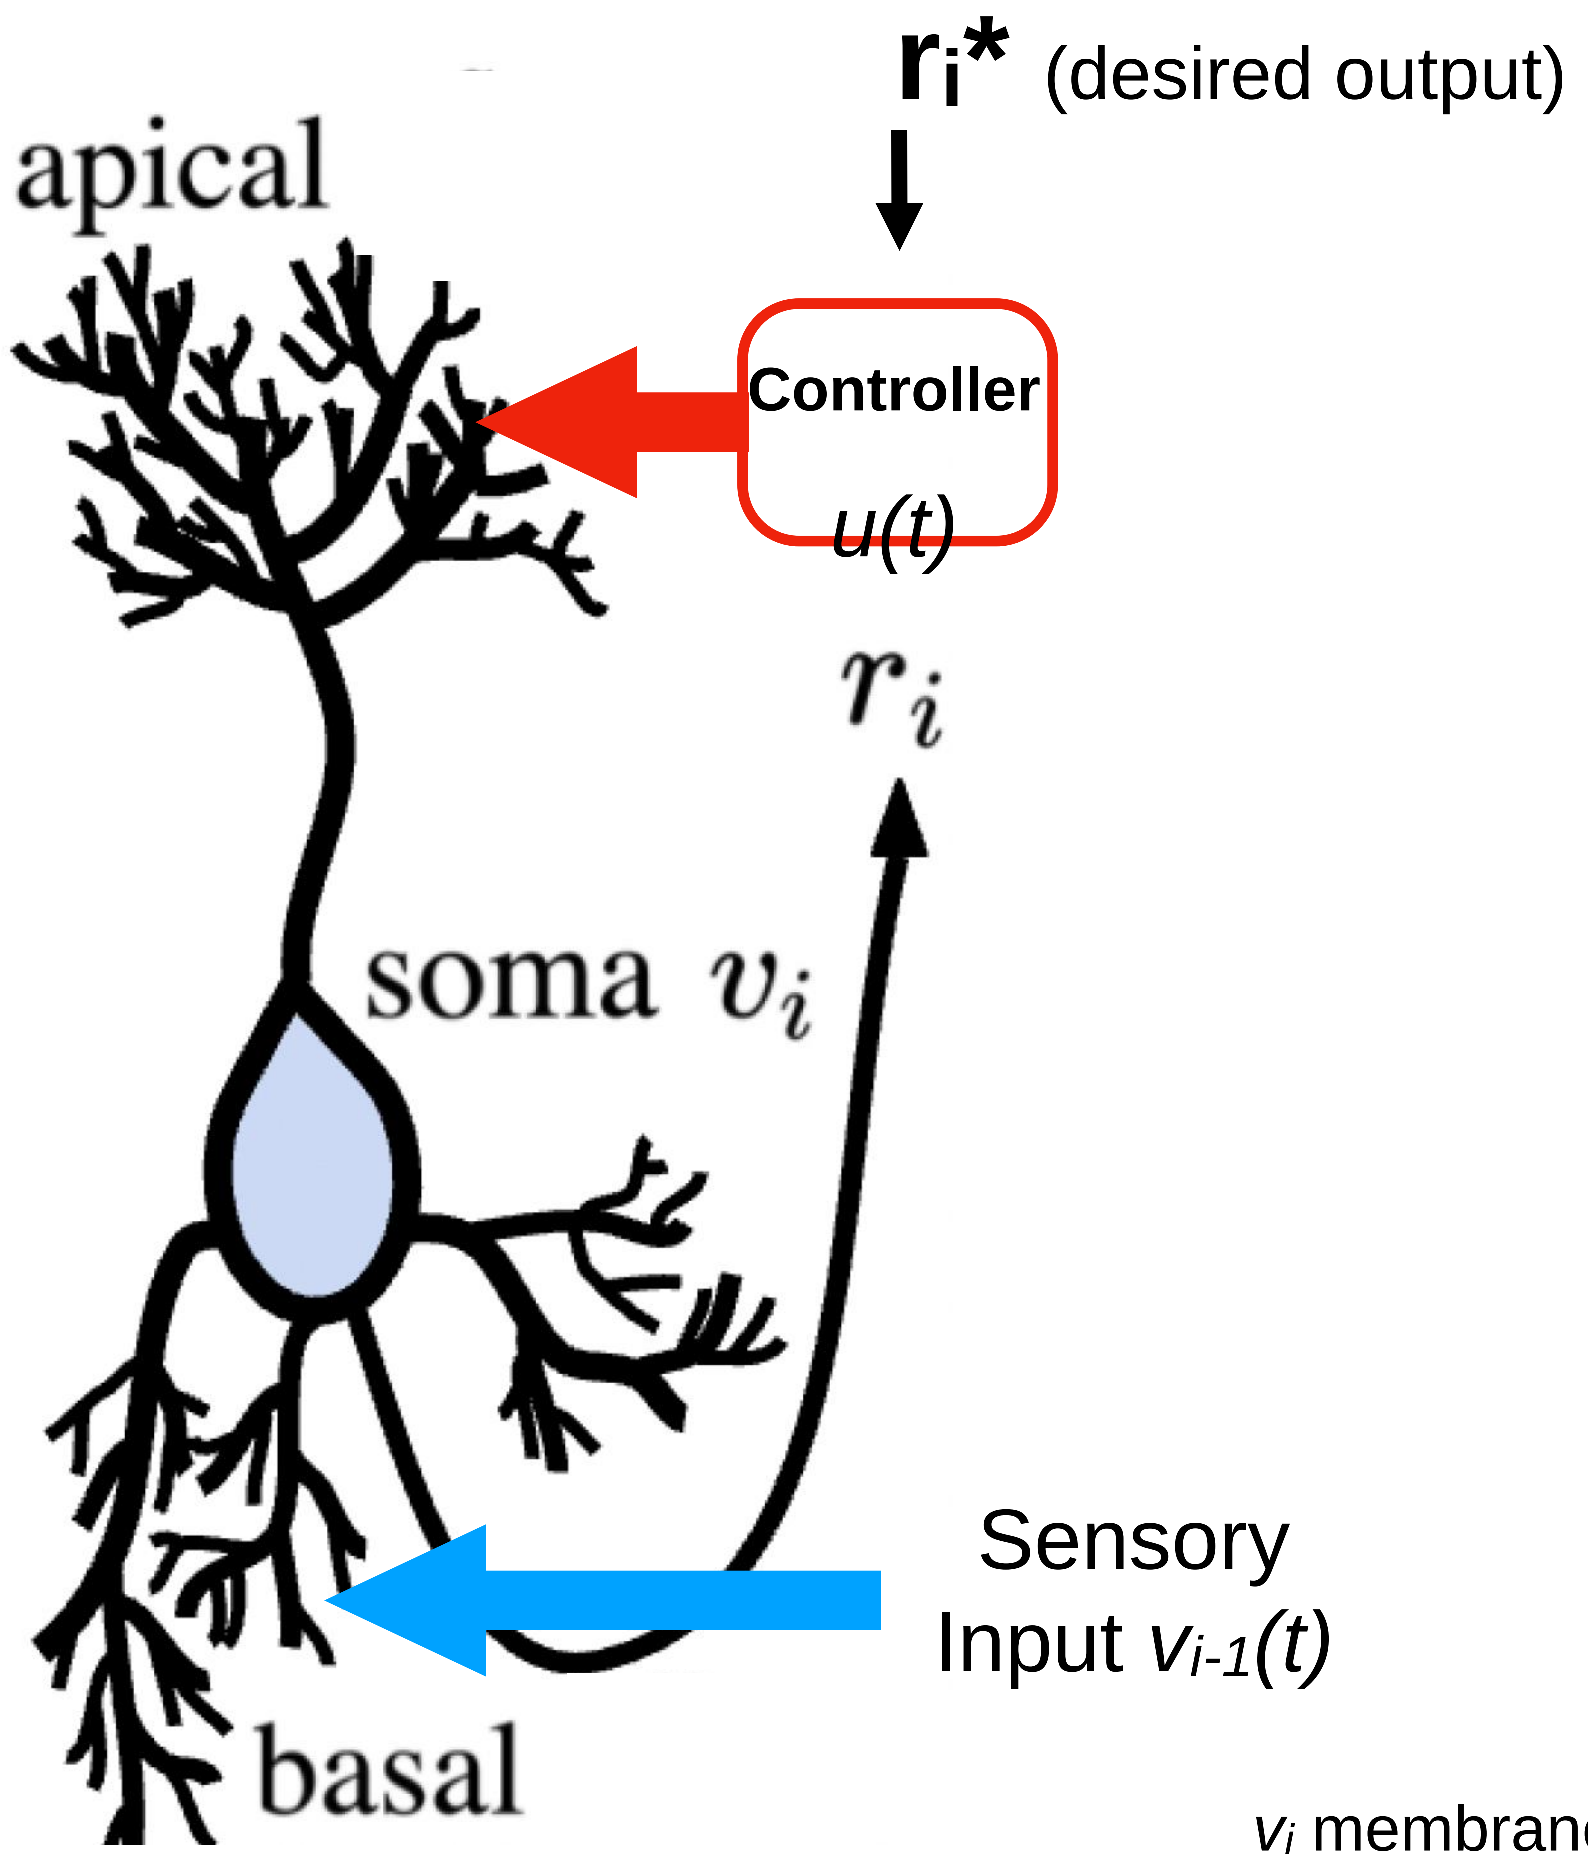

Neuron Output / Firing Rate

$$\mathbf{r}_i = \phi(\mathbf{v}_i)$$

Non. Linearity

Membrane Potential Dynamics

$$\tau_v \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{v}_i(t) = -\mathbf{v}_i(t) + \underbrace{\boxed{W_i\phi\big(\mathbf{v}_{i-1}(t)\big)}}_{\text{Sensory}} + \underbrace{\boxed{Q_i\mathbf{u}(t)}}_{\text{Feedback}}$$

v_i membrane potential
r_i neuron firing rate

$u(t)$ Feedback signal
Q_i Feedback weights

Meulemans et al. 2020, 2021

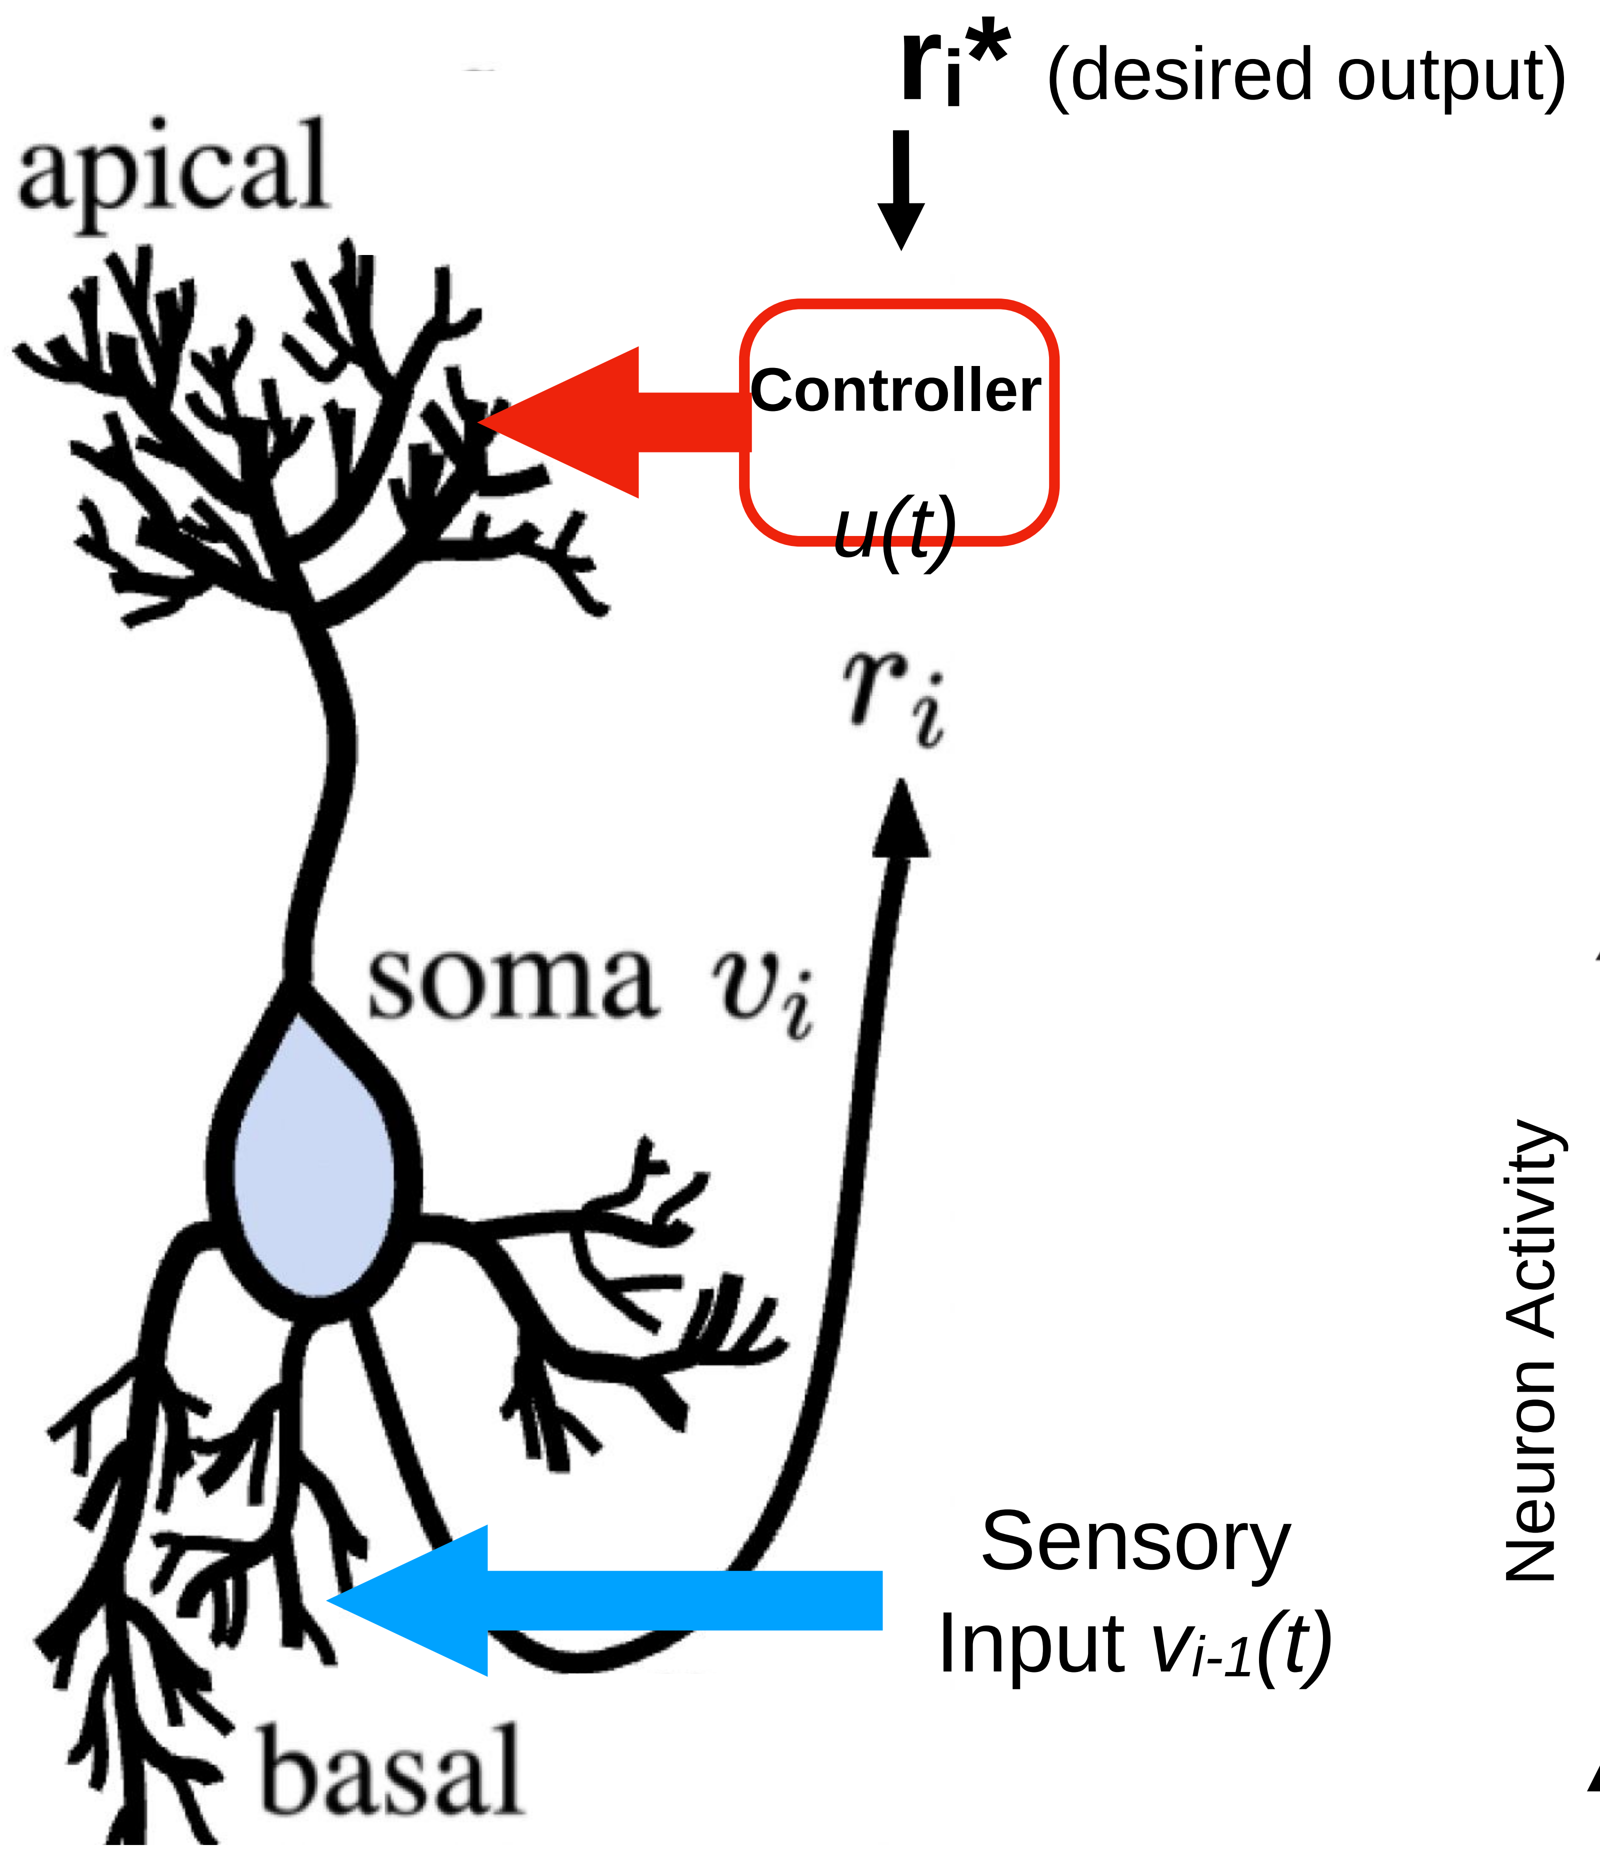

Neuron Output / Firing Rate

$$\mathbf{r}_i = \phi(\mathbf{v}_i)$$

Non. Linearity

Membrane Potential Dynamics

$$\tau_v \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{v}_i(t) = -\mathbf{v}_i(t) + \boxed{W_i\phi\big(\mathbf{v}_{i-1}(t)\big)} + \boxed{Q_i\mathbf{u}(t)}$$

Sensory

Feedback

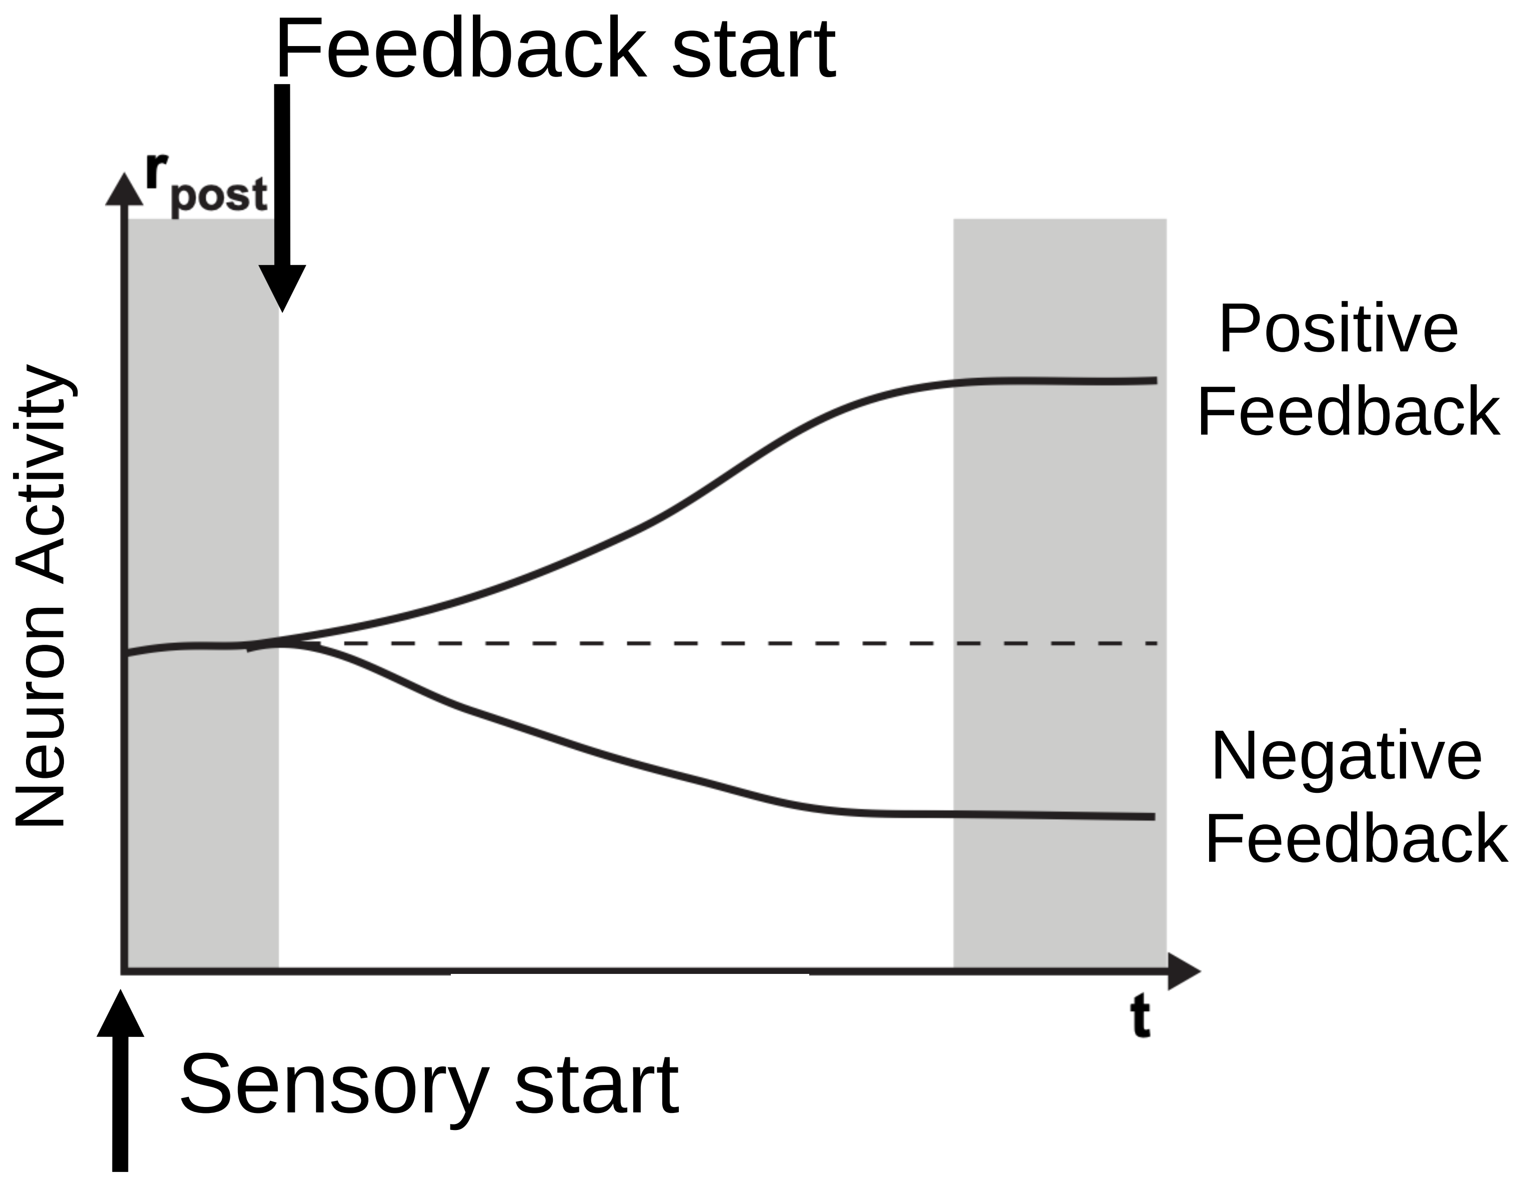

Meulemans et al. 2020, 2021

Deep Learning Through Feedback Control

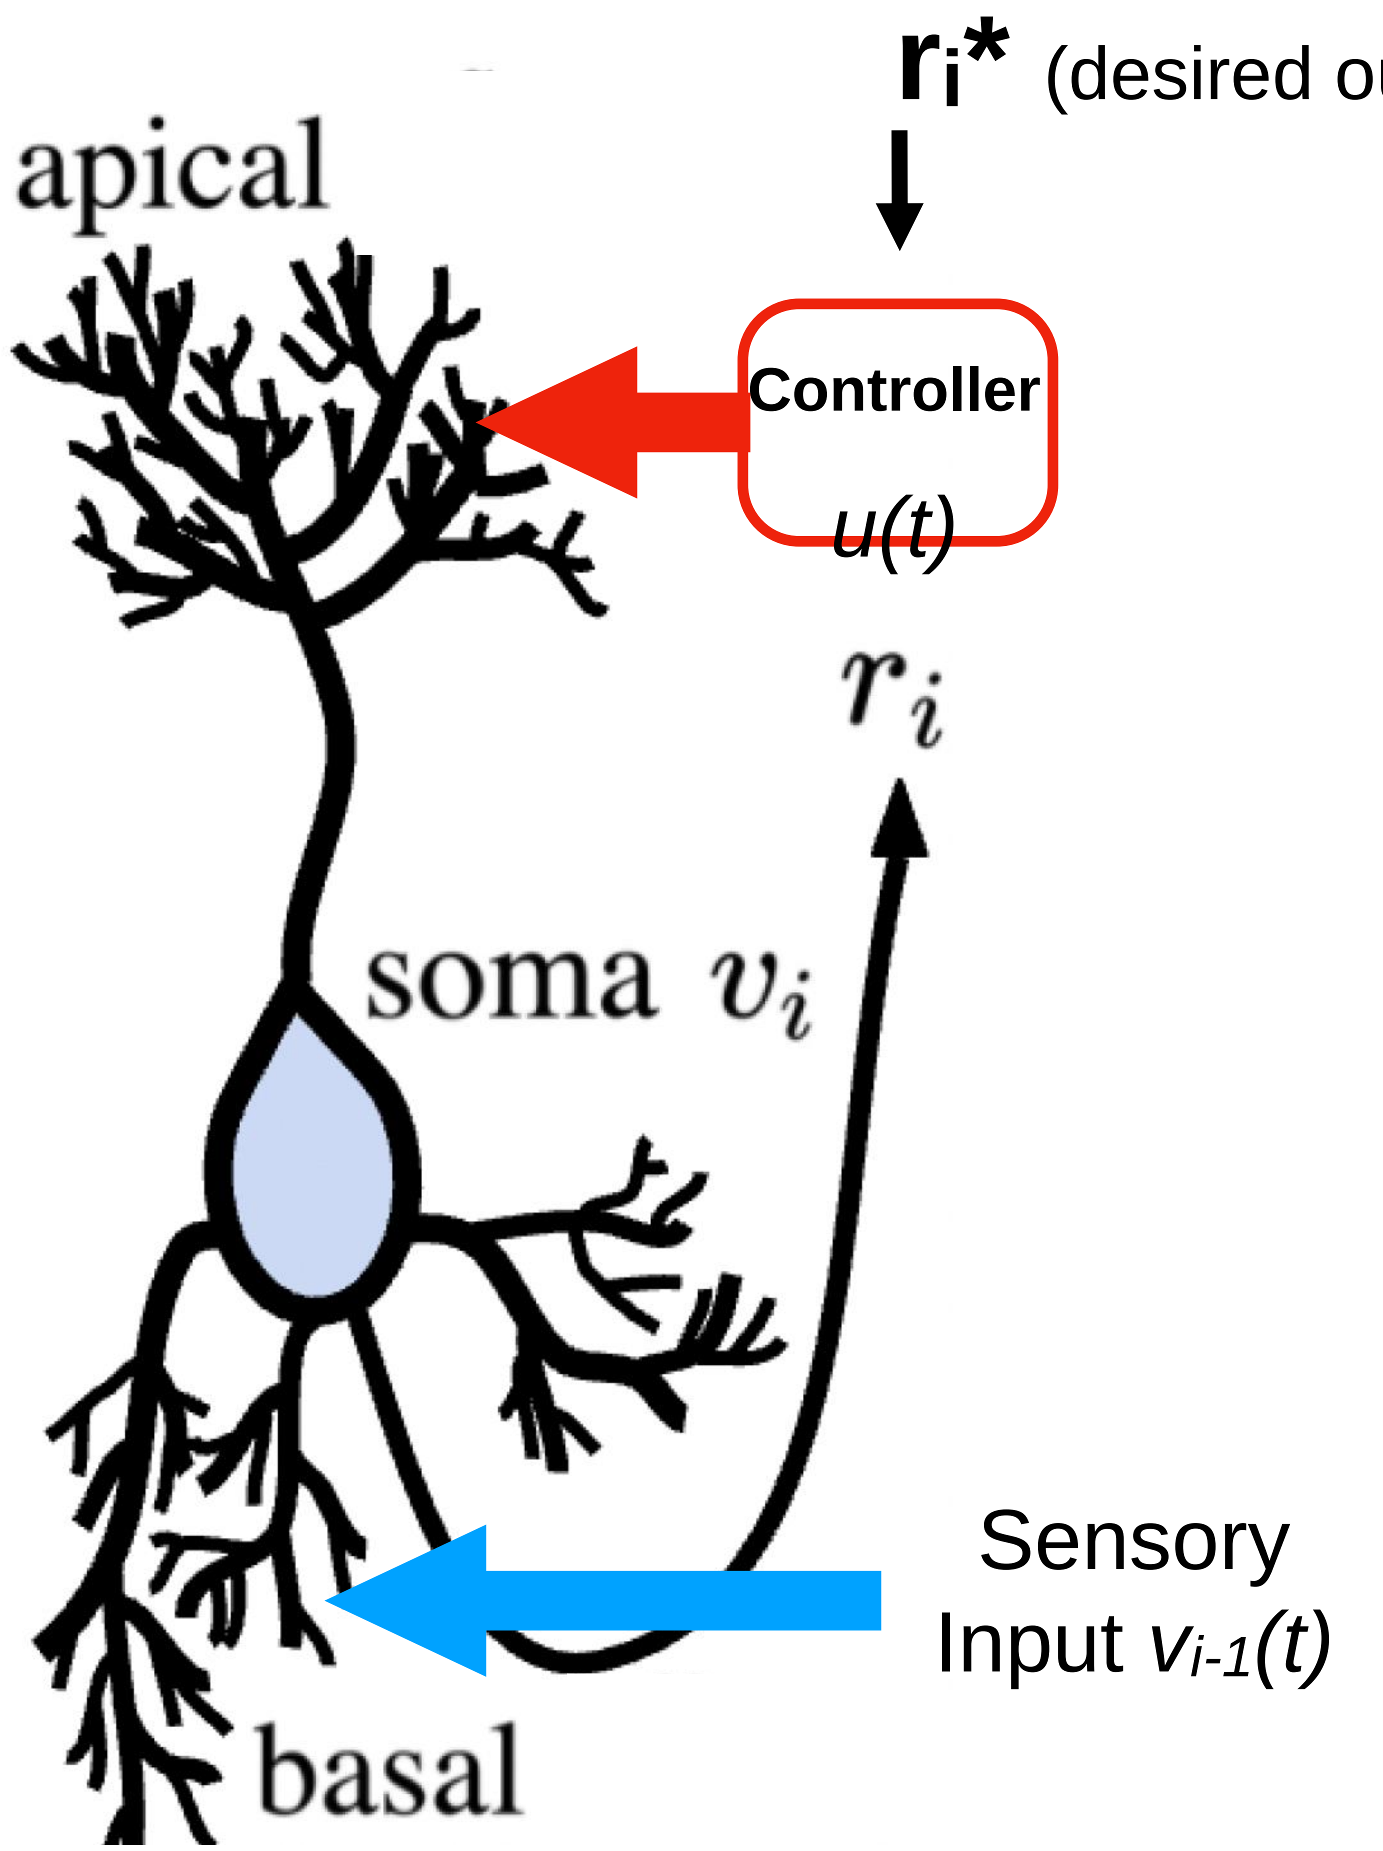

Neuron Output / Firing Rate

$$\mathbf{r}_i = \phi(\mathbf{v}_i)$$

Non. Linearity

Membrane Potential Dynamics

$$\tau_v \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{v}_i(t) = -\mathbf{v}_i(t) + \underbrace{\boxed{W_i\phi\big(\mathbf{v}_{i-1}(t)\big)}}_{\text{Sensory}} + \underbrace{\boxed{Q_i\mathbf{u}(t)}}_{\text{Feedback}}$$

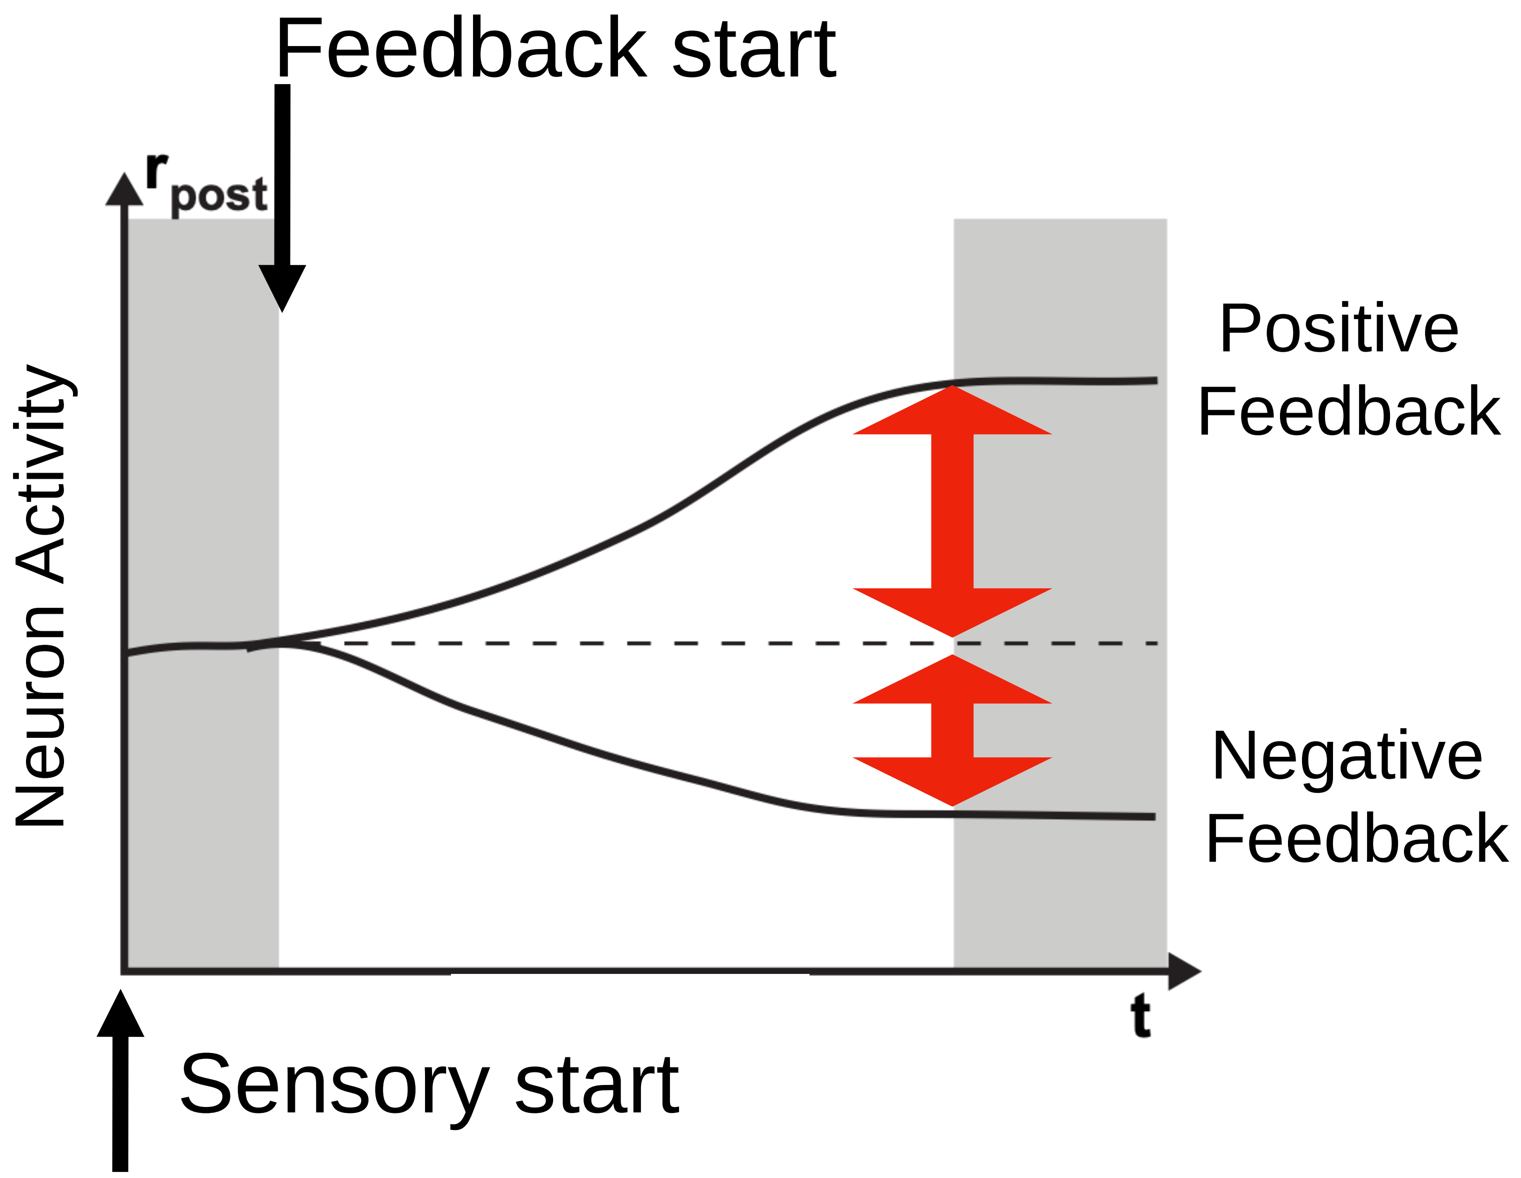

Meulemans et al. 2020, 2021

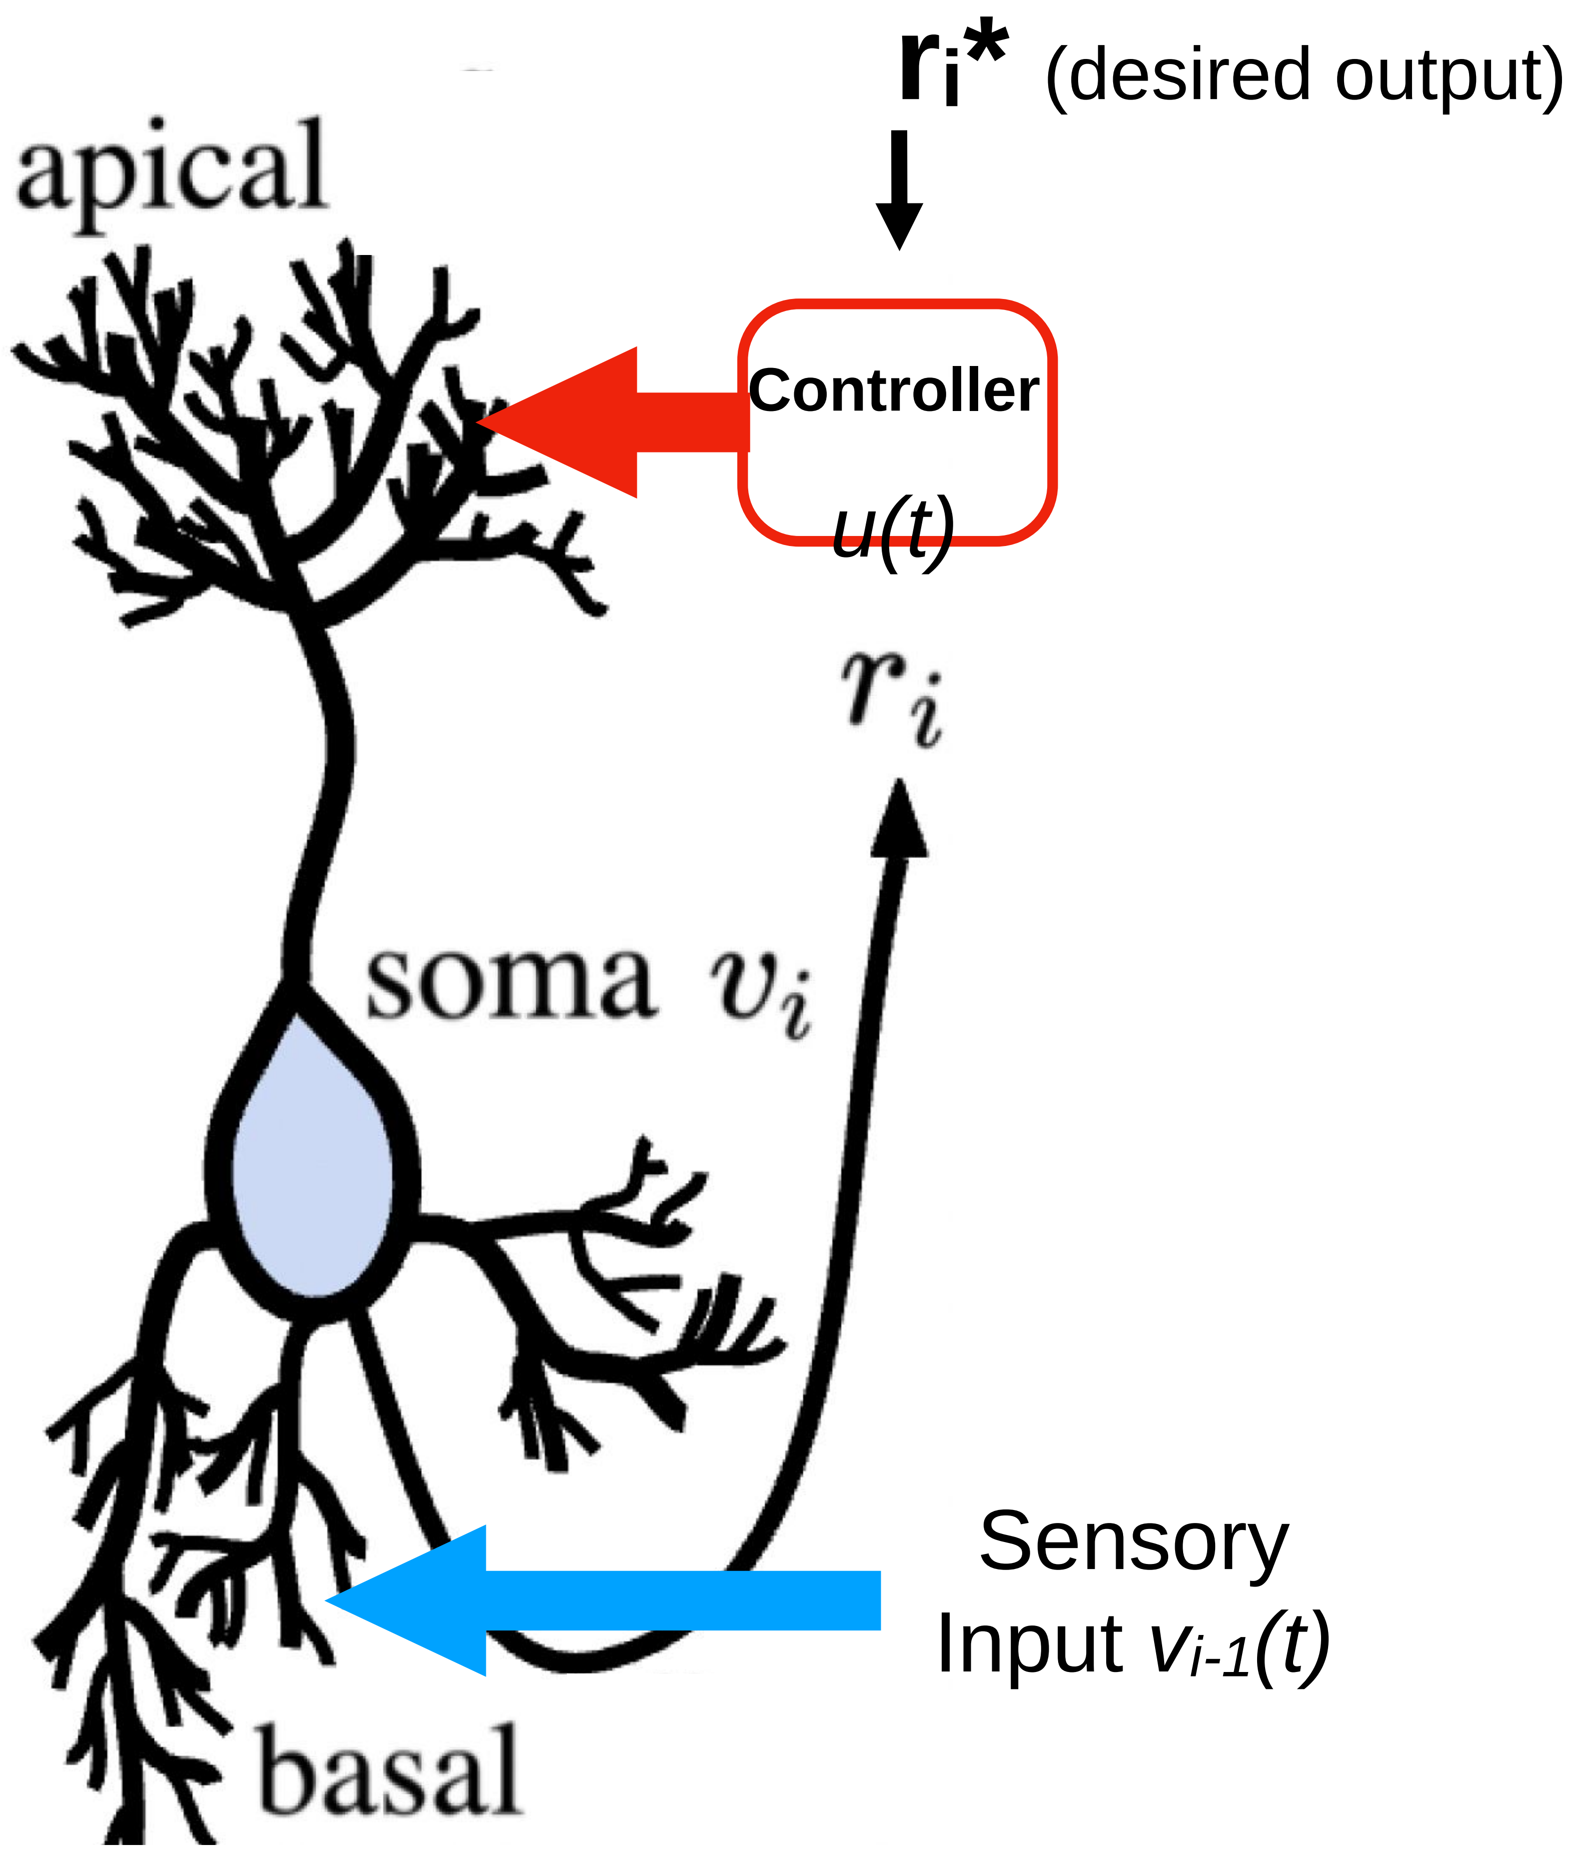

Neuron Output / Firing Rate

$$\mathbf{r}_i = \phi(\mathbf{v}_i)$$

Non. Linearity

Membrane Potential Dynamics

$$\tau_v \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{v}_i(t) = -\mathbf{v}_i(t) + \boxed{W_i\phi\big(\mathbf{v}_{i-1}(t)\big)} + \boxed{Q_i\mathbf{u}(t)}$$

Sensory   Feedback

Synaptic Weight Update

$$\tau_W \frac{\mathrm{d}}{\mathrm{d}t}W_i(t) = \boxed{\big(\phi(\mathbf{v}_i(t))} - \boxed{\phi(W_i\mathbf{r}_{i-1}(t))\big)}\,\boxed{\mathbf{r}_{i-1}(t)^T}$$

Activity w/o Feedback   Activity with Feedback   Presynaptic Term

Meulemans et al. 2020, 2021

ETH zürich

Deep Learning Through Feedback Control

institute of neuroinformatics

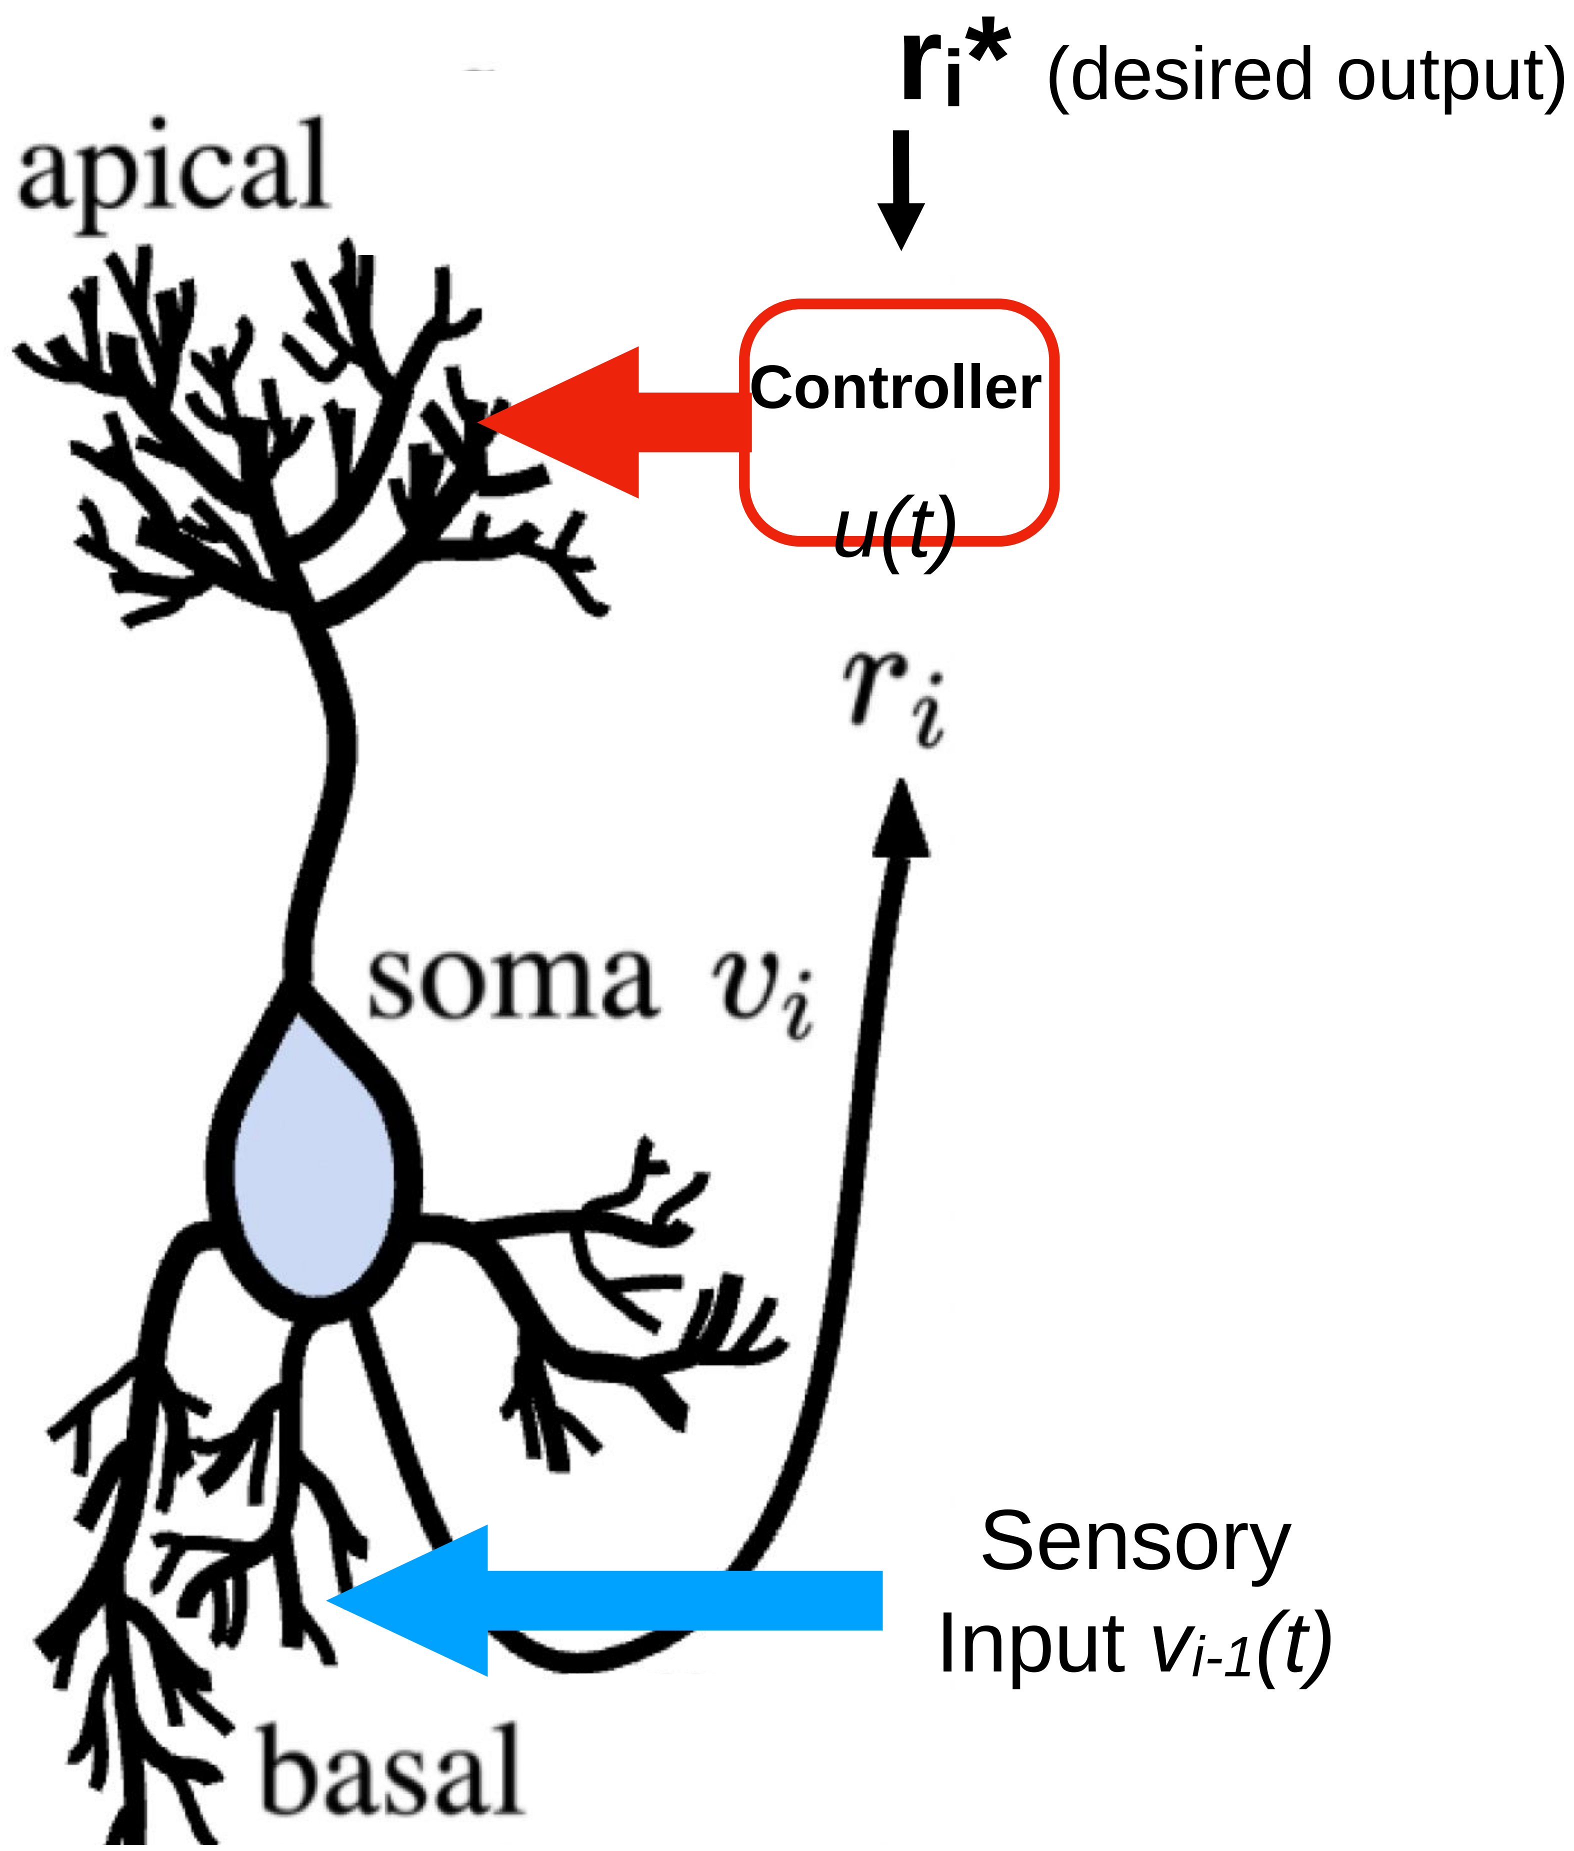

Neuron Output / Firing Rate

$$\mathbf{r}_i = \phi(\mathbf{v}_i)$$

Non. Linearity

Neuron Activity

Membrane Potential Dynamics

$$\tau_v \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{v}_i(t) = -\mathbf{v}_i(t) + \boxed{W_i \phi\big(\mathbf{v}_{i-1}(t)\big)} + \boxed{Q_i \mathbf{u}(t)}$$

Sensory

Feedback

Synaptic Weight Update

Neuron Plasticity

$$\tau_W \frac{\mathrm{d}}{\mathrm{d}t} W_i(t) = \boxed{\big(\phi(\mathbf{v}_i(t))} - \boxed{\phi(W_i \mathbf{r}_{i-1}(t))\big)} \boxed{\mathbf{r}_{i-1}(t)^T}$$

Activity w/o Feedback

Activity with Feedback

Presynaptic Term

Meulemans et al. 2020, 2021

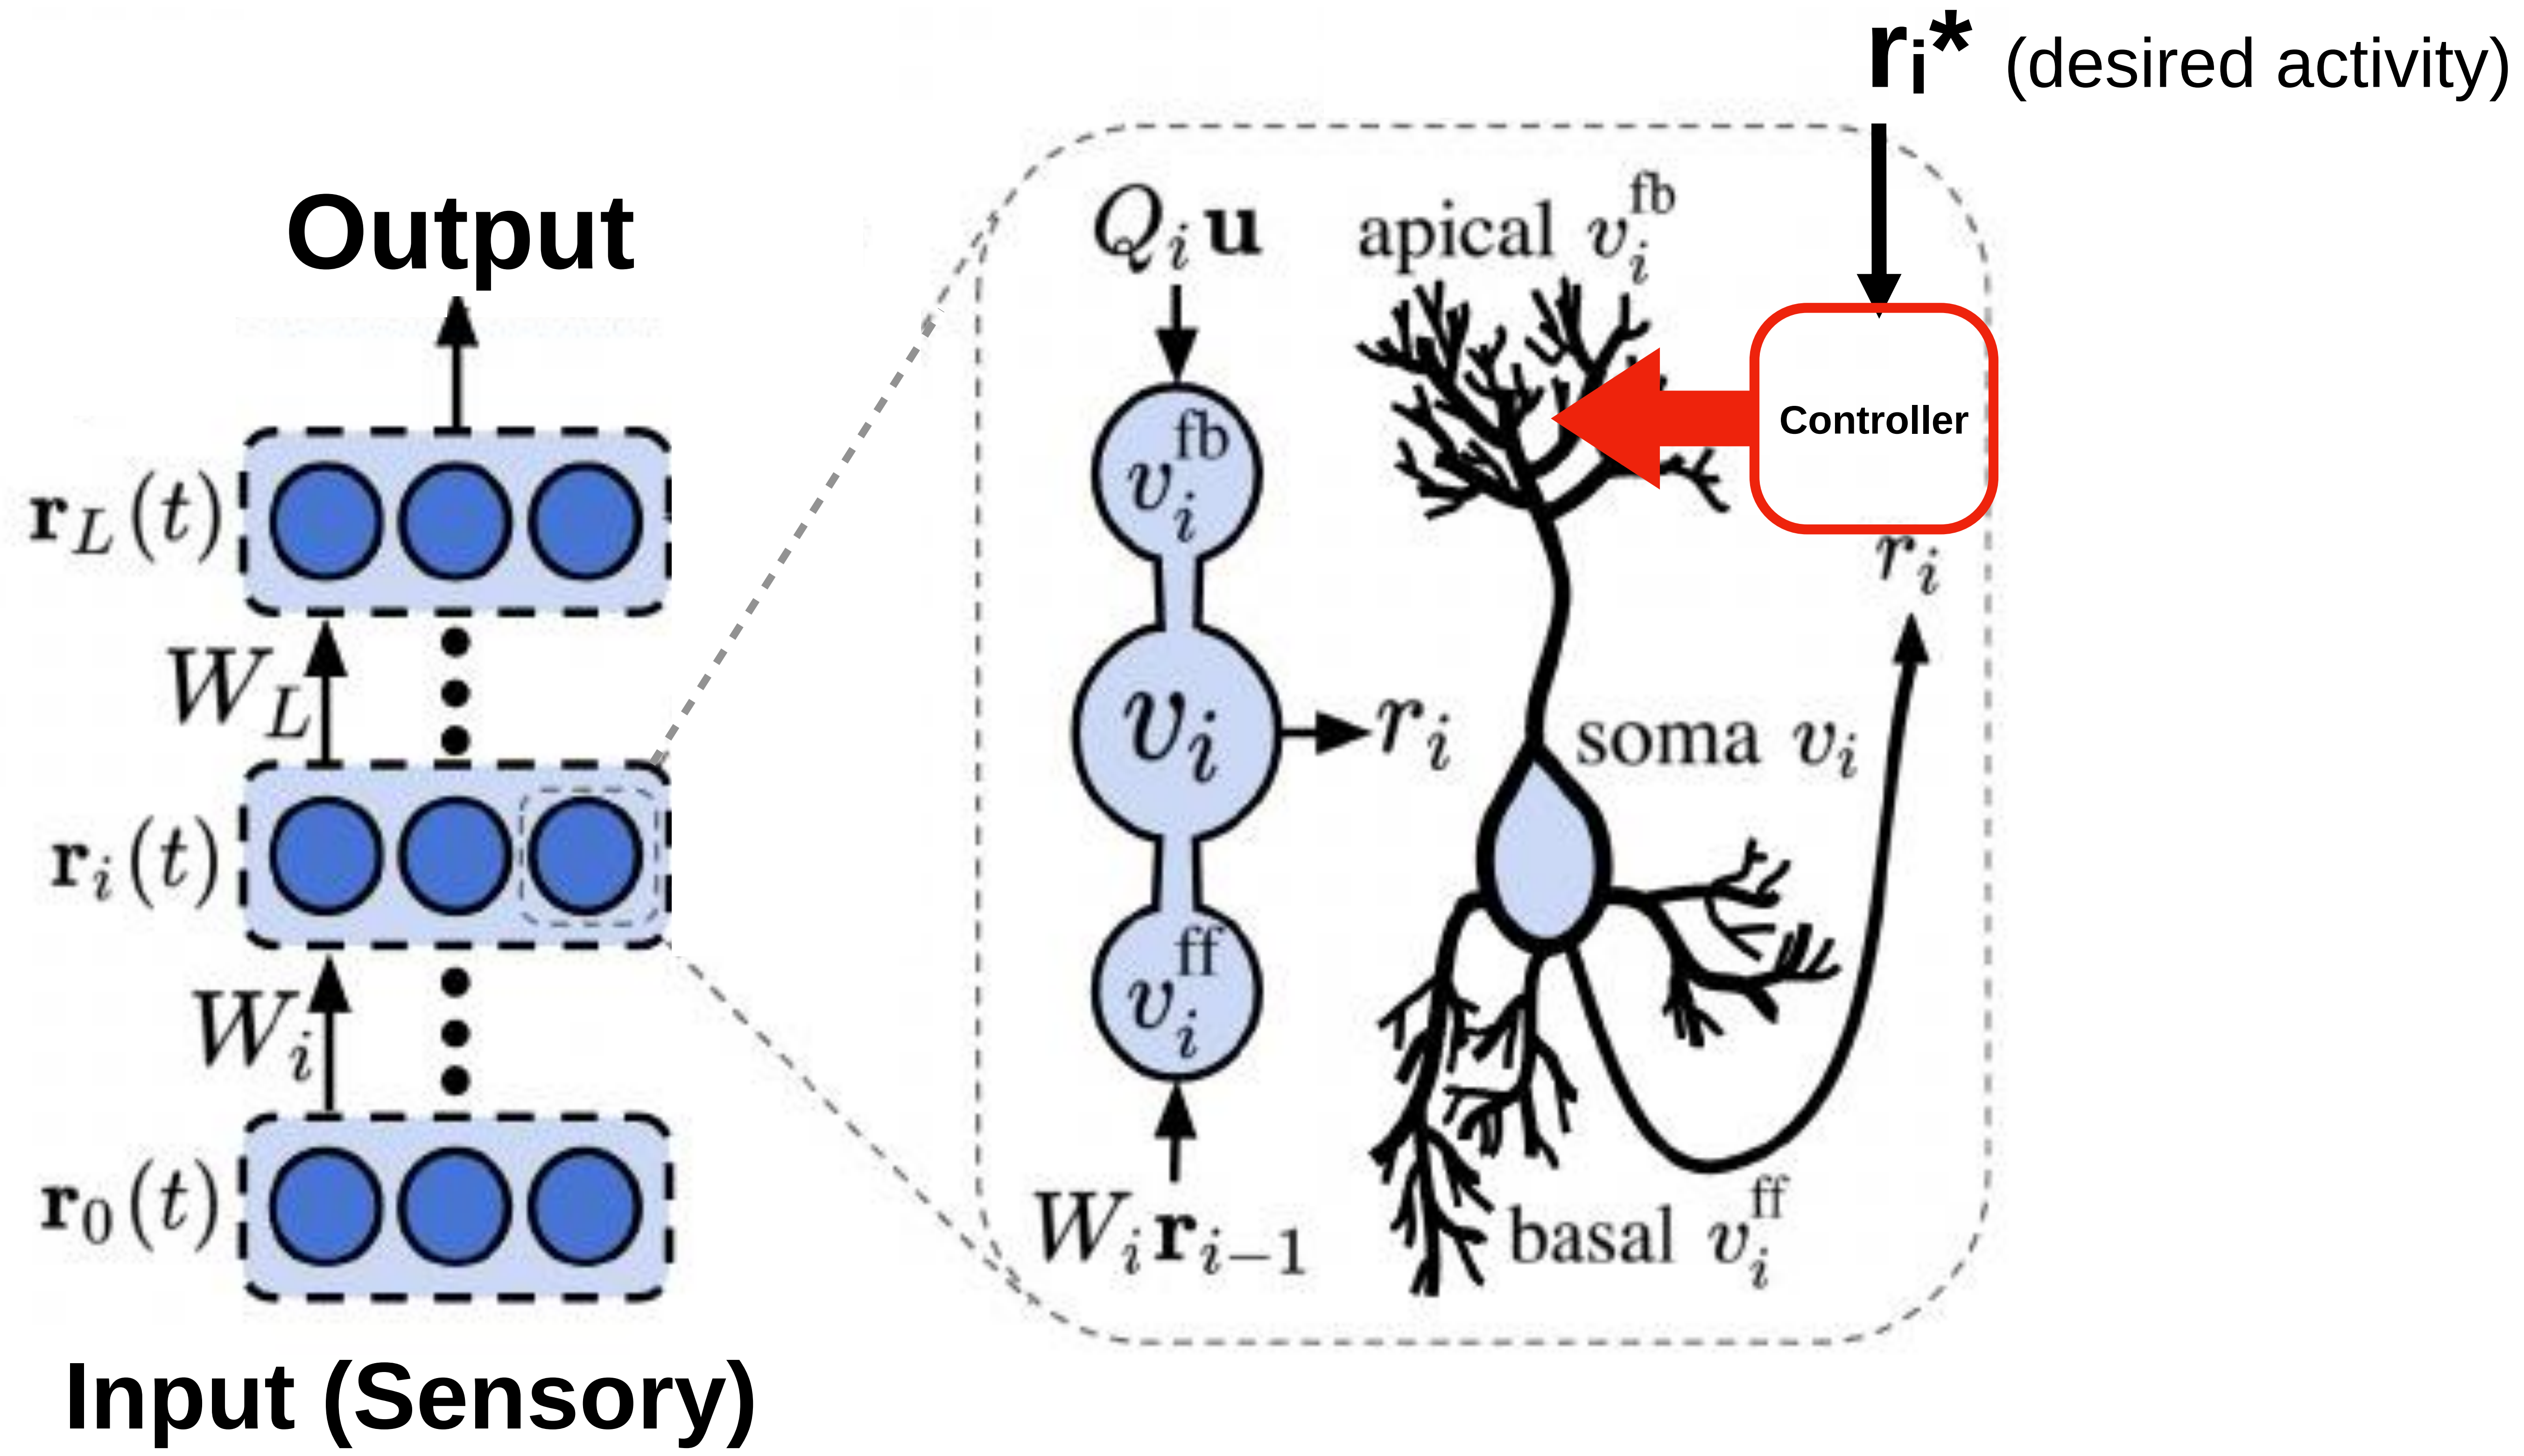

Meulemans et al. 2020, 2021

Deep Learning Through Feedback Control

<u>We solve a complex control problem!</u>

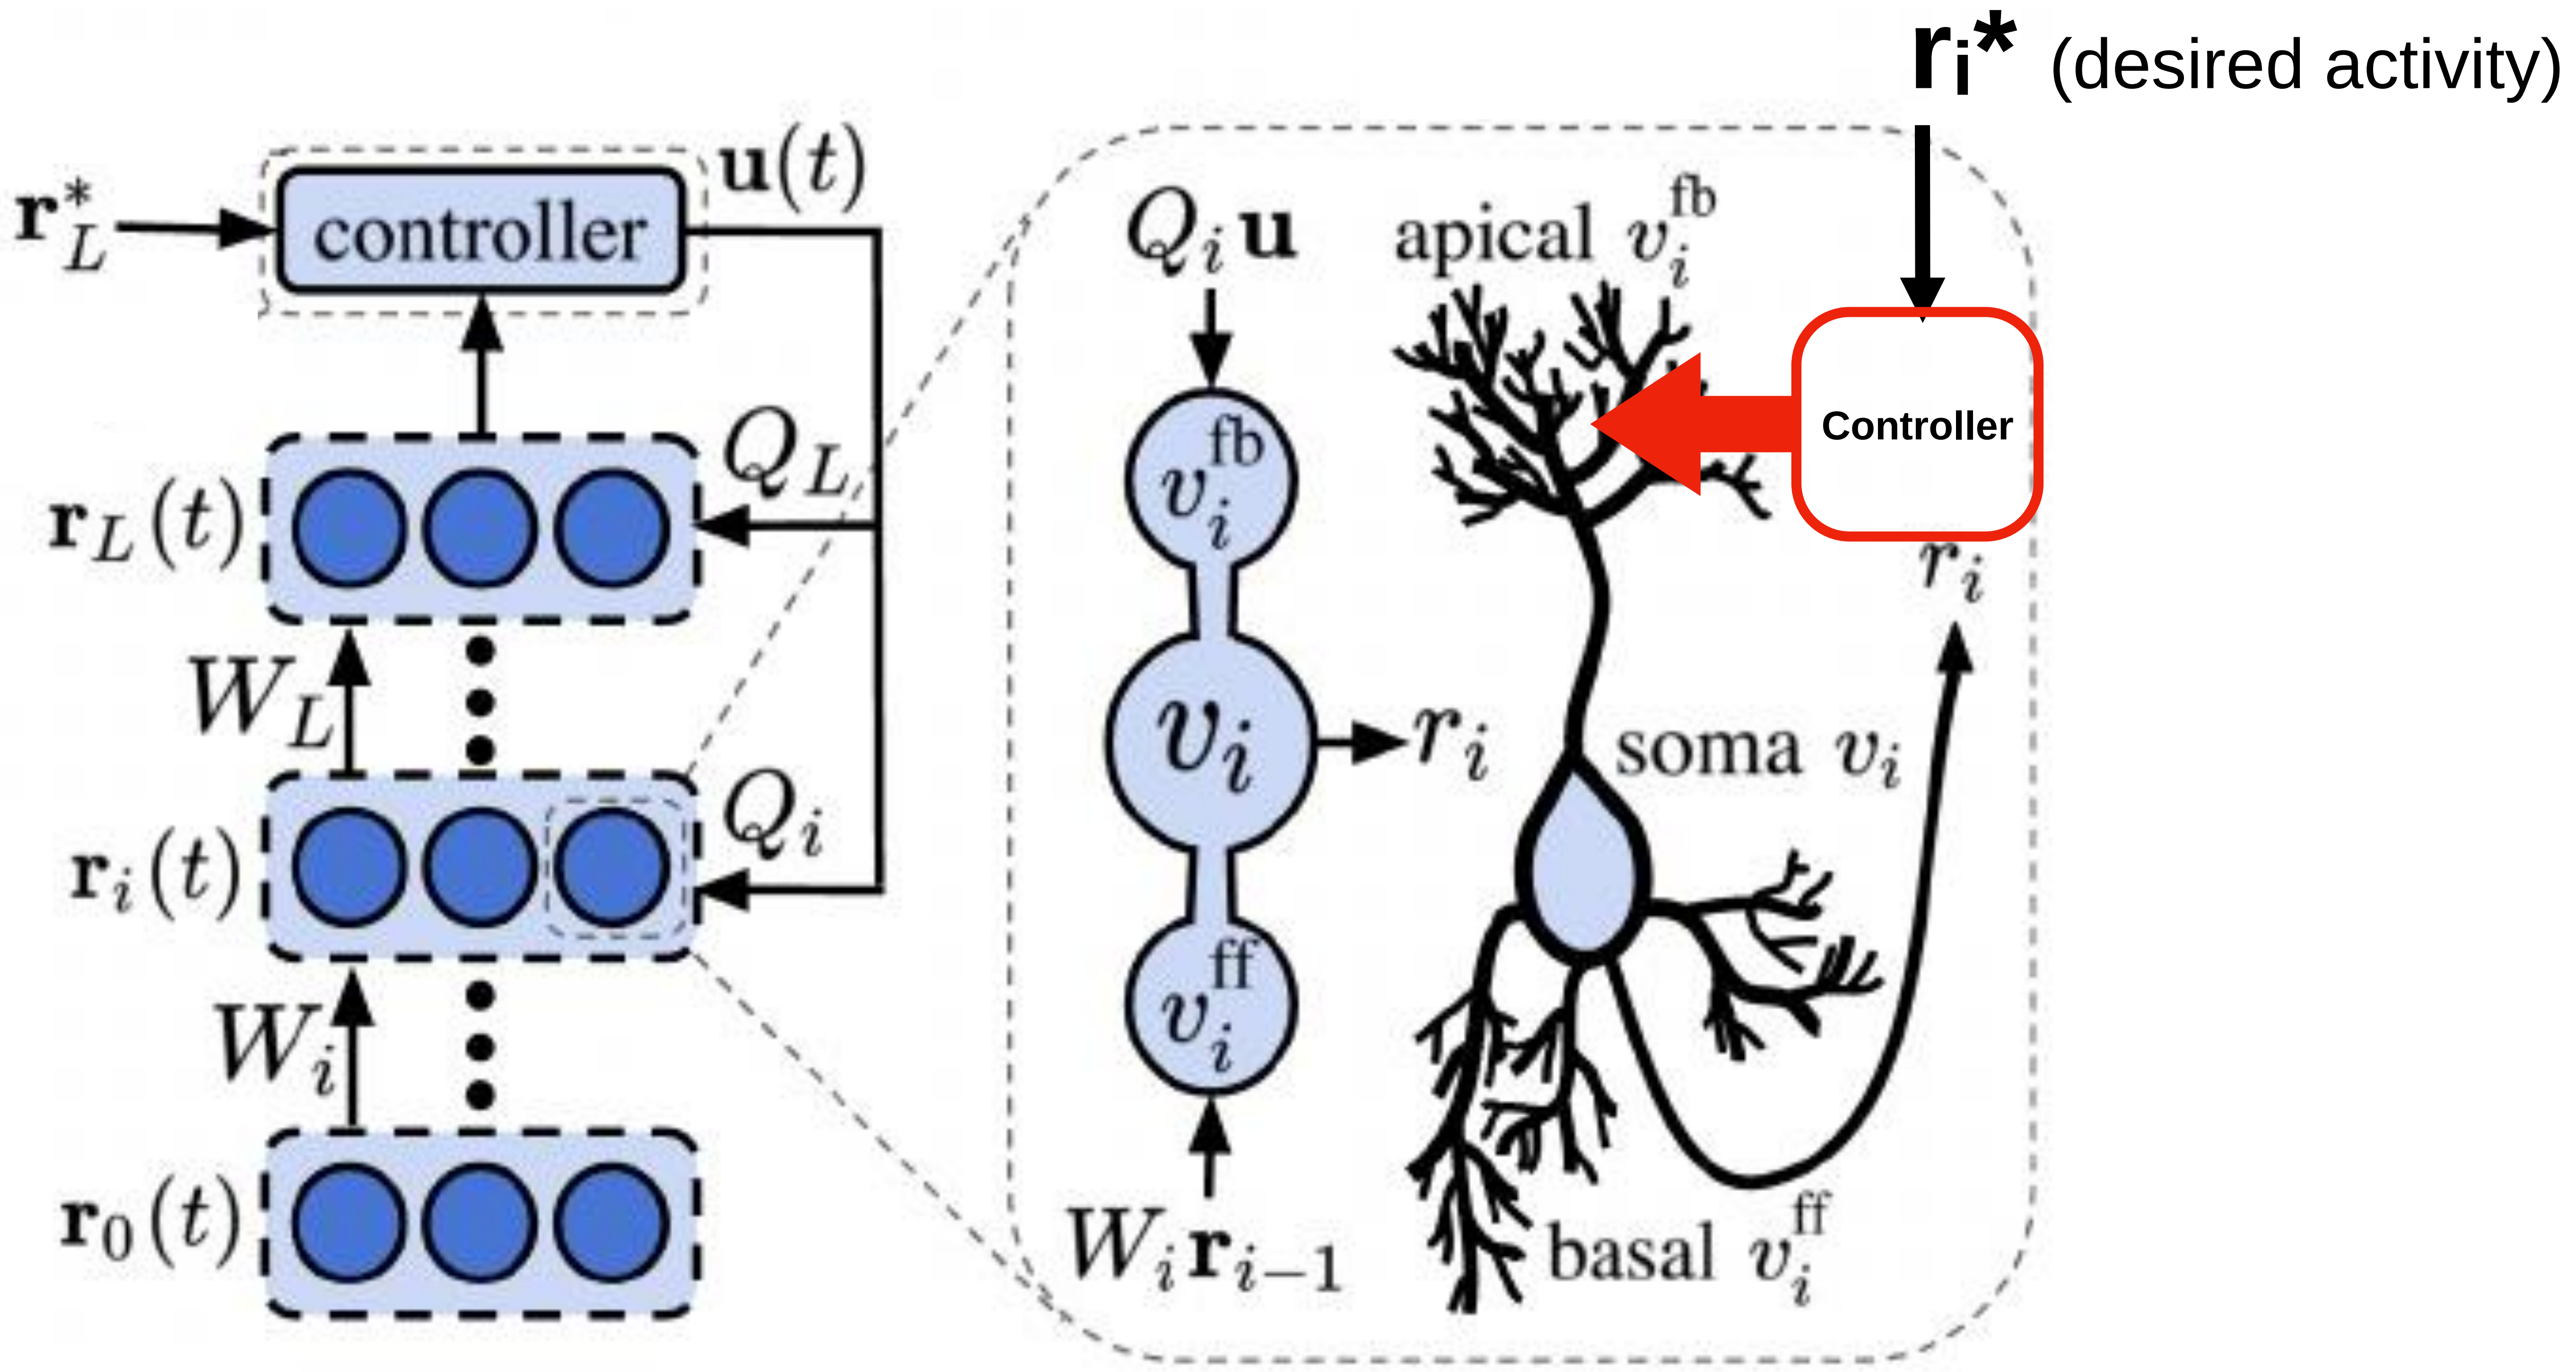

Input (Sensory)

Meulemans et al. 2020, 2021

Minimising Control for Credit Assignment

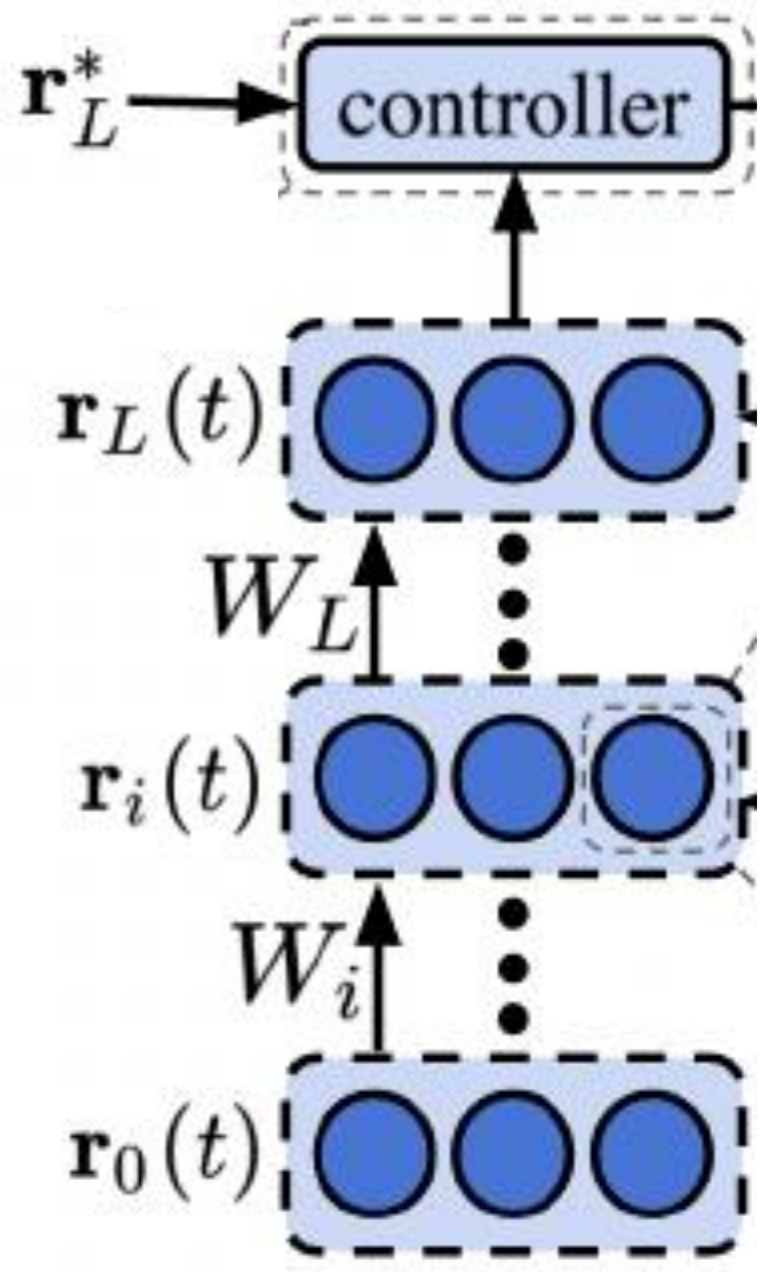

Input (Sensory)

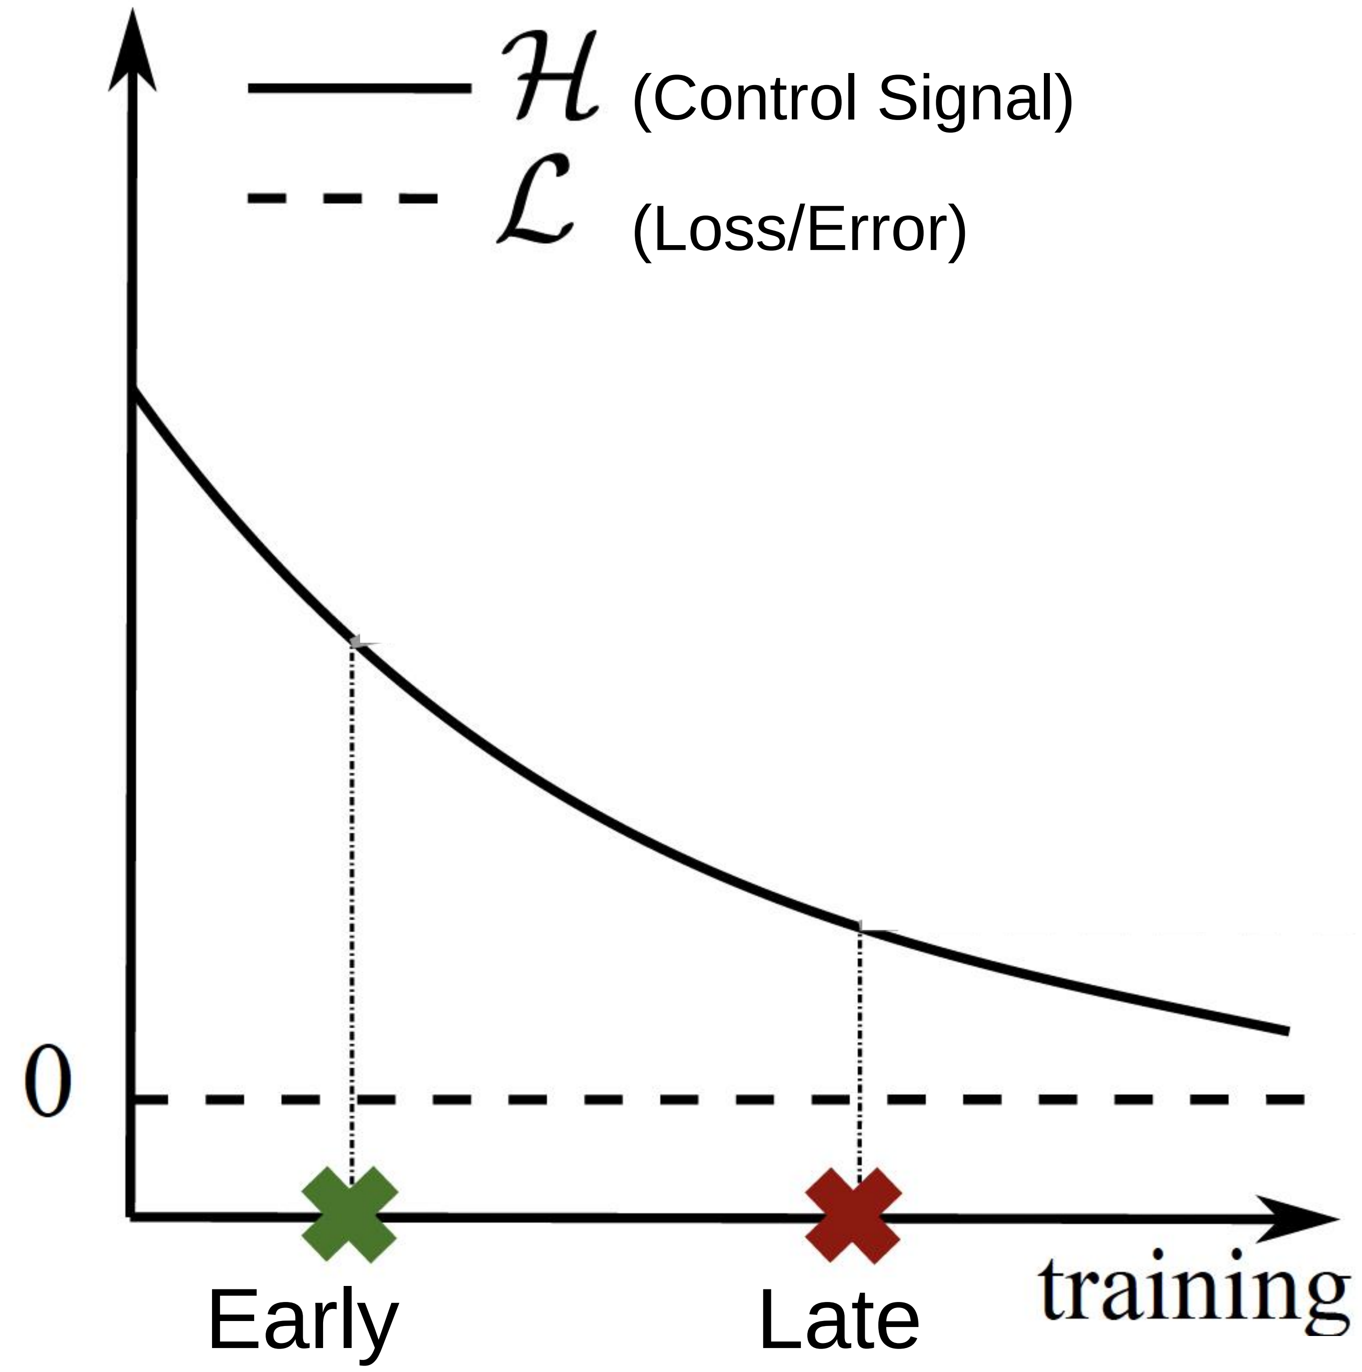

Early Stage

Late Stage

Learning = Reducing Help!

Meulemans et al. 2022

Checkpoint: How Bio-Plausible are we?

1. We **no longer** use separate forward and backward phases.
2. We **don't send** sensory information forward and errors backward.
3. We **don't send** feedback signals though the **SAME** weights.
4. We allow continuous (in time) computation.
5. Our update/plasticity rule still not biologically plausible.

 

Spotlight Awards
(Best 1-2%)
* *
Meulemans et al. 2020, 2021, 2022

Neuronal Plasticity in Biology

Spike Timing Dependent Plasticity (STDP)

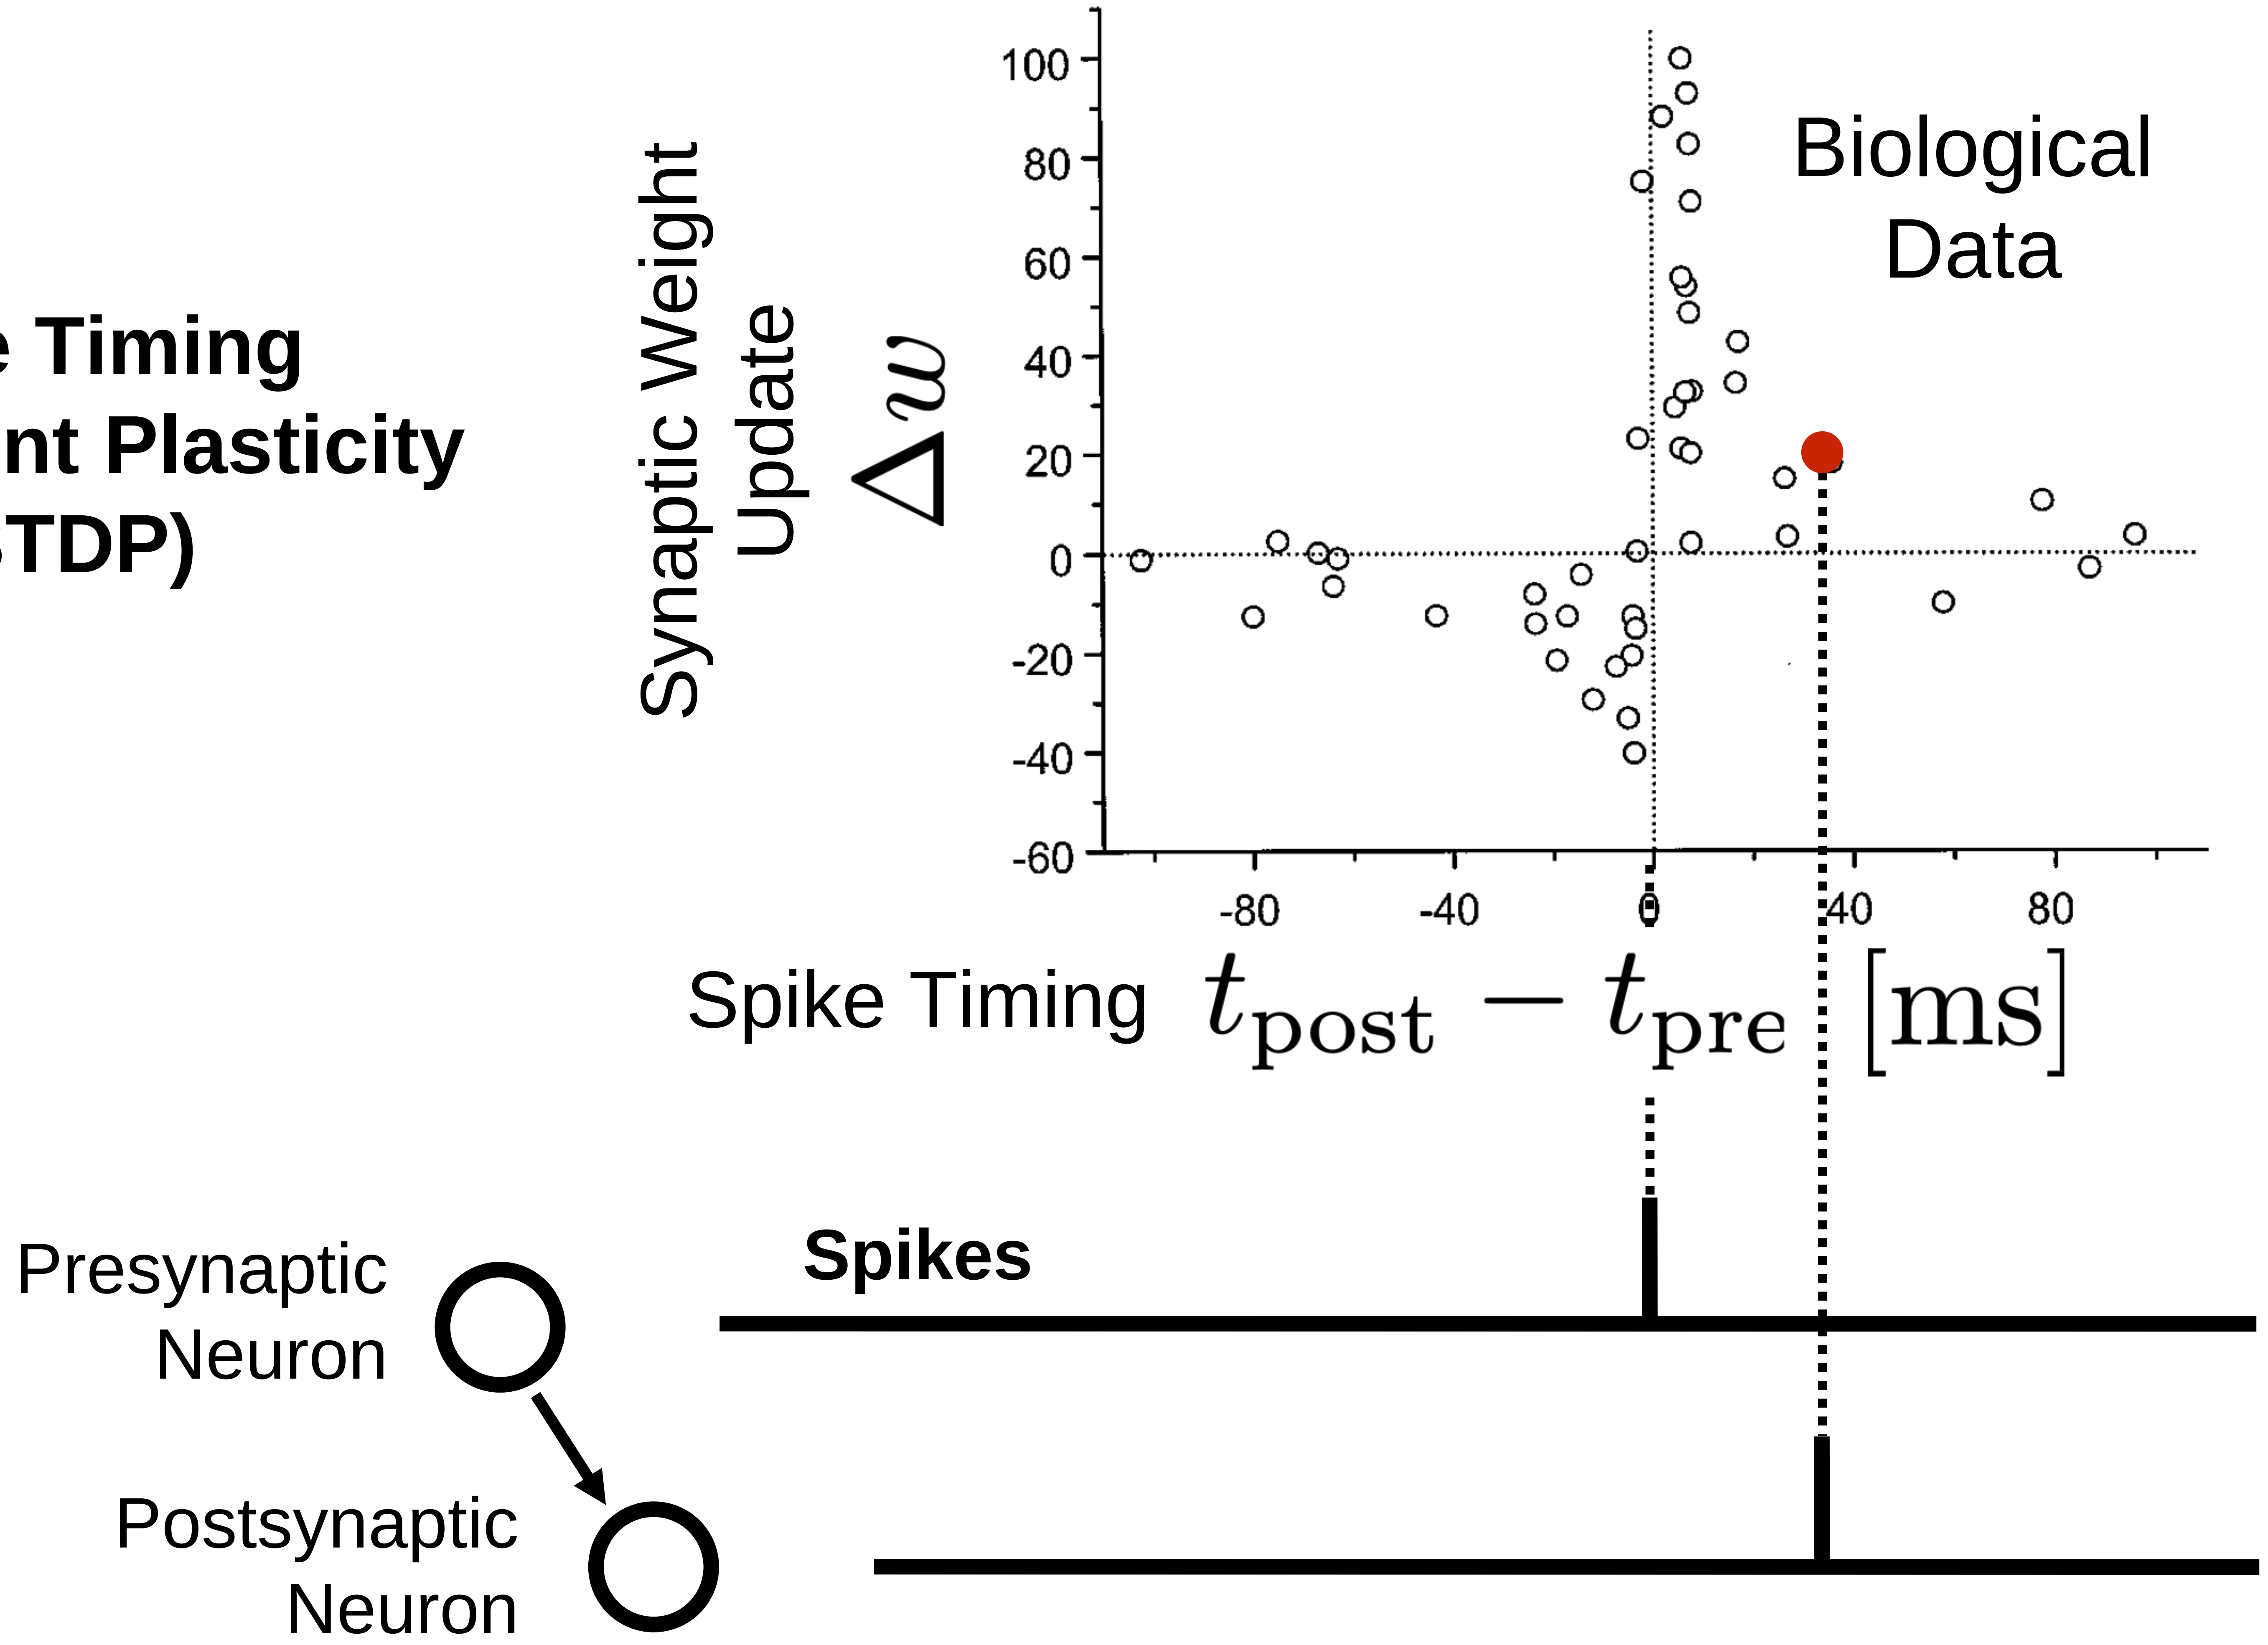

Bi and Poo, 1998

Neuronal Plasticity in Biology

Spike Timing Dependent Plasticity (STDP)

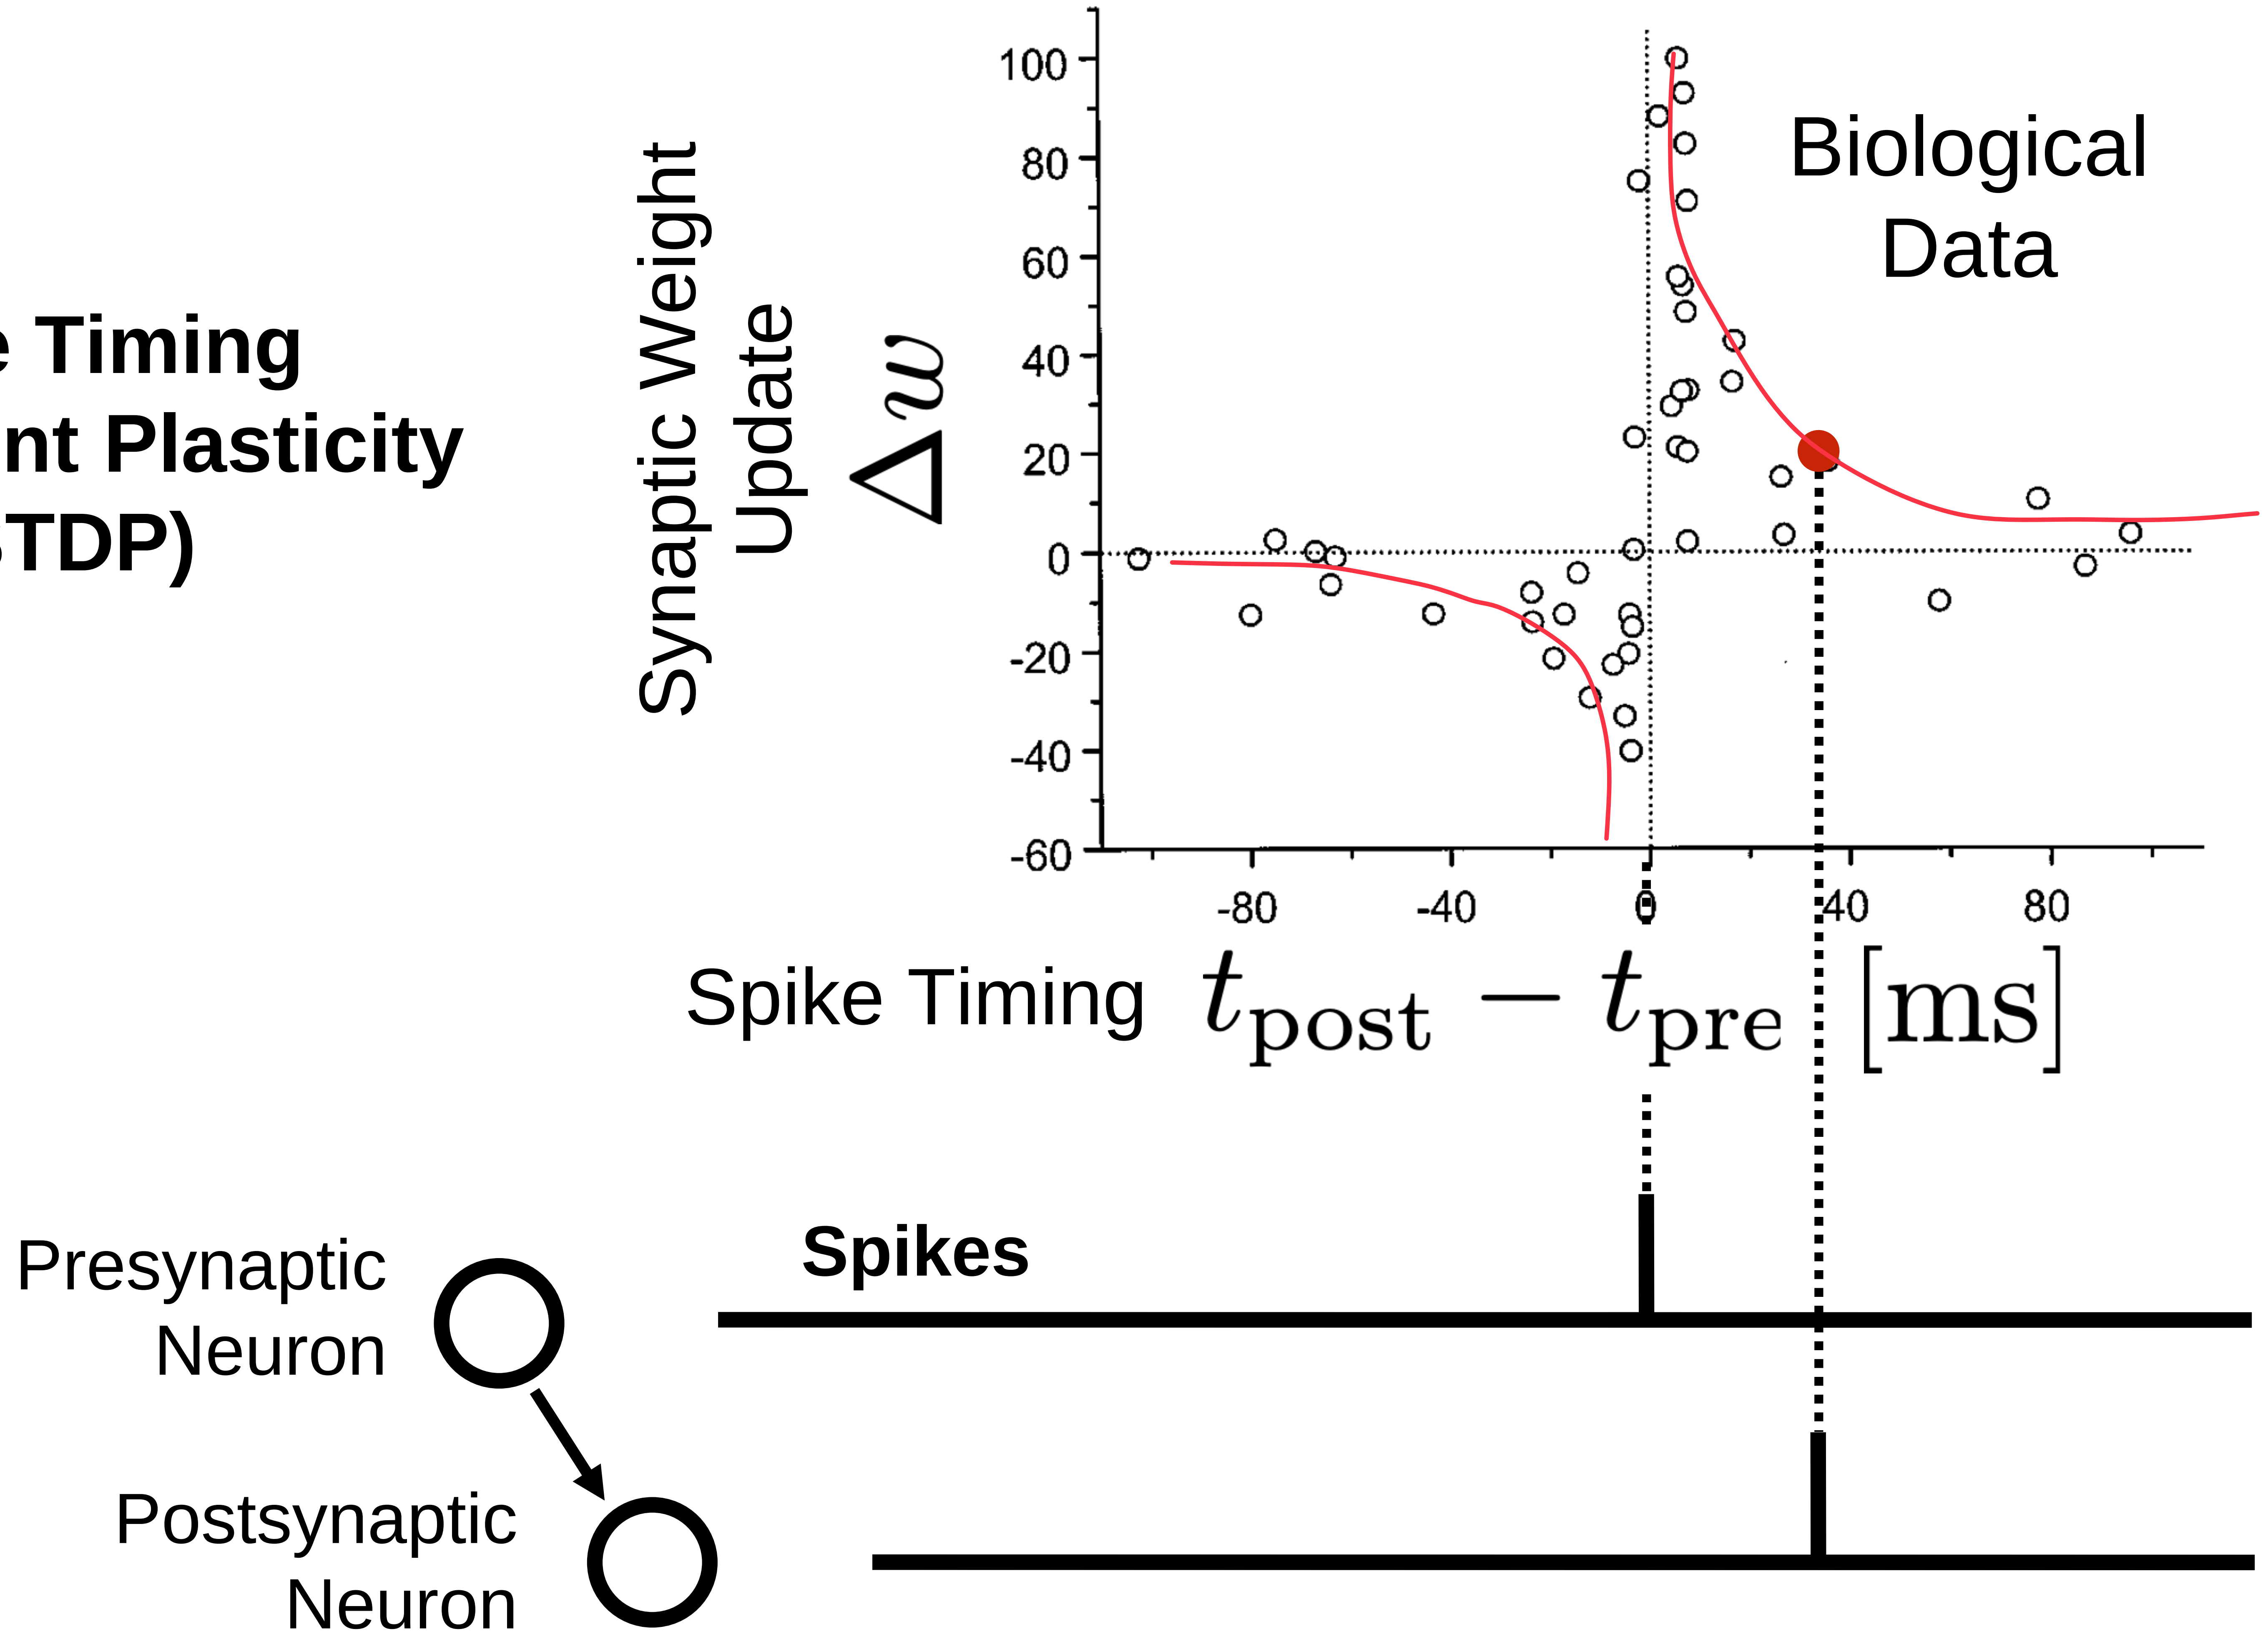

Bi and Poo, 1998

Mimicking Biological Plasticity

Pau Aceituno

<u>Biology:</u>
Spike Timing
Dependent Plasticity
(STDP)

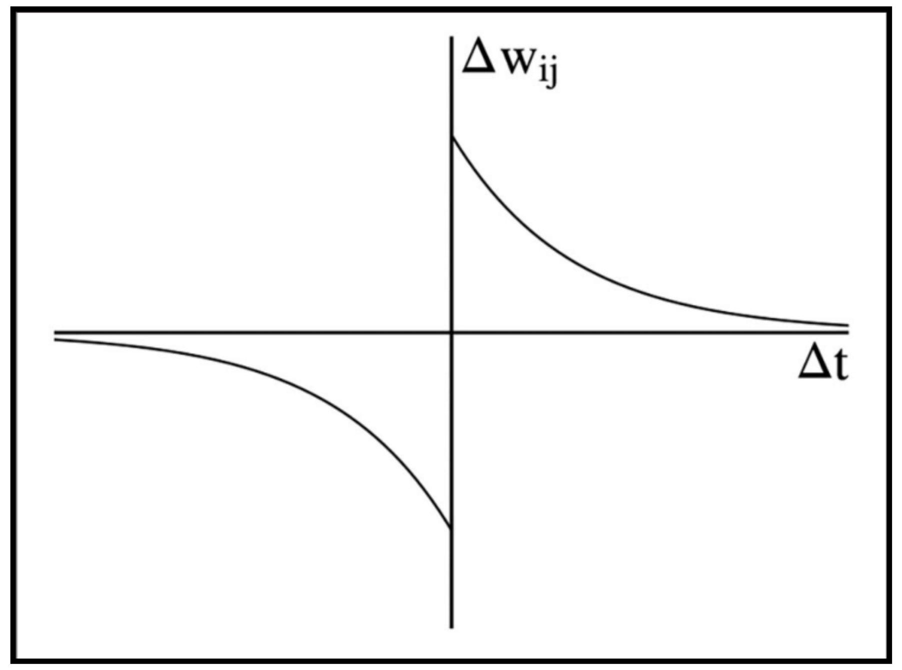

Mimicking Biological Plasticity

Biology:
Spike Timing
Dependent Plasticity
(STDP)

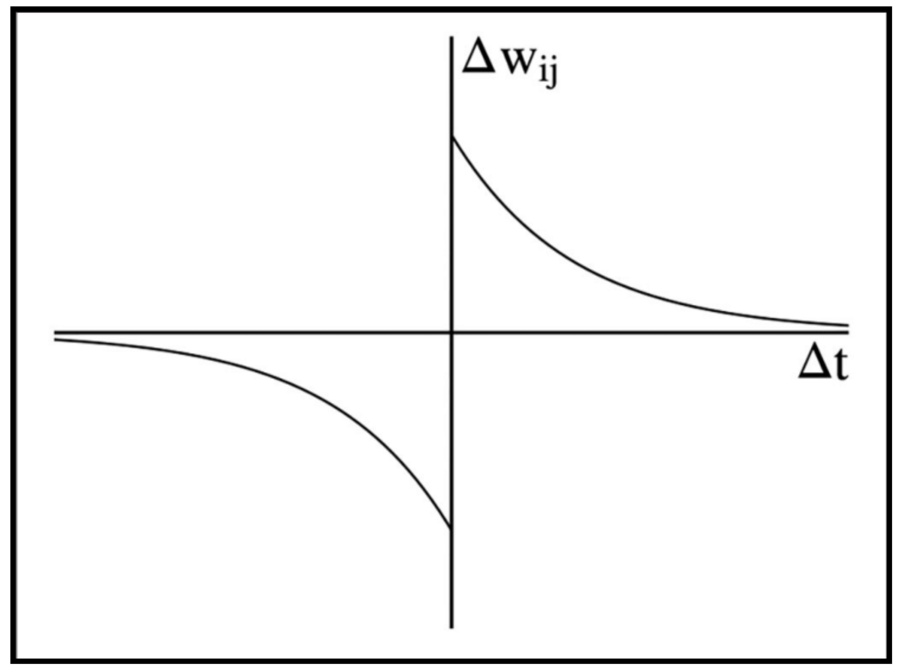

Learning Through Control
Differential Hebbian +
Feedback Control

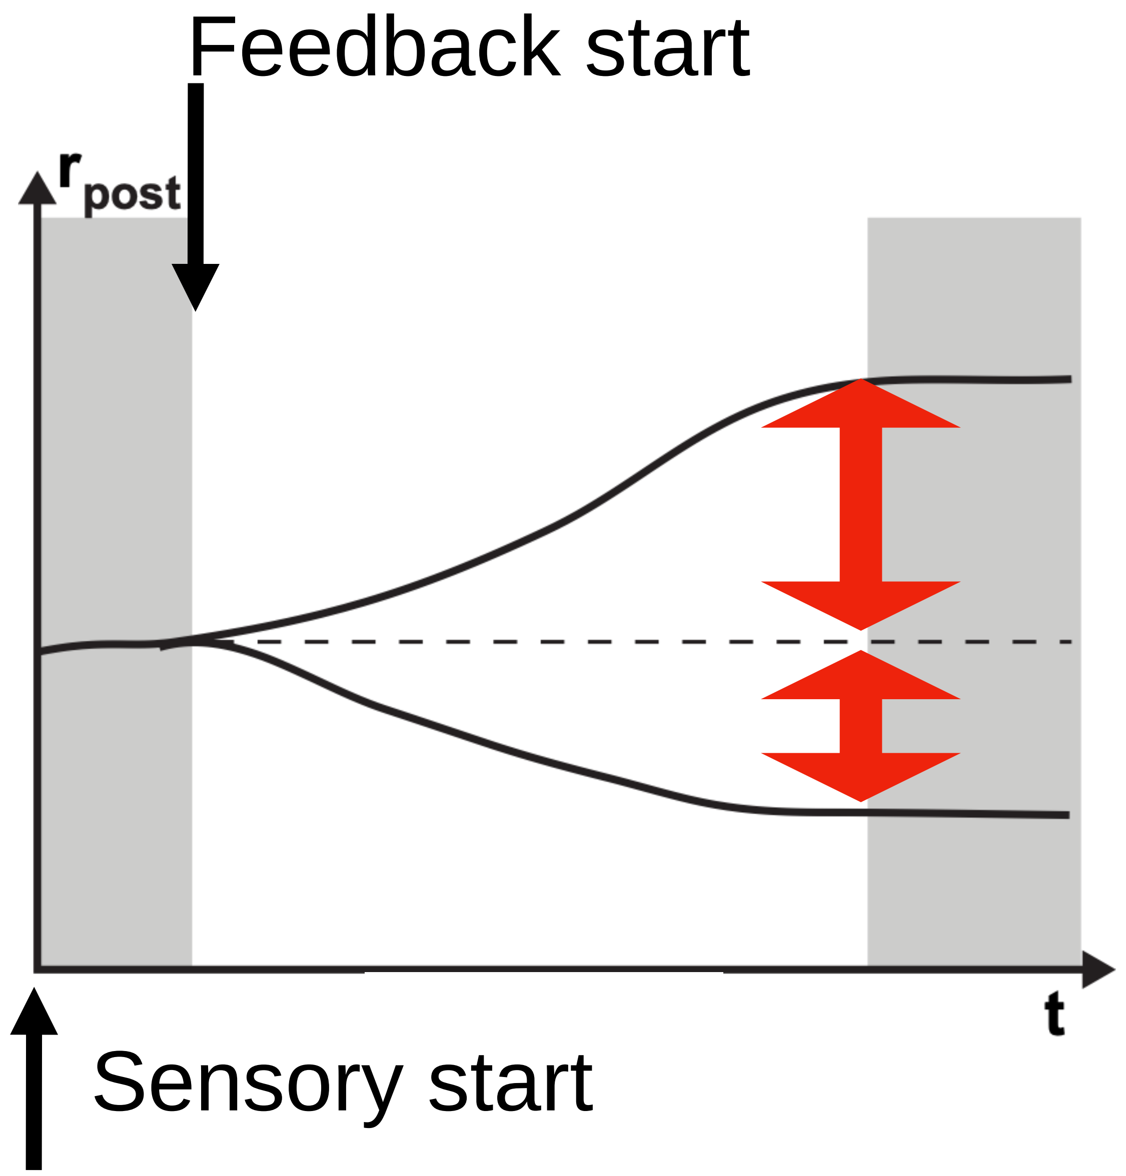

Mimicking Biological Plasticity

<u>Biology:</u>
Spike Timing
Dependent Plasticity
(STDP)

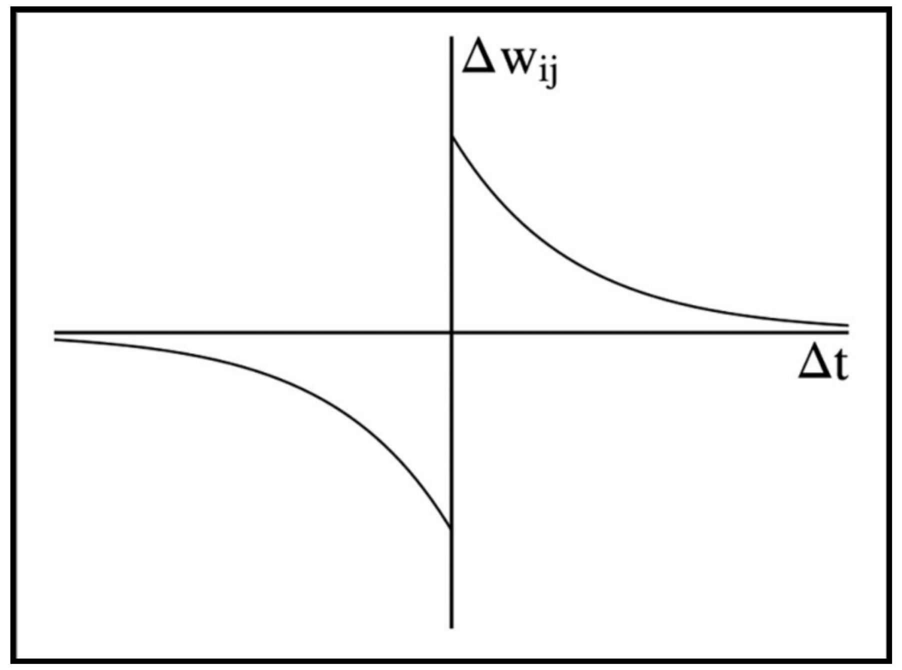

<u>Learning Through Control</u>
Differential Hebbian +
Feedback Control

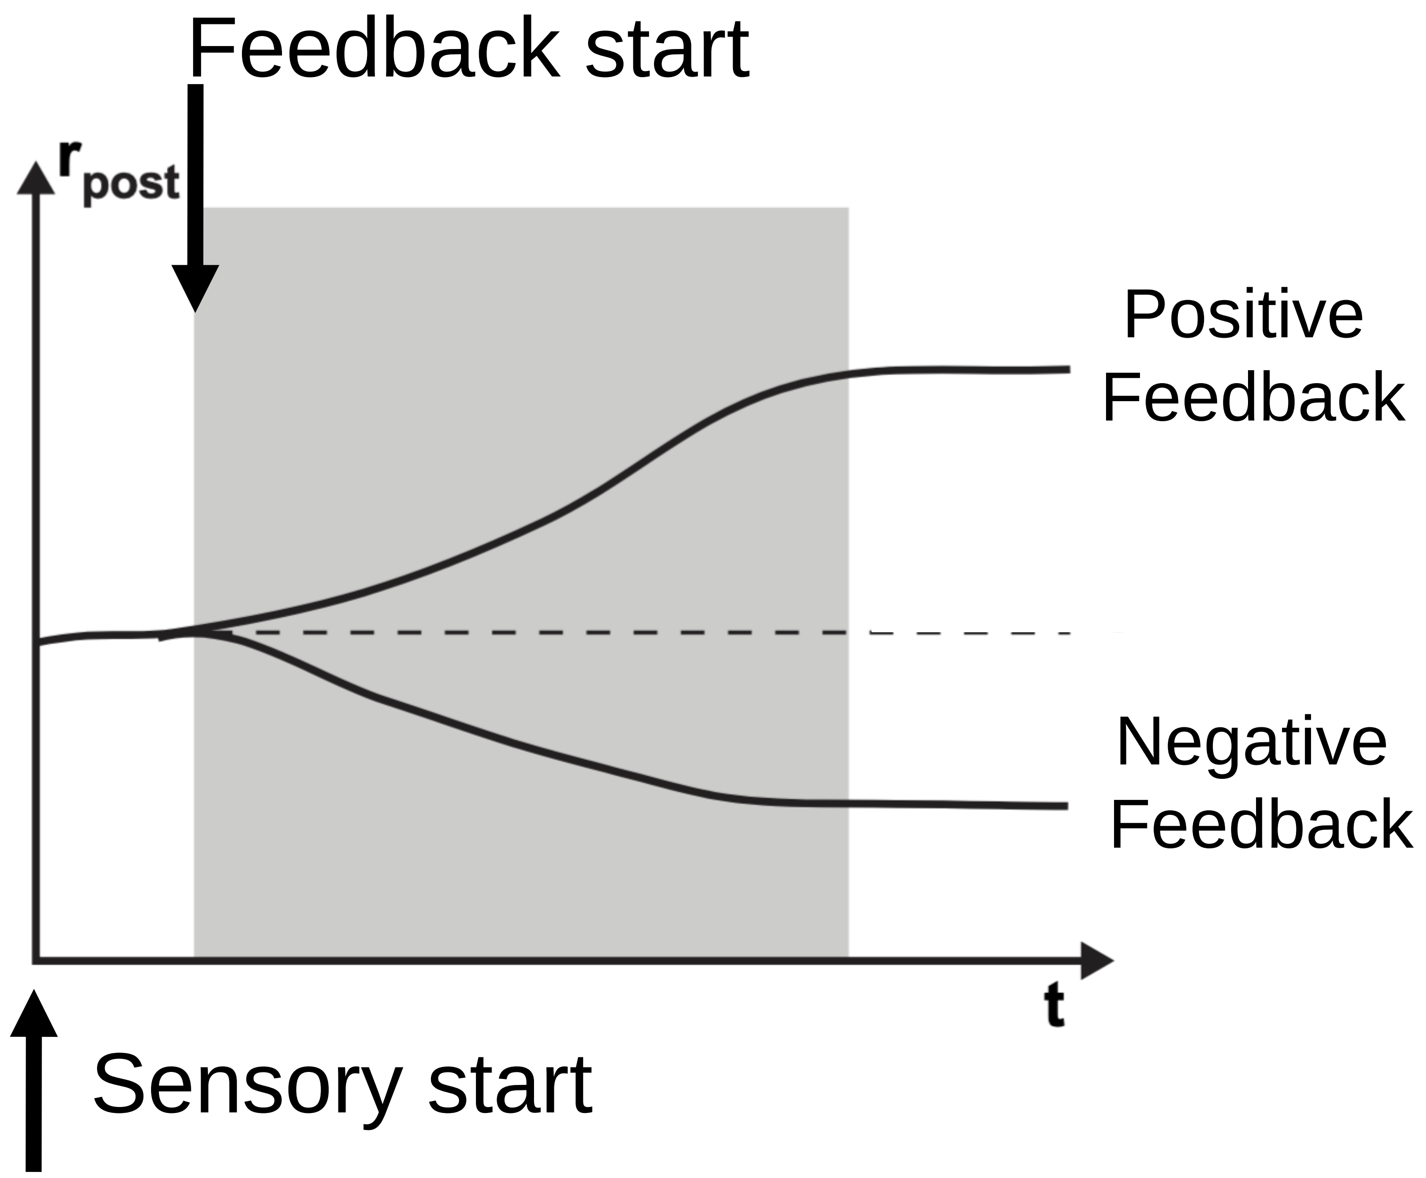

Weight Update

$$\Delta w \propto \int r_{\text{pre}}(t)\dot{r}_{\text{post}}(t)dt$$

Aceituno et al, 2023, Front. Comp. Neurosc

Mimicking Biological Plasticity

Biology:
Spike Timing
Dependent Plasticity
(STDP)

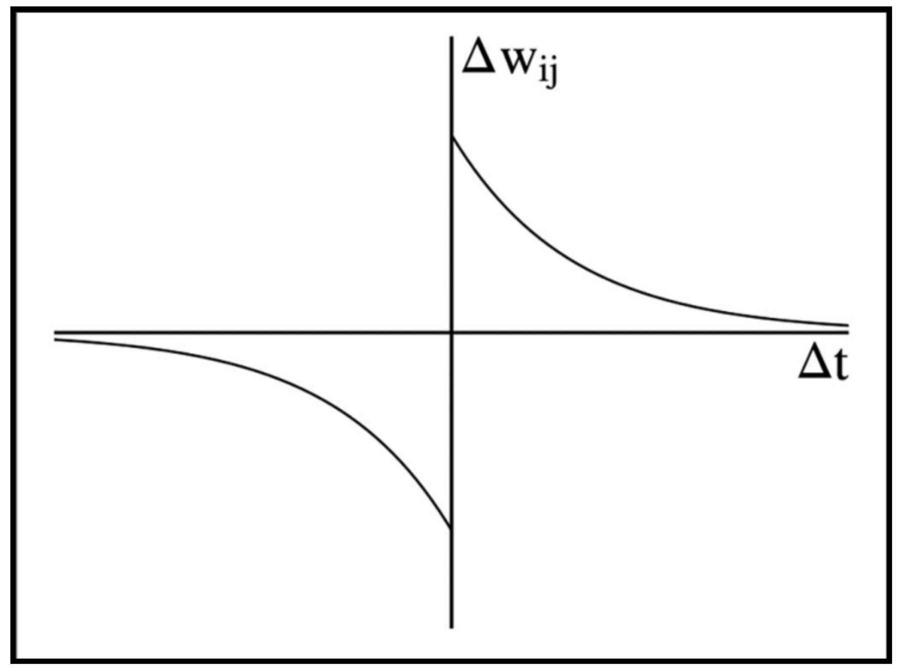

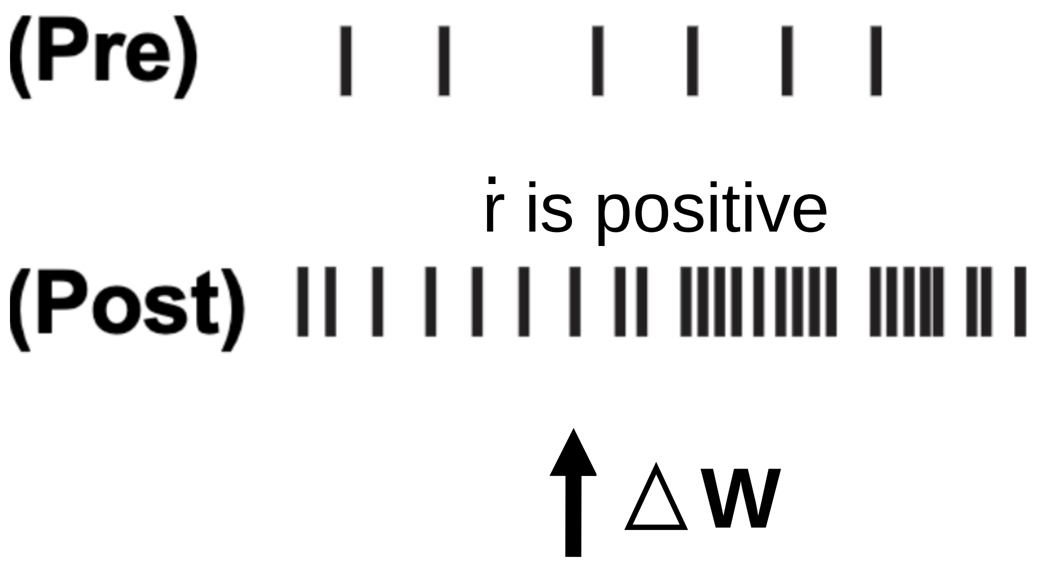

Learning Through Control
Differential Hebbian +
Feedback Control

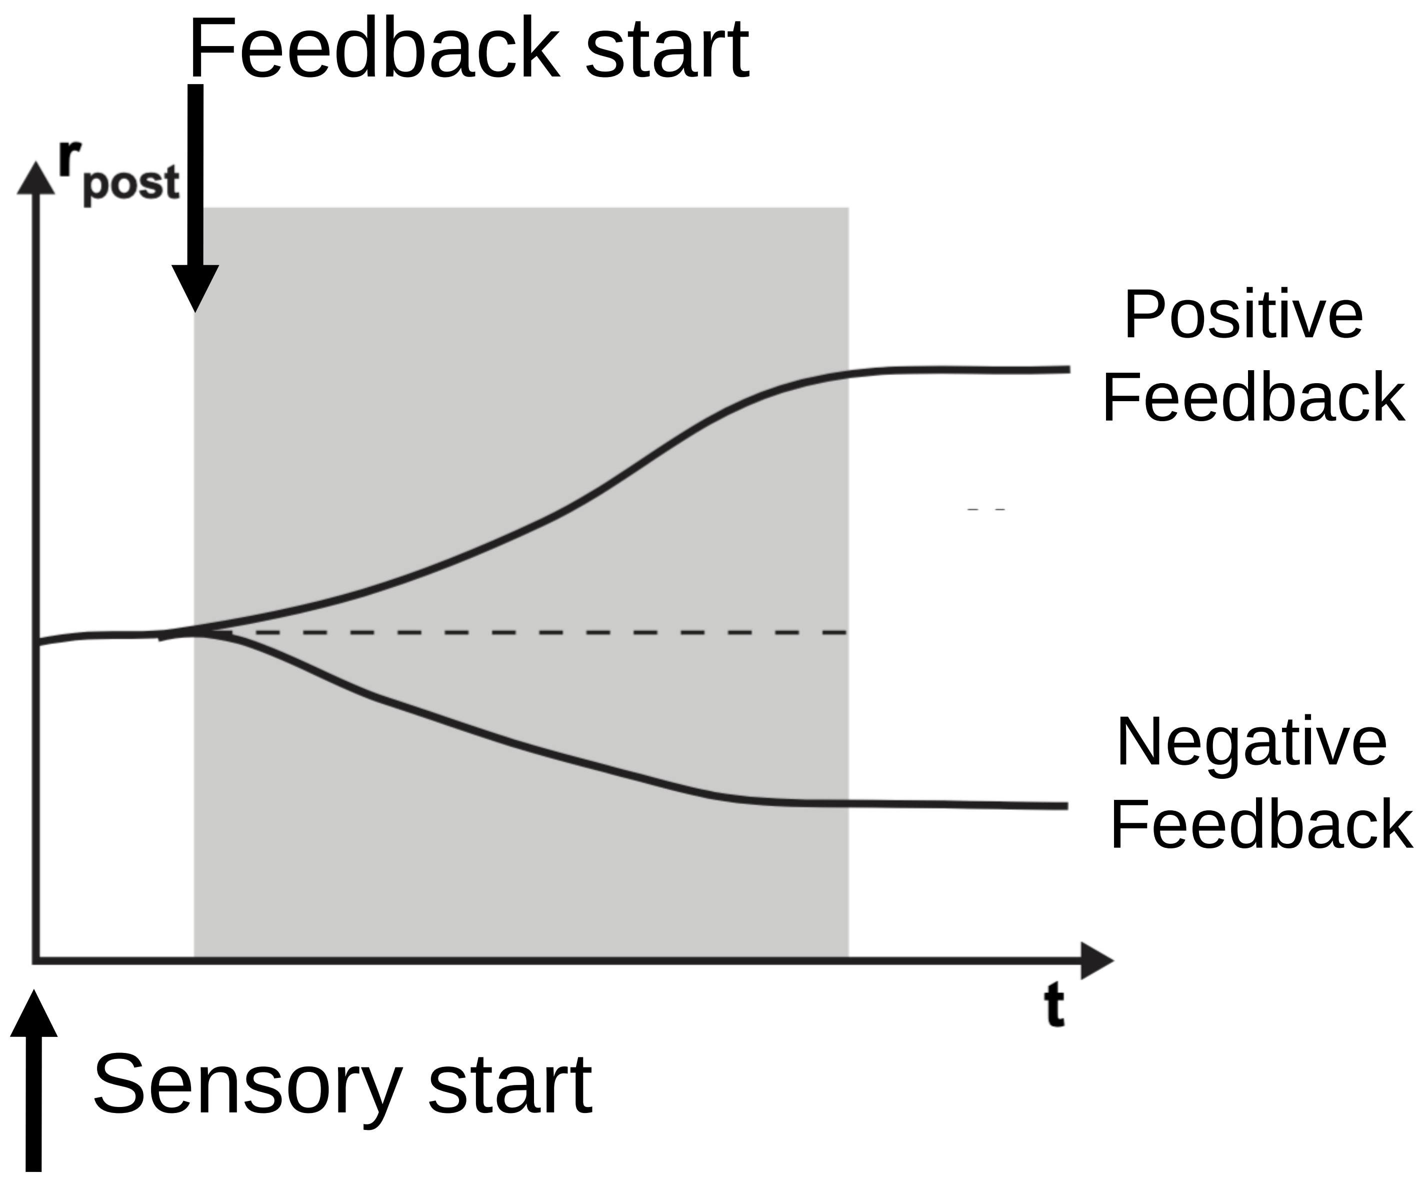

Weight Update

$$\Delta w \propto \int r_{\text{pre}}(t)\dot{r}_{\text{post}}(t)dt$$

Aceituno et al, 2023, Front. Comp. Neurosc

Mimicking Biological Plasticity

Biology:
Spike Timing
Dependent Plasticity
(STDP)

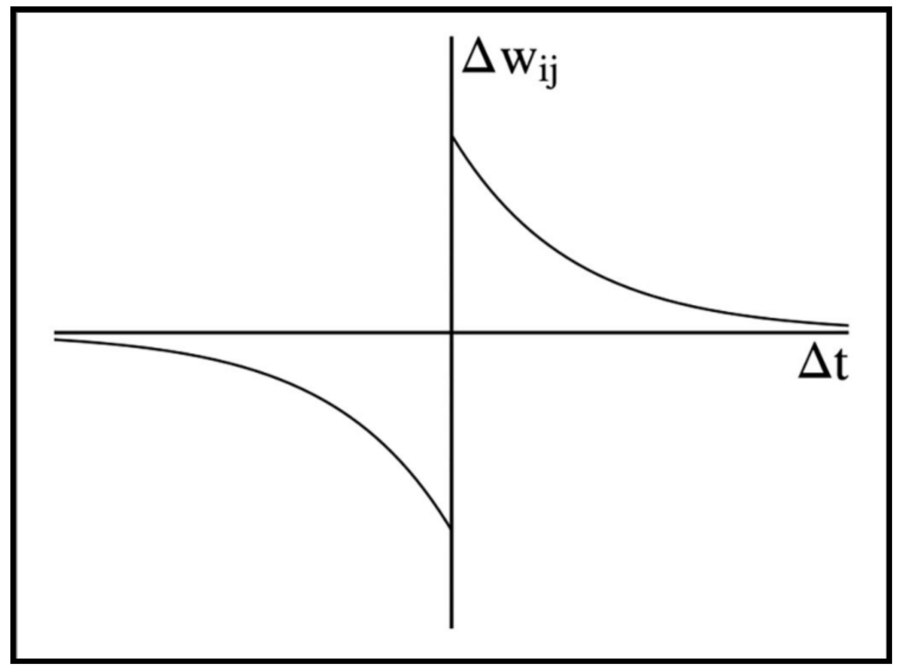

=

Learning Through Control
Differential Hebbian +
Feedback Control

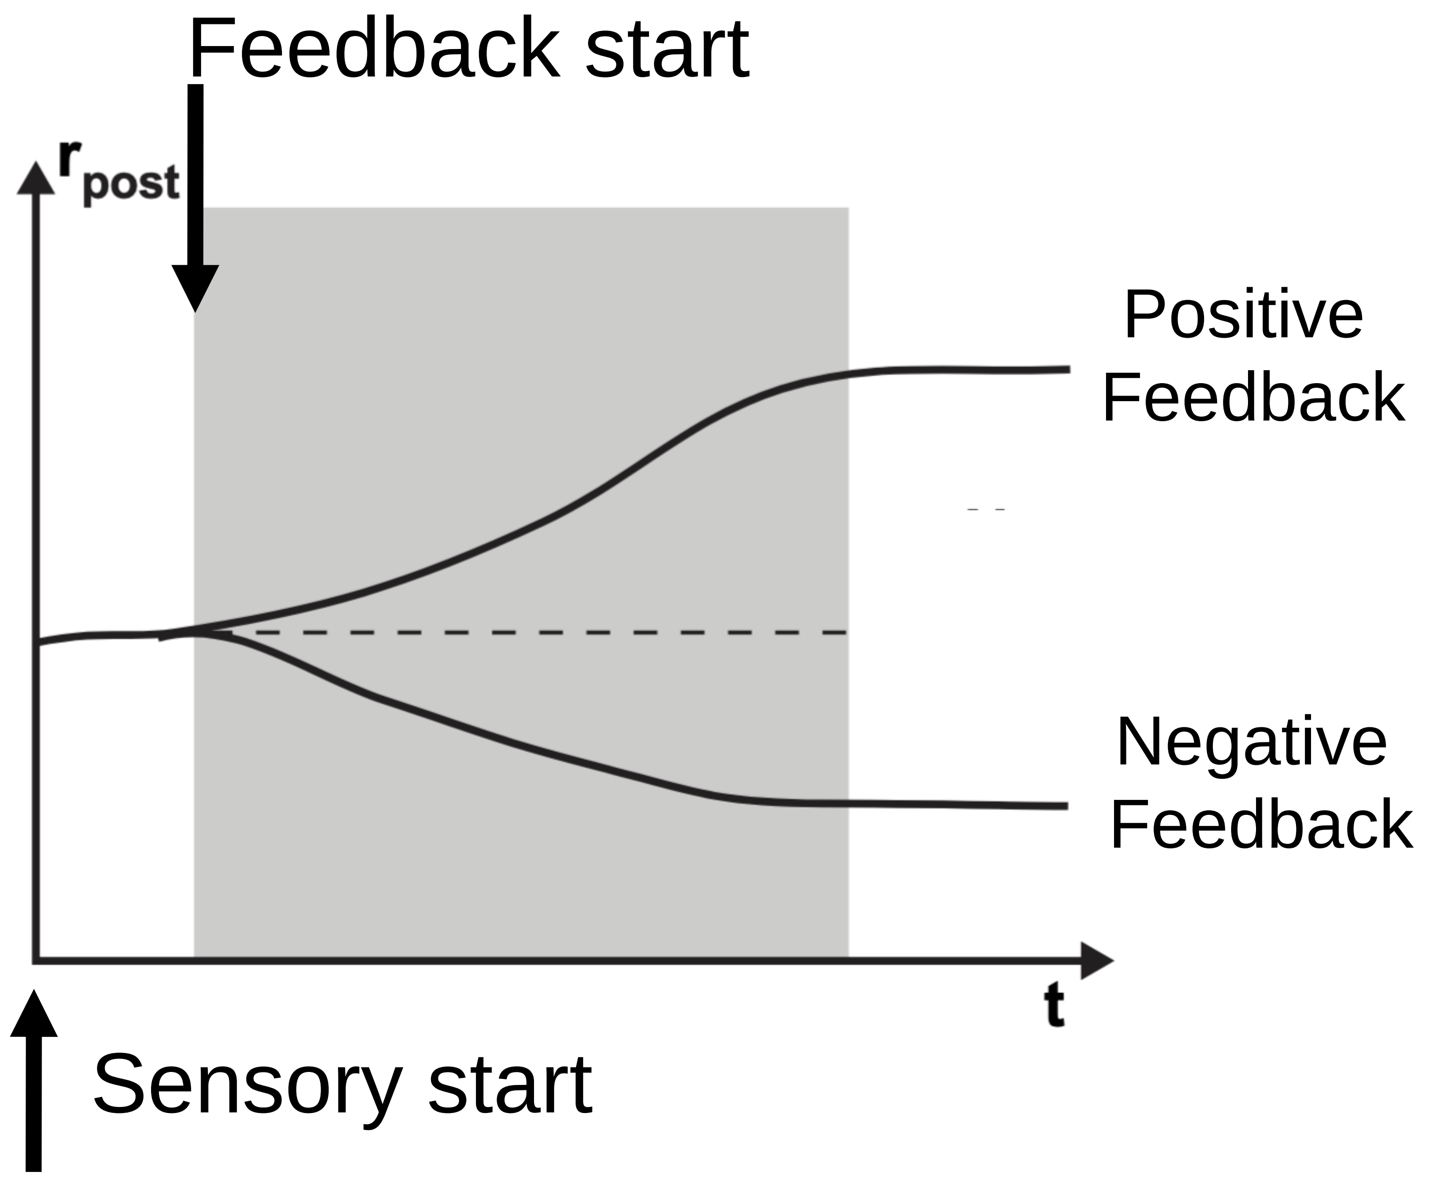

Weight Update

$$\Delta w \propto \int r_{\text{pre}}(t)\dot{r}_{\text{post}}(t)dt$$

Aceituno et al, 2023, Front. Comp. Neurosc

ETH zürich

Temporal Hebbian Updates
Allow to Learn Cortical Hierarchies

institute of neuroinformatics

MNIST Dataset
(standard benchmark for
handwritten digit recognition)

Aceituno et al, 2023, Front. Comp. Neurosc

Temporal Hebbian Updates Allow to Learn Cortical Hierarchies

MNIST Dataset
(standard benchmark for handwritten digit recognition)

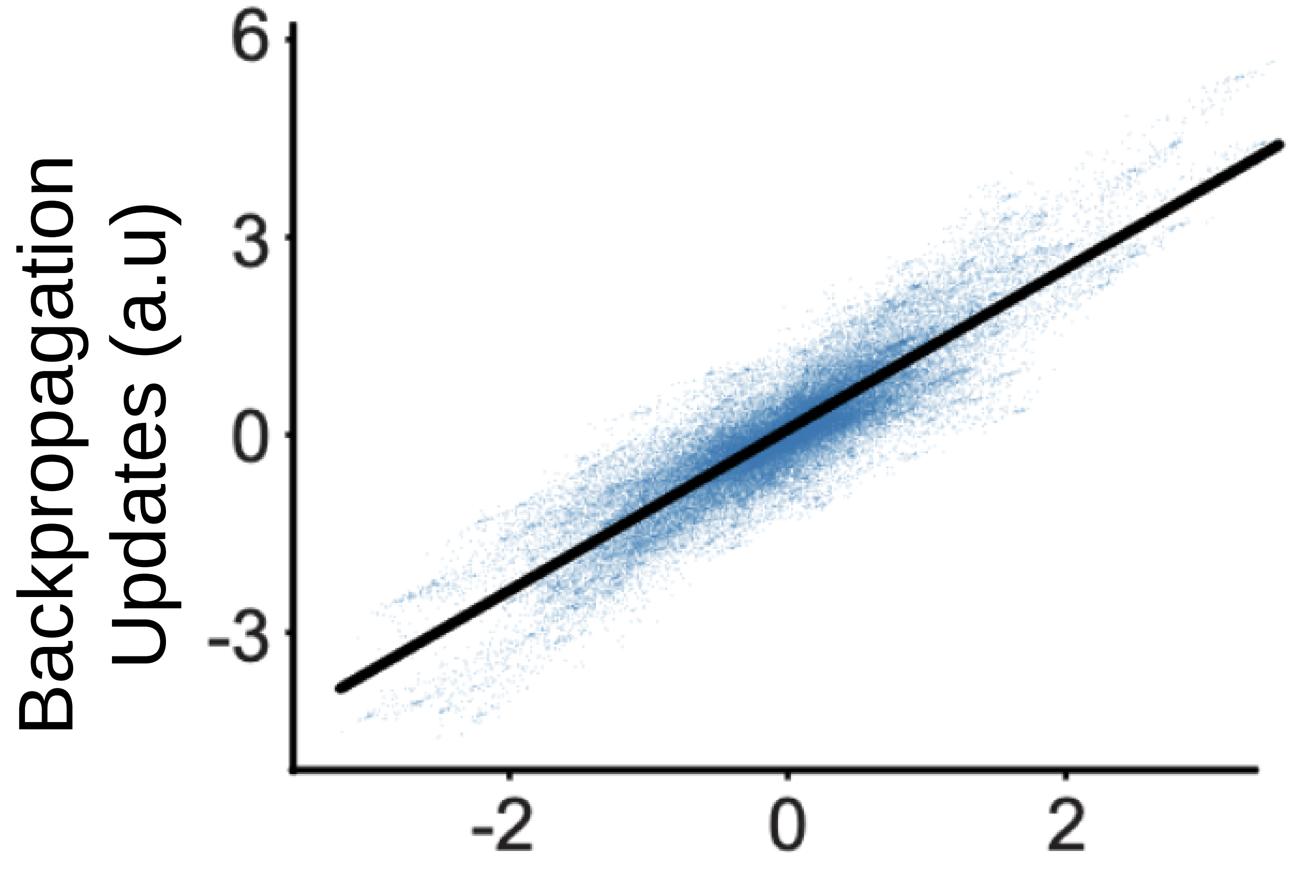

Aceituno et al, 2023, Front. Comp. Neurosc

ETH zürich

institute of neuroinformatics

Temporal Hebbian Updates Allow to Learn Cortical Hierarchies

MNIST Dataset
(standard benchmark for handwritten digit recognition)

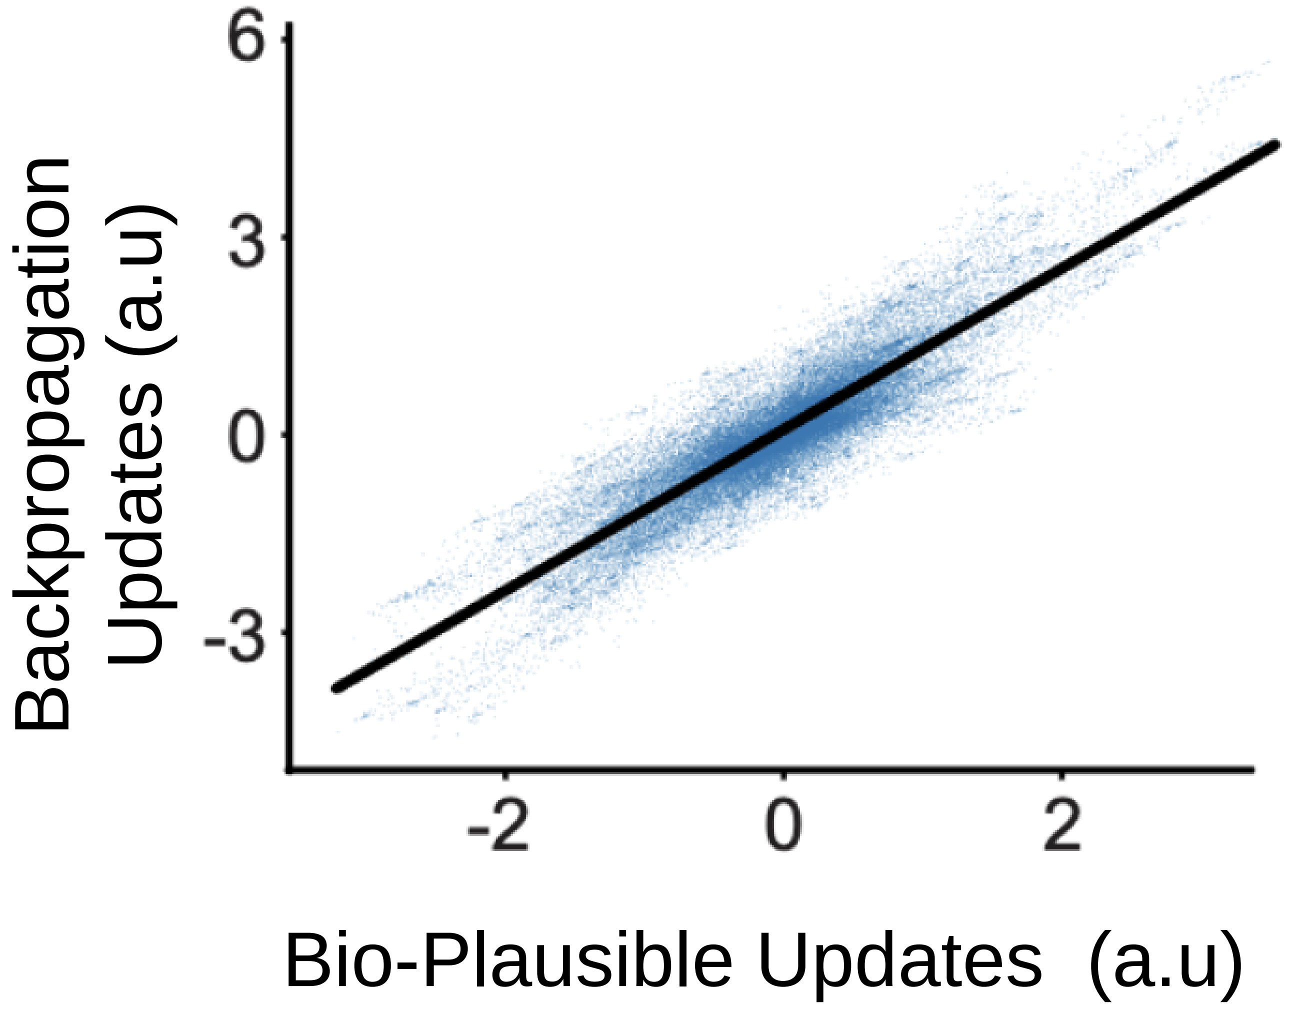

	Classification Error
BP	1.74 ± 0.10%
DFC	1.98 ± 0.05%
Bio-DFC	1.89 ± 0.15%

Aceituno et al, 2023, Front. Comp. Neurosc

Temporal Hebbian Updates Allow to Learn Cortical Hierarchies

MNIST Dataset
(standard benchmark for handwritten digit recognition)

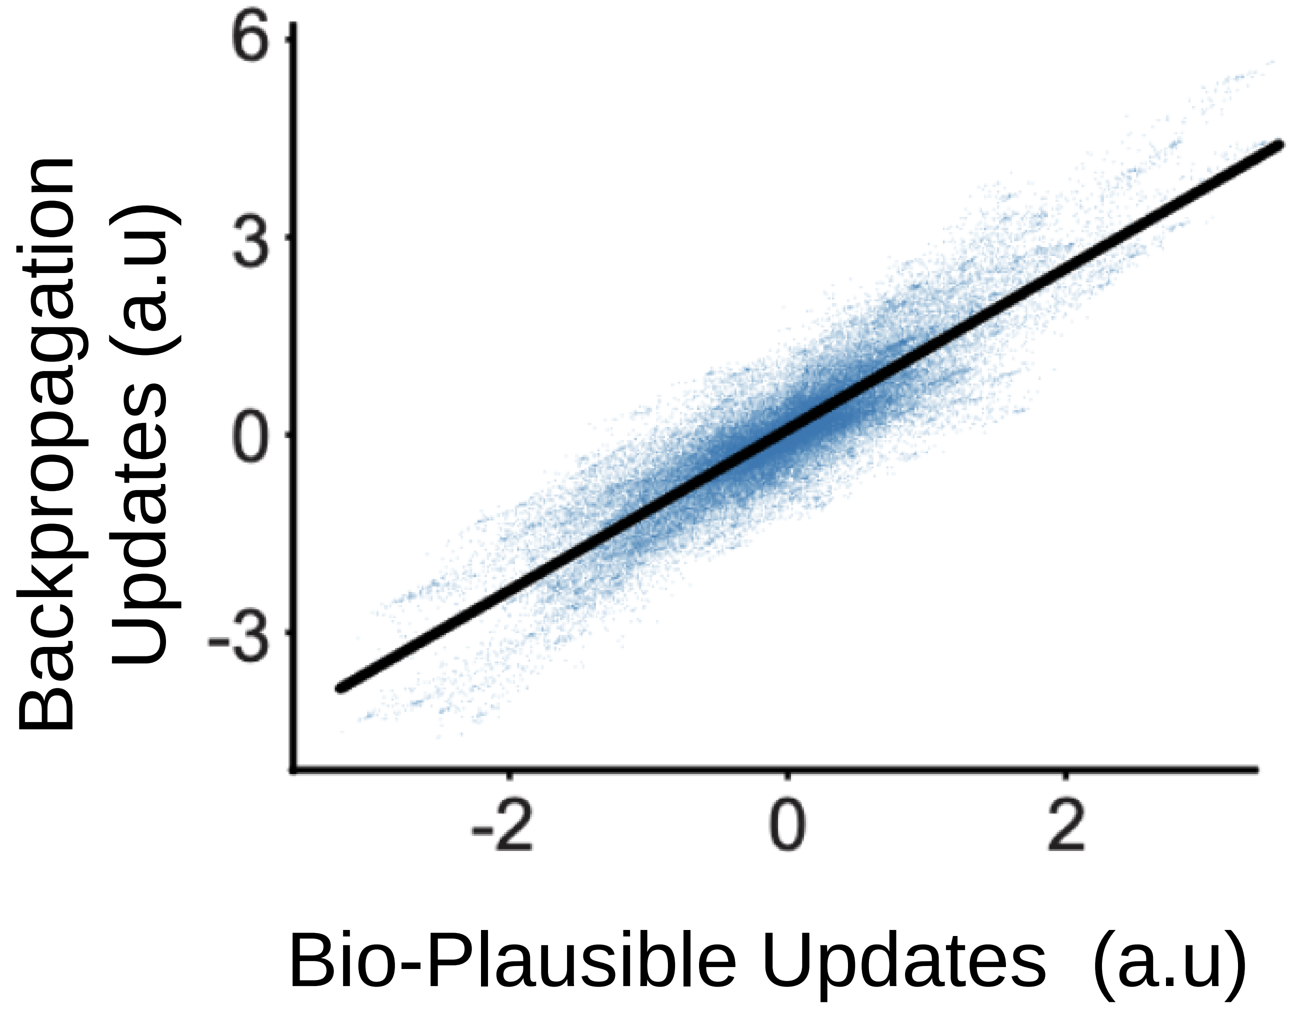

	Classification Error
BP	1.74 ± 0.10%
DFC	1.98 ± 0.05%
Bio-DFC	1.89 ± 0.15%

Classification Error

Aceituno et al, 2023, Front. Comp. Neurosc

Hierarchical (Deep) Learning though Control

Backpropagation of the Error

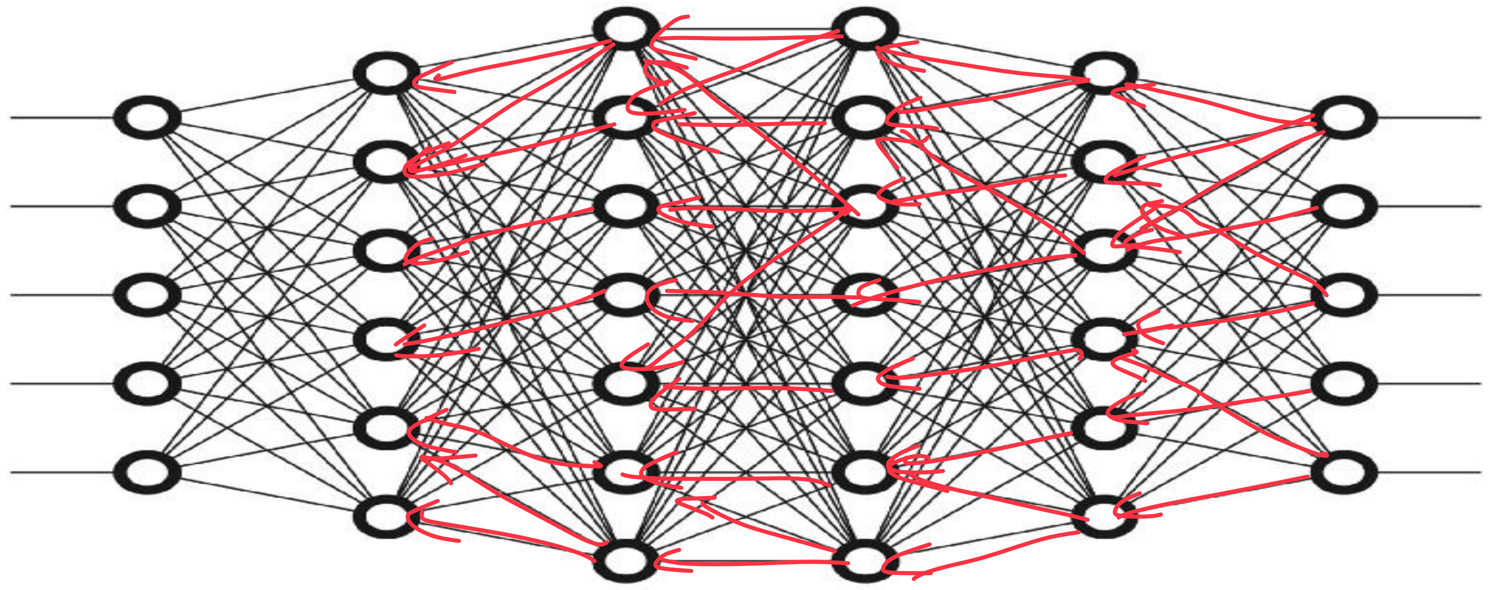

Learning Through Controlling a Complex Dynamic System

Learning = Reducing Control Feedback!

Advantages:

- Continuous / asynchronous updates and data processing.
- Absence of phases or back-propagation errors (+ uses less memory).
- Naturally works with time series data.
- Very simple learning rule that is local in space and time.

Aceituno et al, 2023, Front. Comp. Neurosc

Hierarchical (Deep) Network Learning though Control

- **Enables Continual Learning when neuronal activity is sparse.**

(Deep) Learning though Control in Hardware

Indiveri Group @ INI

Matteo Saponati
(Grewe lab)

Neuromorphic Devices

- Ultra low-power
- Scalable to 3D

<u>Learning by Controlling a complex Dynamic System</u>

<u>Building a Neuromorphic Device for low-power Deep Learning</u>

Can do Continual Learning when neuronal activity is sparse.

- **<u>Is ideally suited for low-power deep learning on neuromorphic processors.</u>**

(Deep) Learning though Control in Hardware

Indiveri Group @ INI

Matteo Saponati
(Grewe lab)

Neuromorphic Devices
- Ultra low-power
- Scalable to 3D

Learning by Controlling a complex Dynamic System → **Building a Neuromorphic Device for low-power Deep Learning**

Cartiglia et al, 2022, IEEE

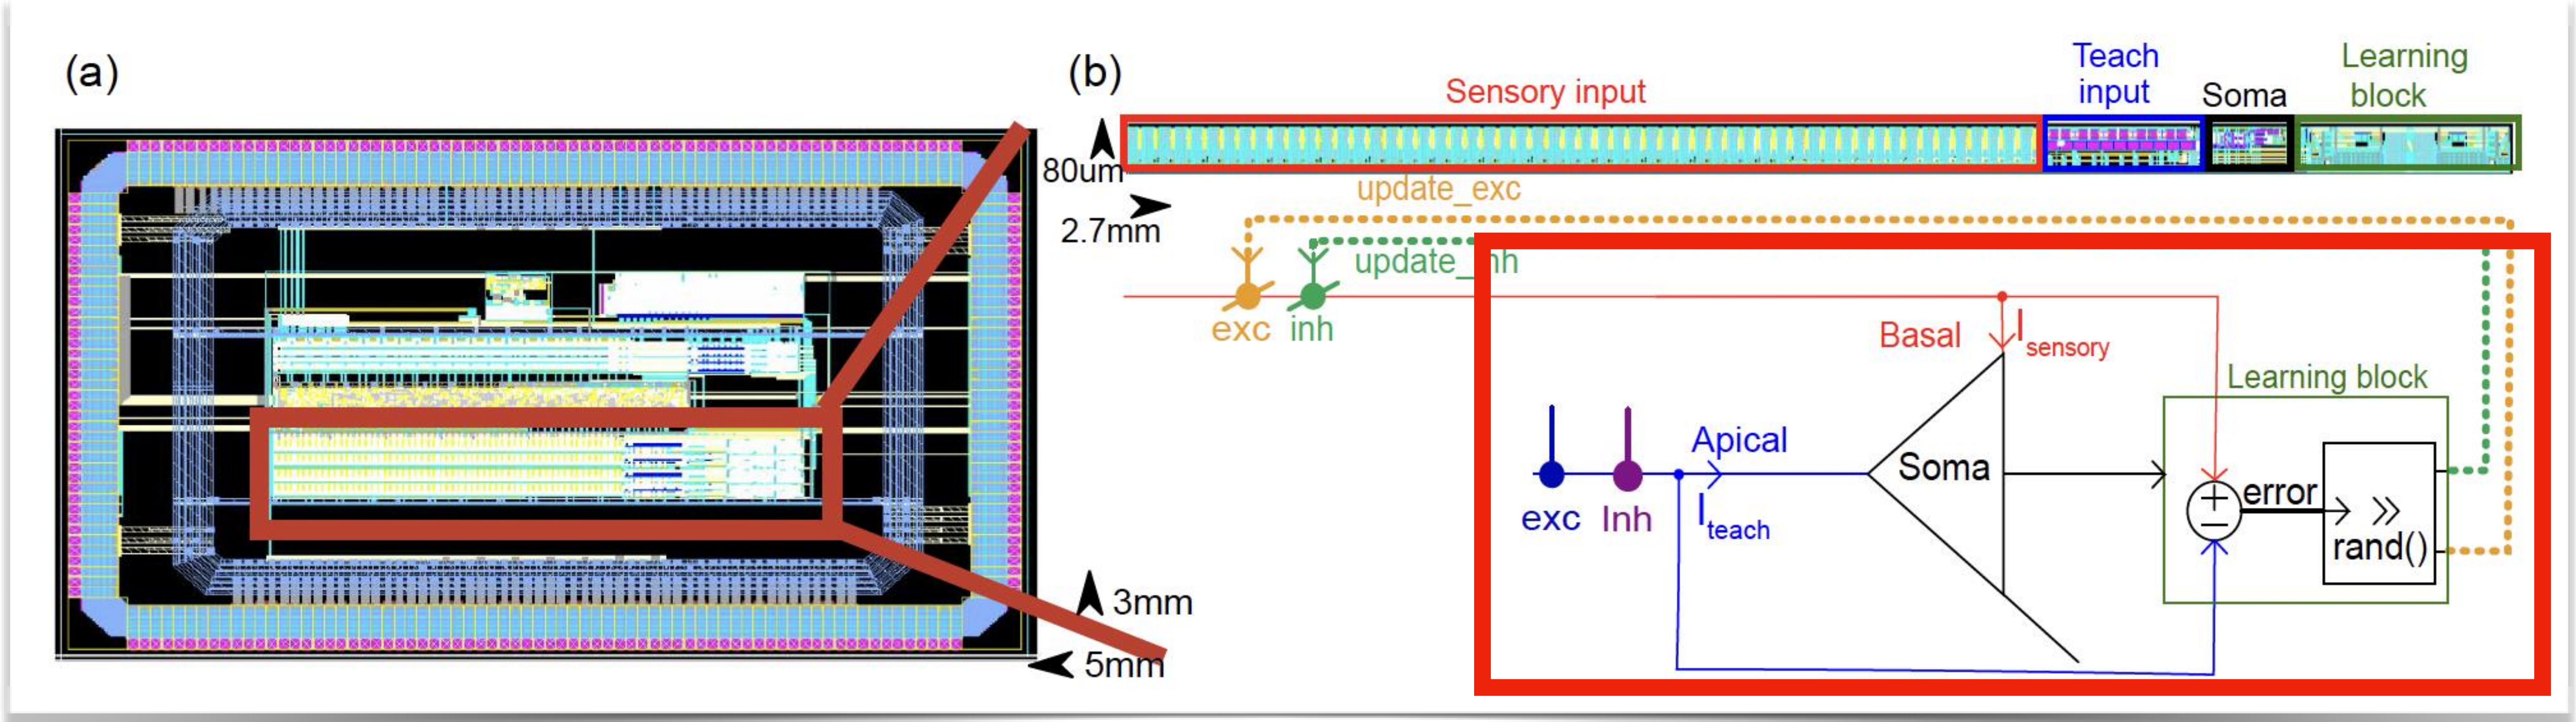

Pyramidal Neuron In Silico

- Can do Continual Learning when neuronal activity is sparse.
- **Is ideally suited for low-power deep learning on neuromorphic processors.**

(Deep) Learning though Control in Hardware

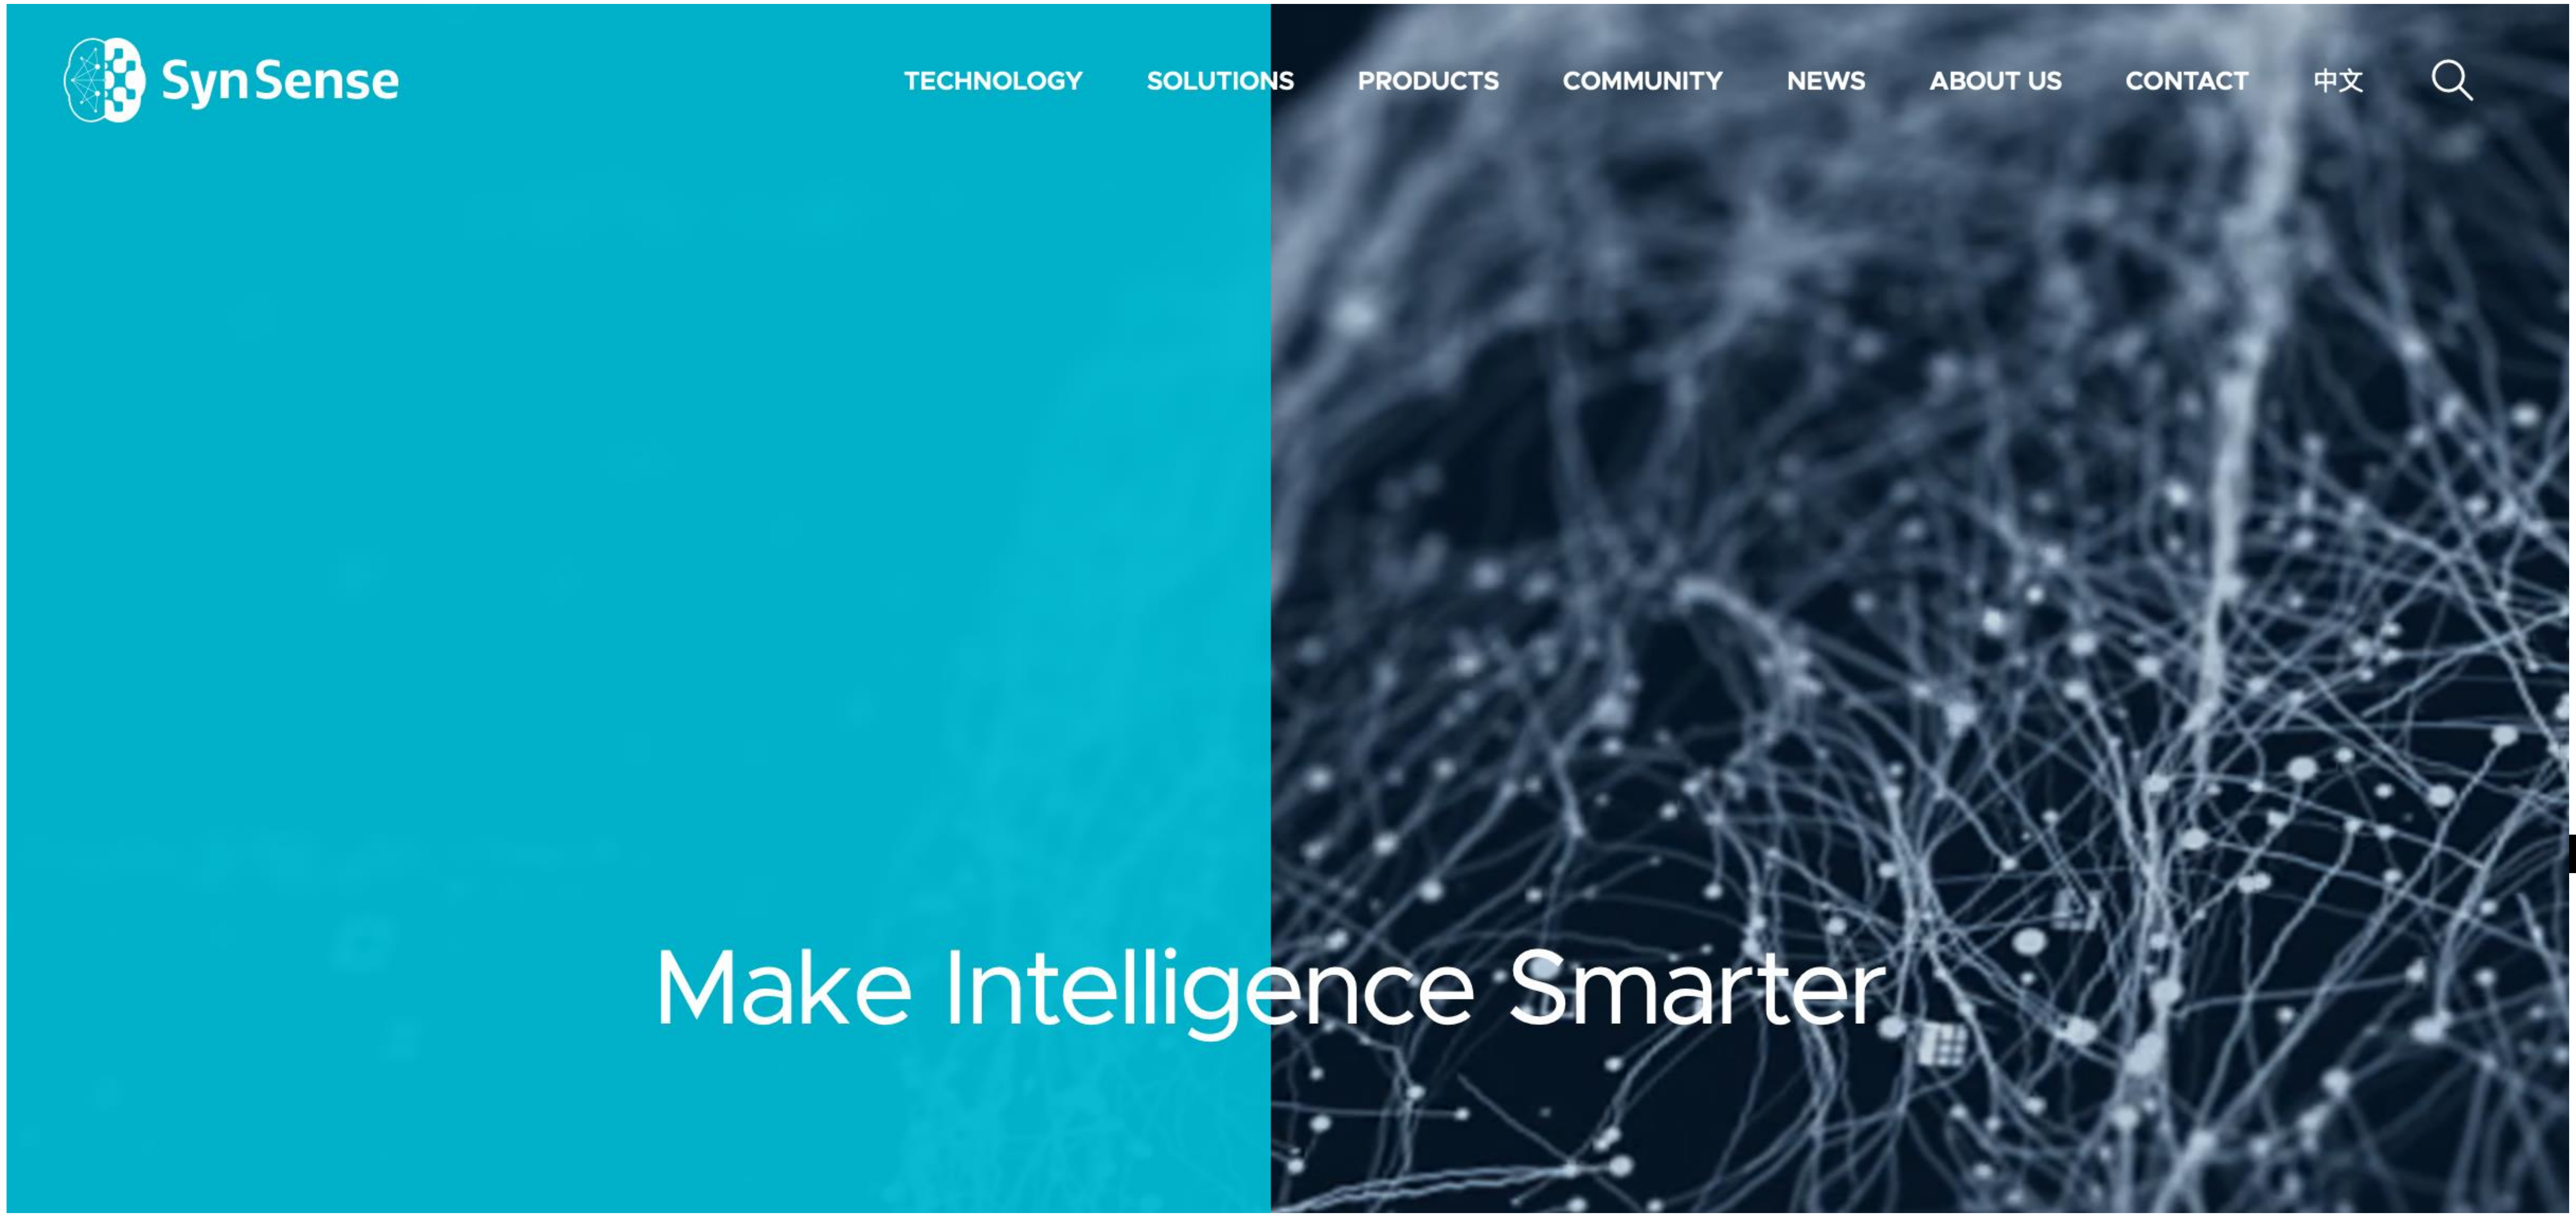

Building a Neuromorphic Device for Deep Learning

- Can do Continual Learning when neuronal activity is sparse.
- **<u>Is ideally suited for low-power neuromorphic processors.</u>**

ETH zürich

AI → Neuroscience

institute of neuroinformatics

Weight Update

$$\Delta w \propto \int r_{\text{pre}}(t)\dot{r}_{\text{post}}(t)dt$$

Plasticity is only 'ON' when Feedback is active!

Weight Update

$$\Delta w \propto \int r_{\text{pre}}(t)\dot{r}_{\text{post}}(t)dt$$

Plasticity is only 'ON' when Feedback is active!

In Vitro Measurement of
Neuronal Activity & Plasticity

We are testing this theoretical prediction in biological neurons!

ETH zürich

institute of neuroinformatics

AI → Neuroscience

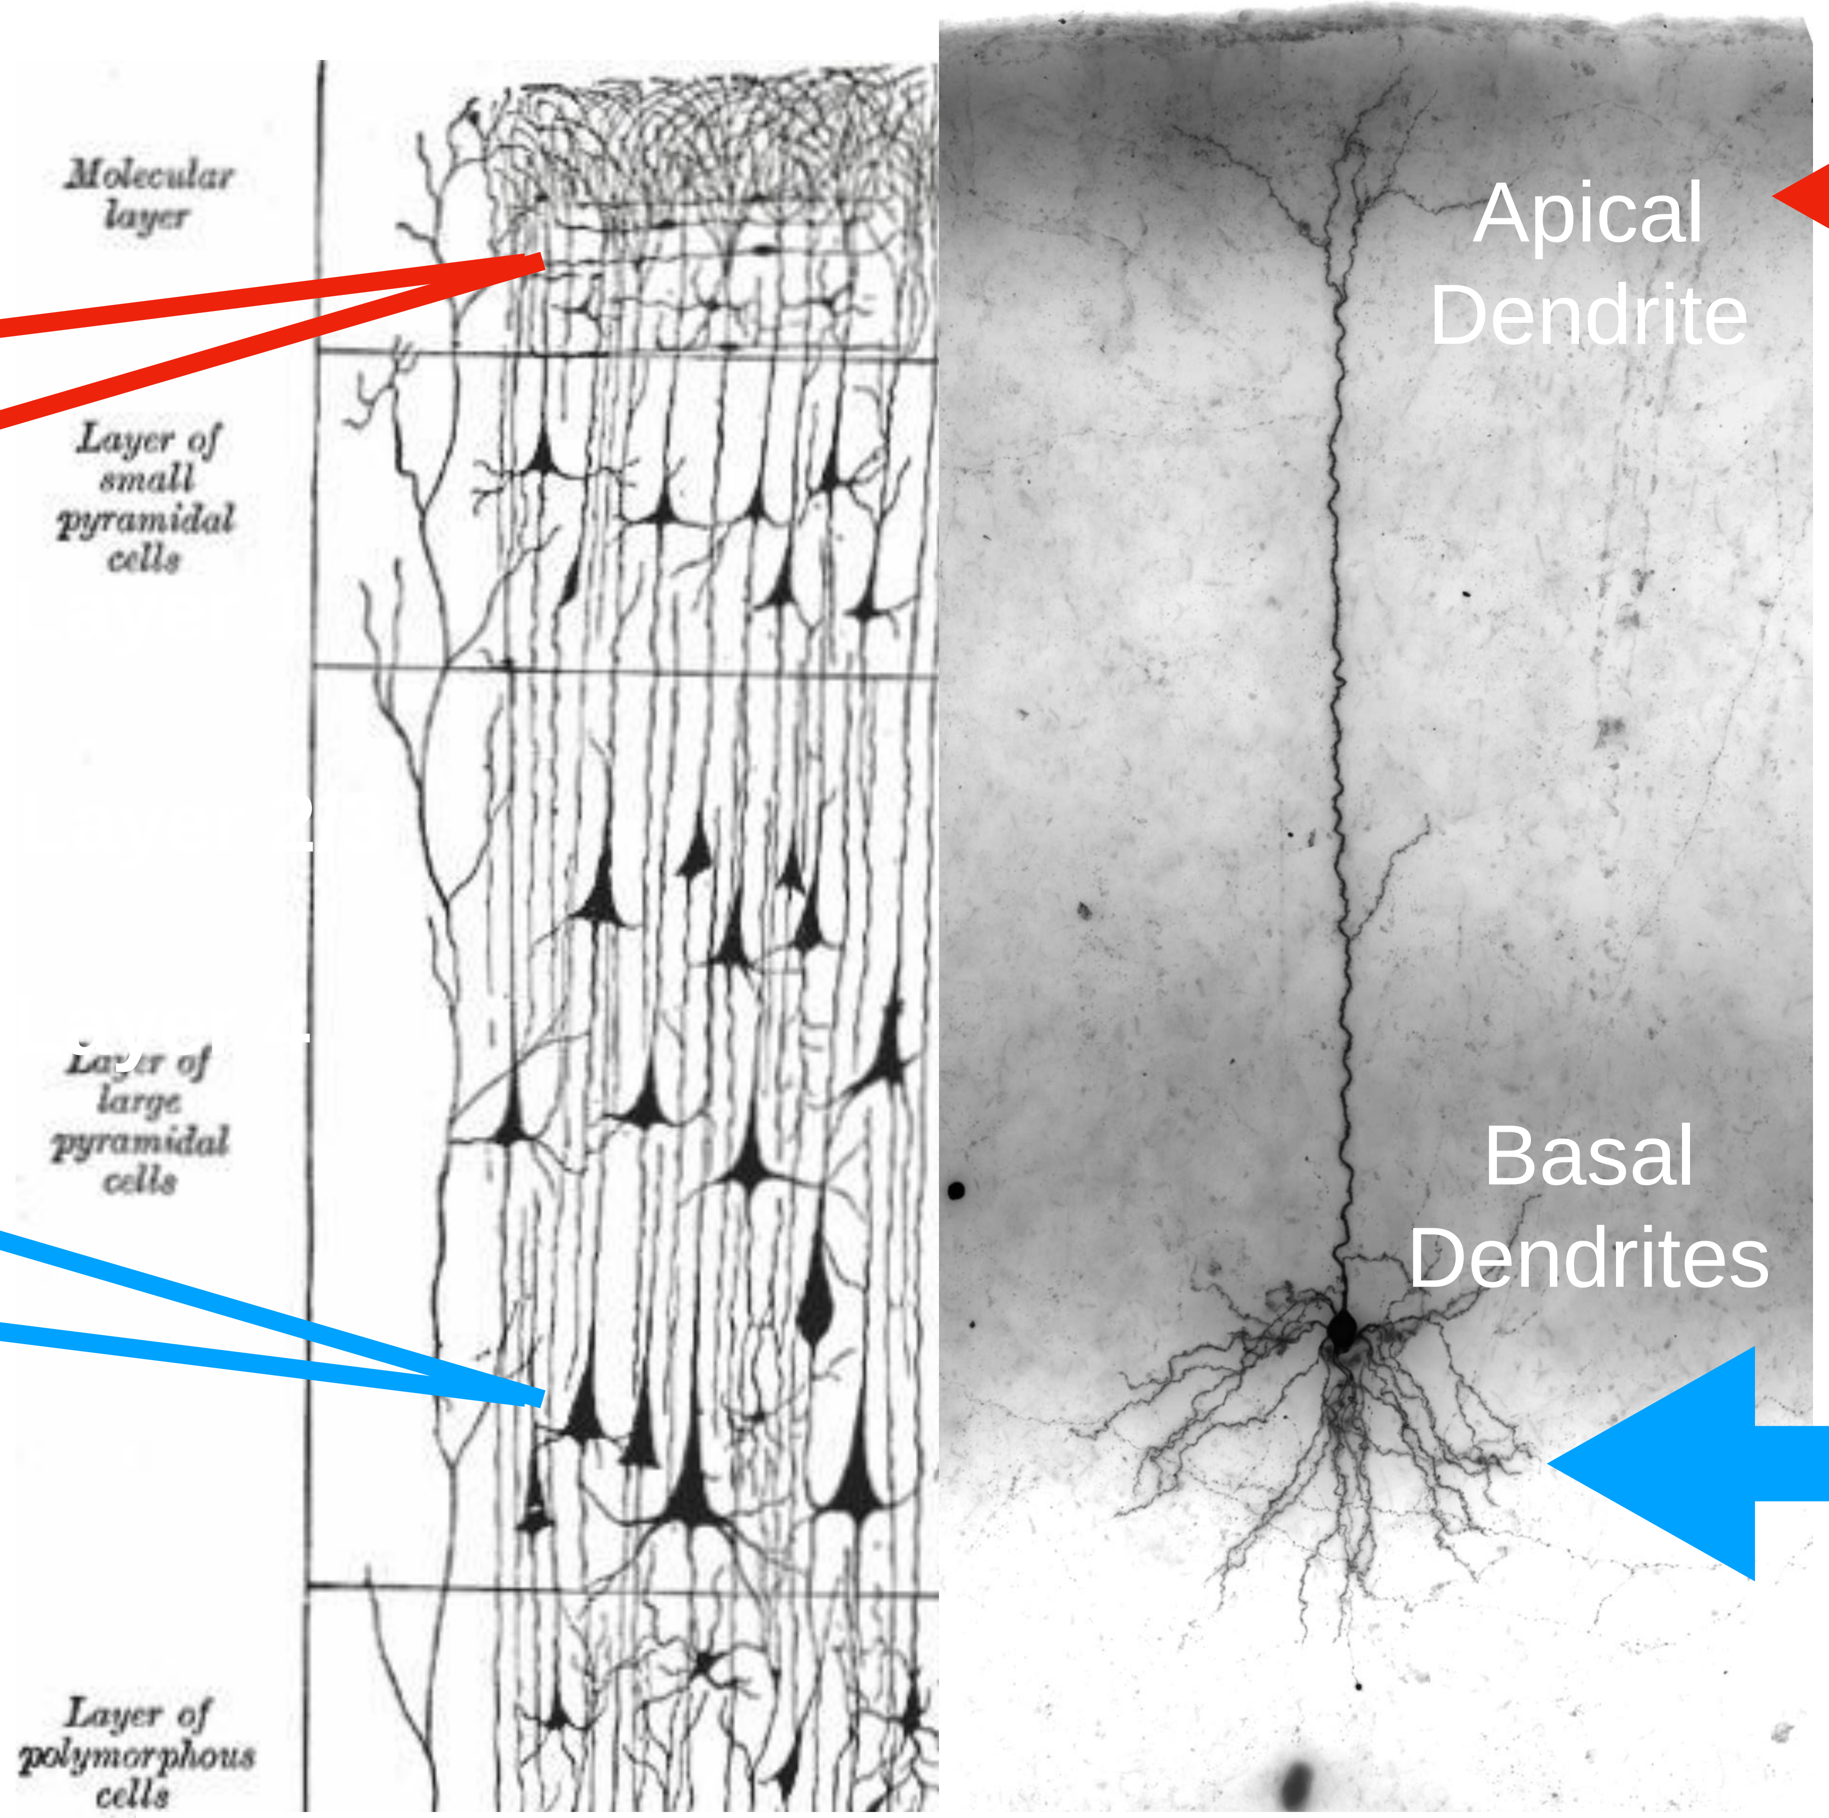

Top-Down Feedback (Error)

Electrical Stimulation Pipette

Sensory Input

ETH zürich

AI → Neuroscience

institute of neuroinformatics

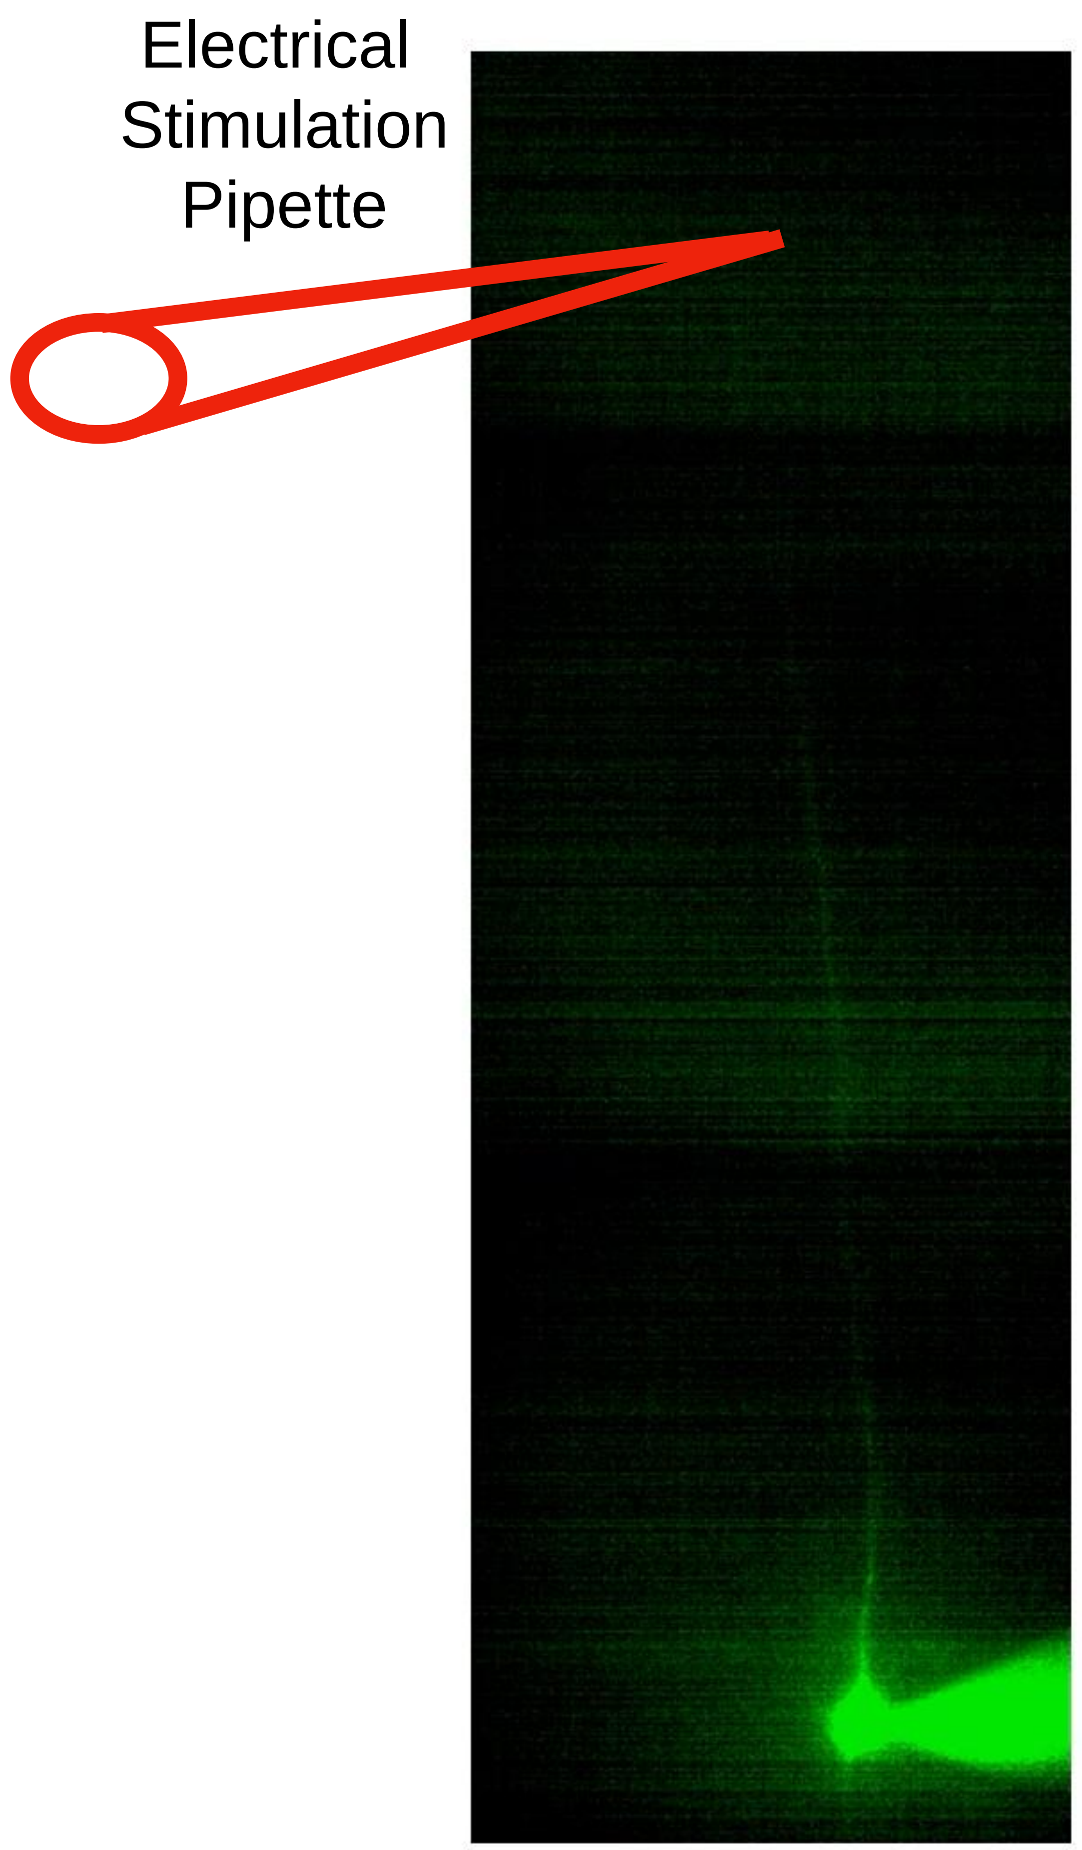

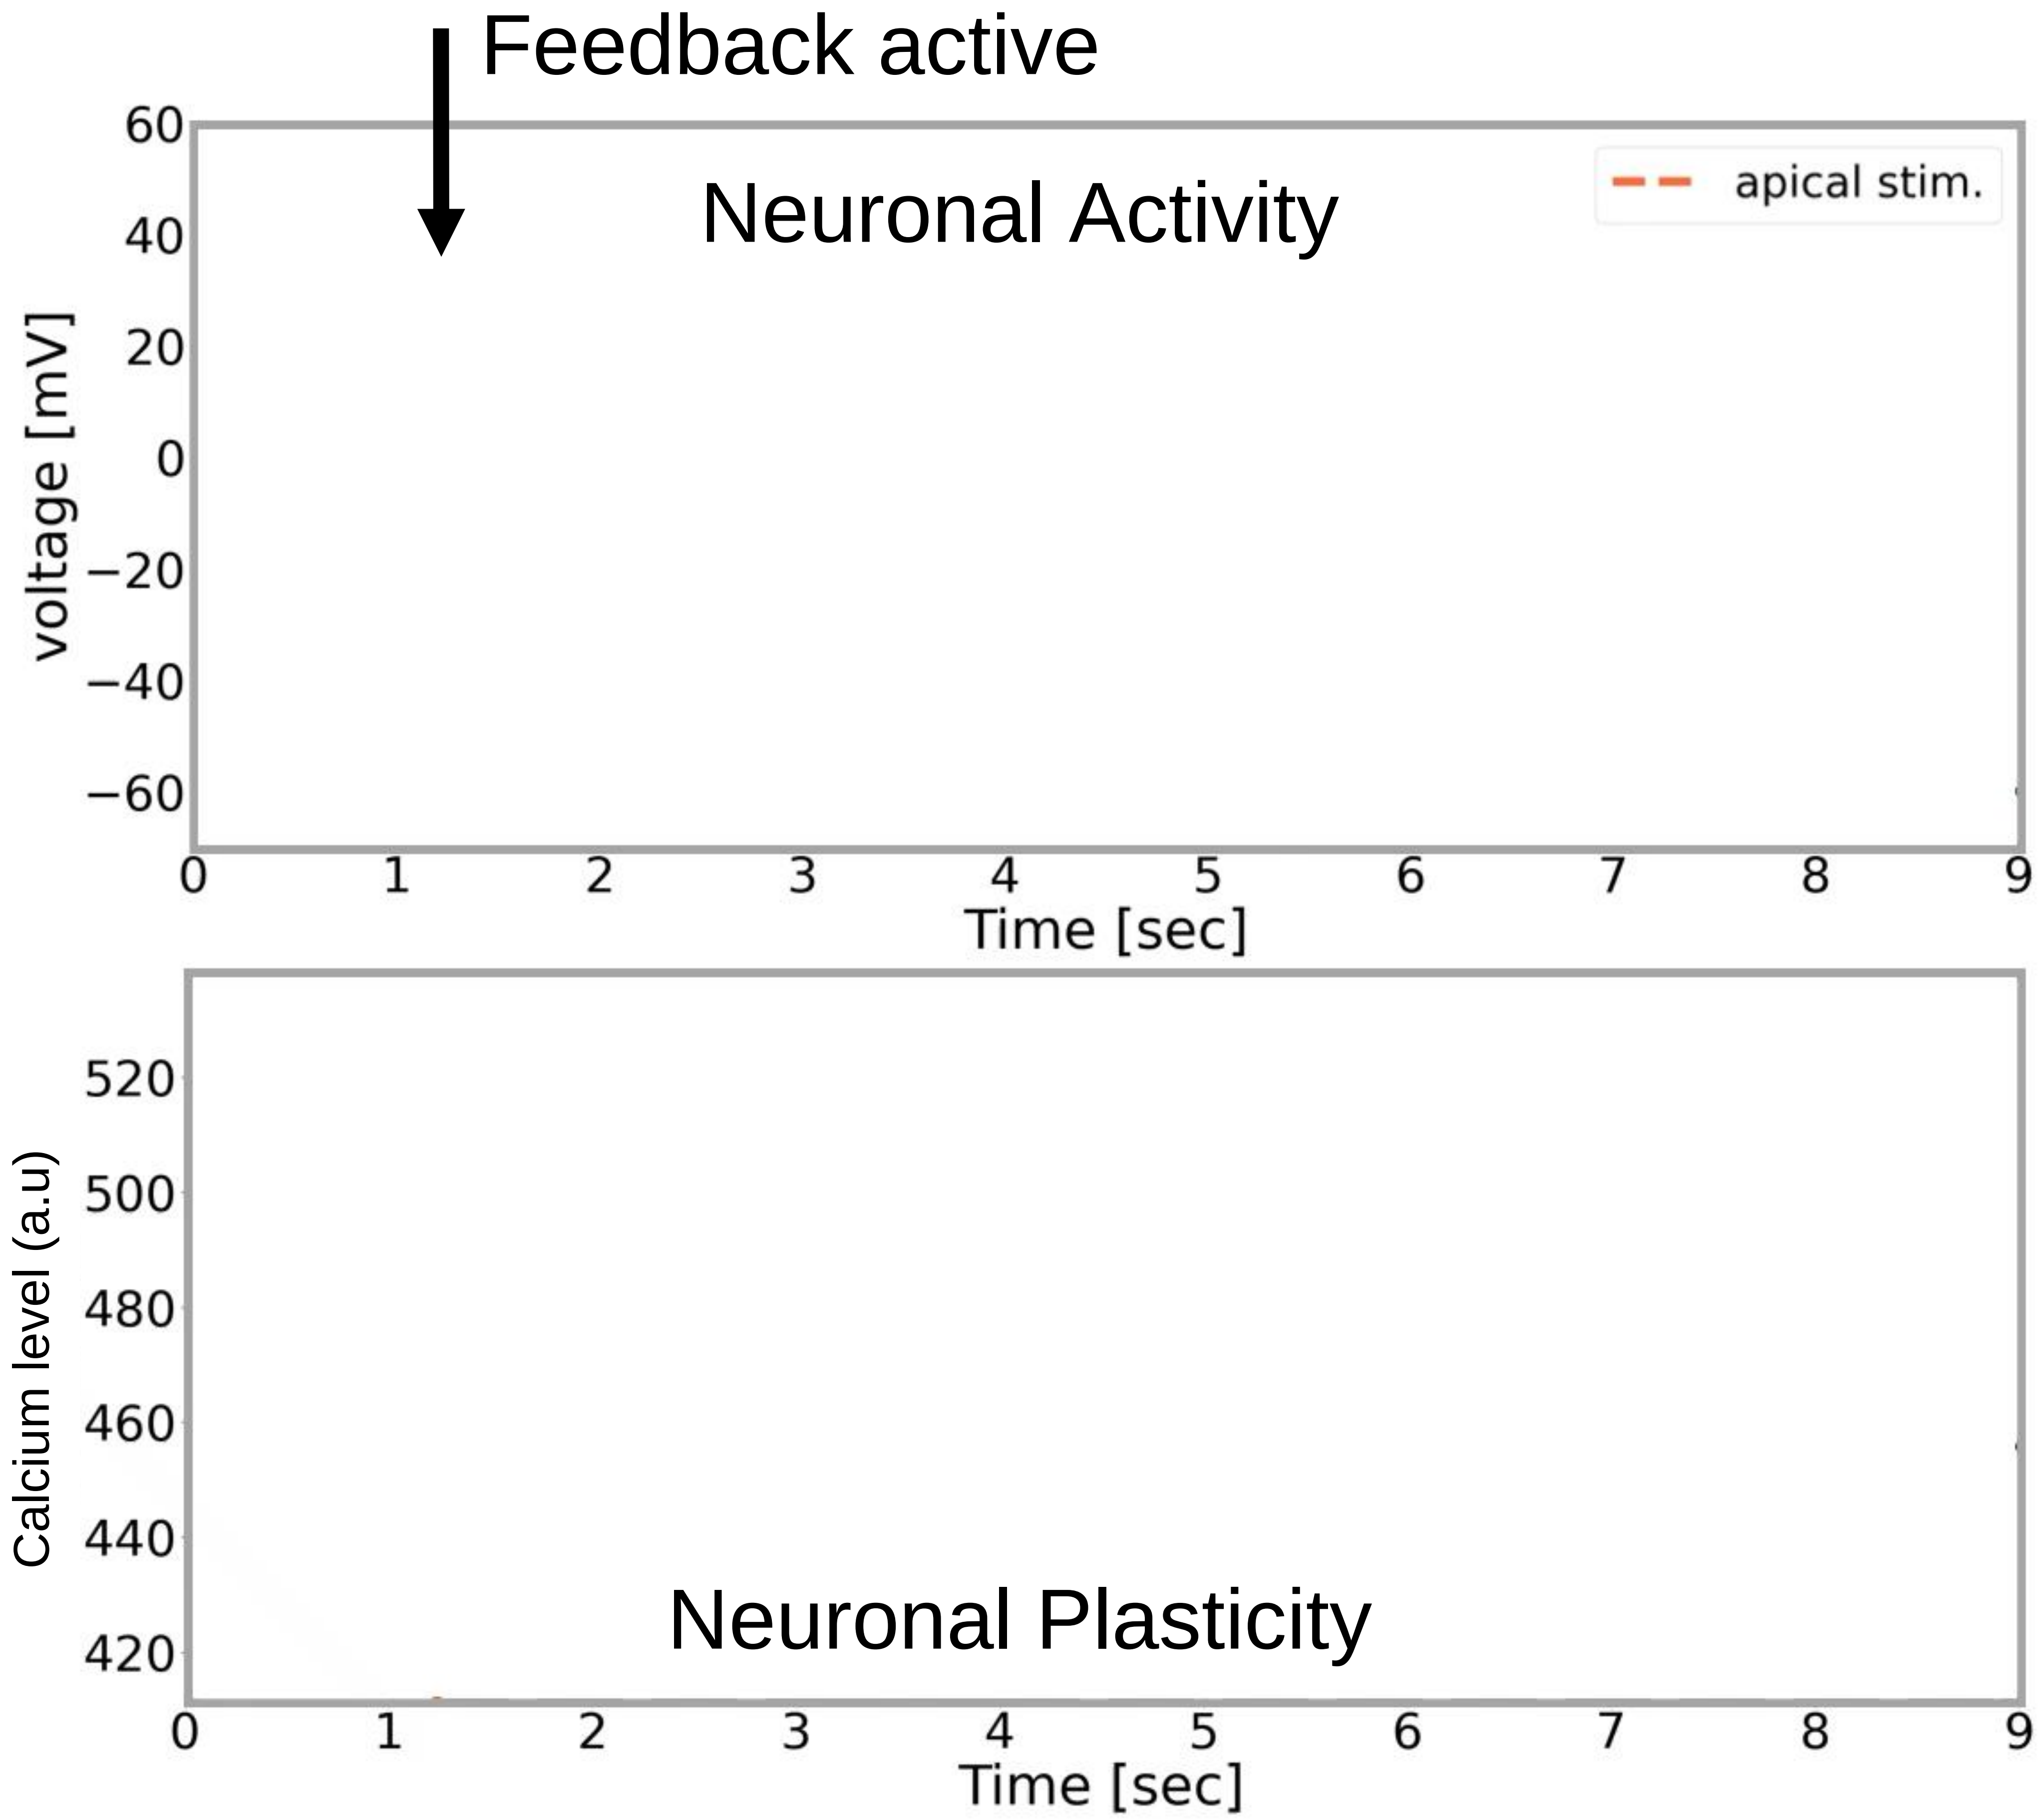

Video Credit: R. Loidl

ETH zürich

Summary Part 1:

institute of neuroinformatics

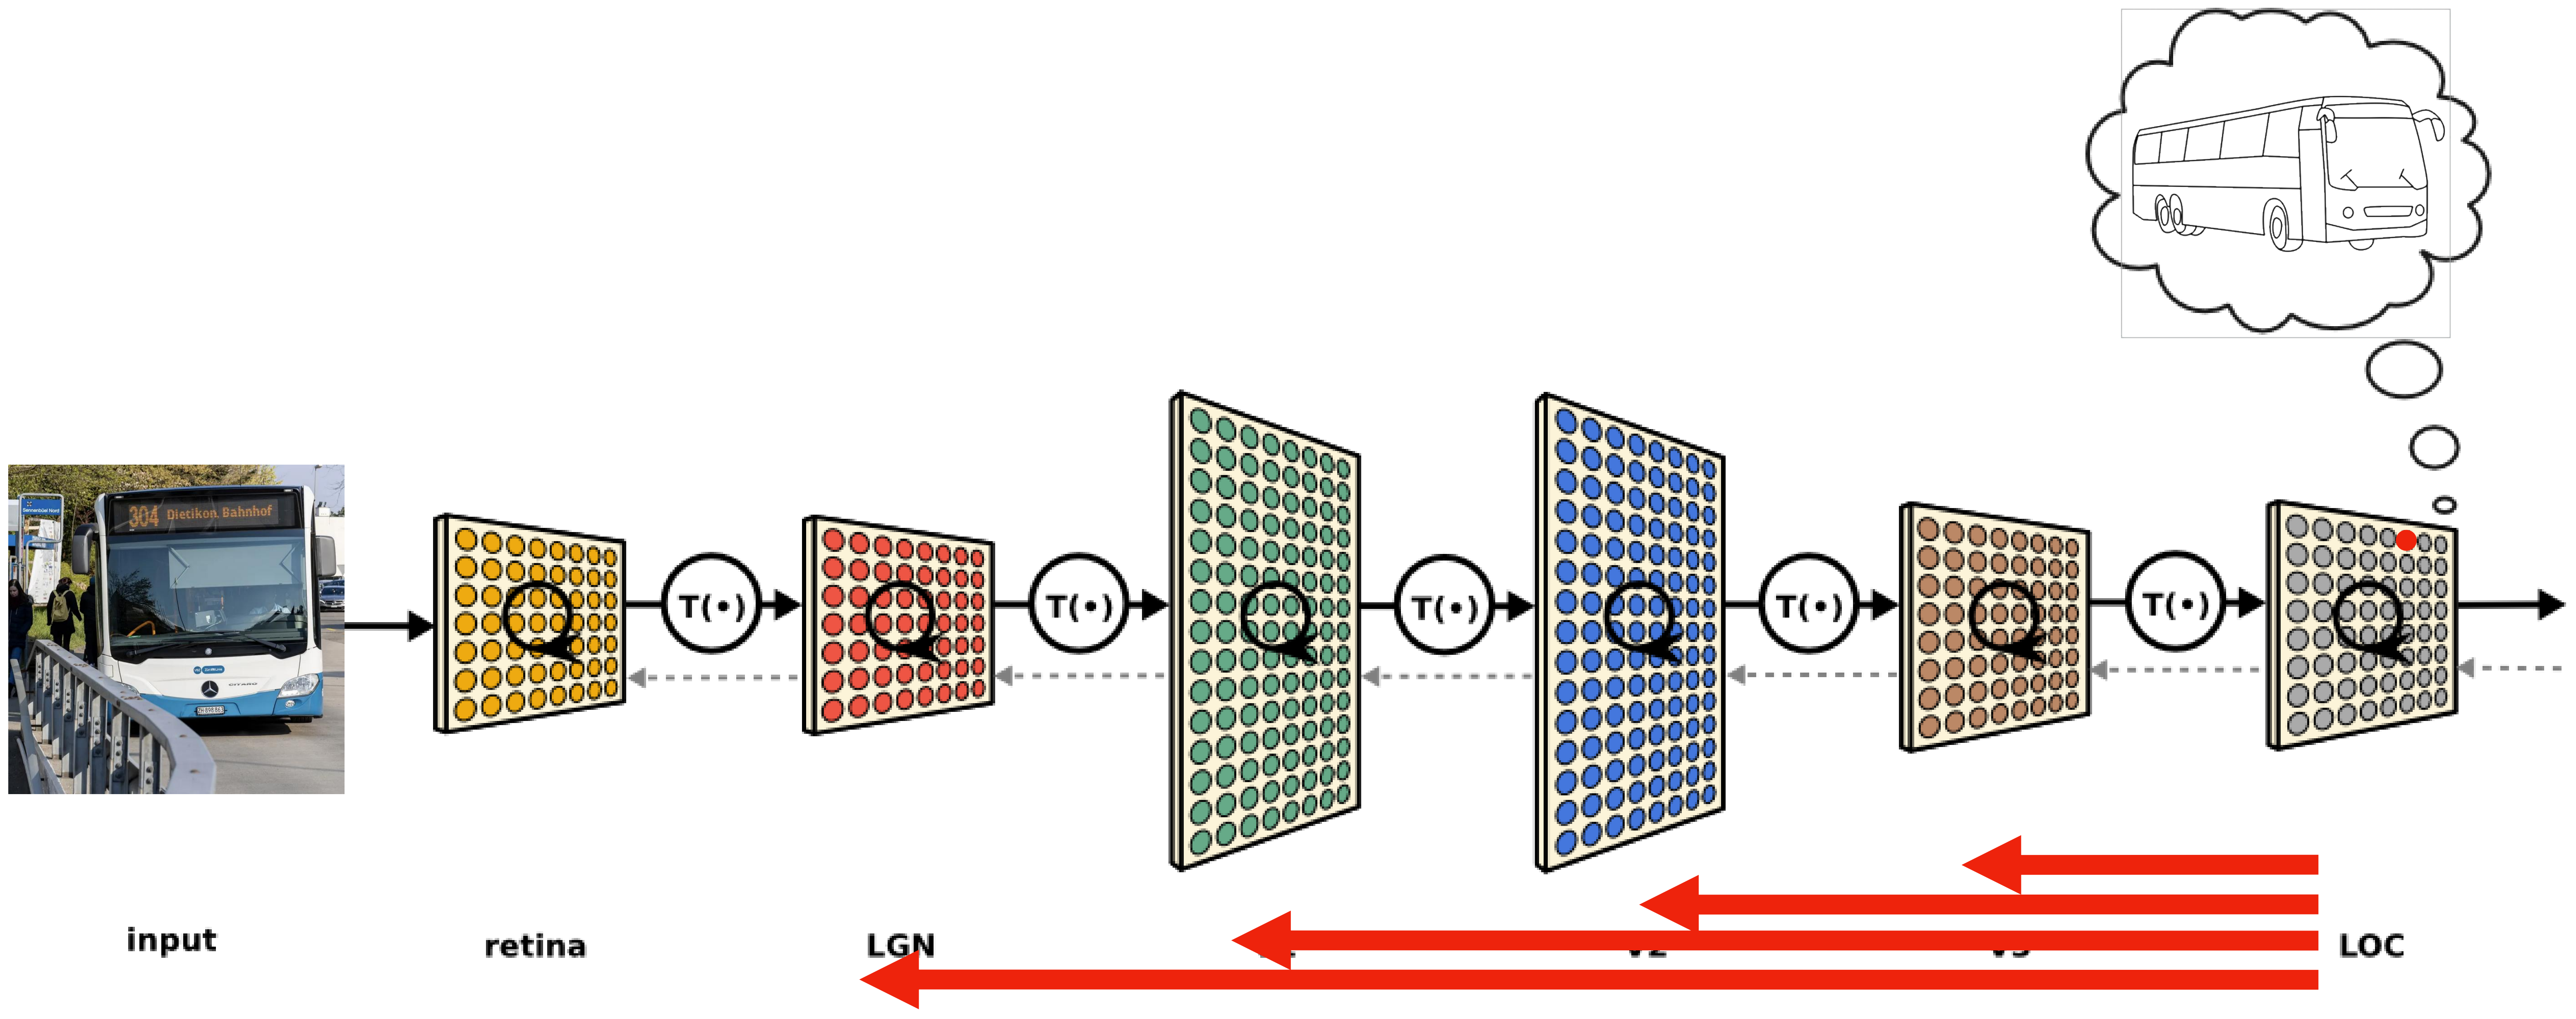

Summary Part 1:

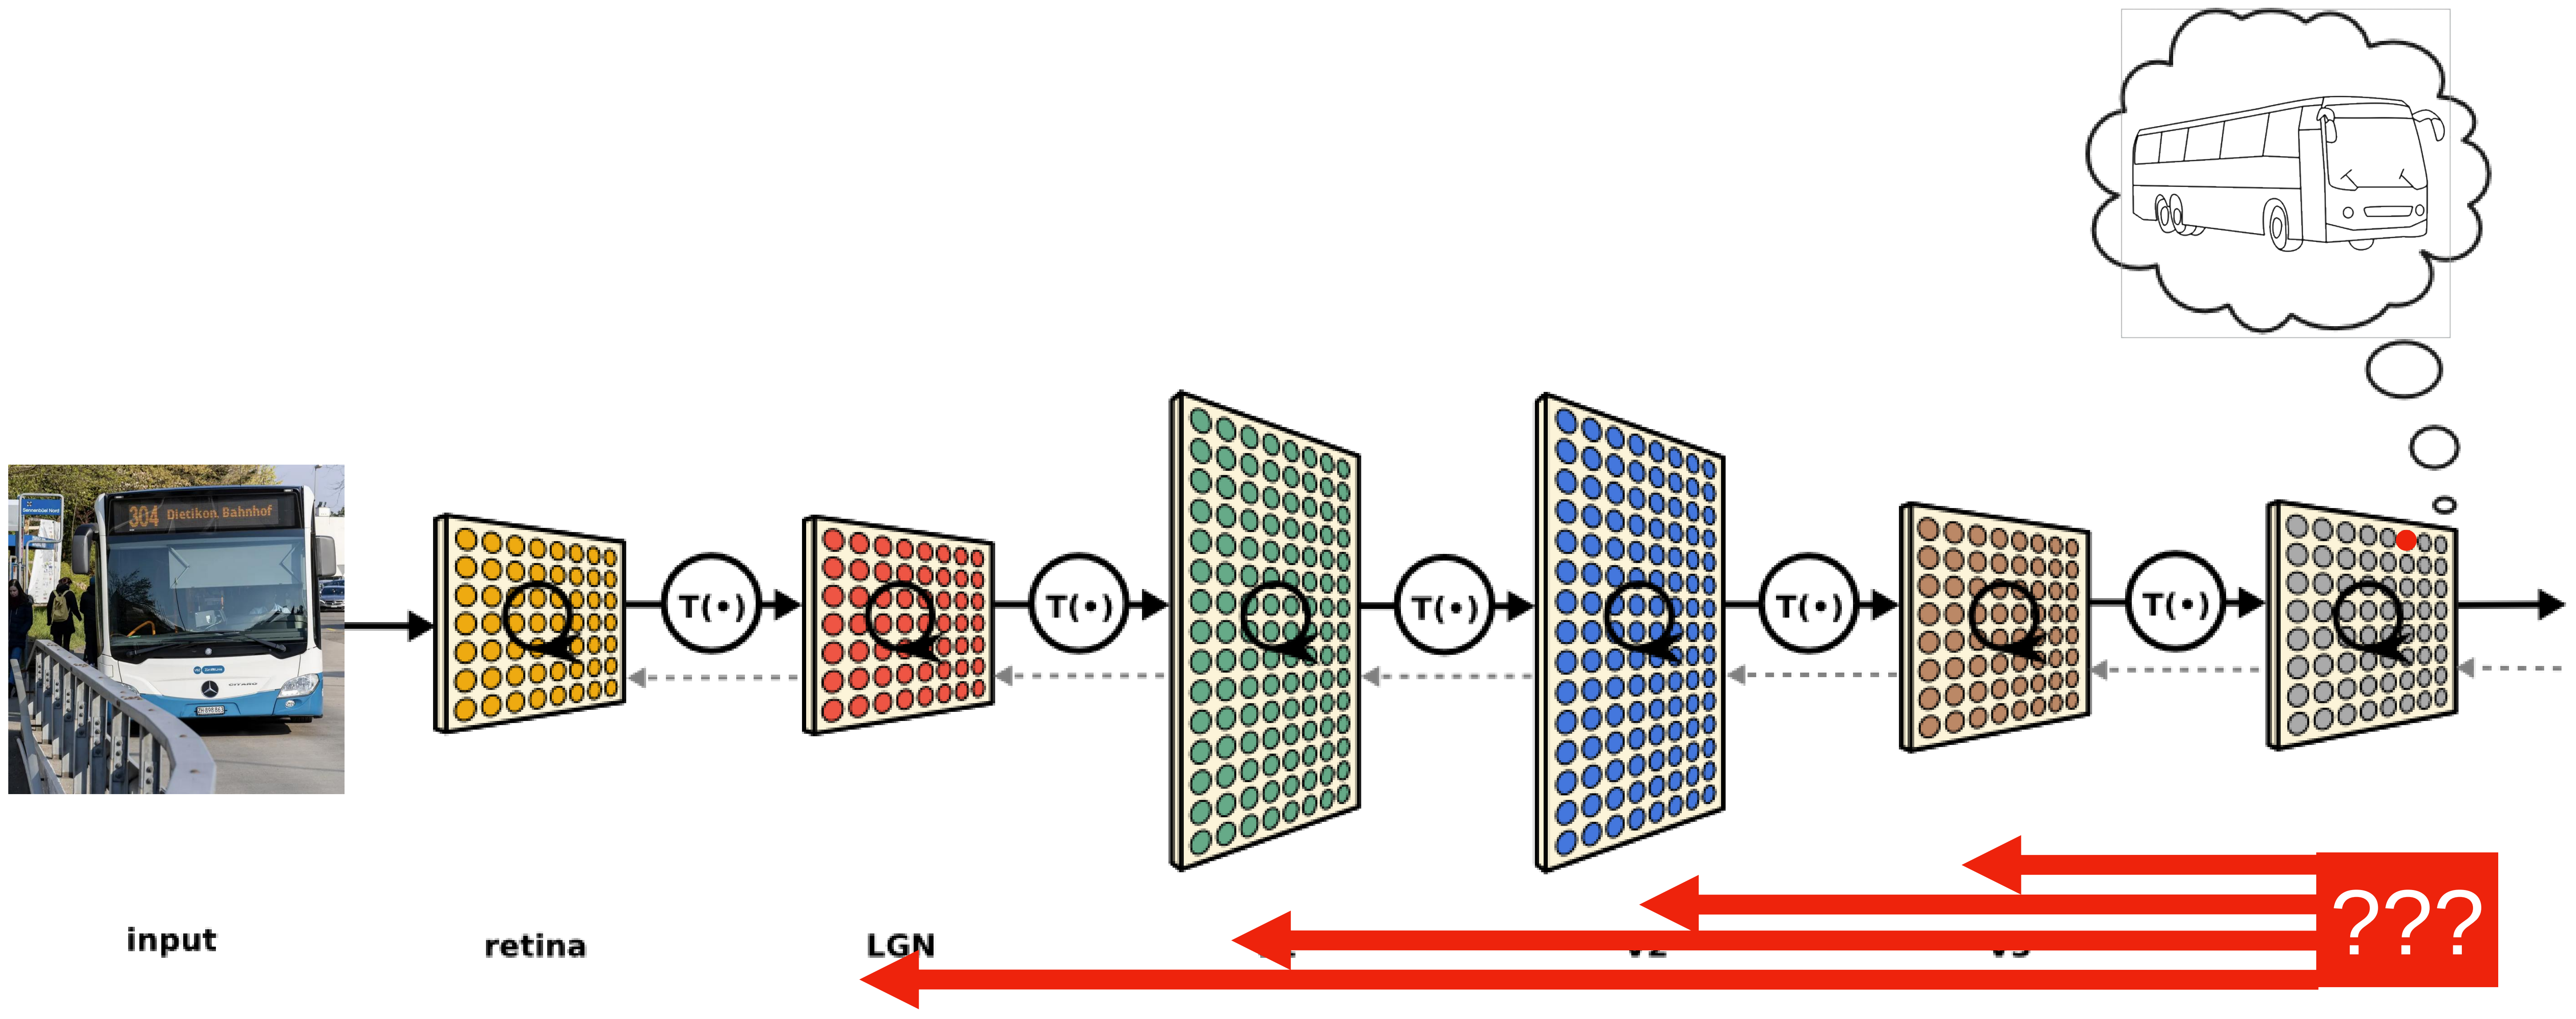

ETH zürich

institute of neuroinformatics

Summary Part 1:

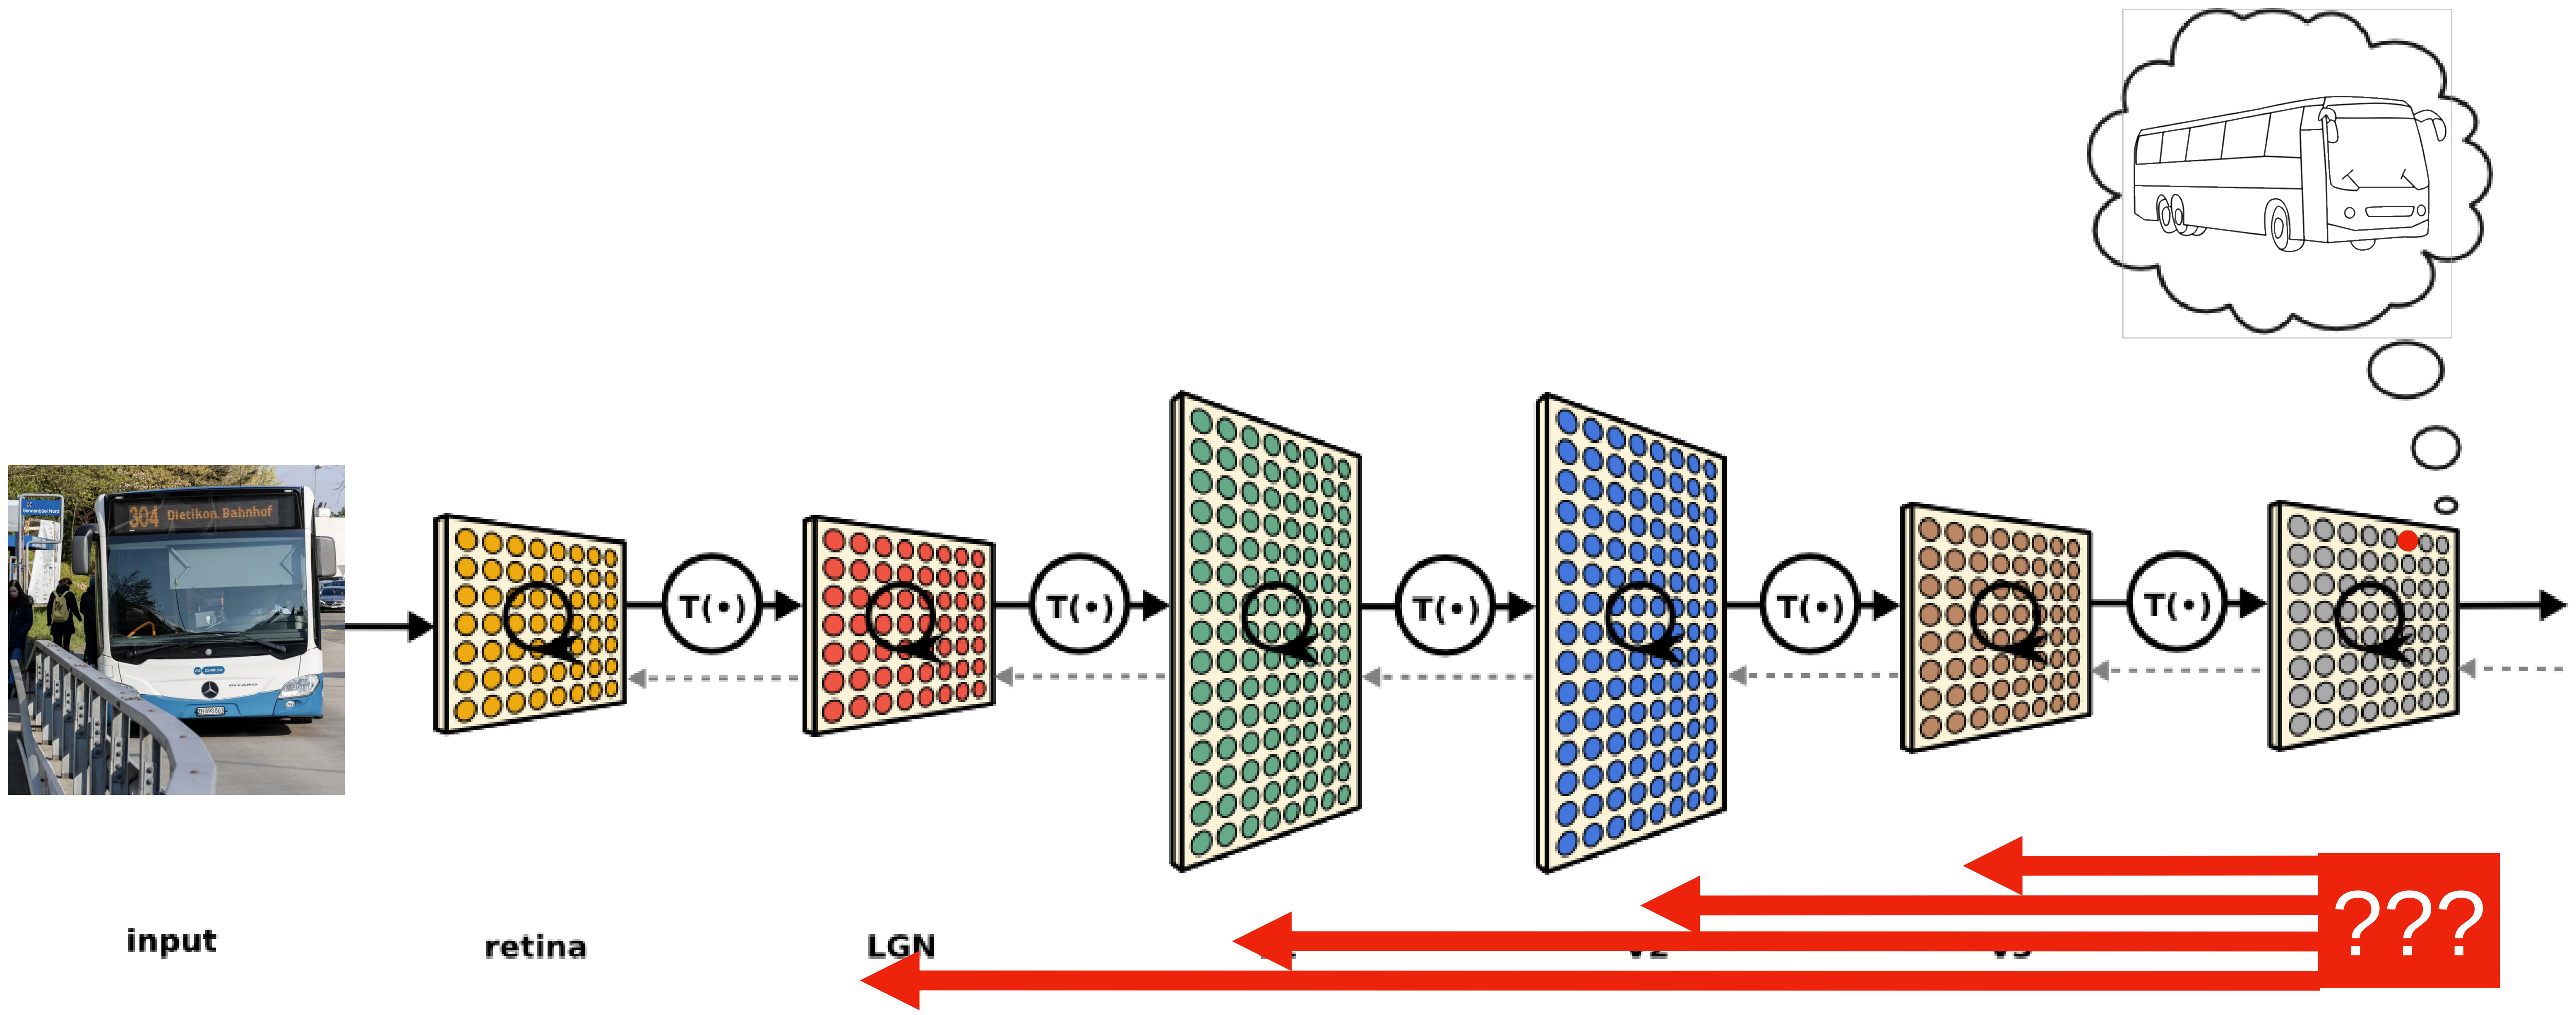

Part II: Understanding Hierarchical Neuronal Representations in Brain

AI

Neuroscience

Scientific Question:
What are the neuronal representations of the sensory input (e.g. image of bus) that allows our brain to generate goal directed actions?

Background and Motivation - Part II

Brain Visual Pathway

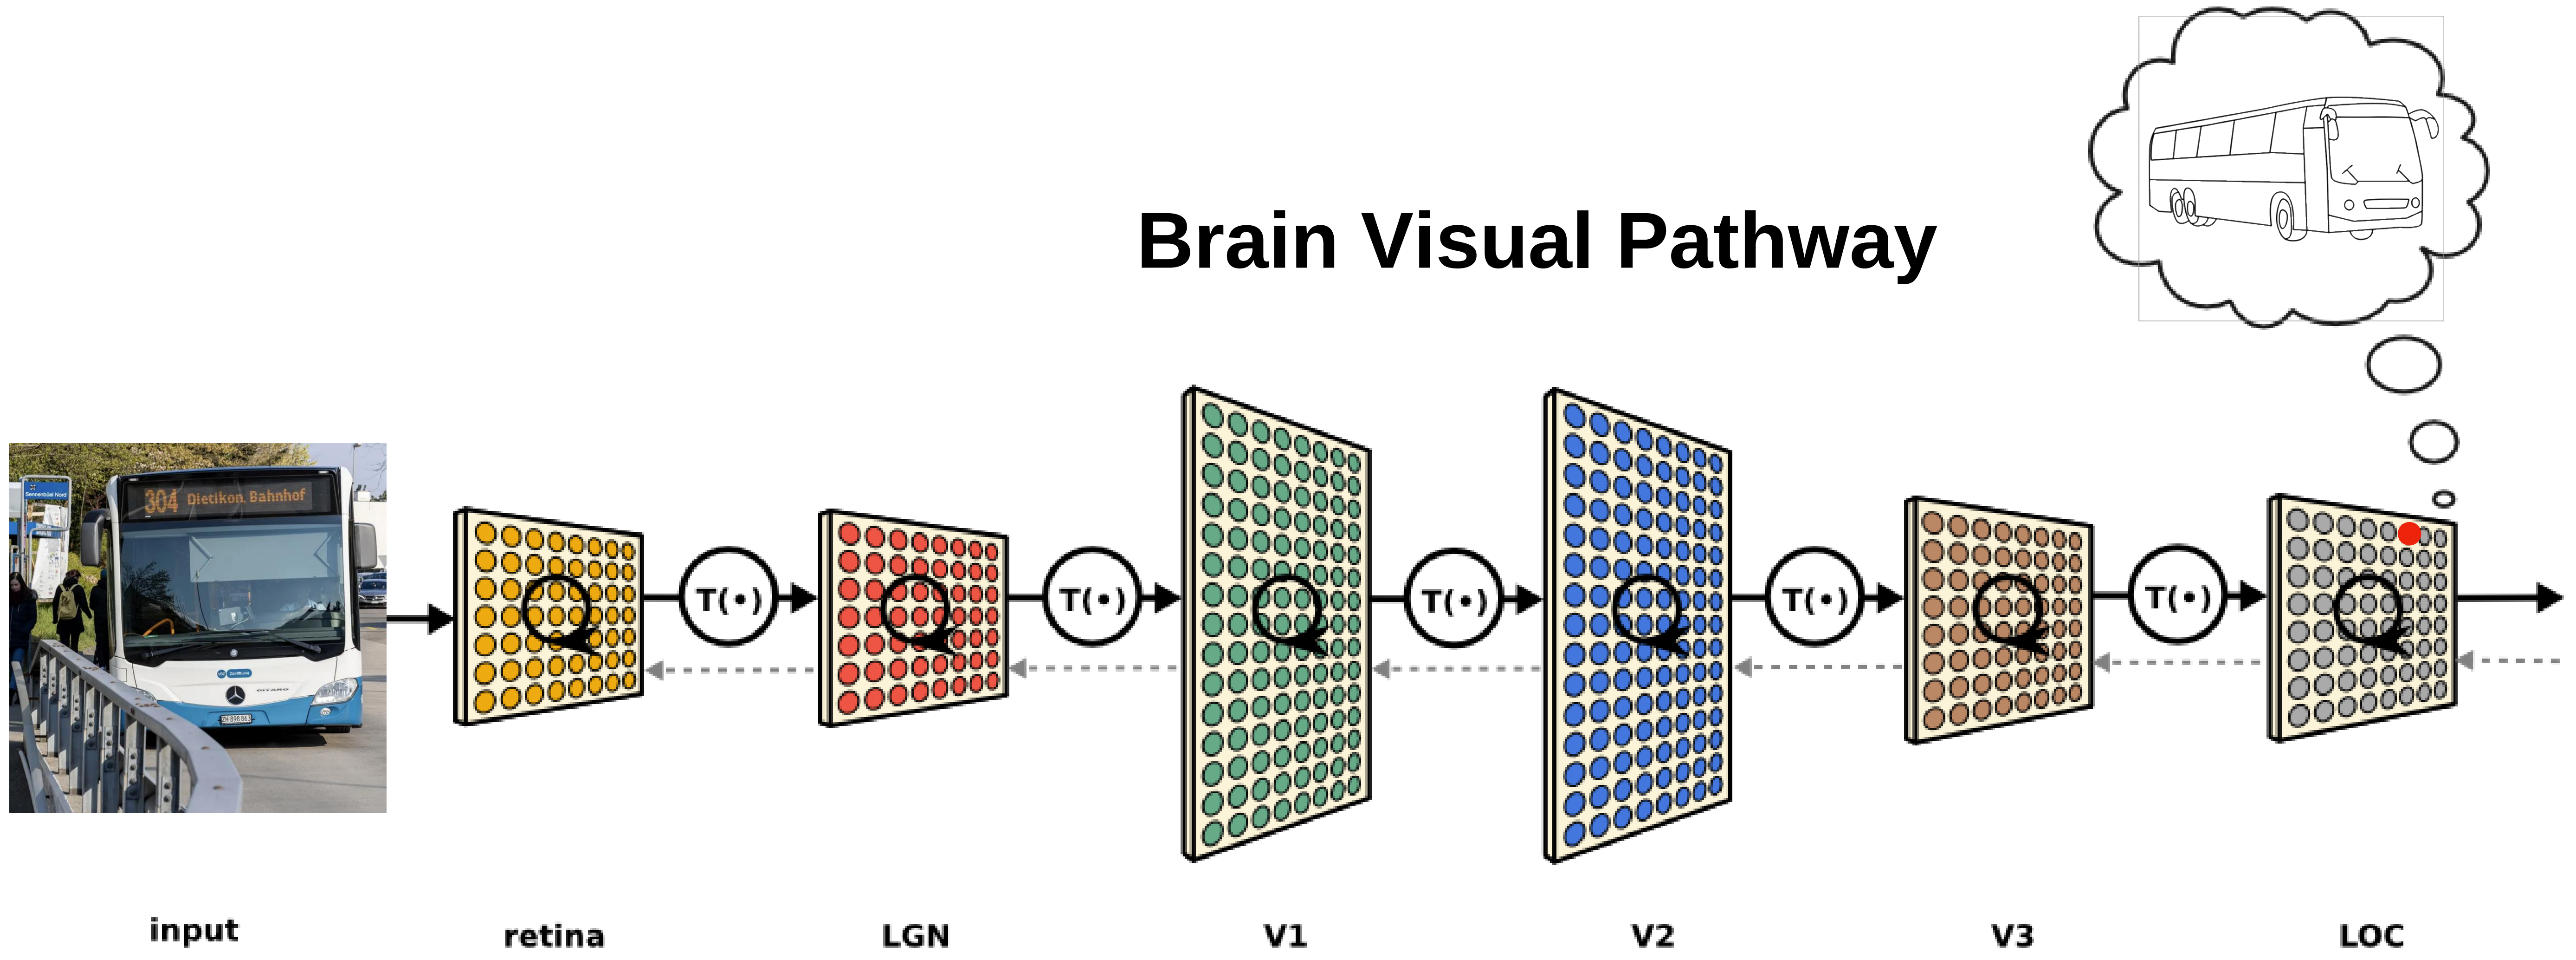

Deep Neuronal Network

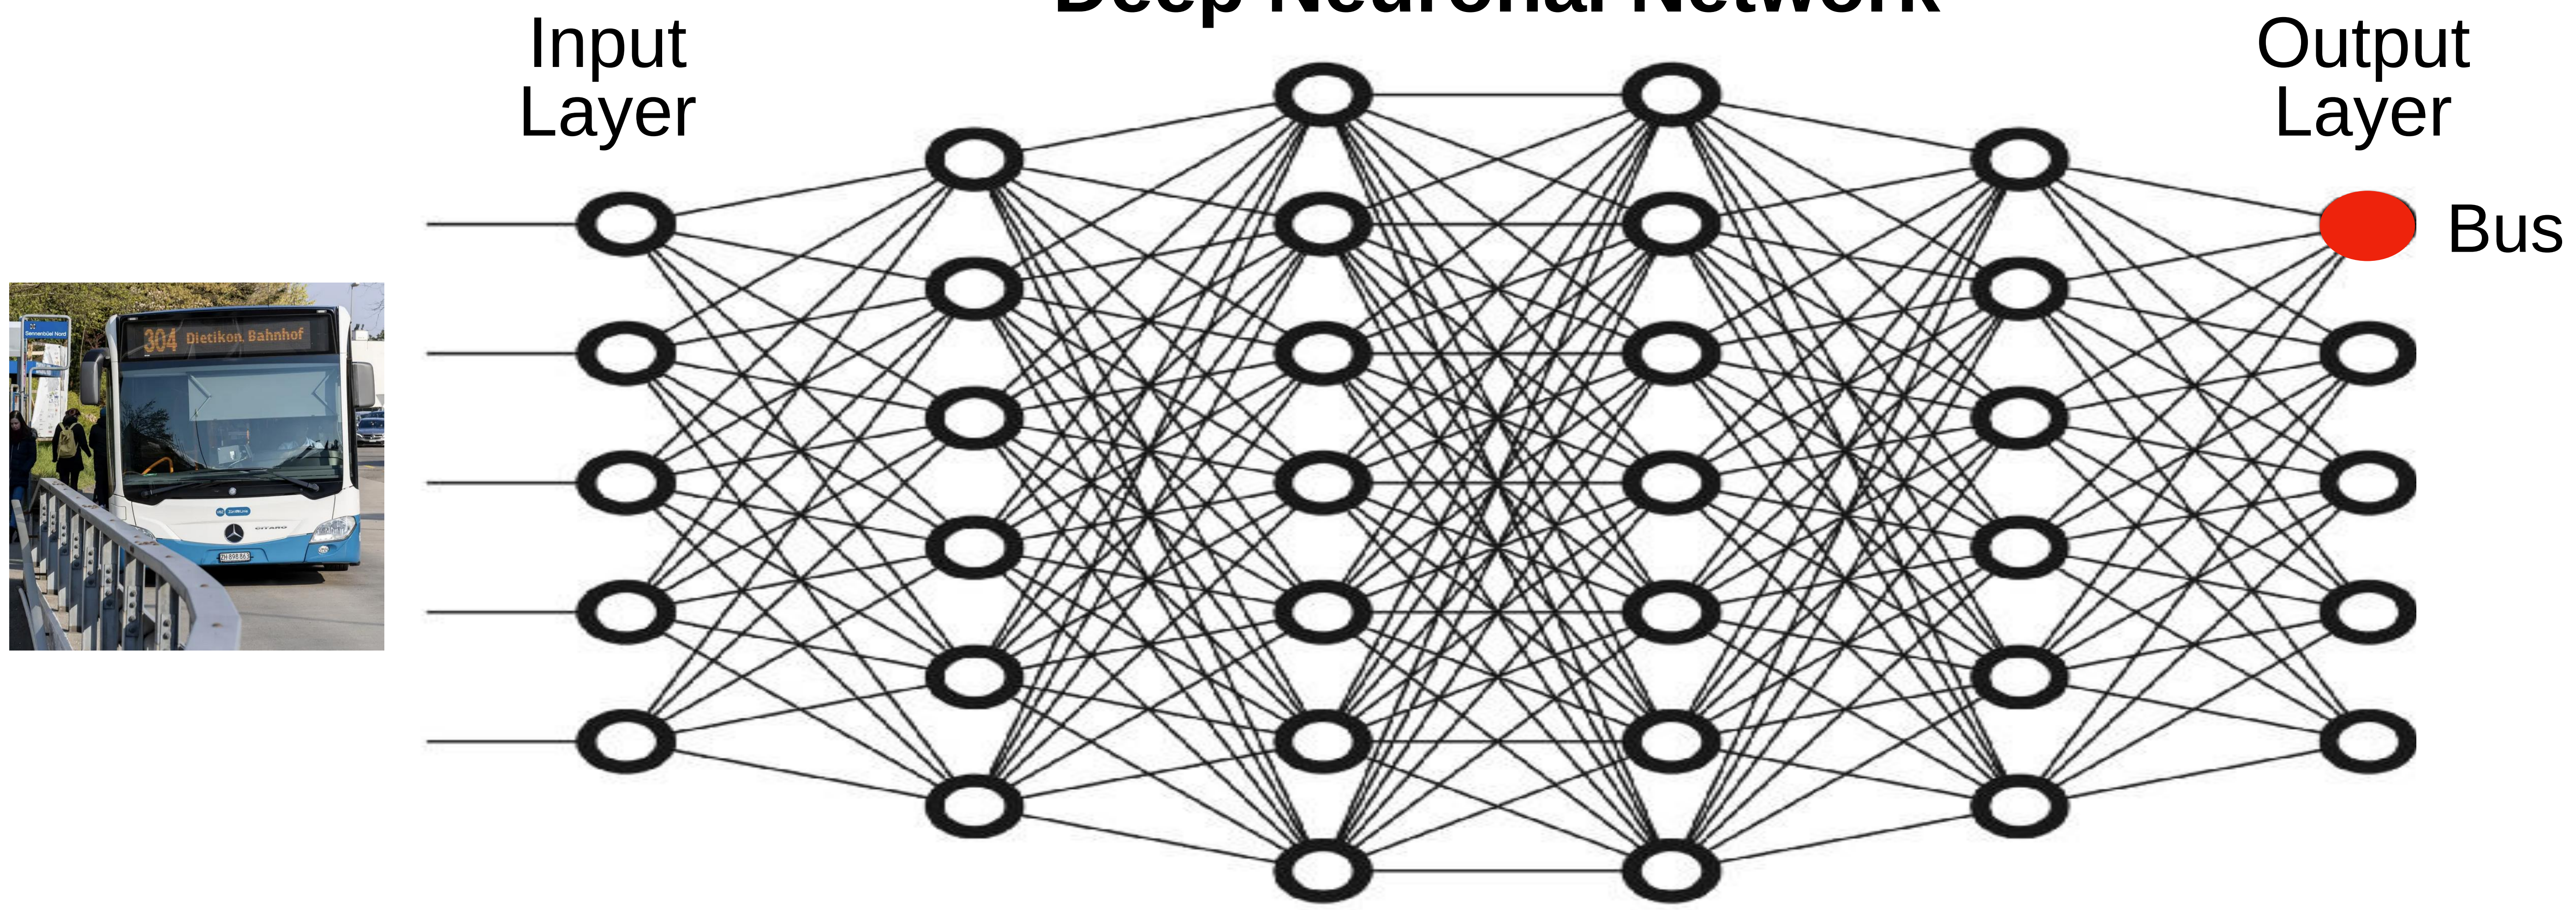

Background and Motivation - Part II

Brain Visual Pathway

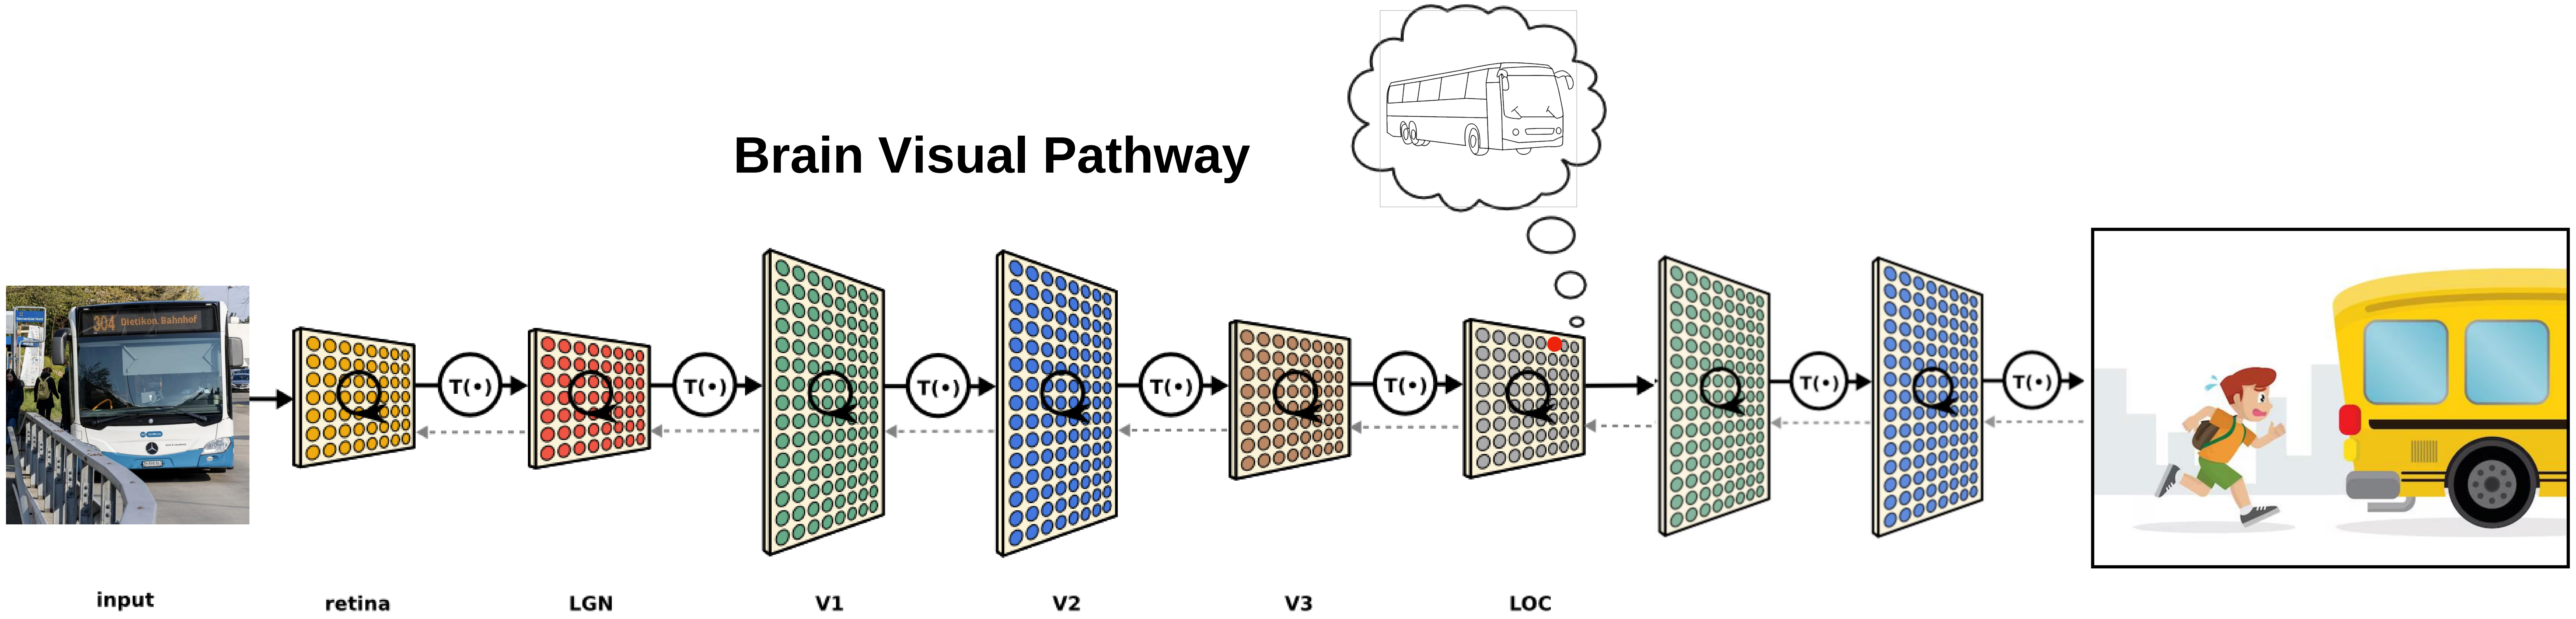

Deep Neuronal Network

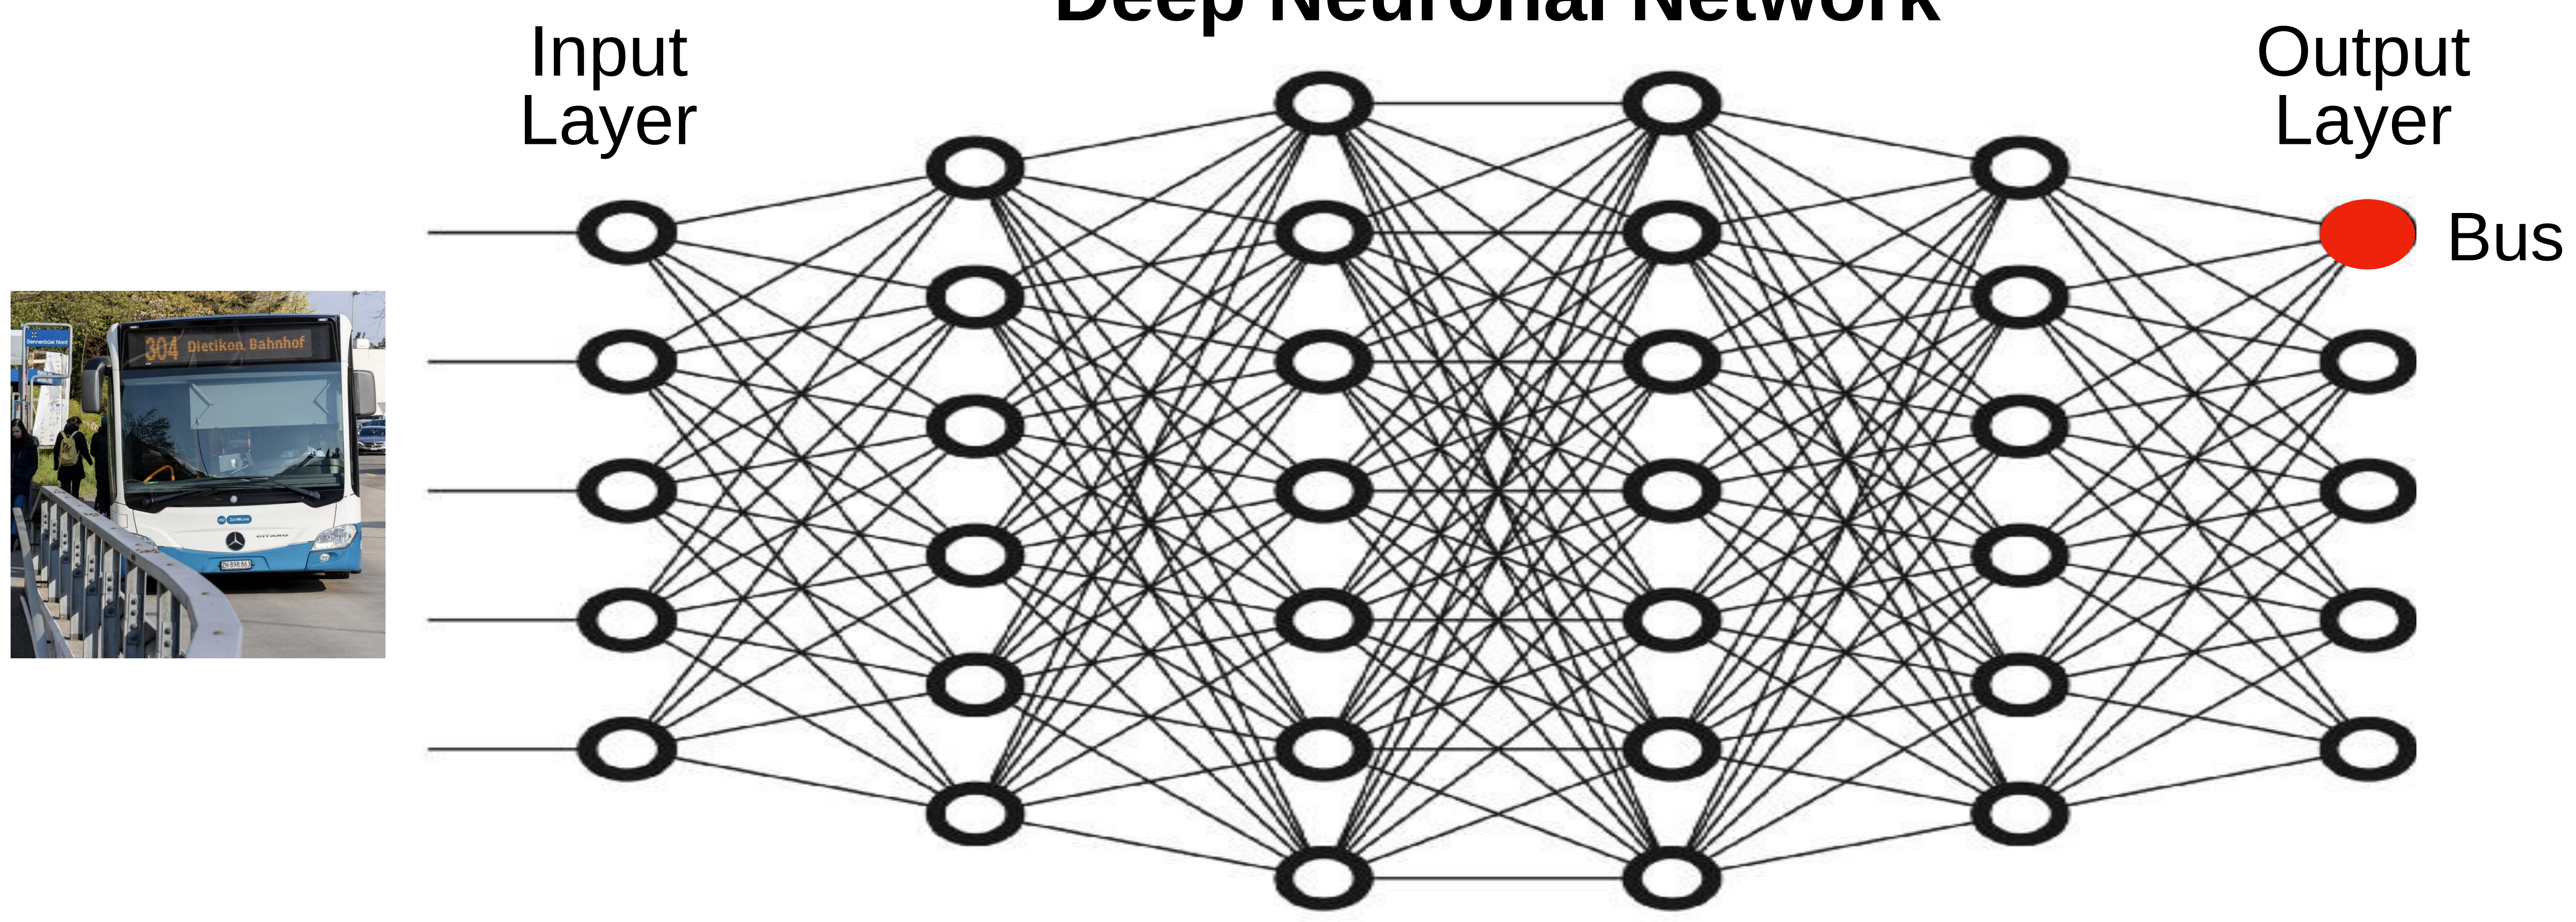

Background and Motivation - Part II

Brain Visual Pathway

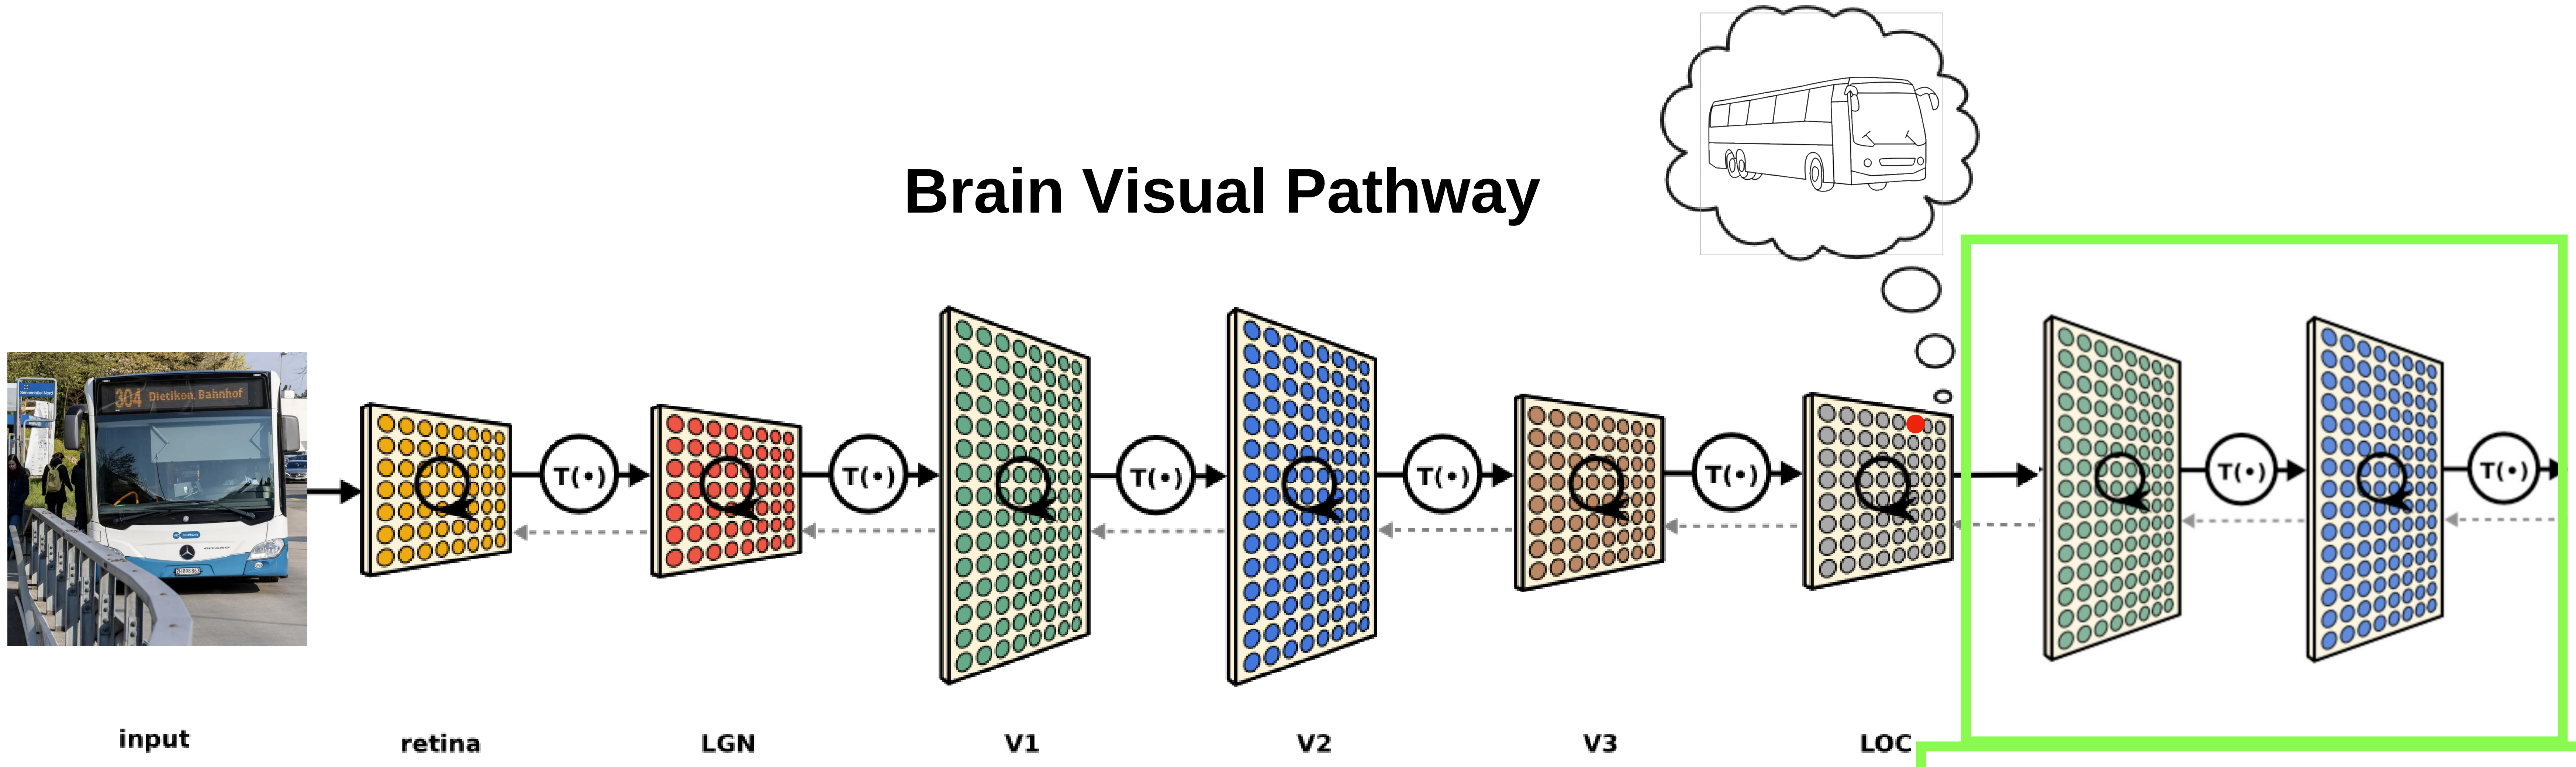

Generates Intelligent Behavior

Deep Neuronal Network

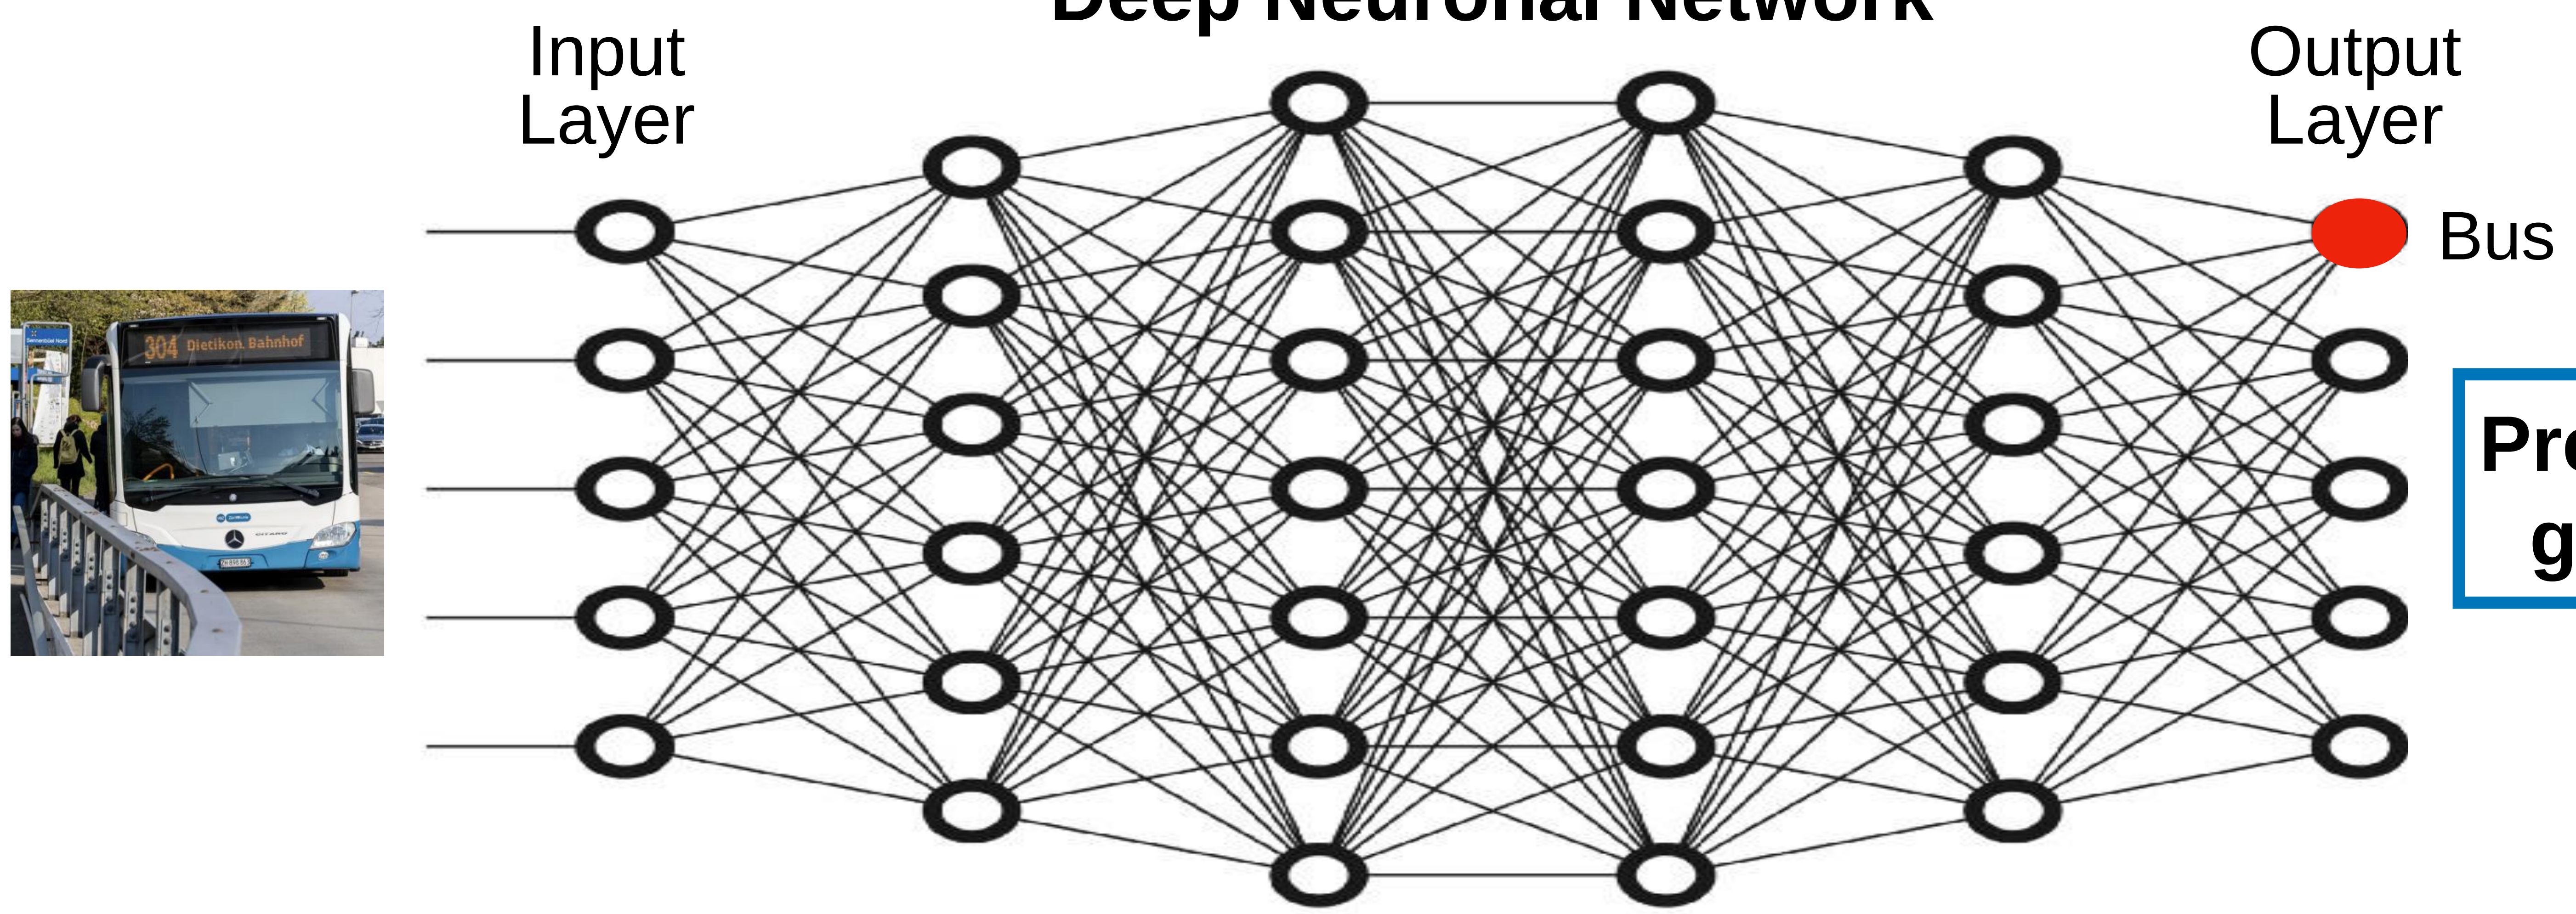

Predicts Labels / Tokens generated by Humans

Background and Motivation - Part II

Brain Visual Pathway

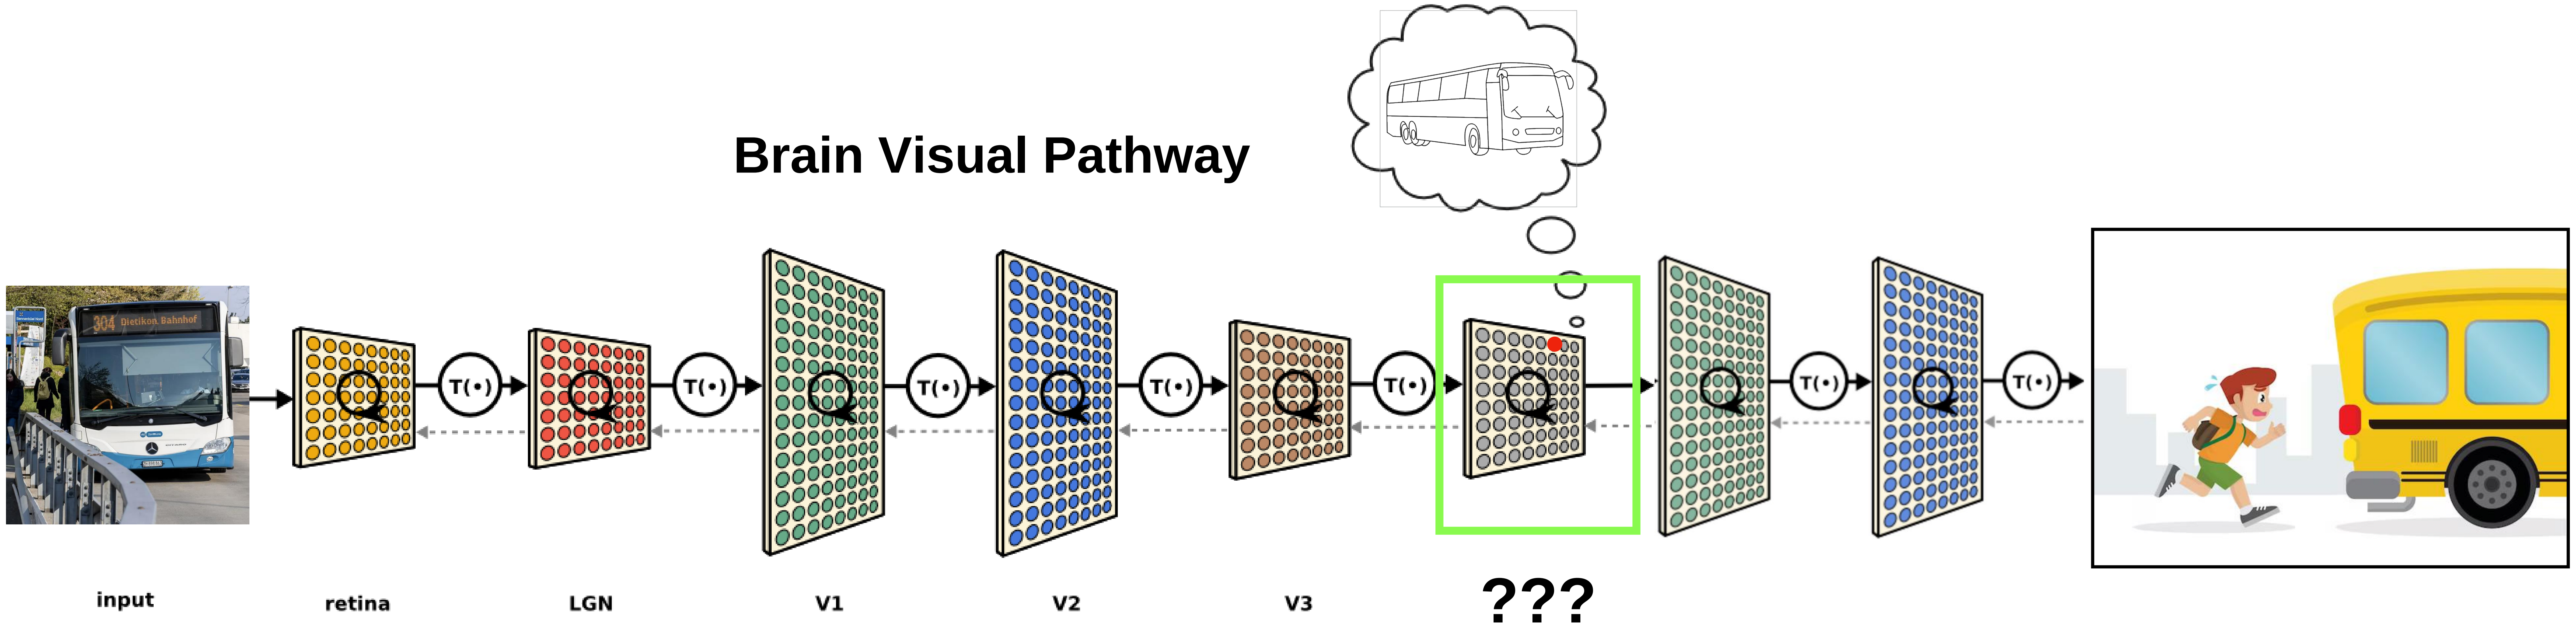

Scientific Question:
What are the neuronal representations of the sensory input (e.g. image of bus) that allows our brain to generate goal directed actions?

ETH zürich

Background and Motivation - Part II

institute of neuroinformatics

Brain Visual Pathway

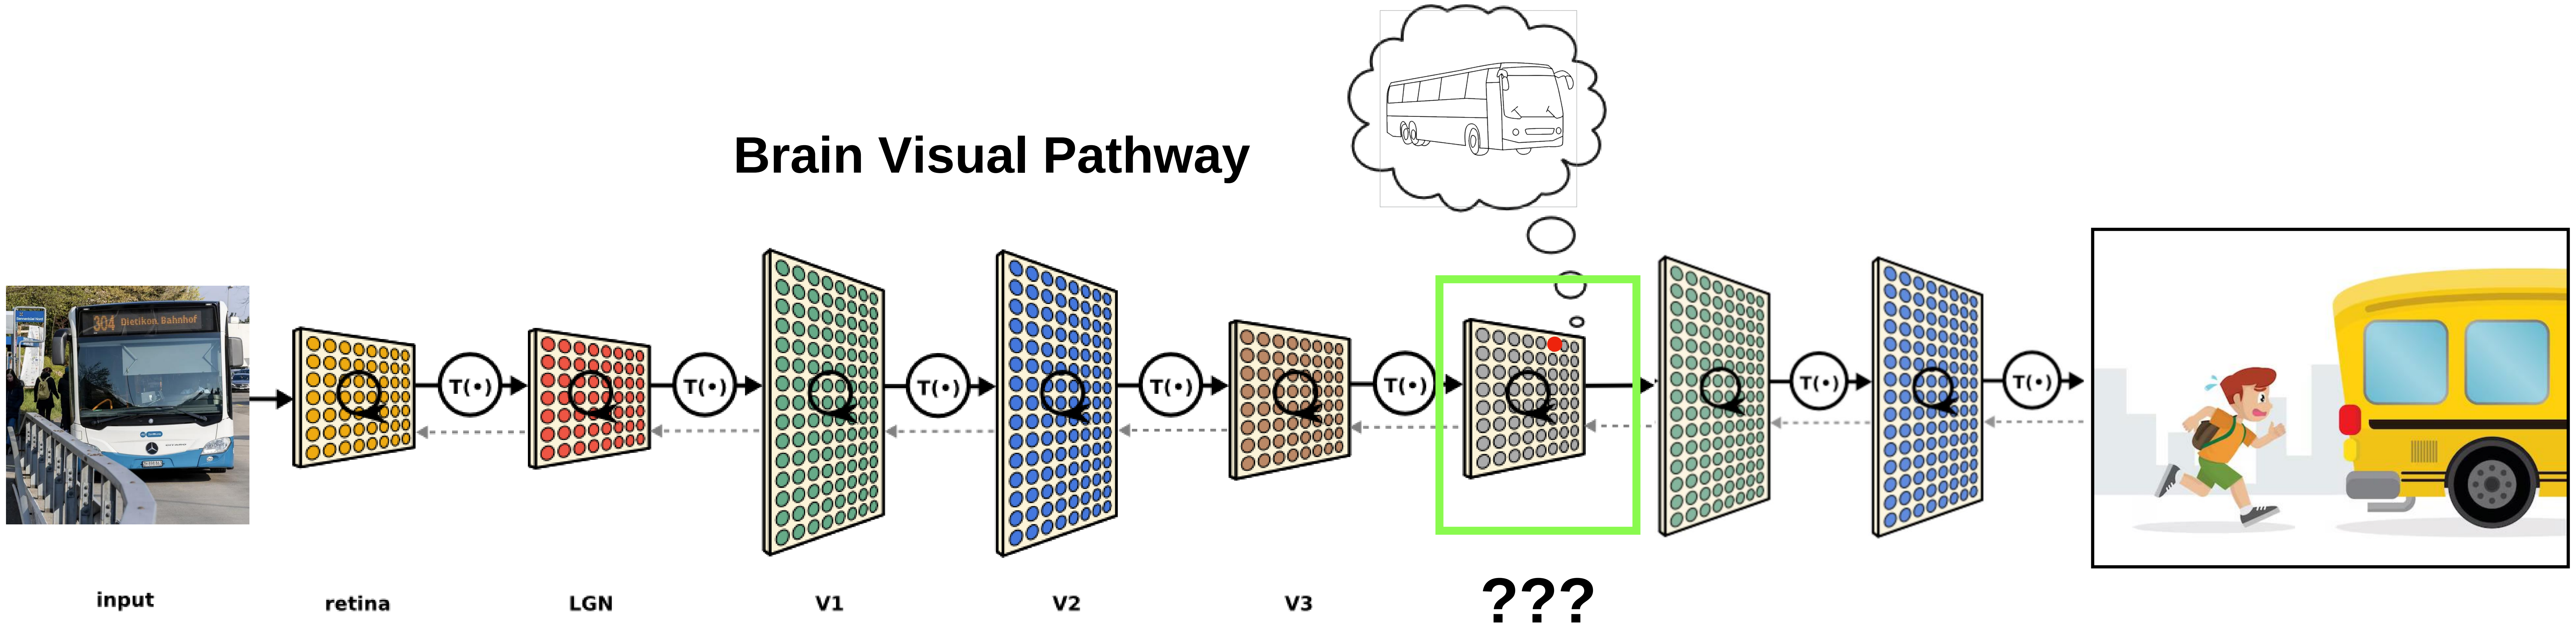

Prof. Dr. Jean Piaget
Swiss psychologist
Neuchatel, 1896-1980

UNIVERSITÉ DE GENÈVE

Université de Paris

The Concept of Affordance in Psychology.

Affordance alludes to the qualities of an object or situation that define its possible use or make clear how it can or should be used.

Affordance adheres to the idea that perception and action are inseparable (Principles of Genetic Epistemology, Jean Piaget).

ETH zürich

institute of neuroinformatics

Investigating Abstract Stimulus Representations in mPFC

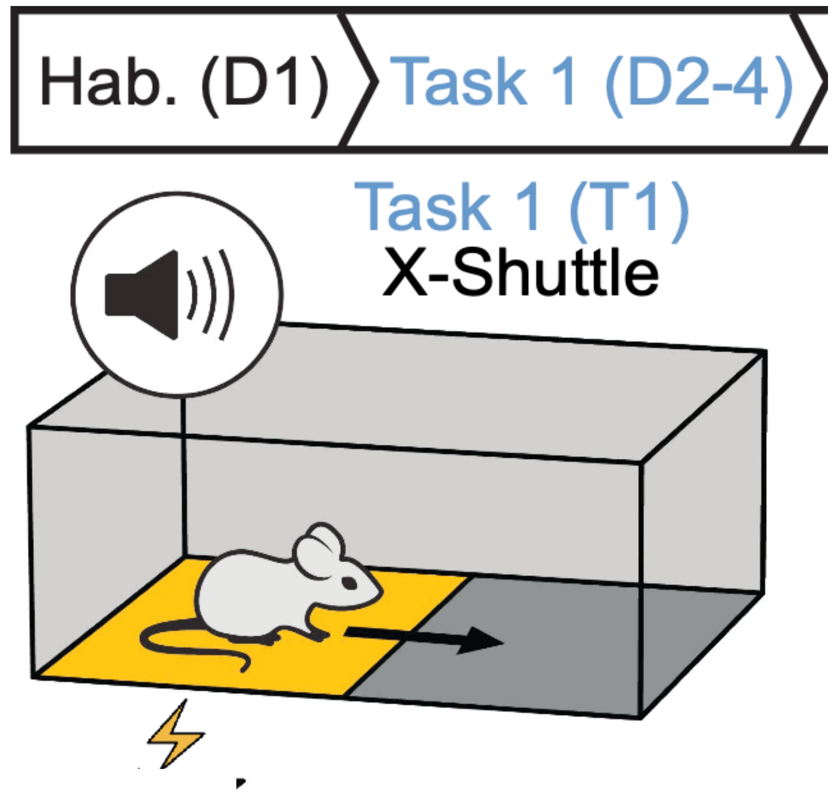

Task 1 (T1)
X-Shuttle

Task 1 (T1)
X-Shuttle

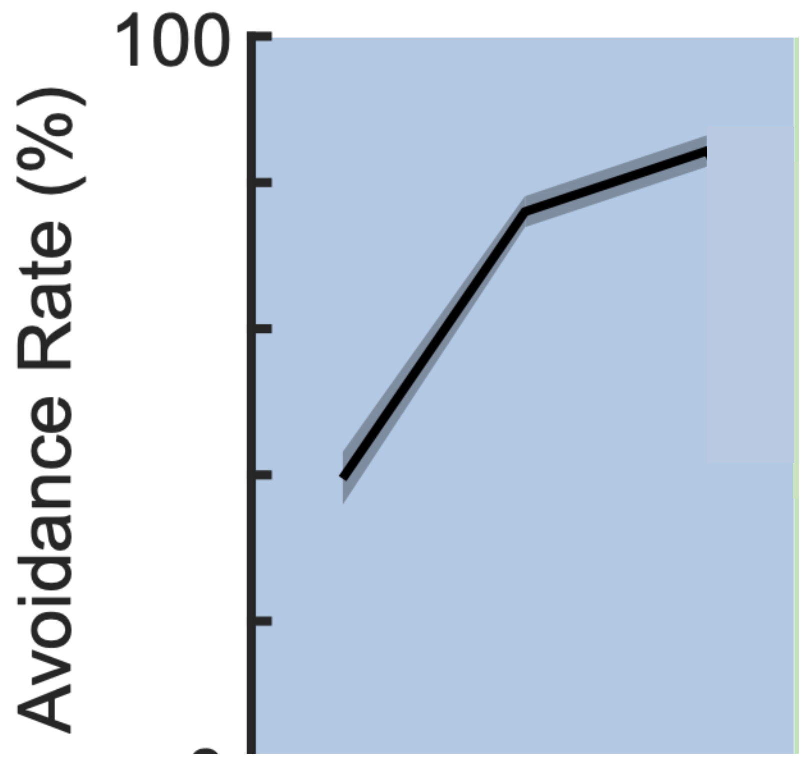

ETH zürich

institute of neuroinformatics

Investigating Abstract Stimulus Representations in mPFC

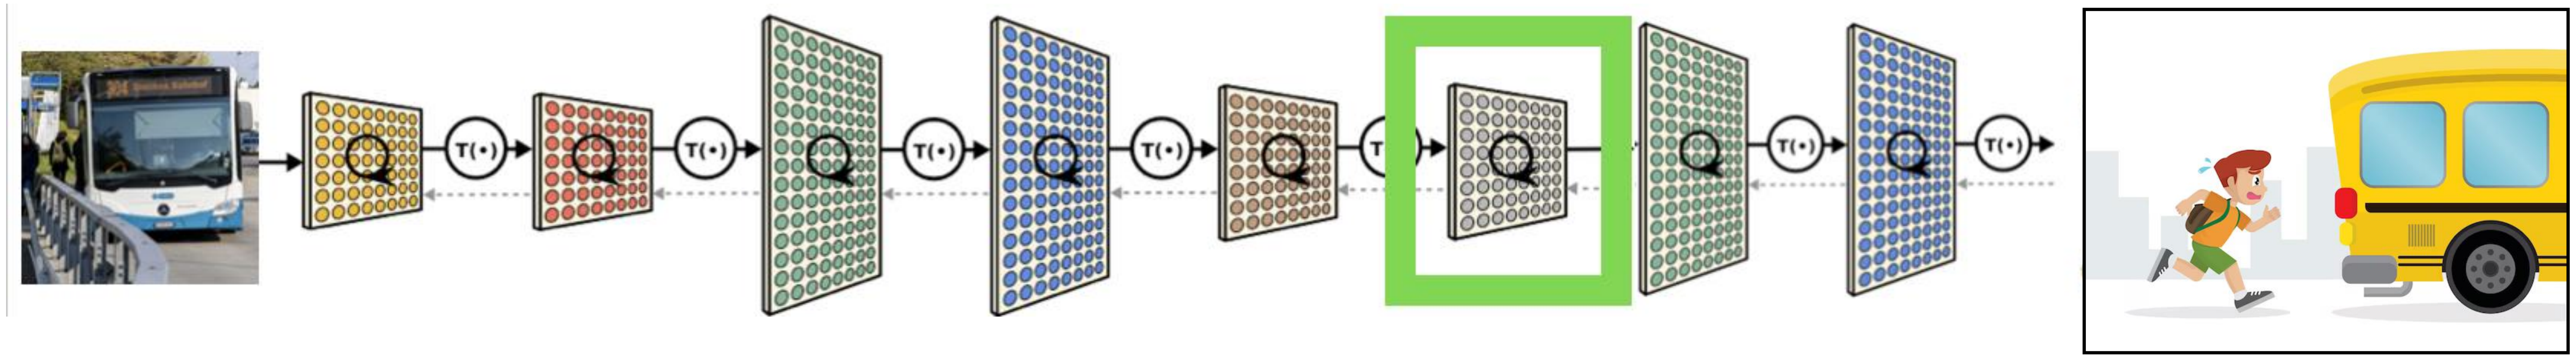

Medial Prefrontal Cortex (mPFC)

Stimulus

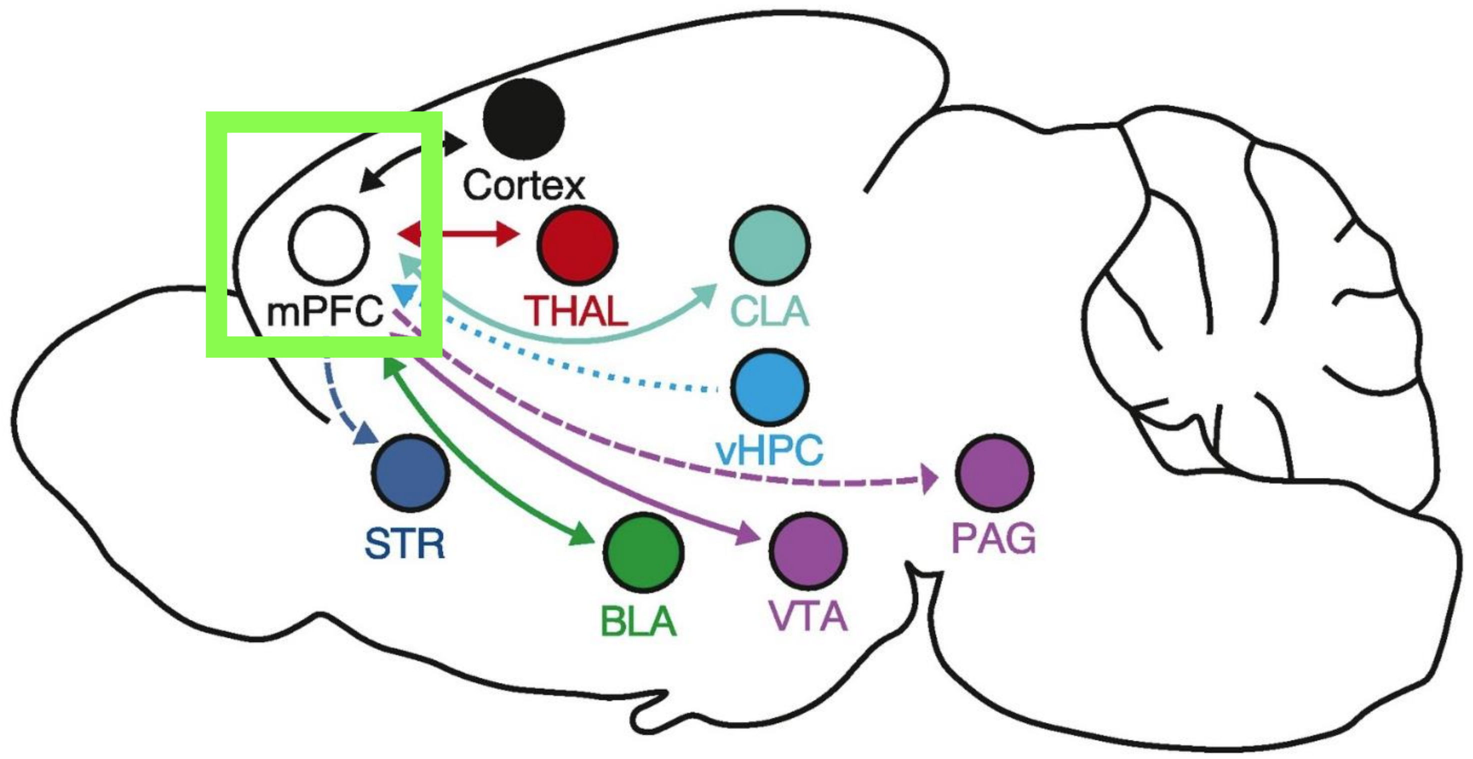

Behavior

Investigating Abstract Stimulus Representations in mPFC

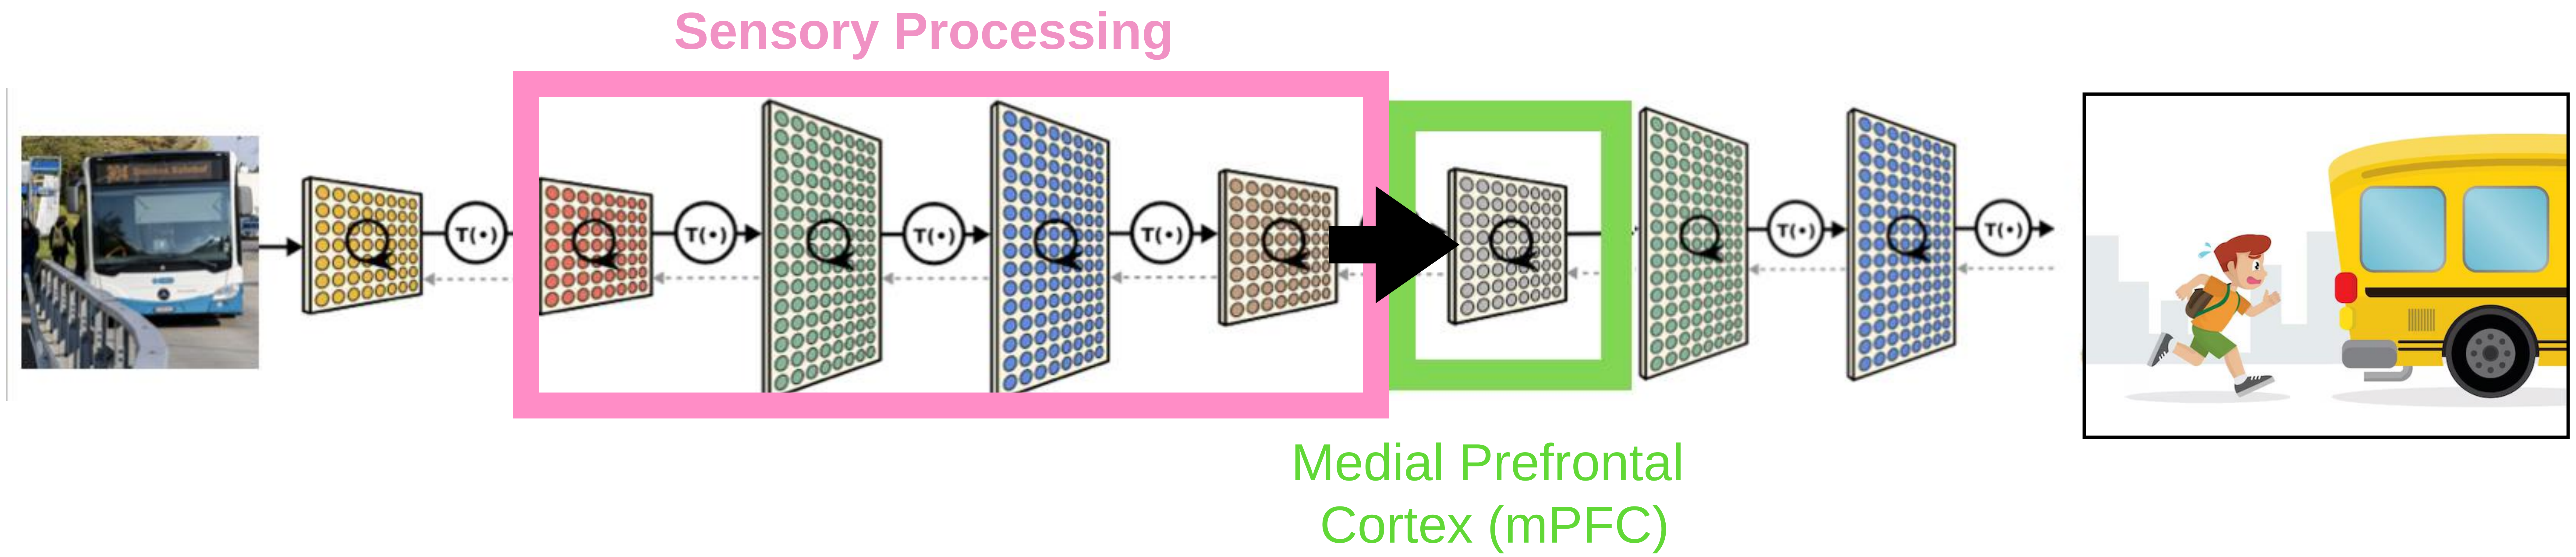

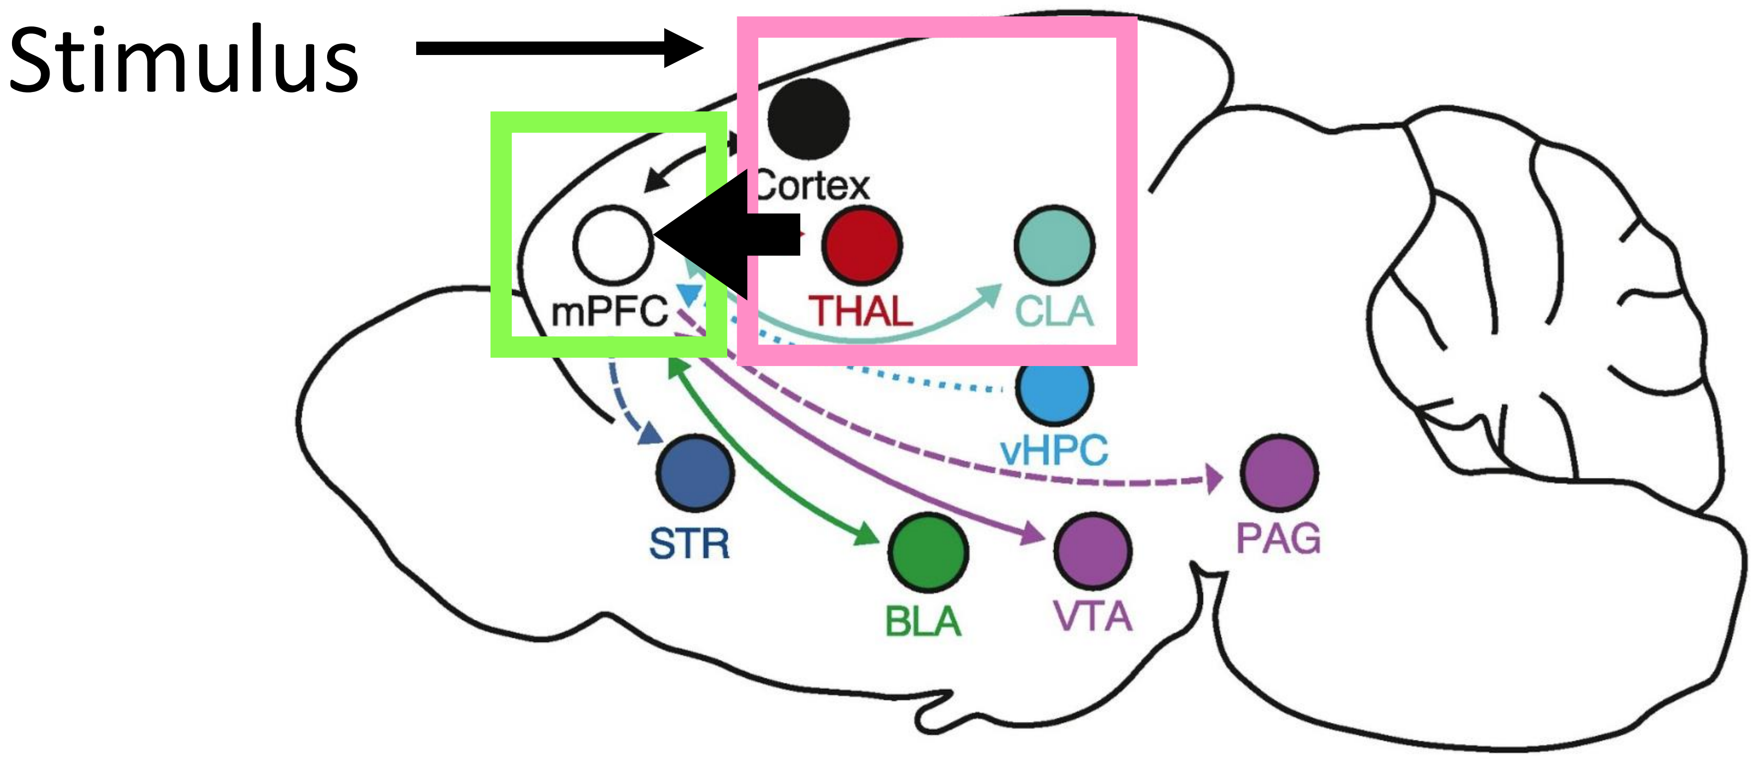

Investigating Abstract Stimulus Representations in mPFC

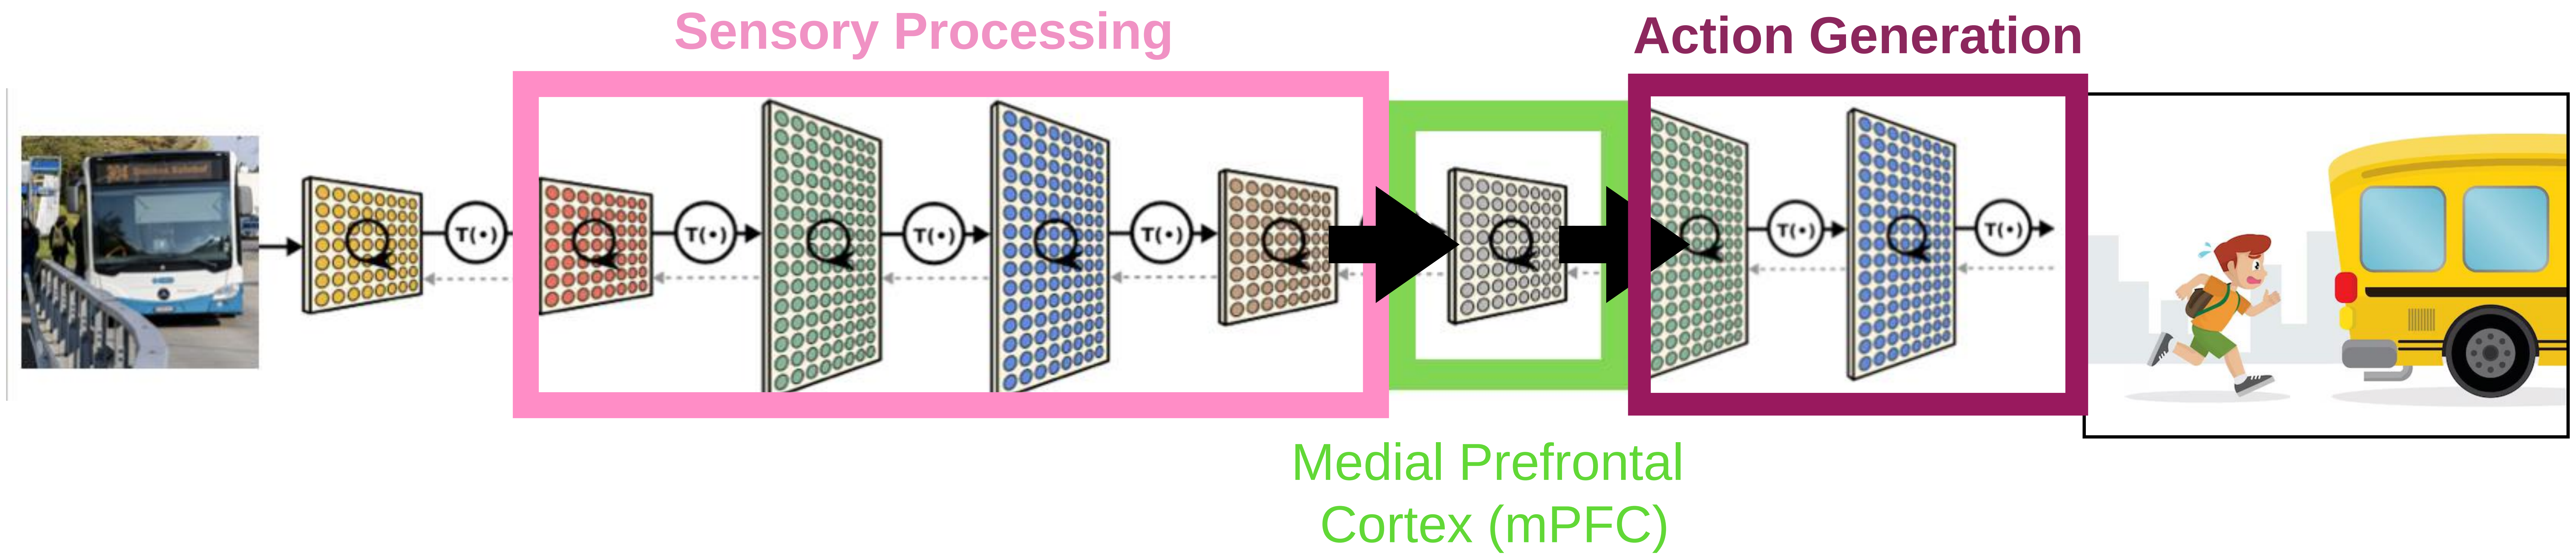

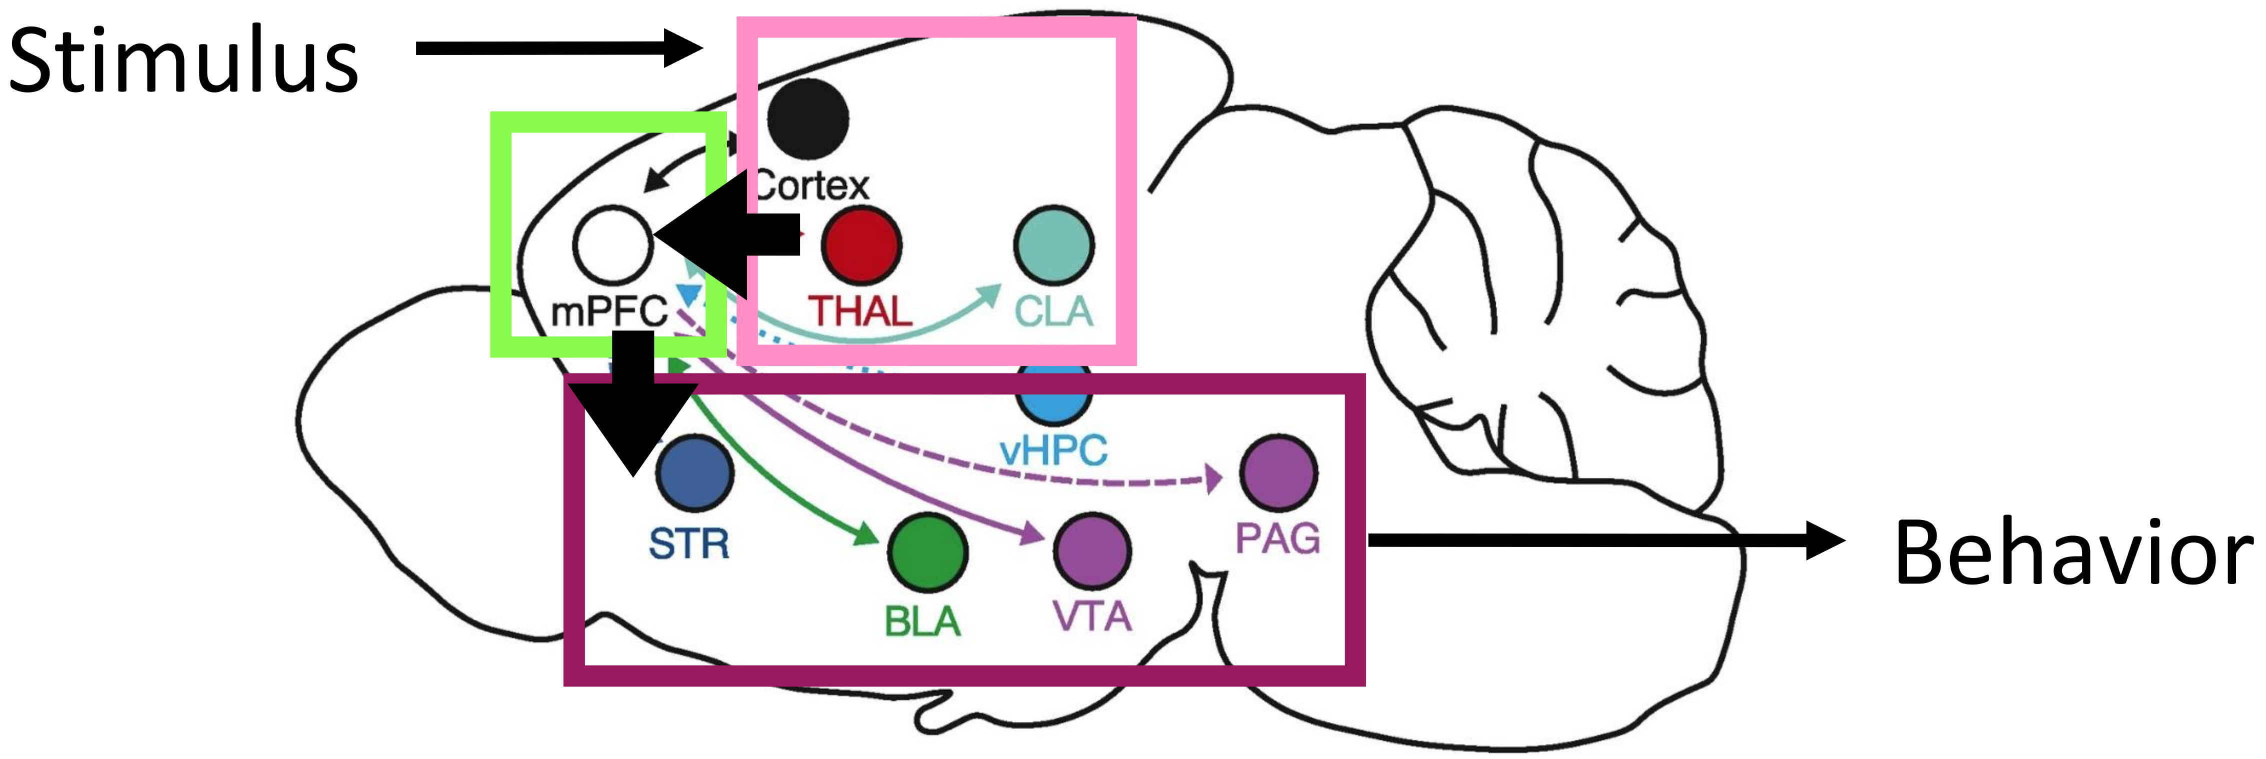

Investigating Abstract Stimulus Representations in mPFC

Sensory Processing

Action Generation

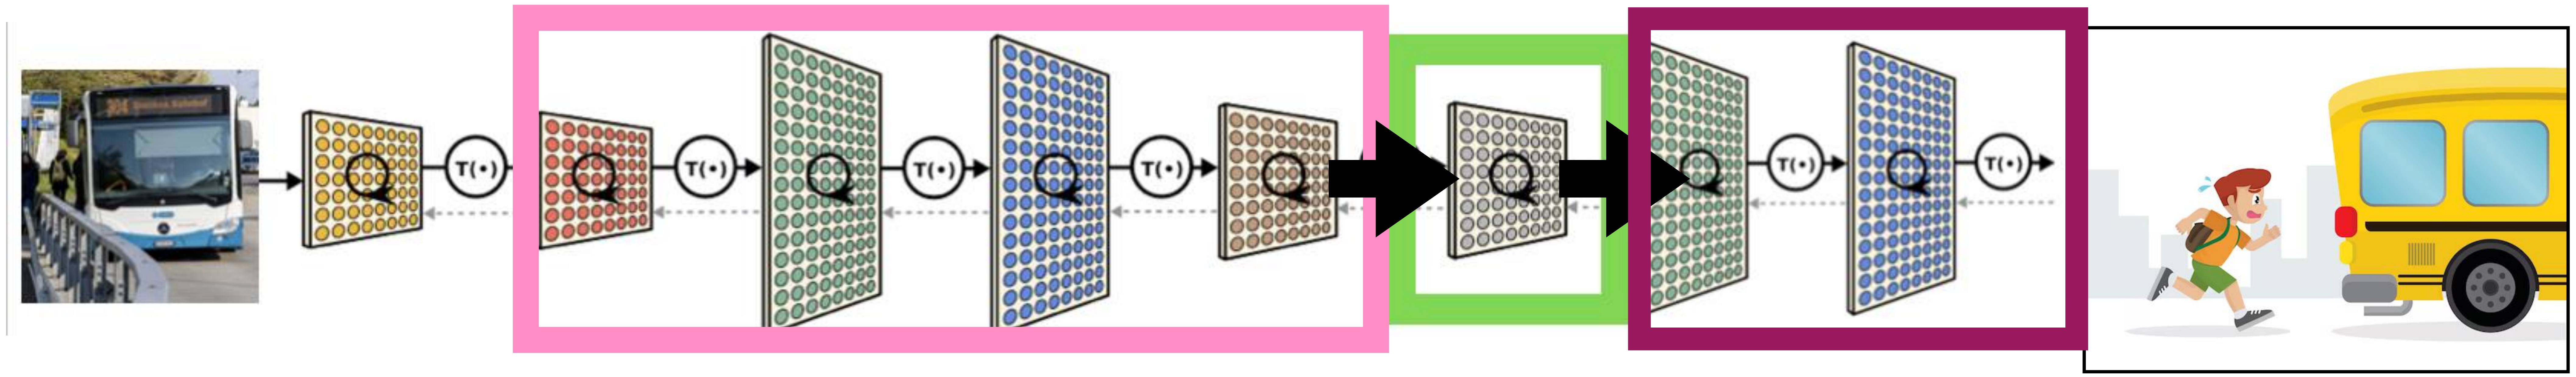

Medial Prefrontal Cortex (mPFC)

‘Learning induces representations of behaviorally relevant stimuli’

Burgos-Robles et al. 2009
Le Merre et al. 2018
Otis et al. 2017

Stimulus

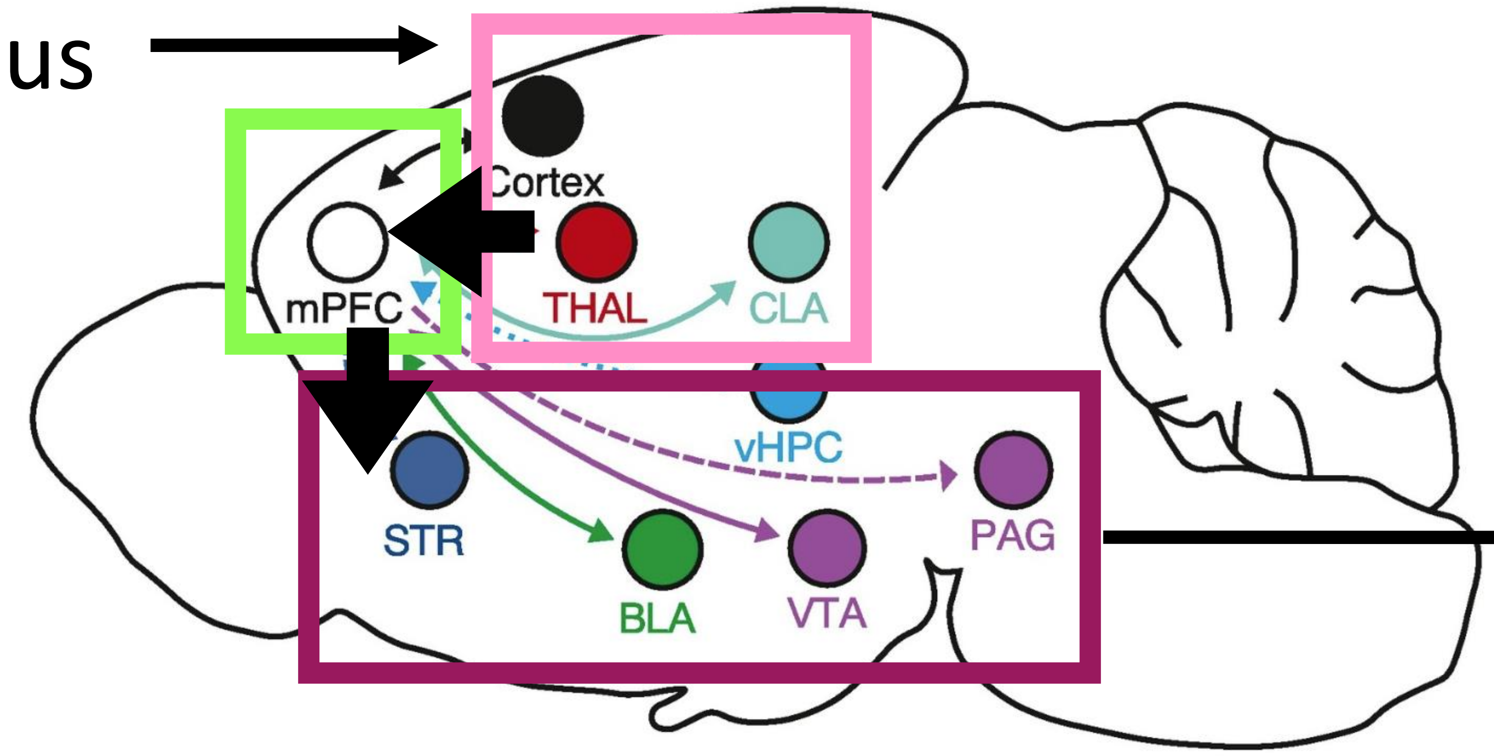

Behavior

ETH zürich

institute of neuroinformatics

Investigating Abstract Stimulus Representations in mPFC

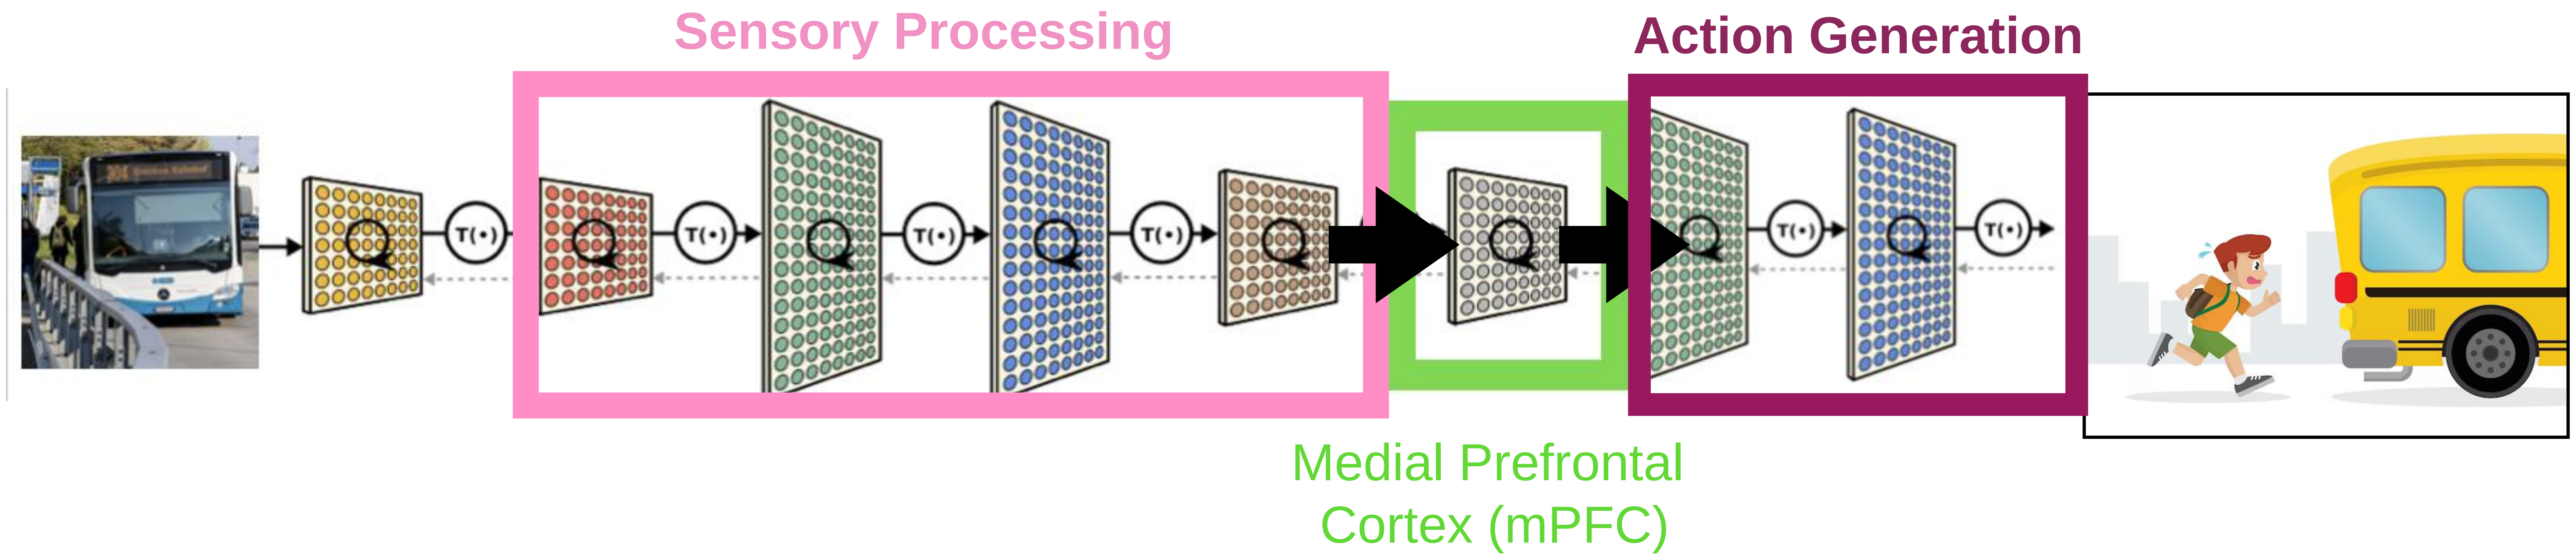

‘Learning induces representations of behaviorally relevant stimuli’

Burgos-Robles et al. 2009
Le Merre et al. 2018
Otis et al. 2017

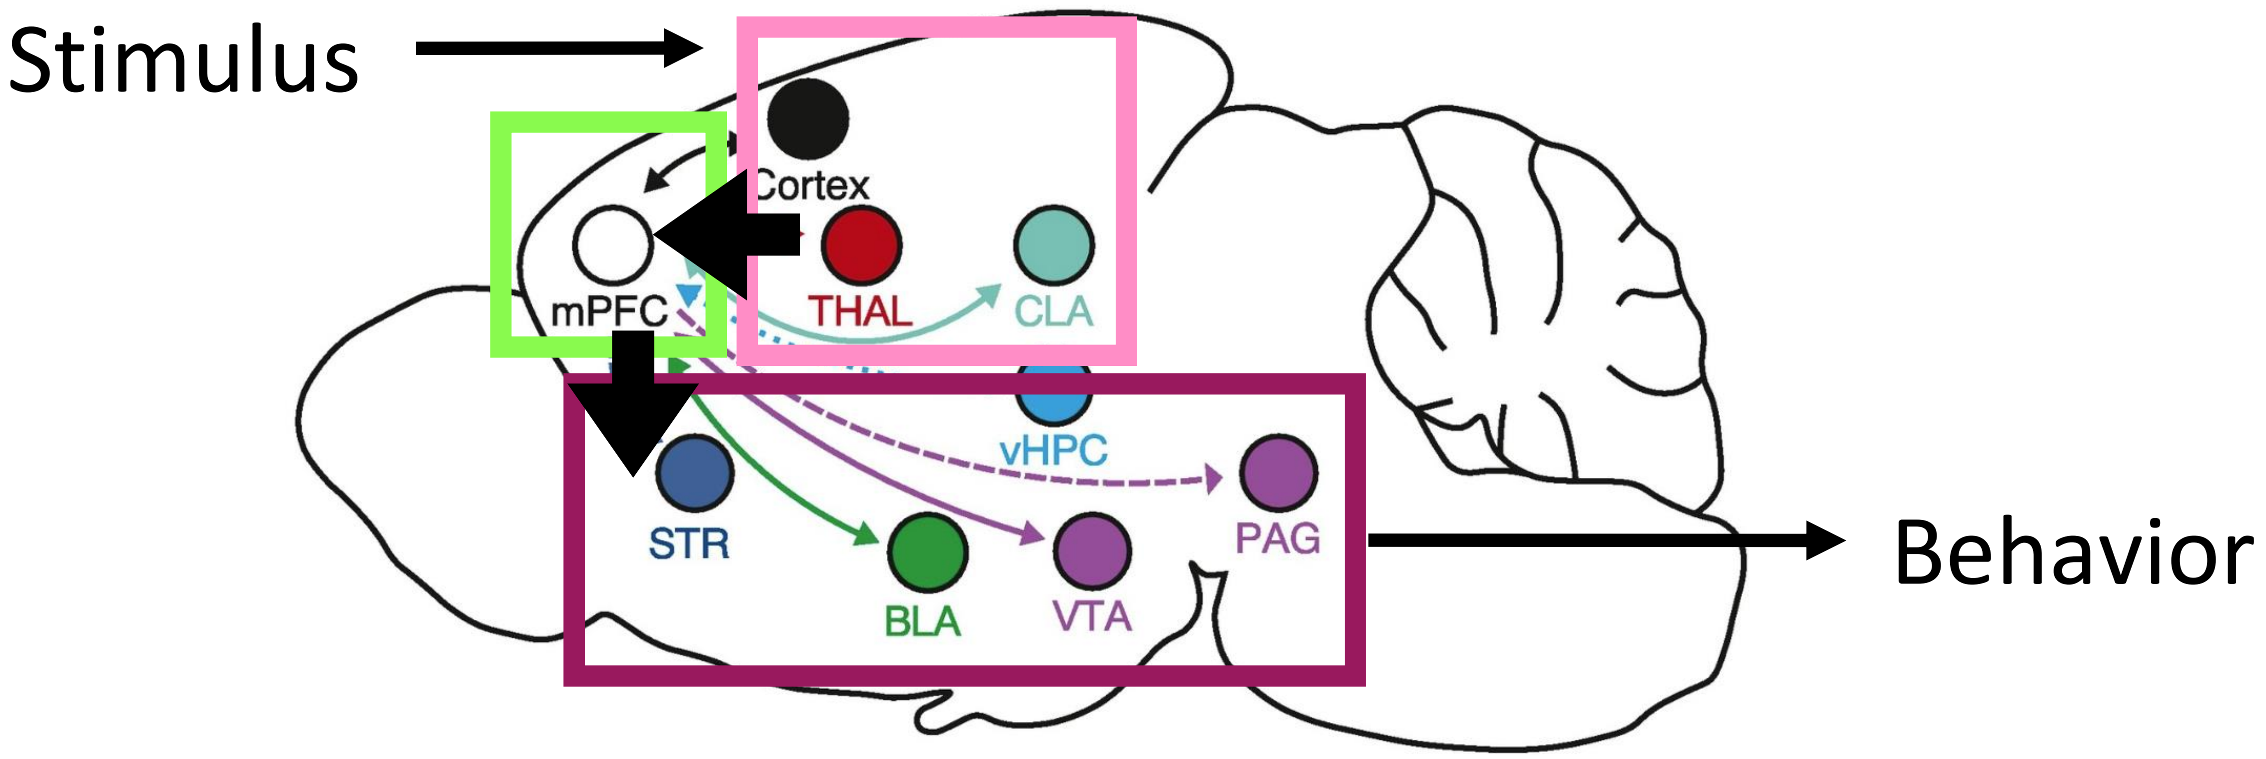

‘mPFC activity guides/ alters behavior’

Murugan et al. 2017
Otis et al. 2017
Rozeske et al. 2018
Diehl et al. 2020

ETH zürich

institute of neuroinformatics

Investigating Abstract Stimulus Representations in mPFC

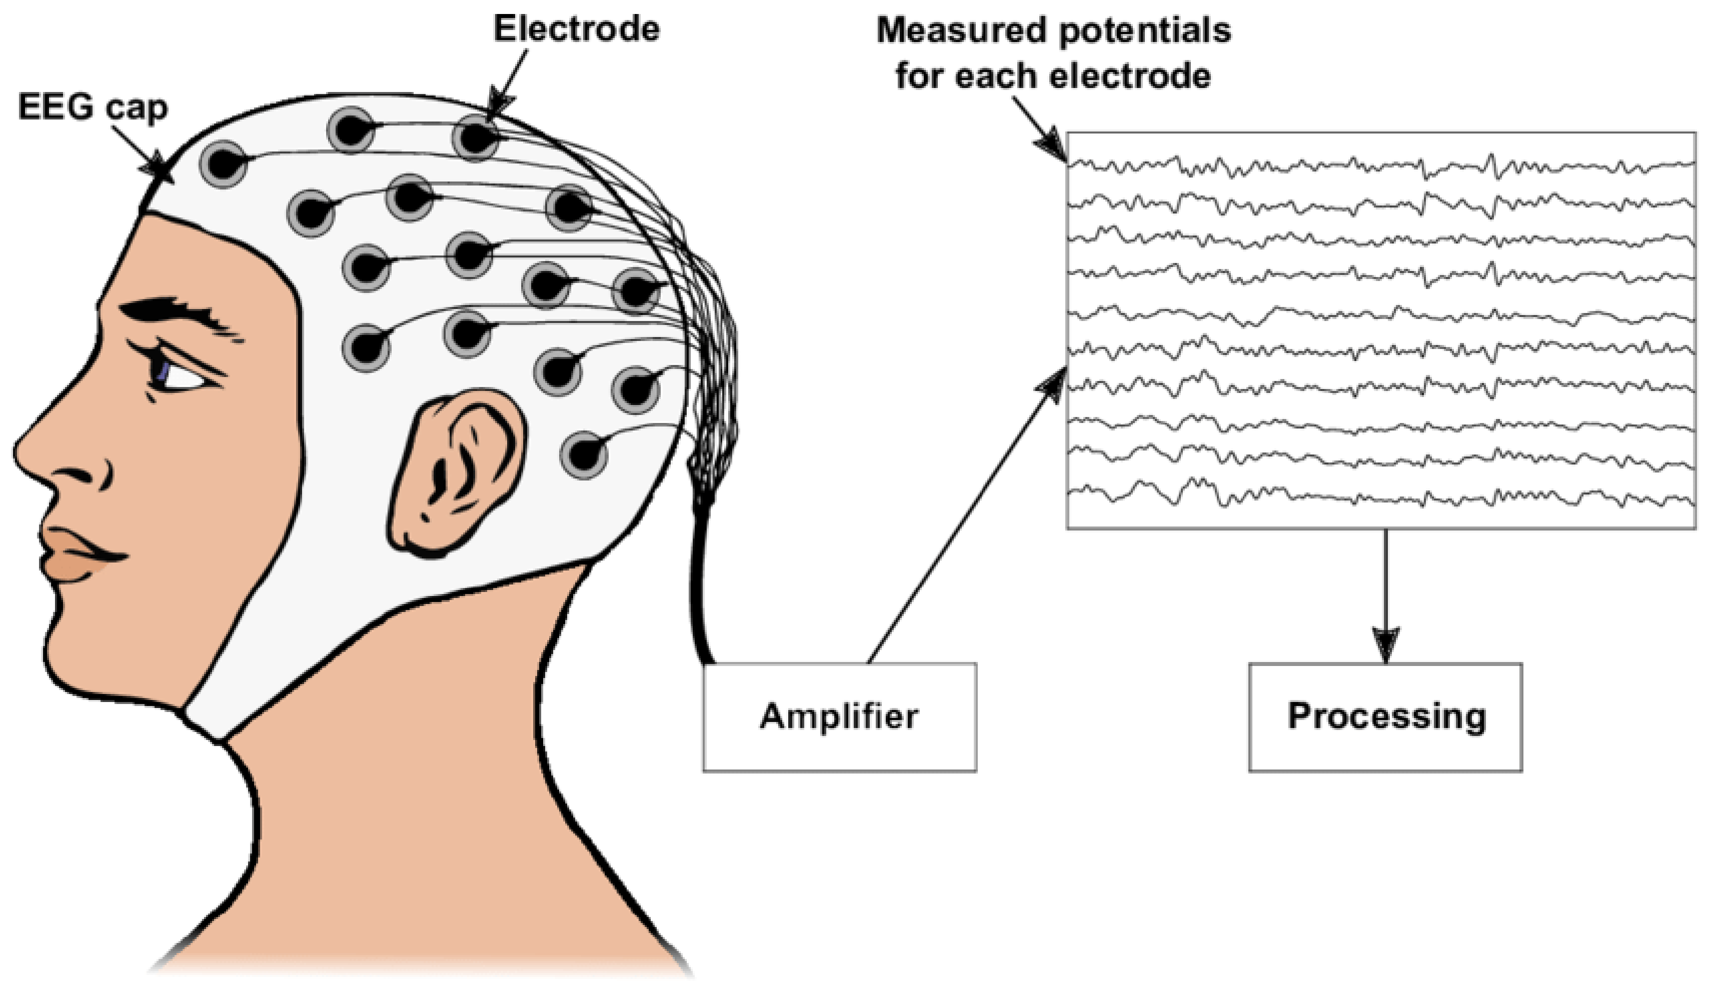

Measuring Neuronal Activity in the Mouse Brain

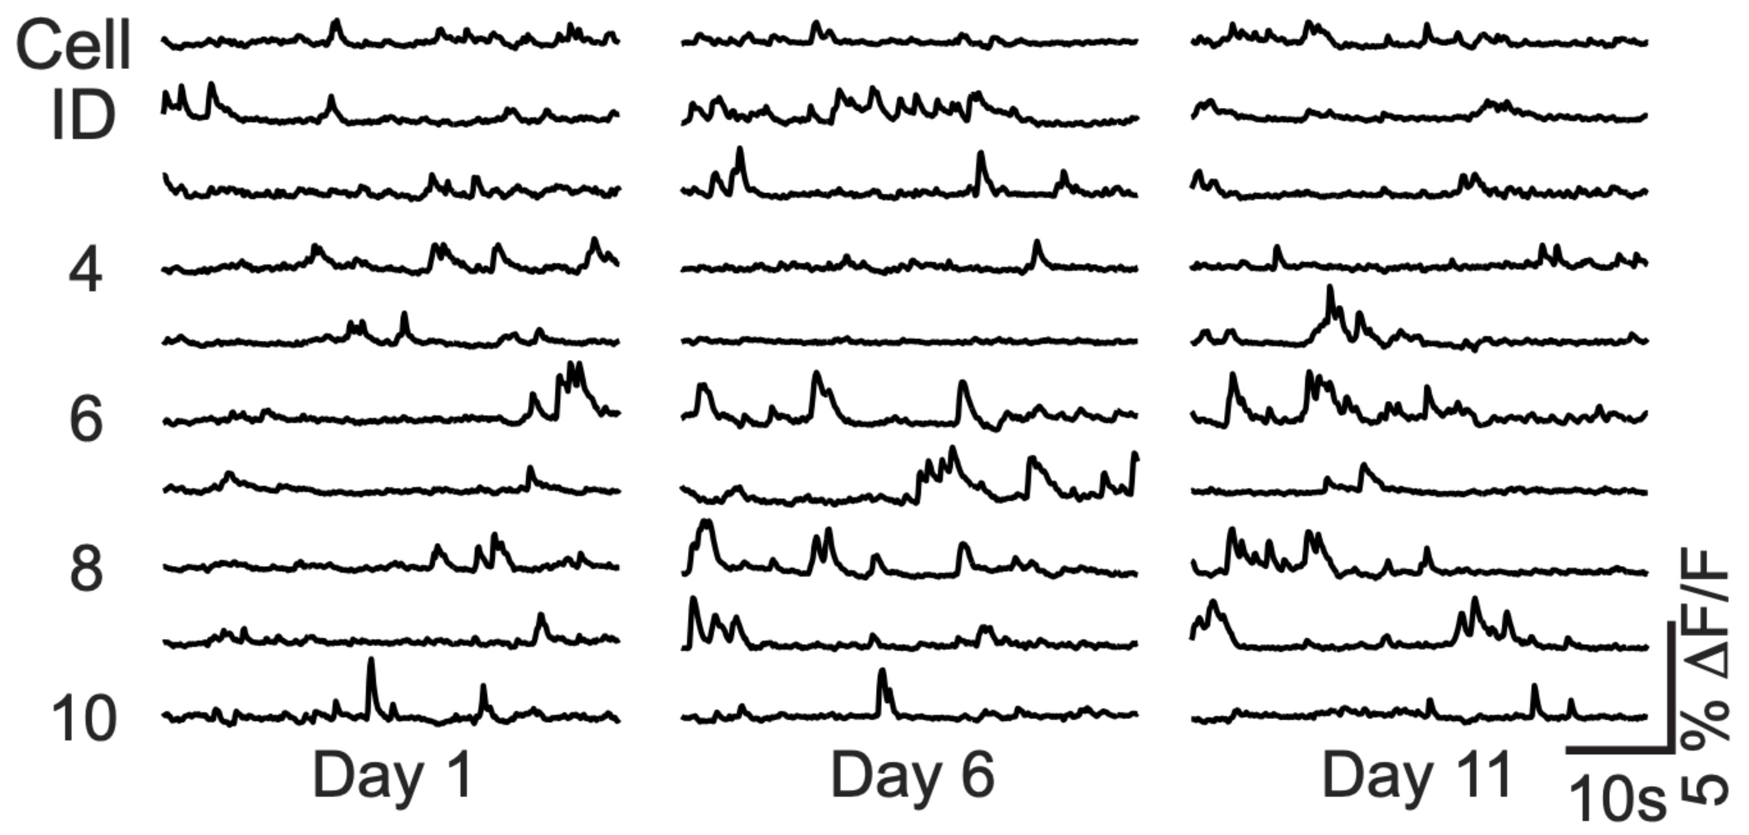

Investigating Abstract Stimulus Representations in mPFC

Data

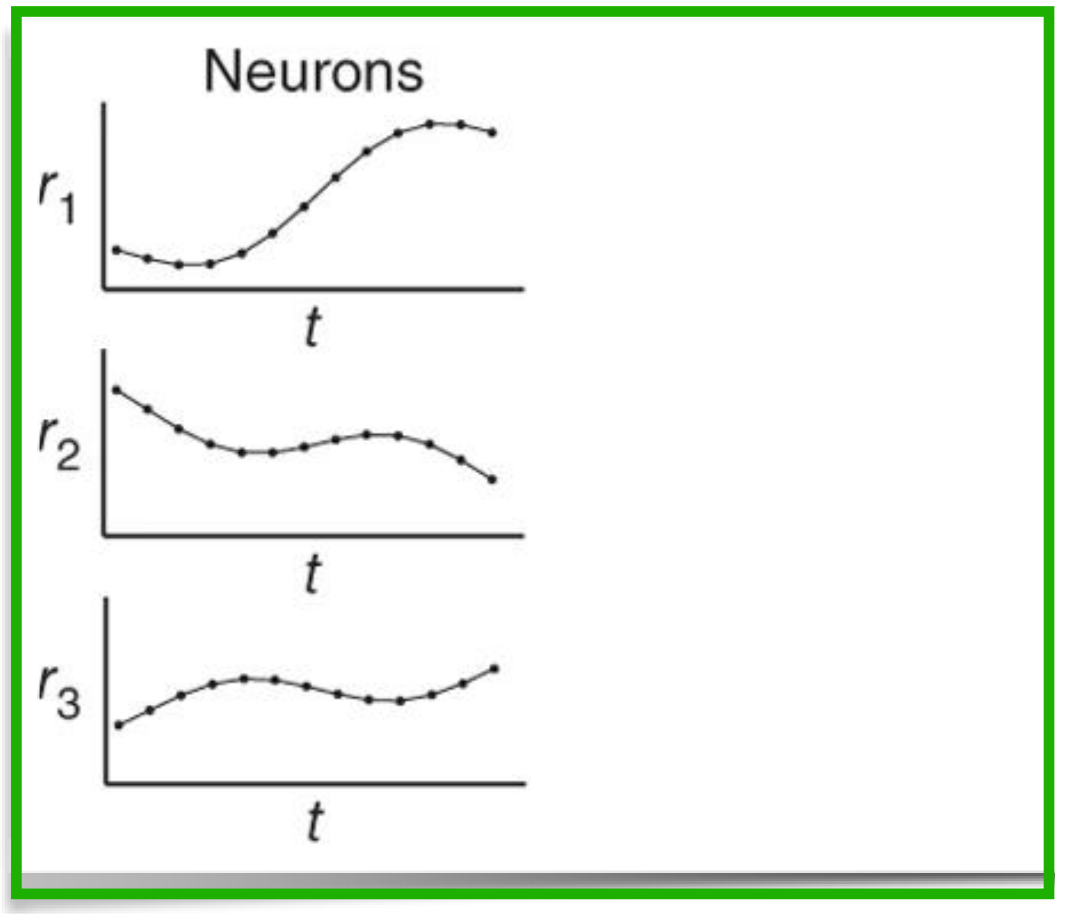

n=12 mice,
3395 neurons

The bus neuron is actually 100 or 1000 neuron's. The bus is encoded as a pattern of activity.

ETH zürich

institute of neuroinformatics

Investigating Abstract Stimulus Representations in mPFC

Data

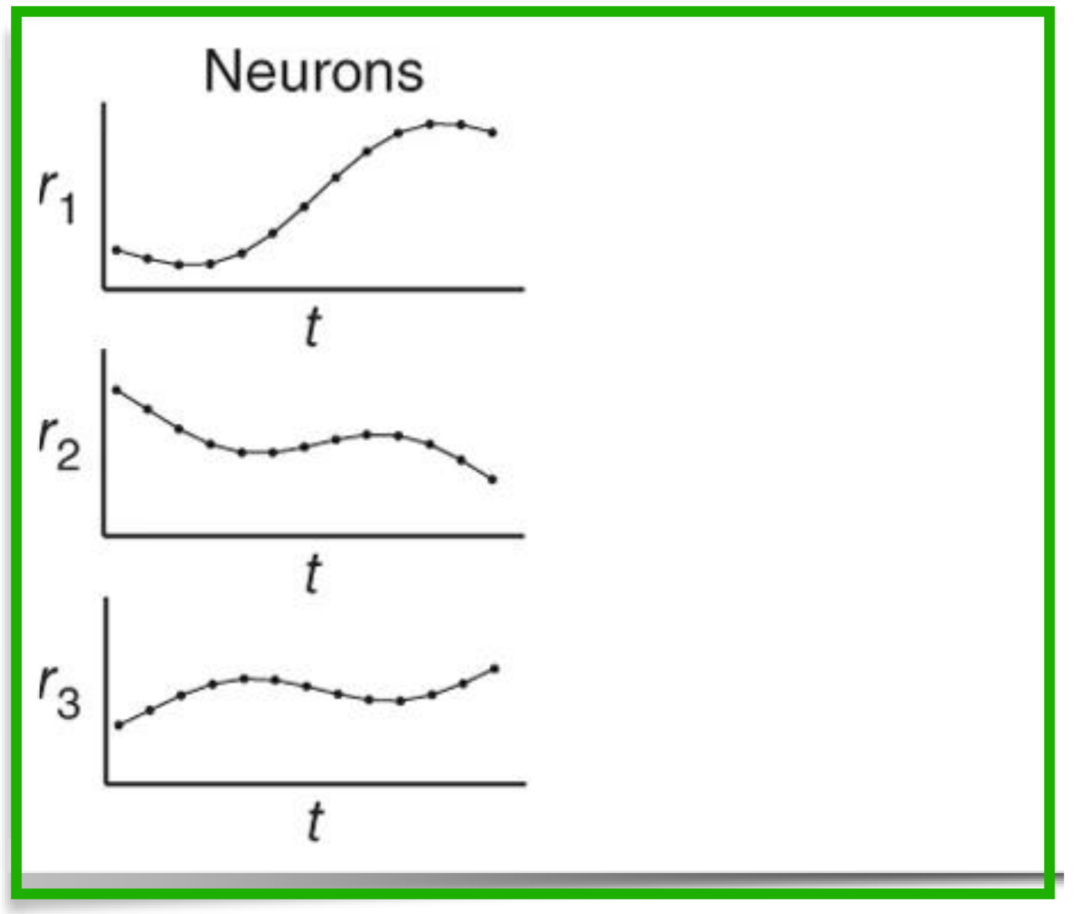

n=12 mice,
3395 neurons

The bus neuron is actually 100 or 1000 neuron's. The bus is encoded as a pattern of activity.

Task Variables

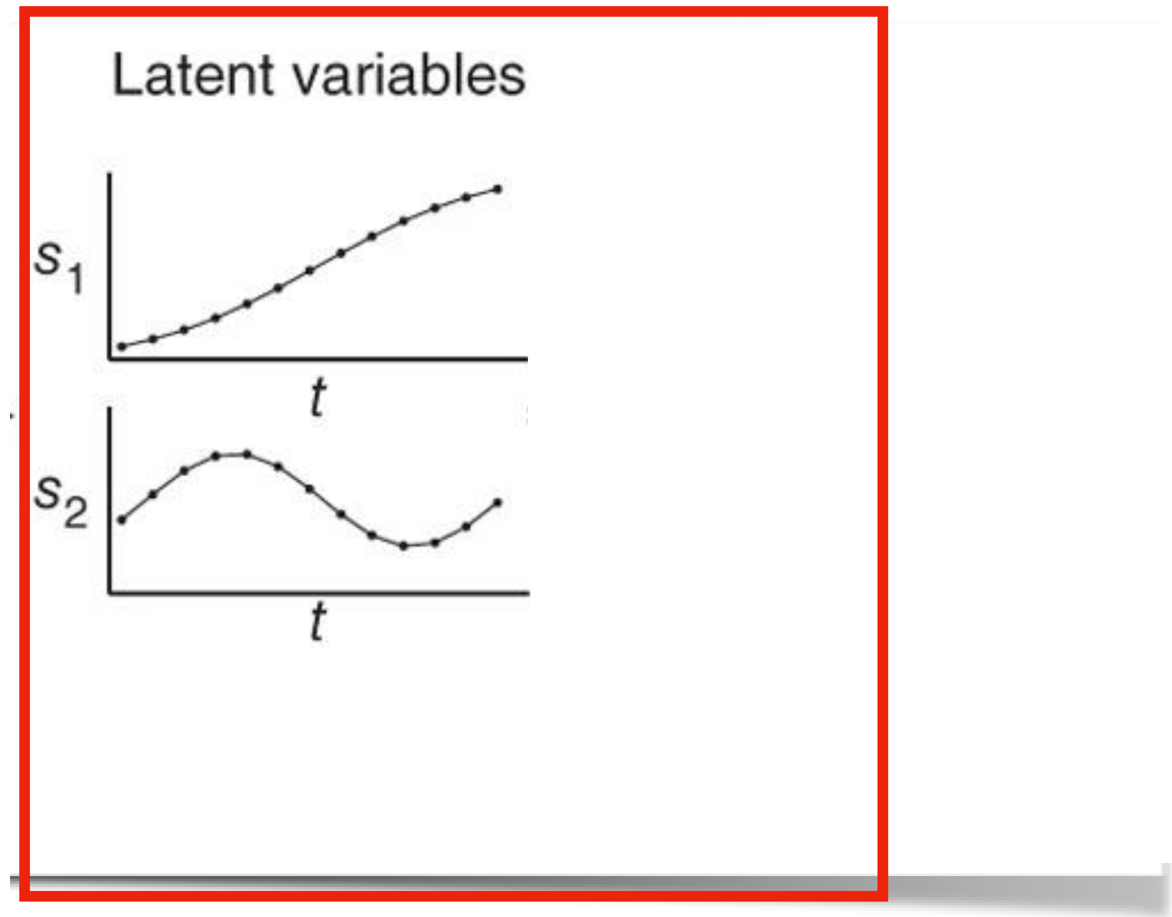

Cunningham and Yu, 2014

<u>Interesting latent/task variables for us:</u>
- Shuttle motion
- Auditory tone stimulus
- Direction of shuttle motion

Investigating Abstract Stimulus Representations in mPFC

Data **Data Analysis Method** **Task Variables**

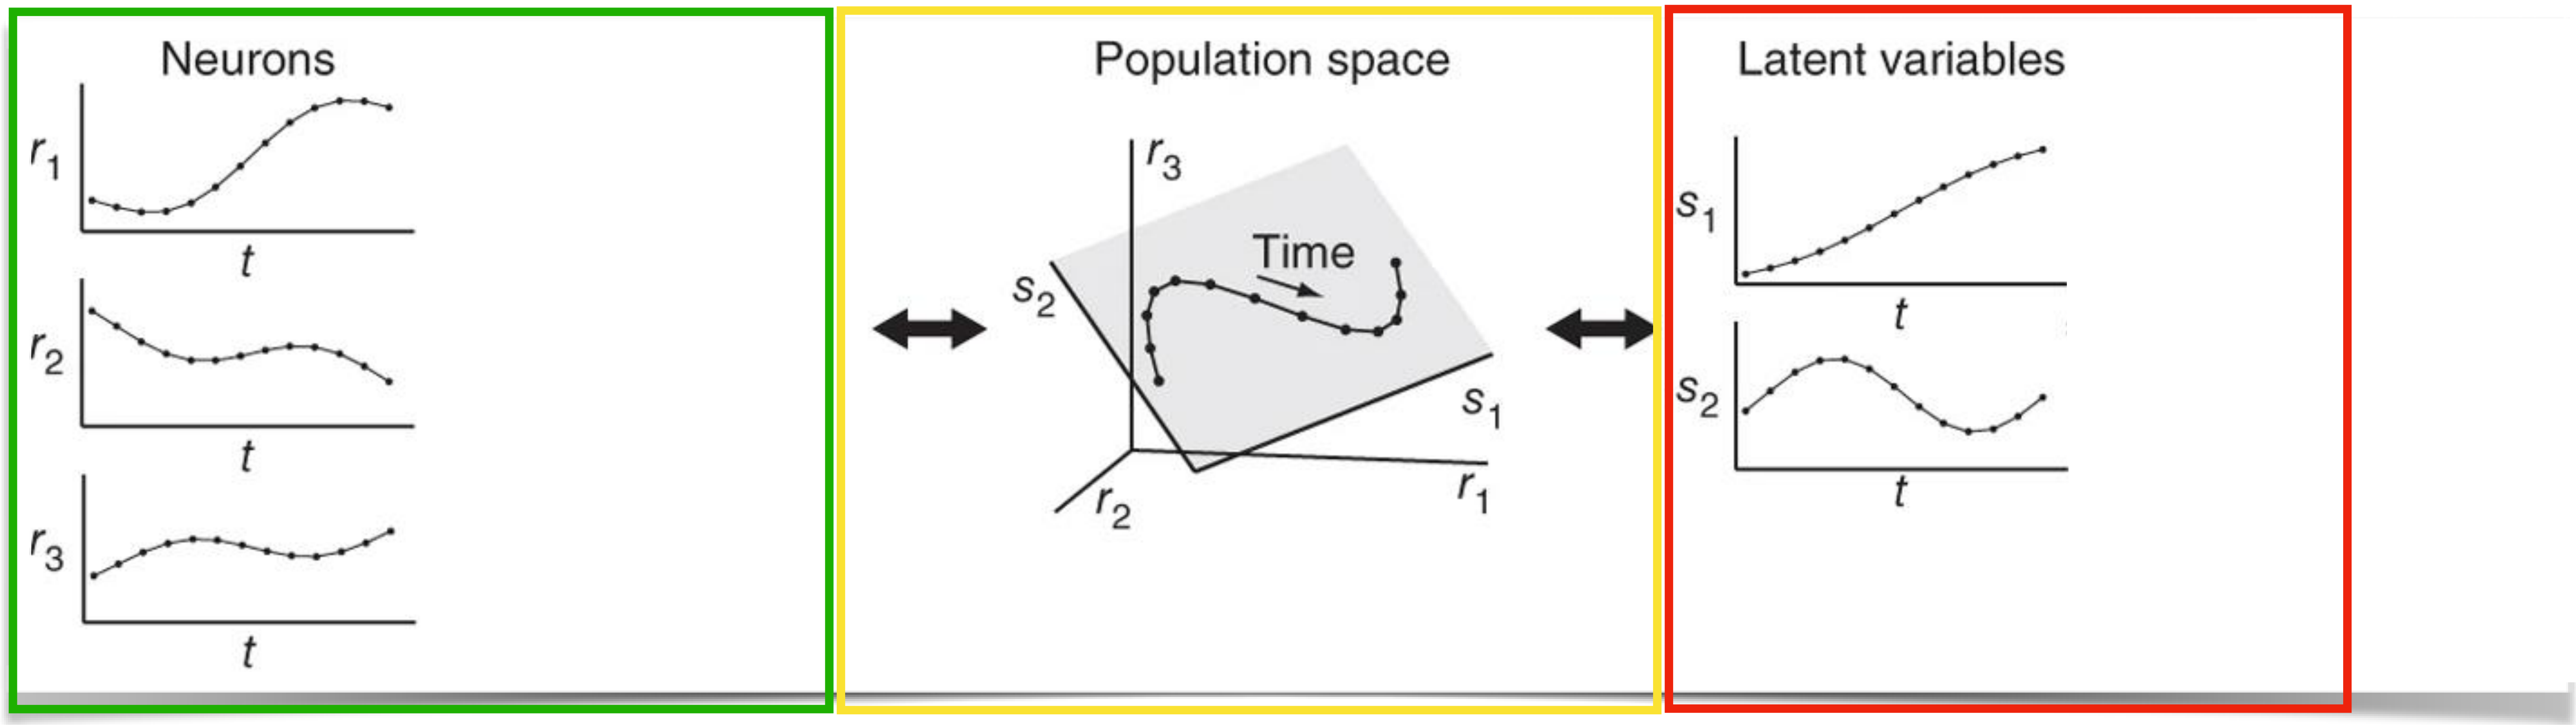

Cunningham and Yu, 2014

n=12 mice,
3395 neurons

The bus neuron is actually 100 or 1000 neuron's. The bus is encoded as a pattern of activity.

<u>Interesting latent/task variables for us:</u>

- Shuttle motion
- Auditory tone stimulus
- Direction of shuttle motion

Stimulus Representations in mPFC

Joint Subspace

n=12 mice,
3395 neurons

Ehret et al. final revisions @ Nature Neurosc.

Stimulus Representations in mPFC

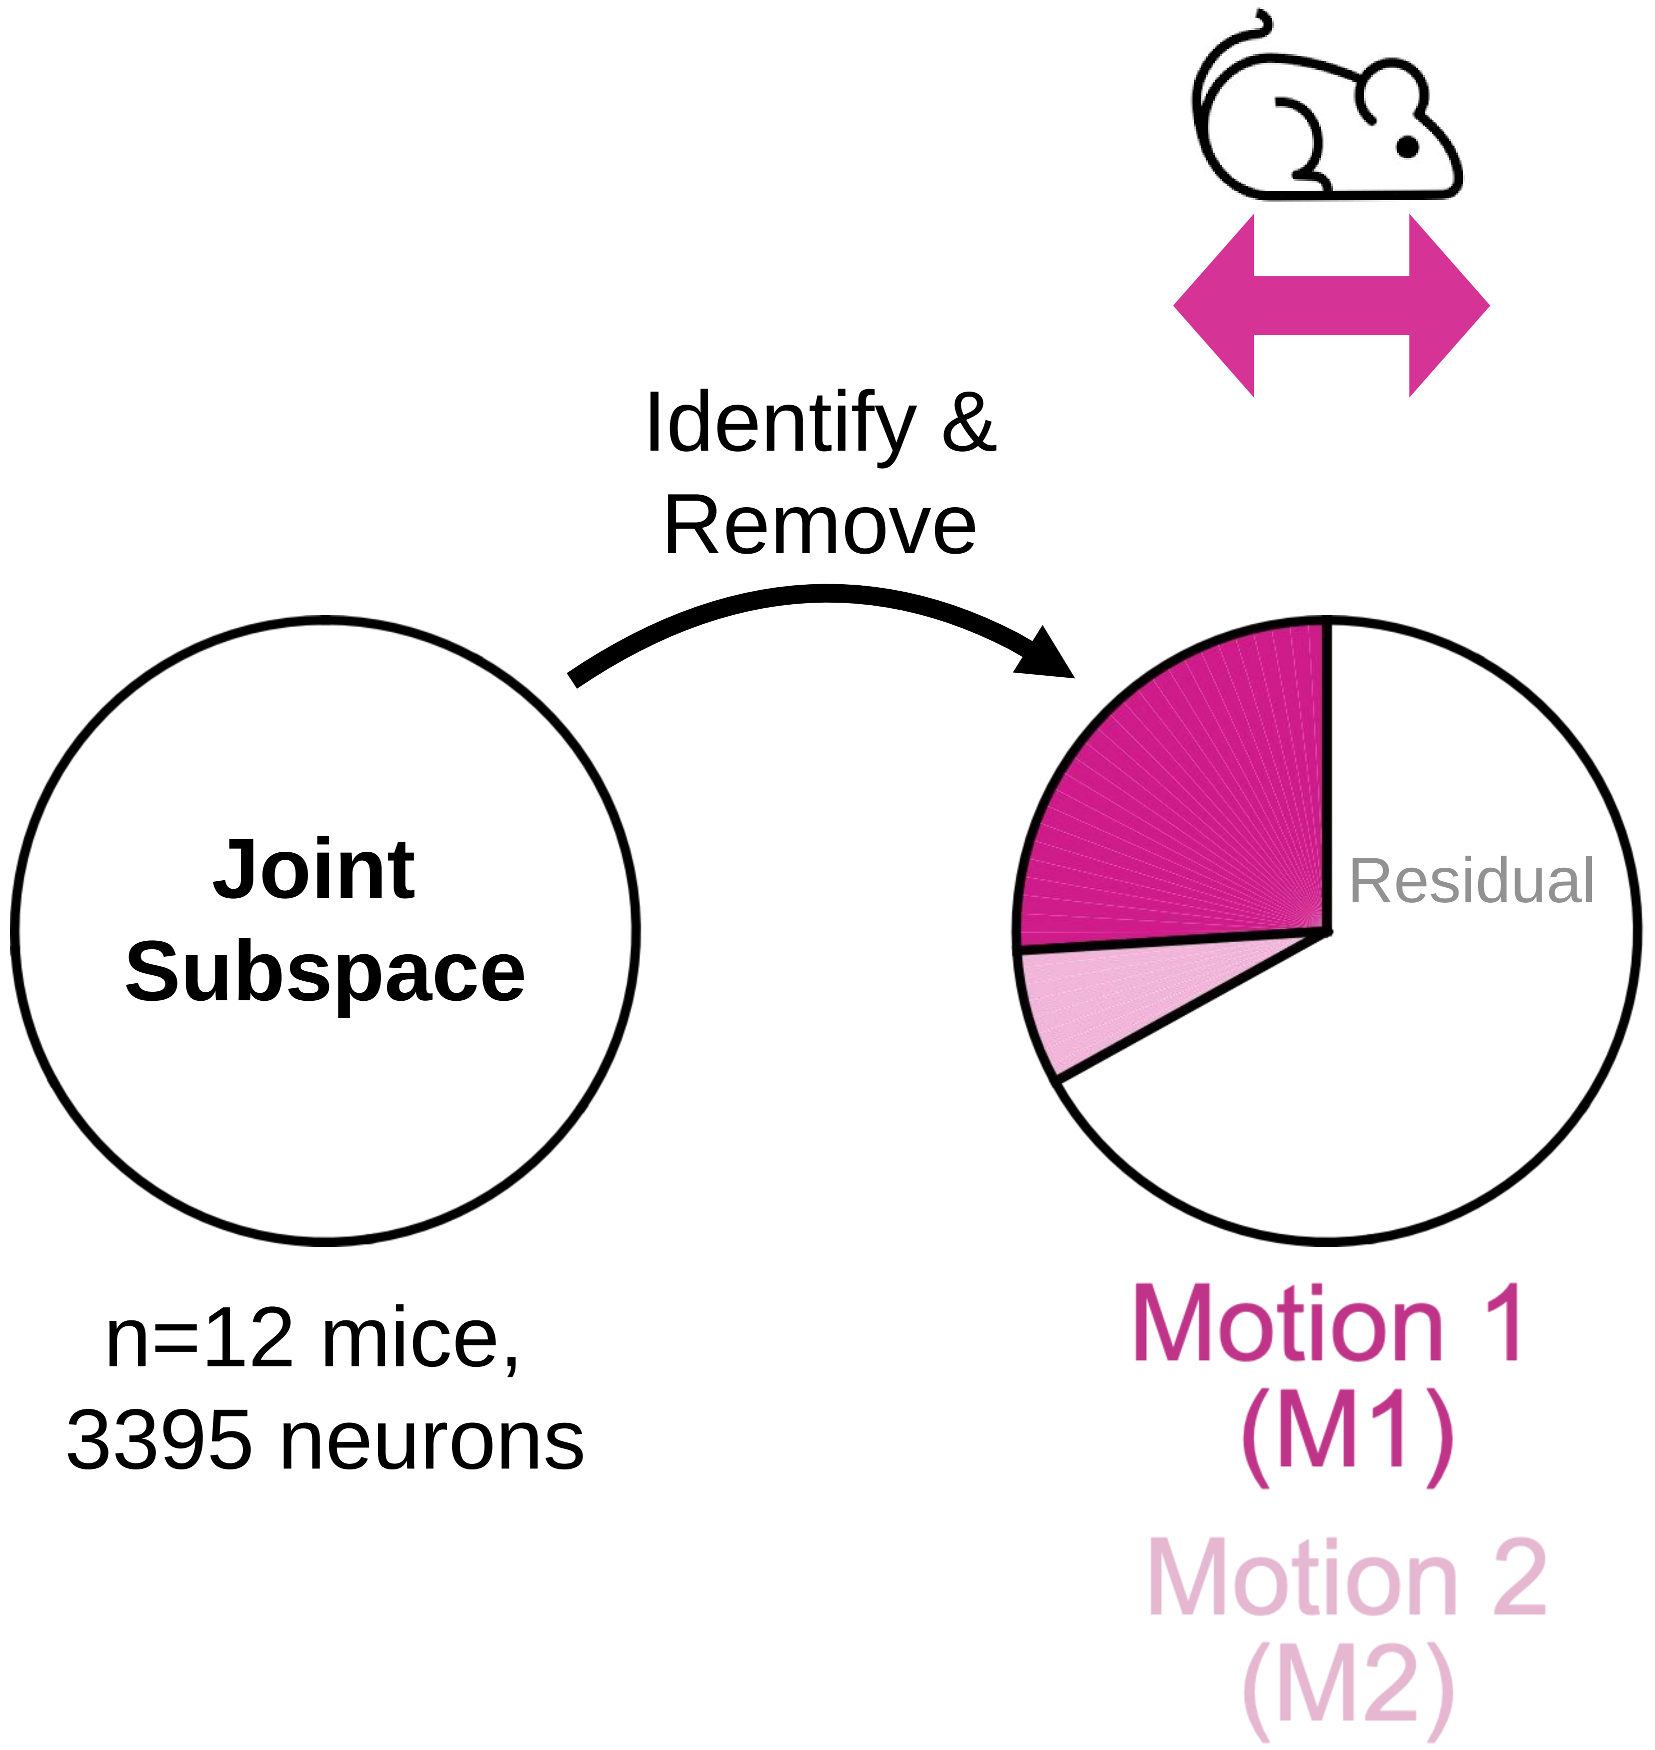

Ehret et al. final revisions @ Nature Neurosc.

Stimulus Representations in mPFC

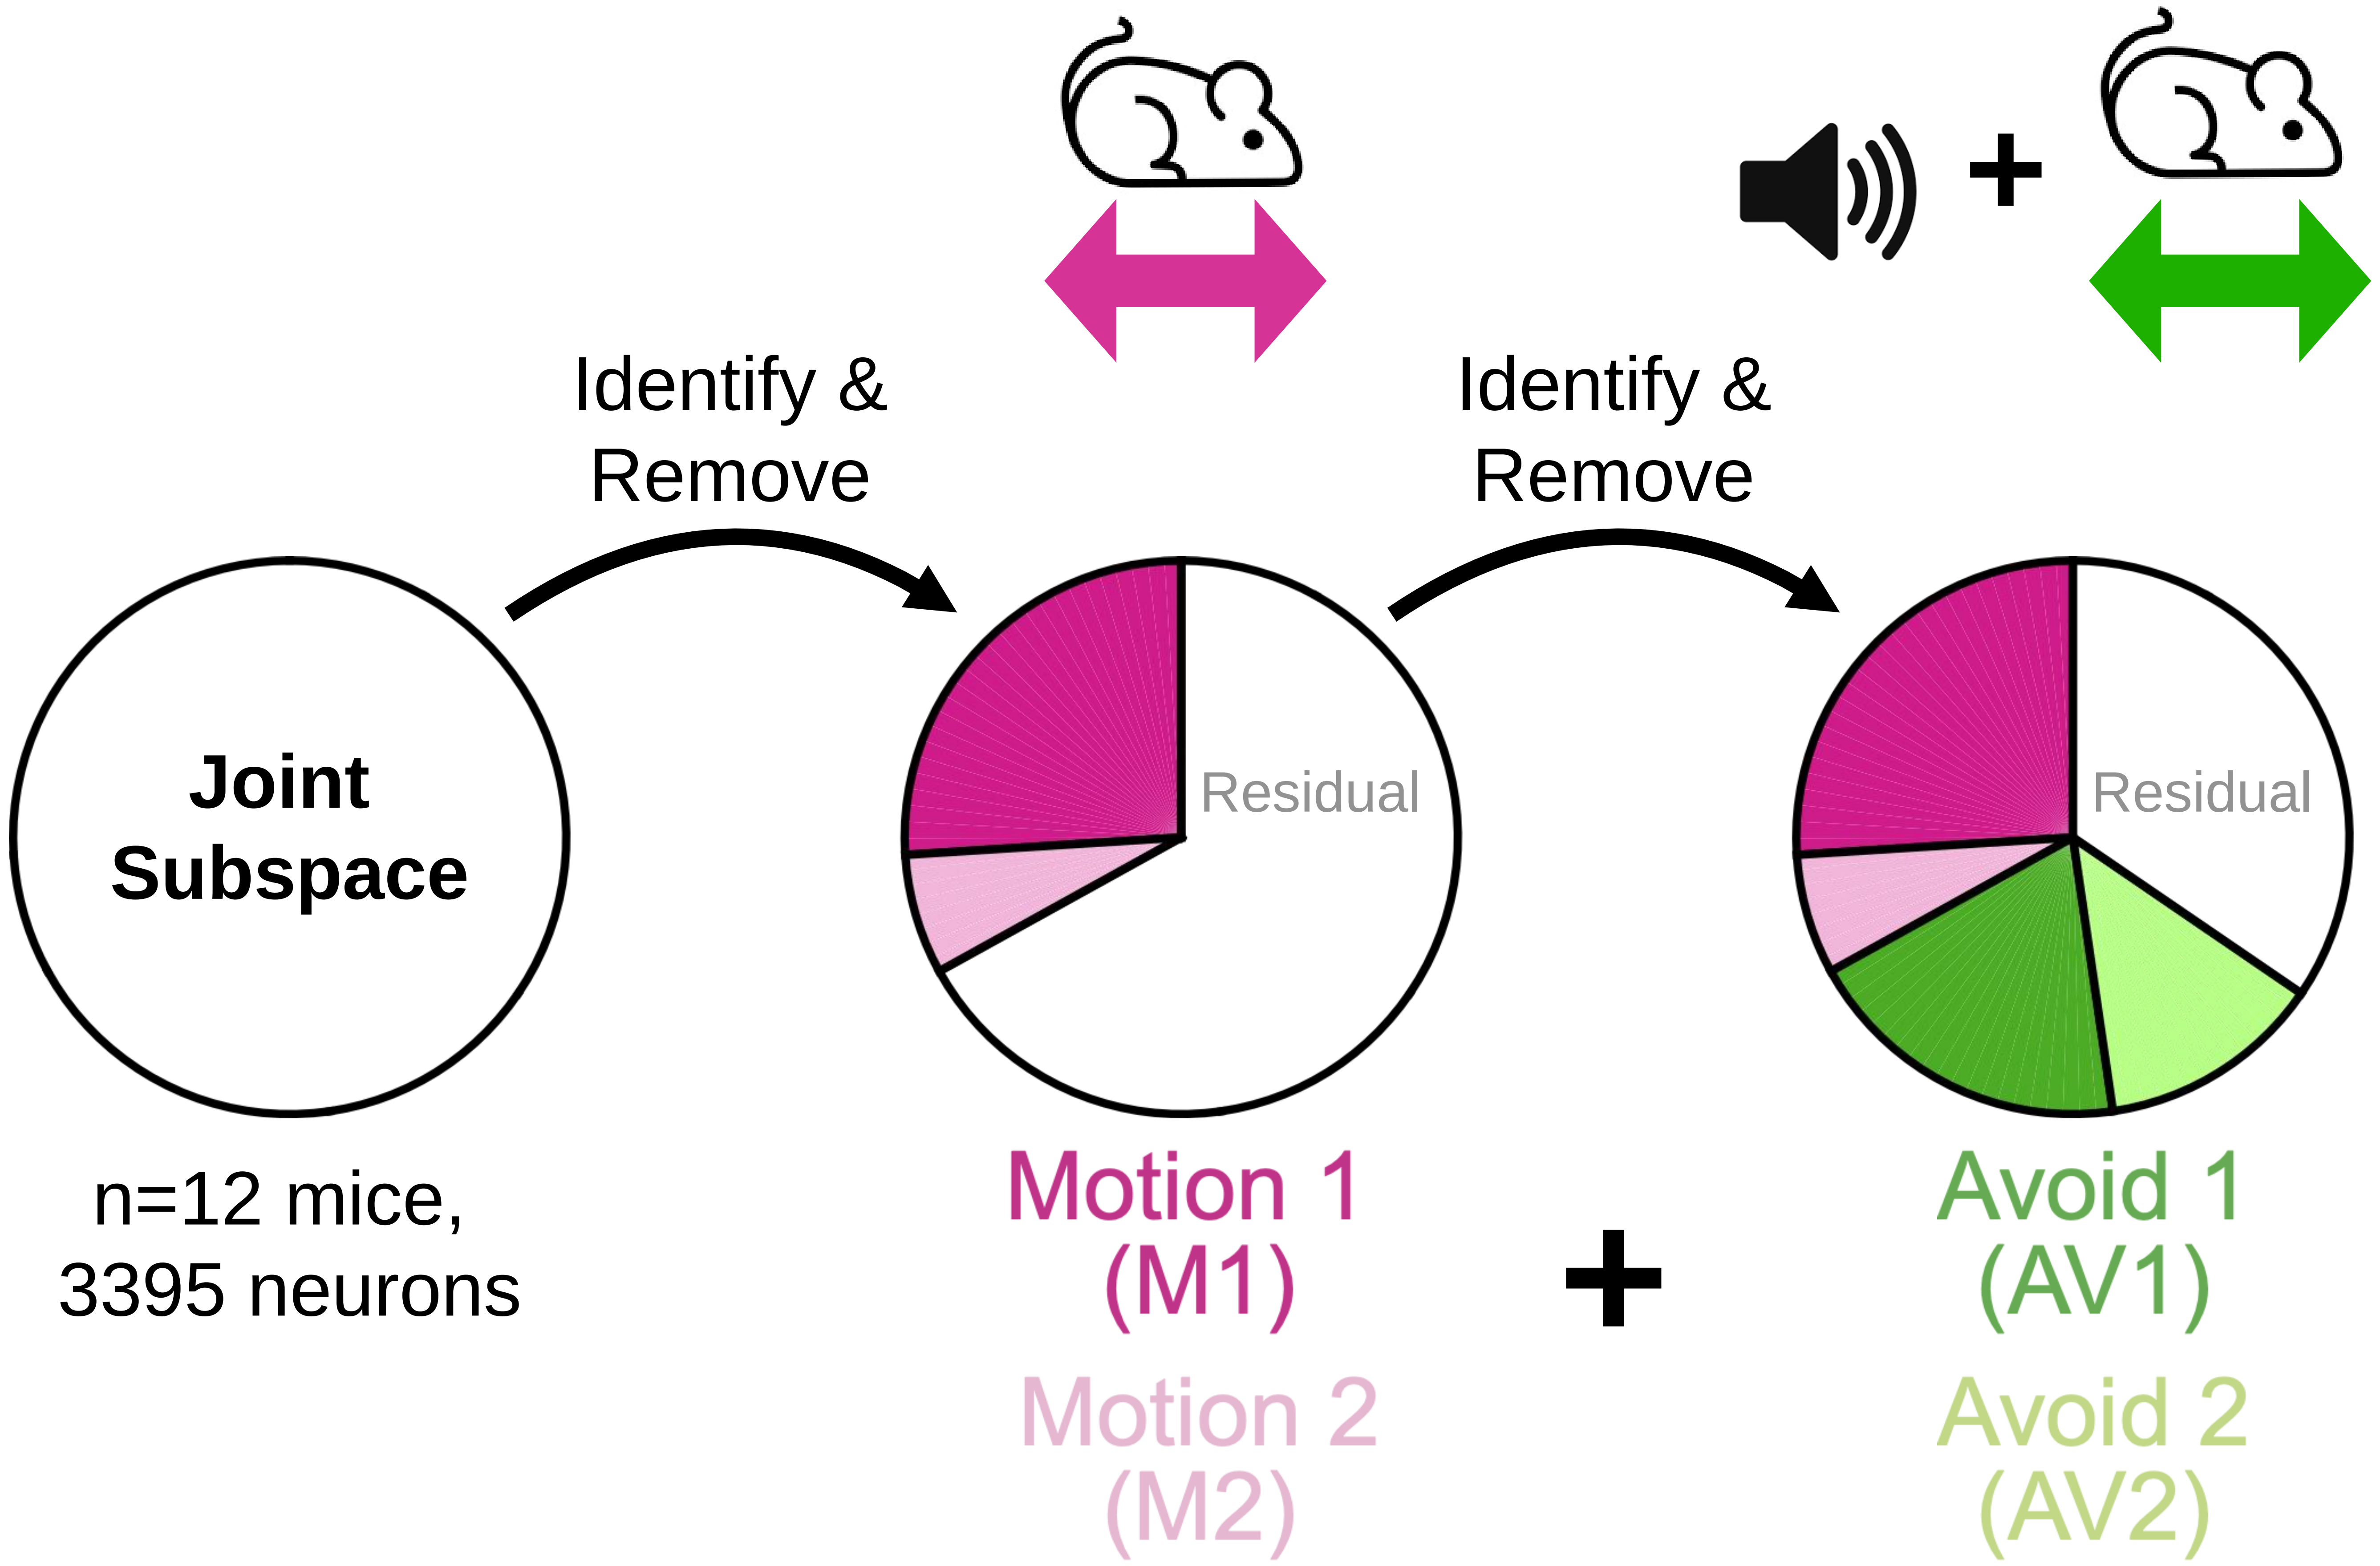

Ehret et al. final revisions @ Nature Neurosc.

Stimulus Representations in mPFC

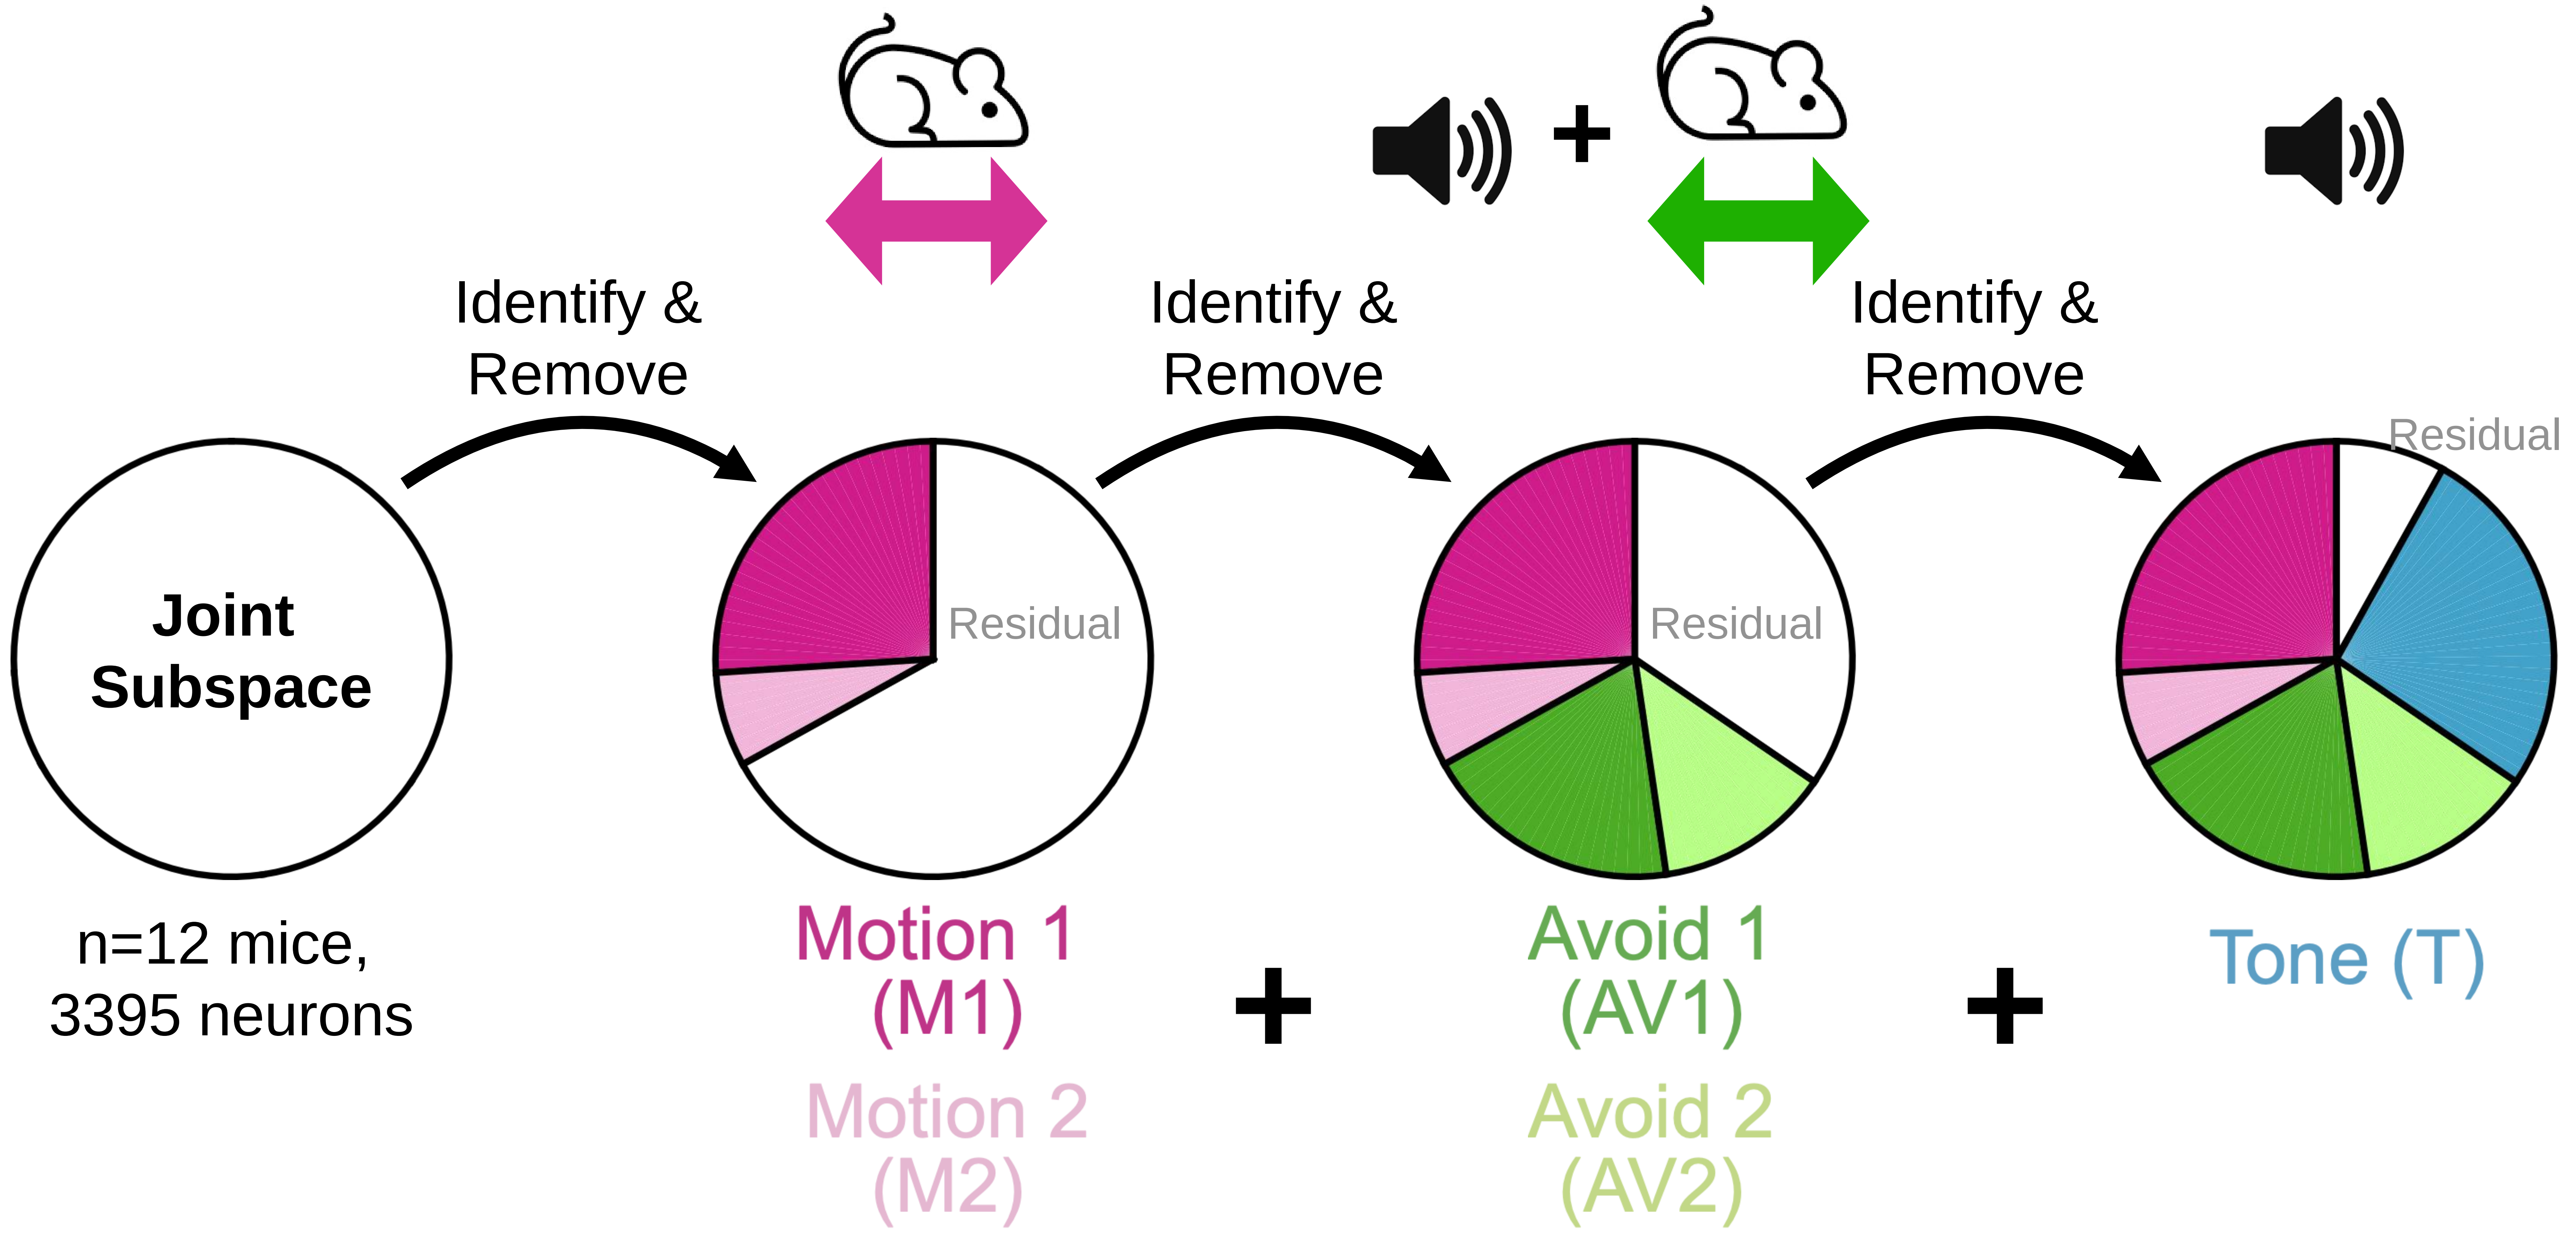

Ehret et al. final revisions @ Nature Neurosc.

Stimulus Representations in mPFC Encode Task(sensory)-Specific Motor Plans

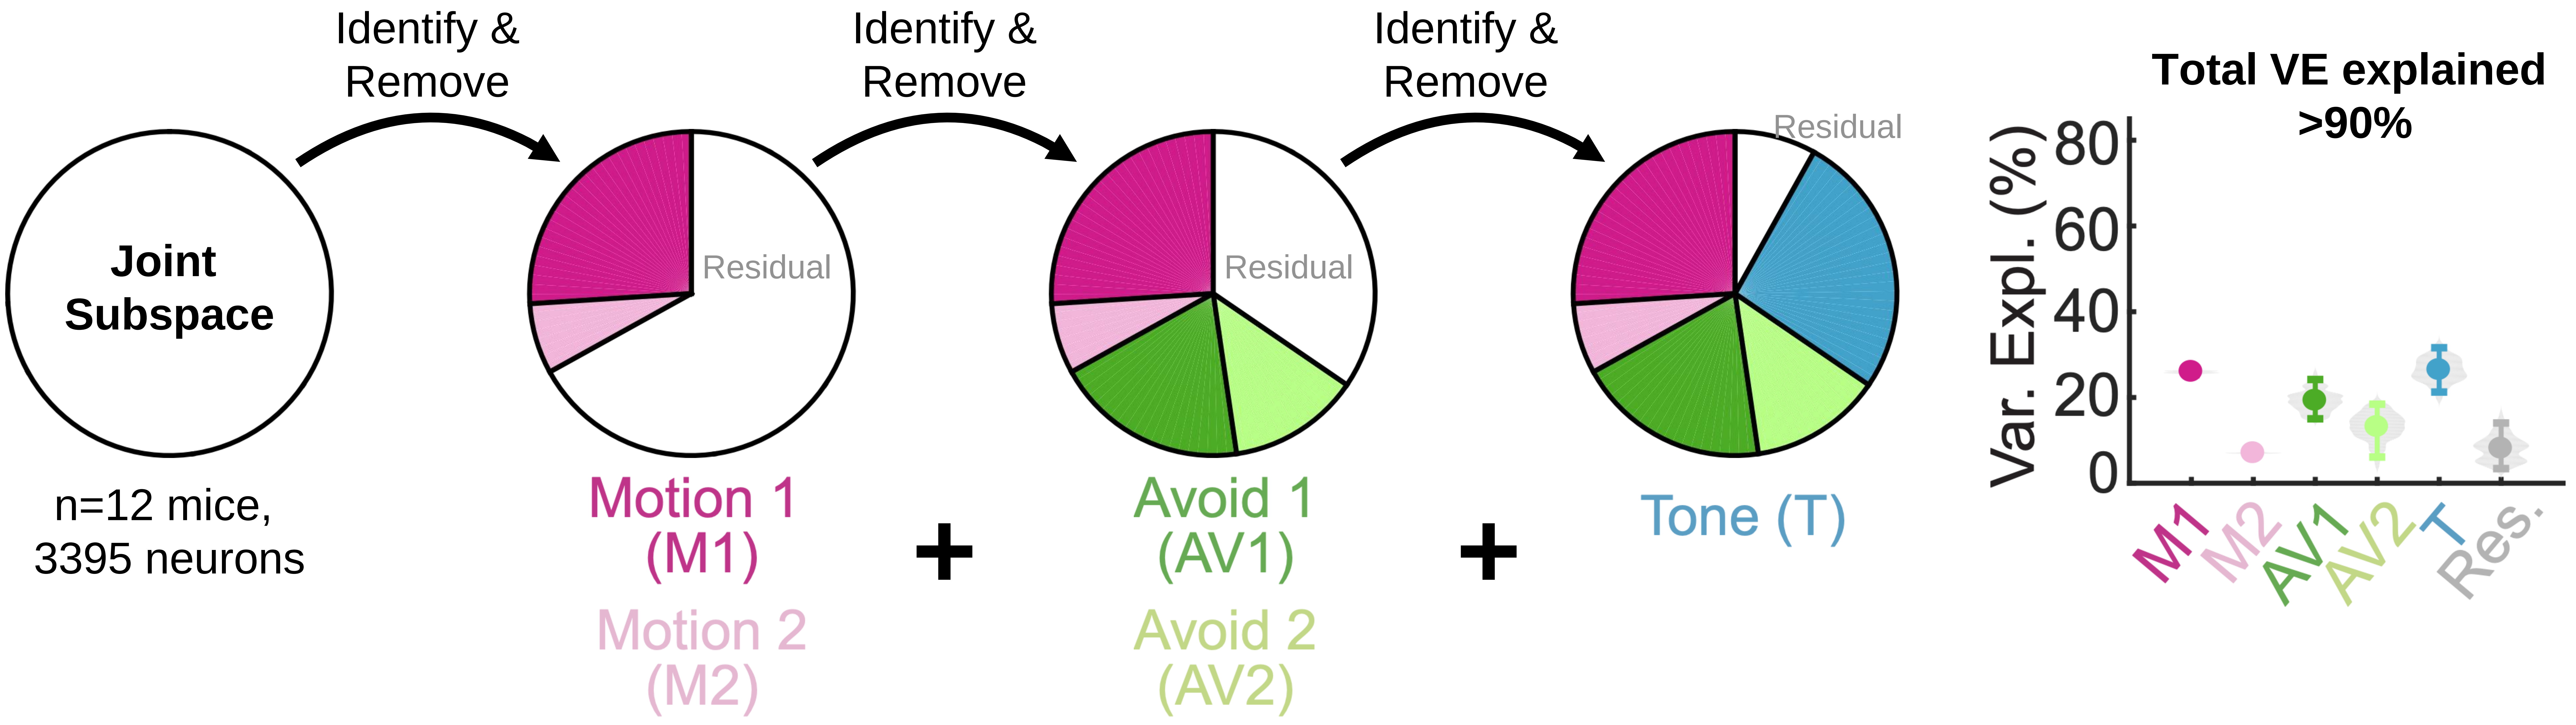

Ehret et al. final revisions @ Nature Neurosc.

Stimulus Representations in mPFC Encode Task(sensory)-Specific Motor Plans

ETH zürich

institute of neuroinformatics

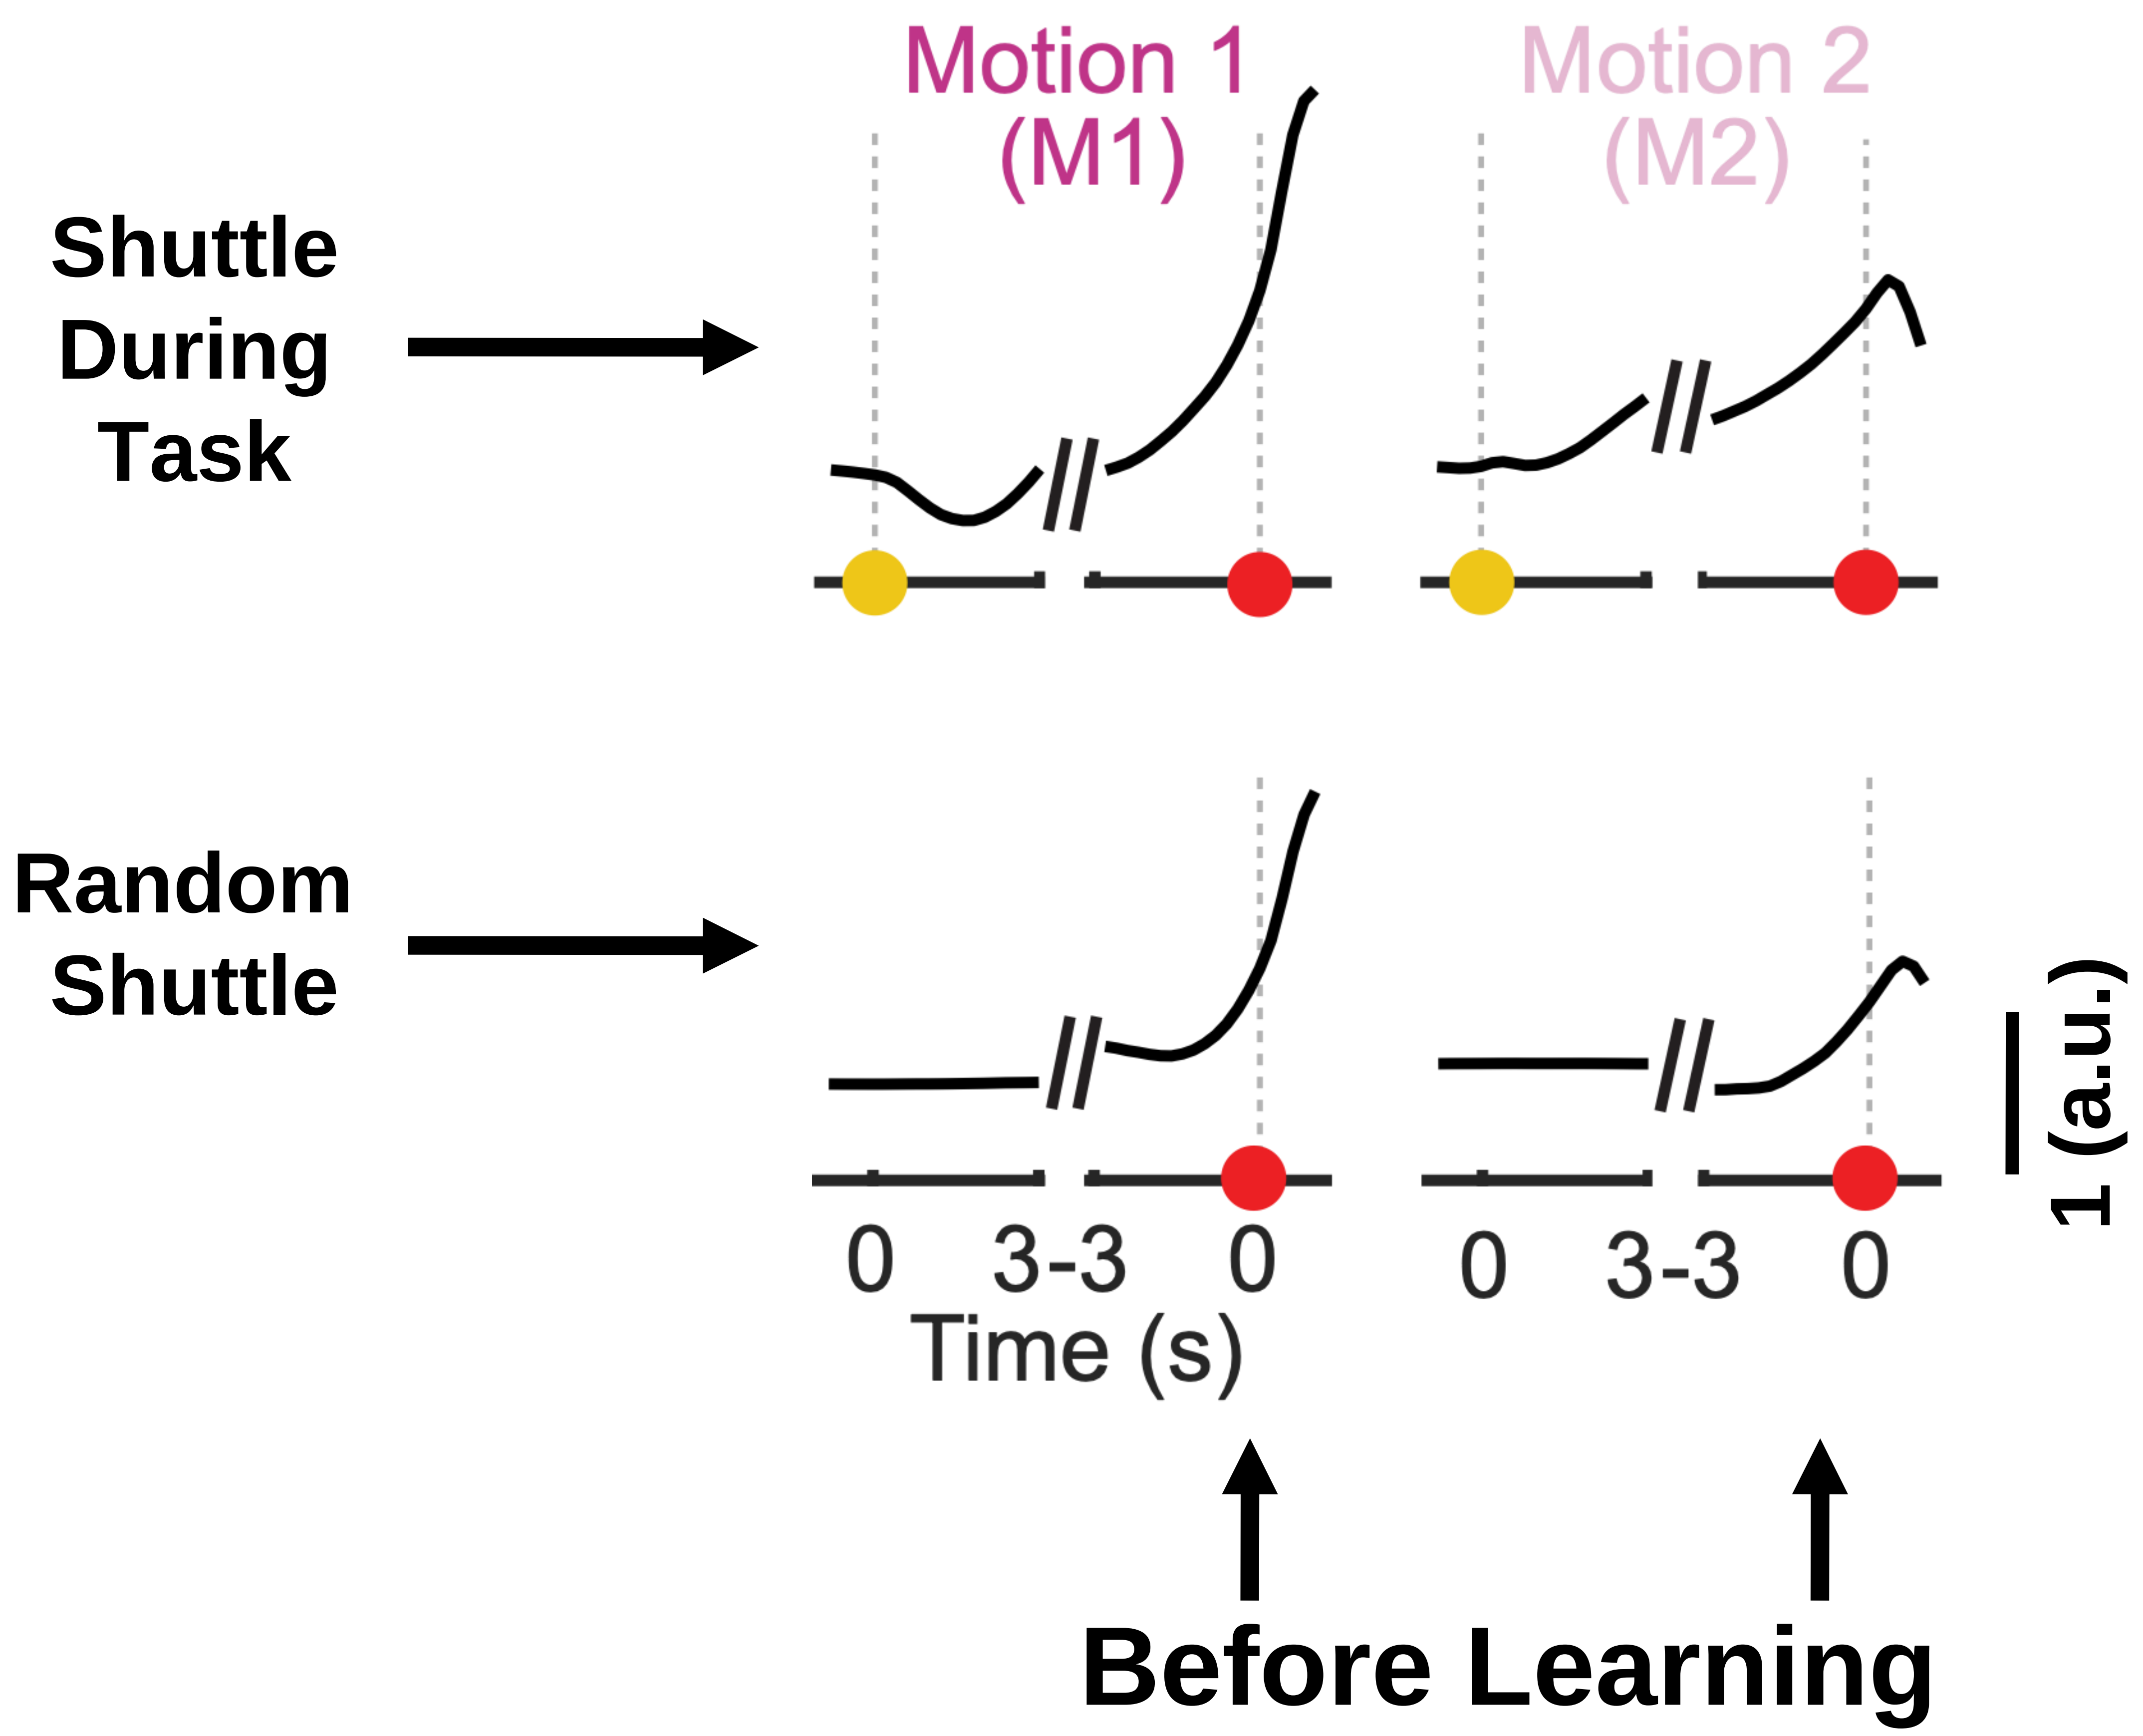

Tone Start
Avoid Action Start

Ehret et al. final revisions @ Nature Neurosc.

Stimulus Representations in mPFC Encode Task(sensory)-Specific Motor Plans

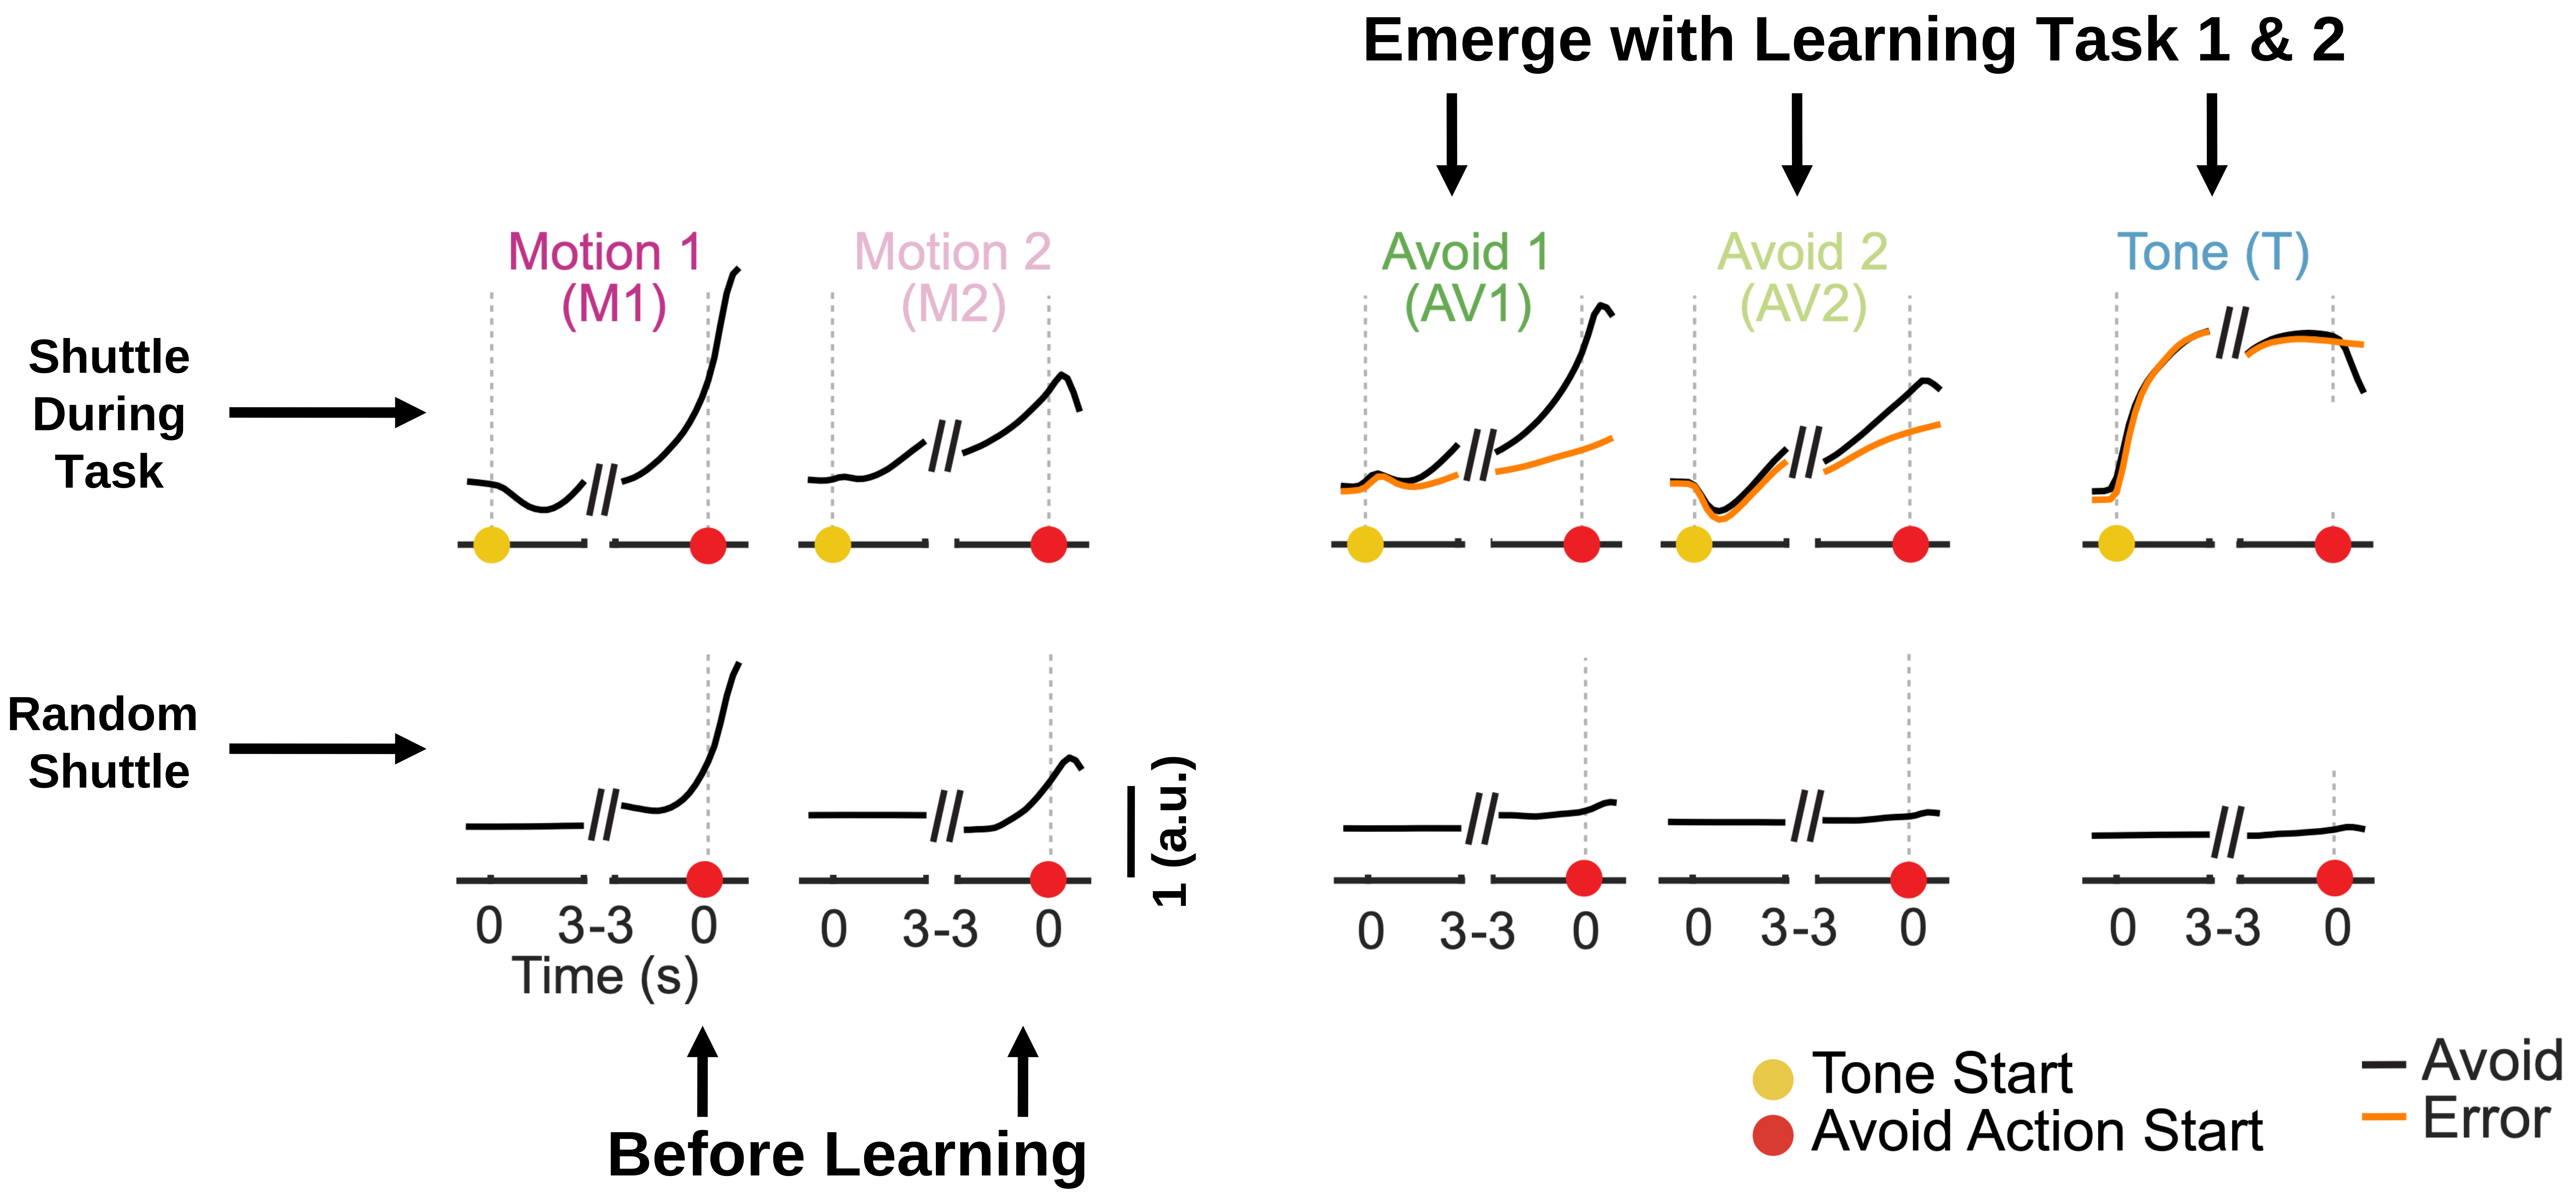

Ehret et al. final revisions @ Nature Neurosc.

Stimulus Representations in mPFC Encode Task(sensory)-Specific Motor Plans

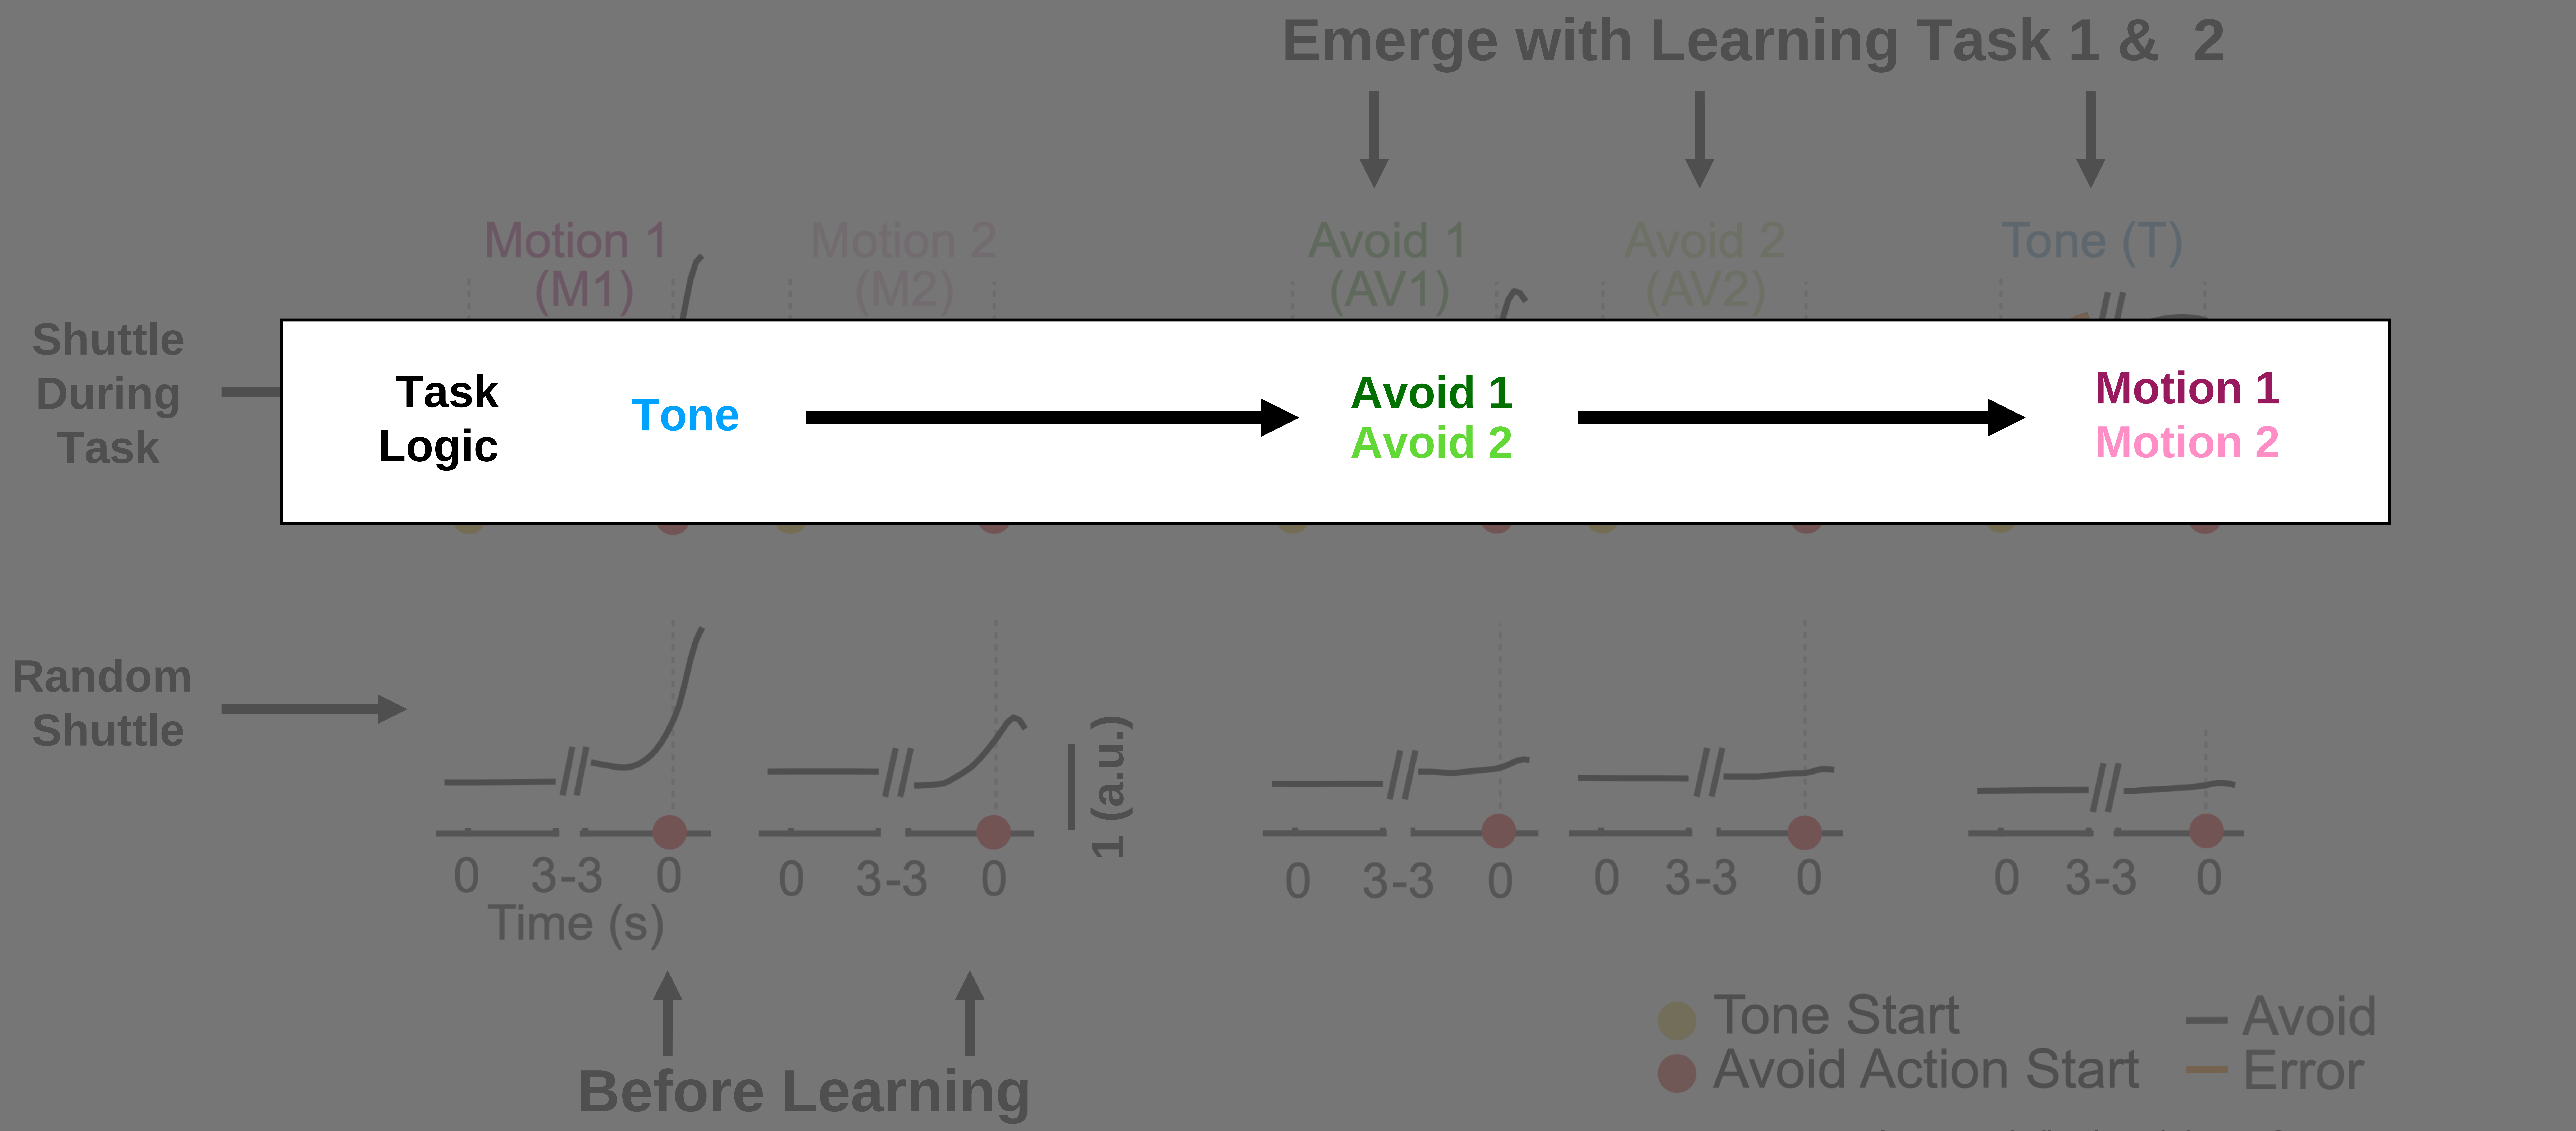

Ehret et al. final revisions @ Nature Neurosc.

Stimulus Representations in mPFC Encode Task(sensory)-Specific Motor Plans

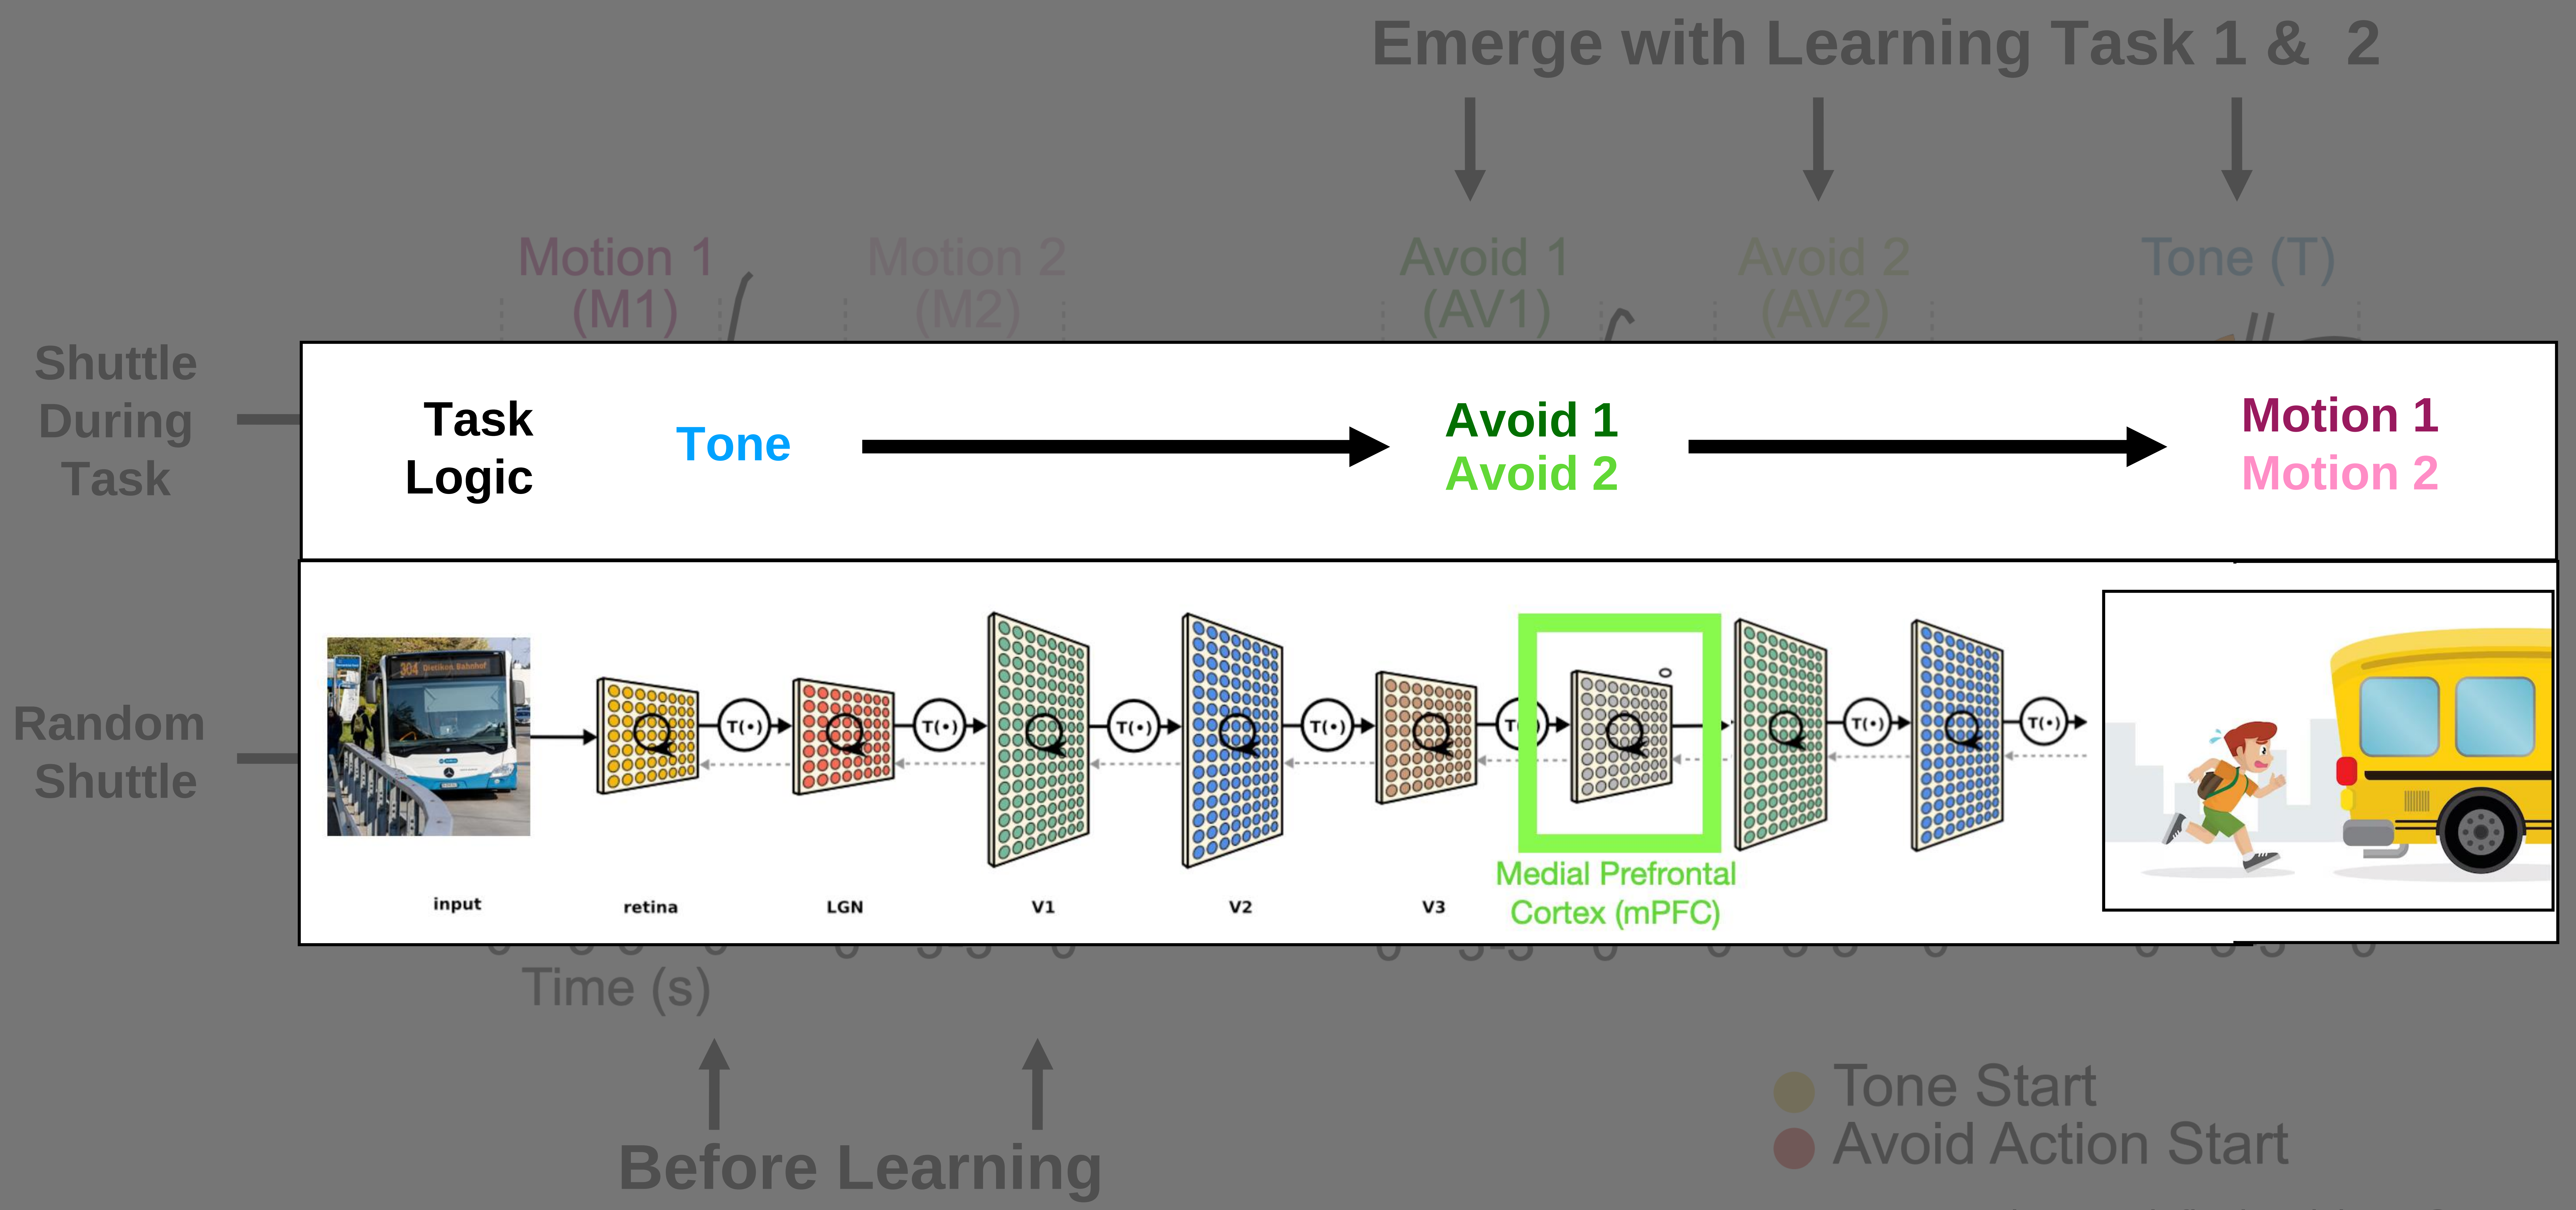

Ehret et al. final revisions @ Nature Neurosc.

A candidate neural representation for 'Affordance'.

Prof. Dr. Jean Piaget
Swiss psychologist and pioneer
Neuchatel, 1896-1980

UNIVERSITÉ DE GENÈVE
Université de Paris

The Concept of 'Affordance' in Psychology

Affordance alludes to the qualities of an object or situation that define its possible use or make clear how it can or should be used.

Affordance adheres to the idea that perception and action are inseparable (Principles of Genetic Epistemology, Piaget).

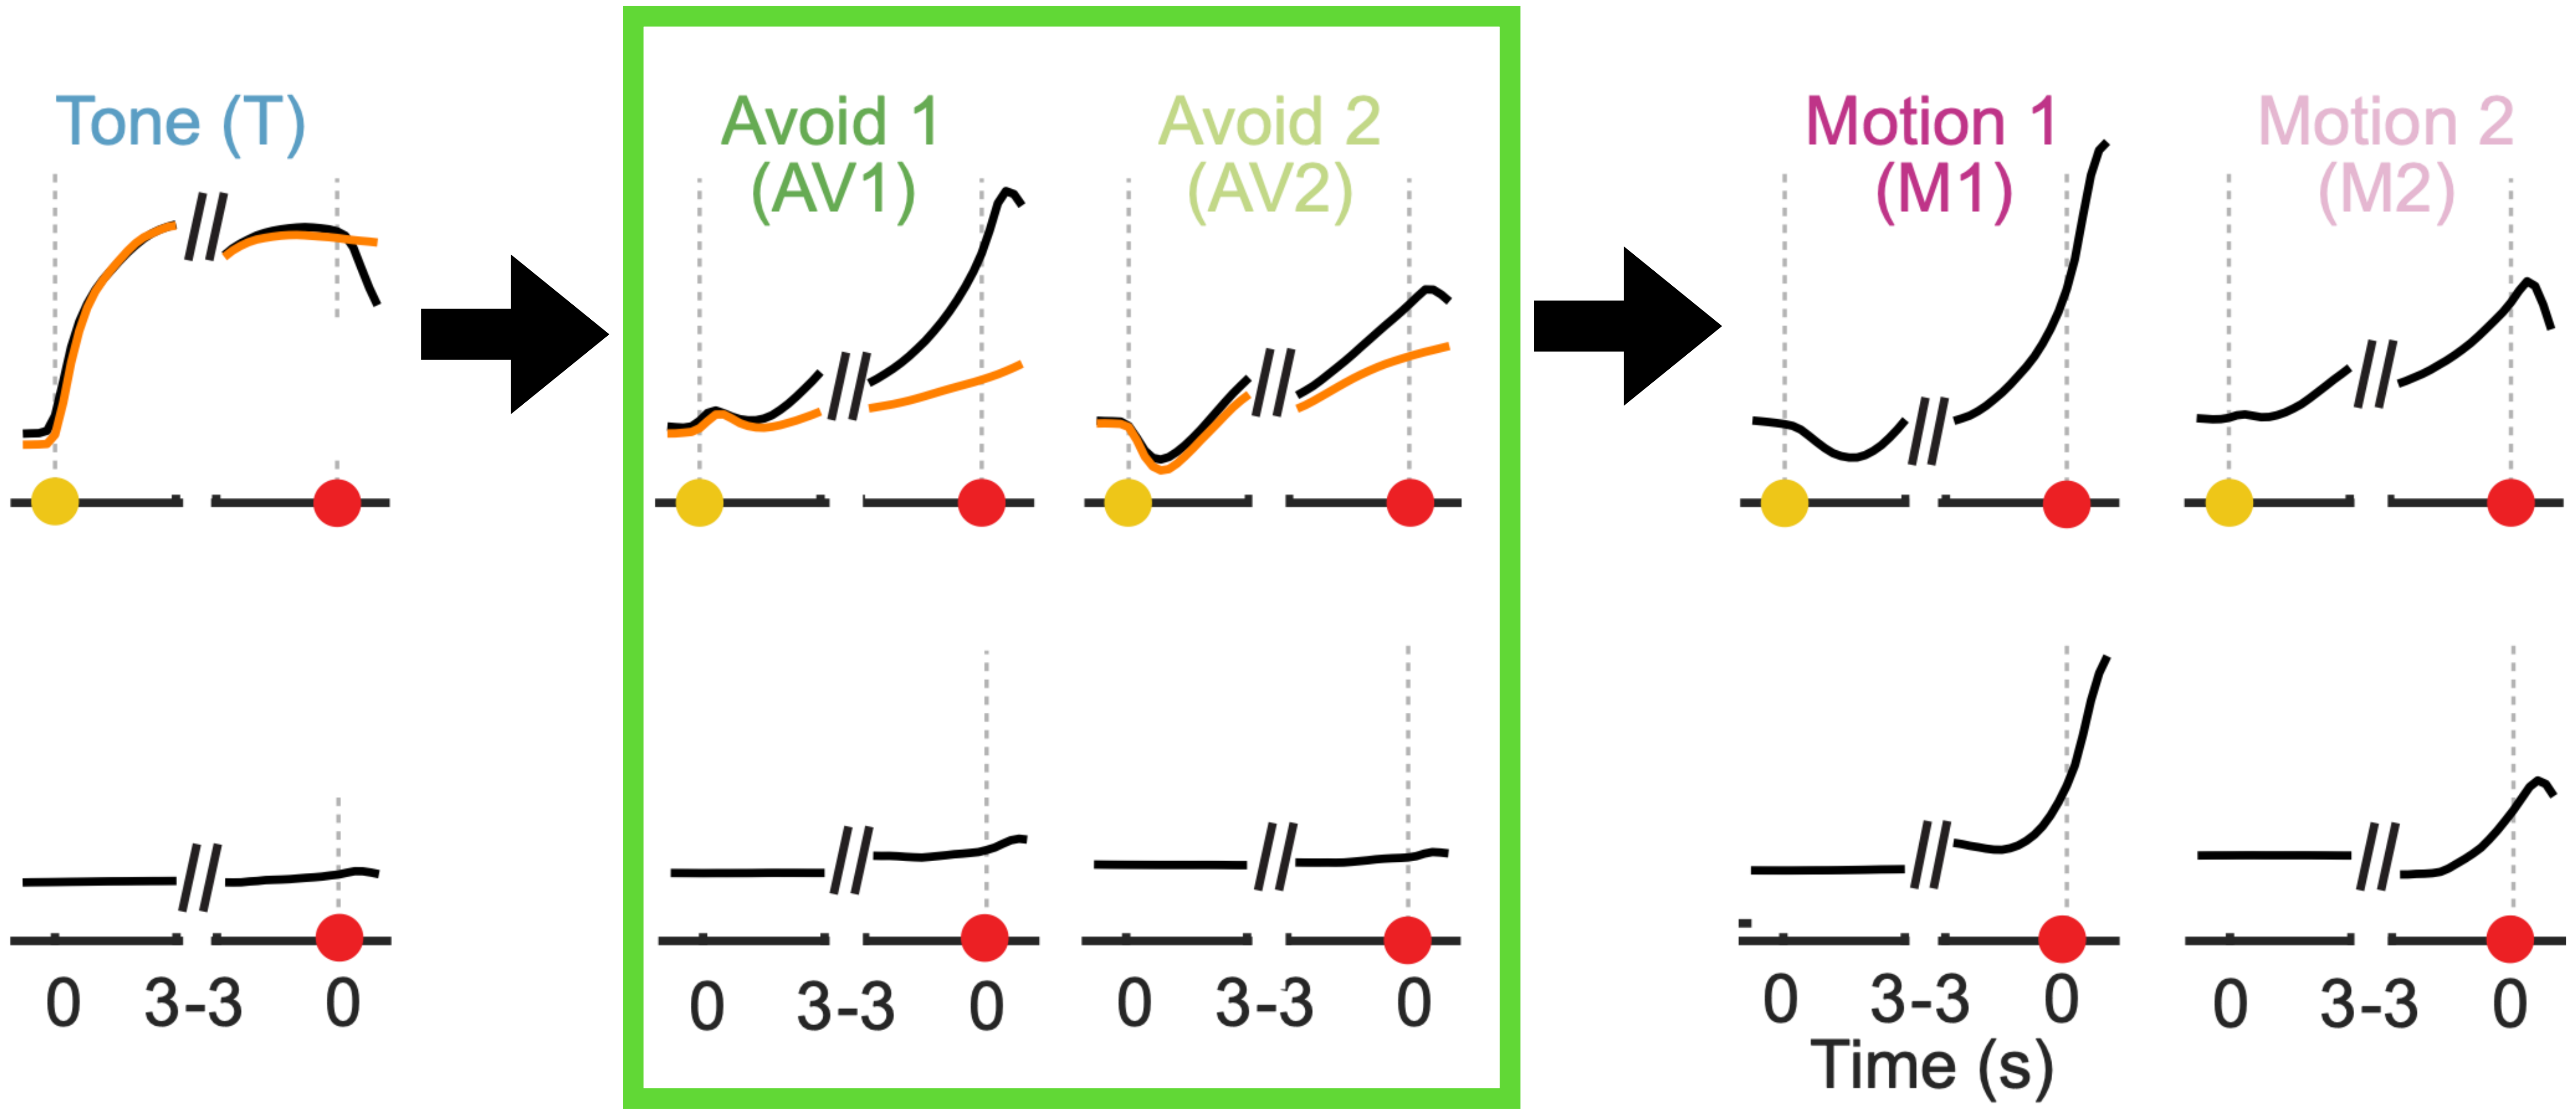

Conclusions Part II: The Nature of Semantic Representations

Brain Visual Pathway

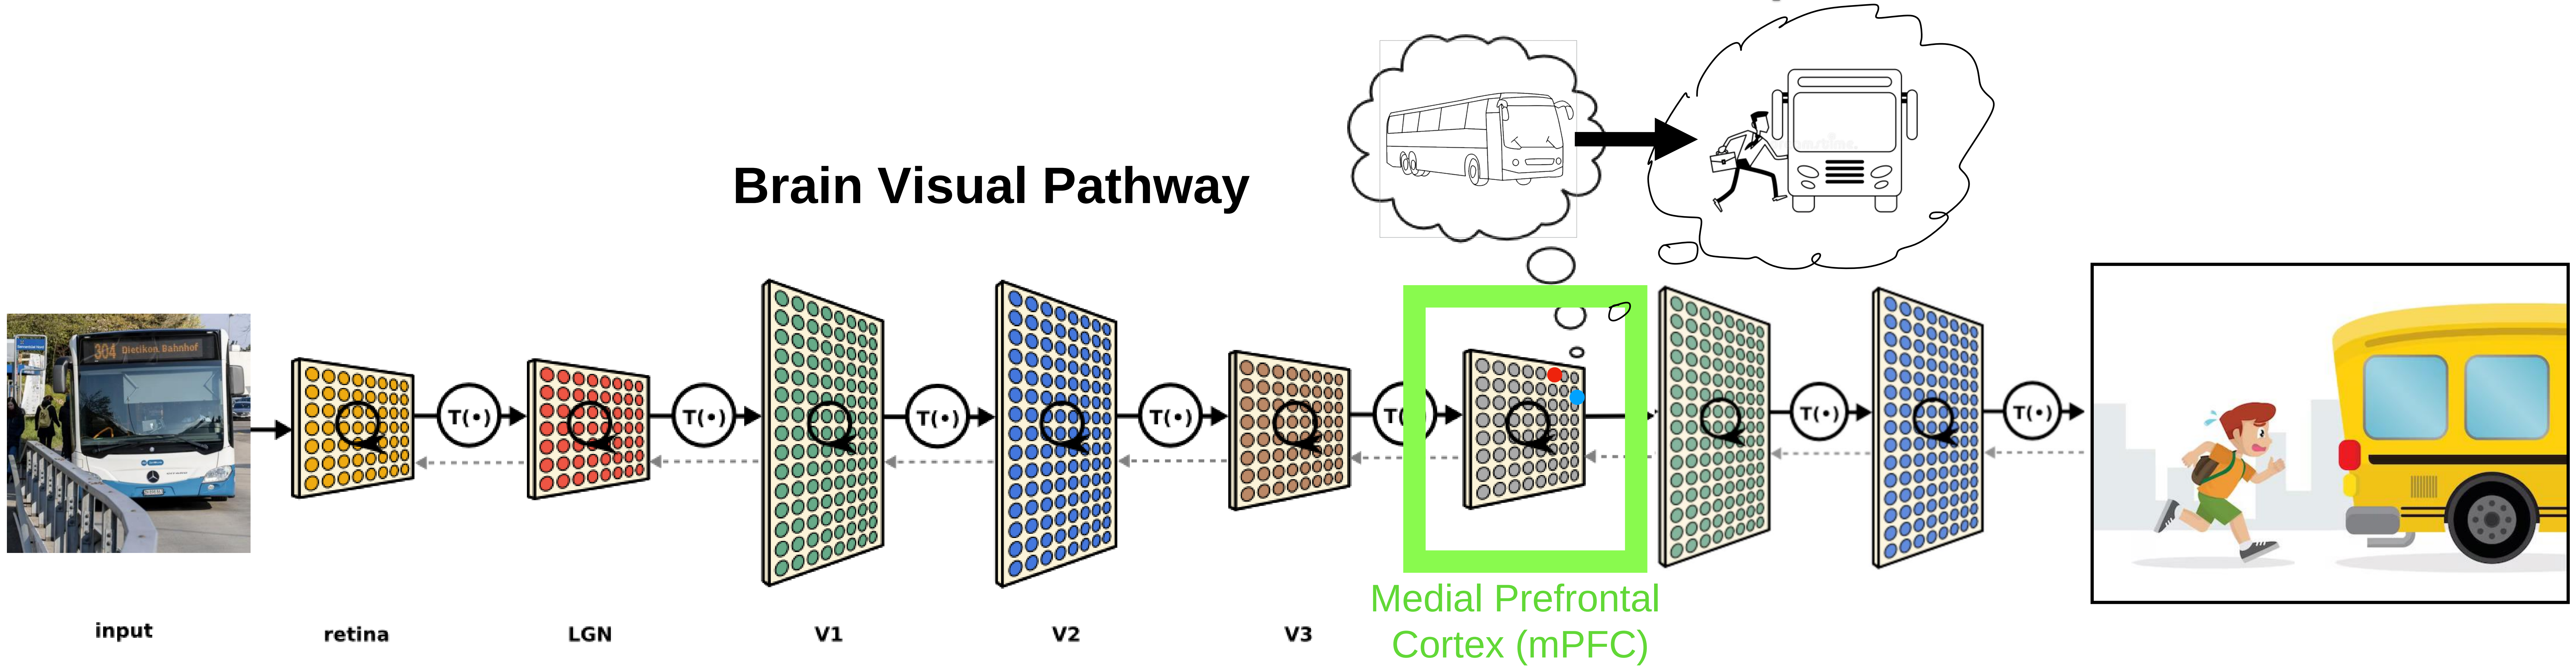

Deep Neuronal Network

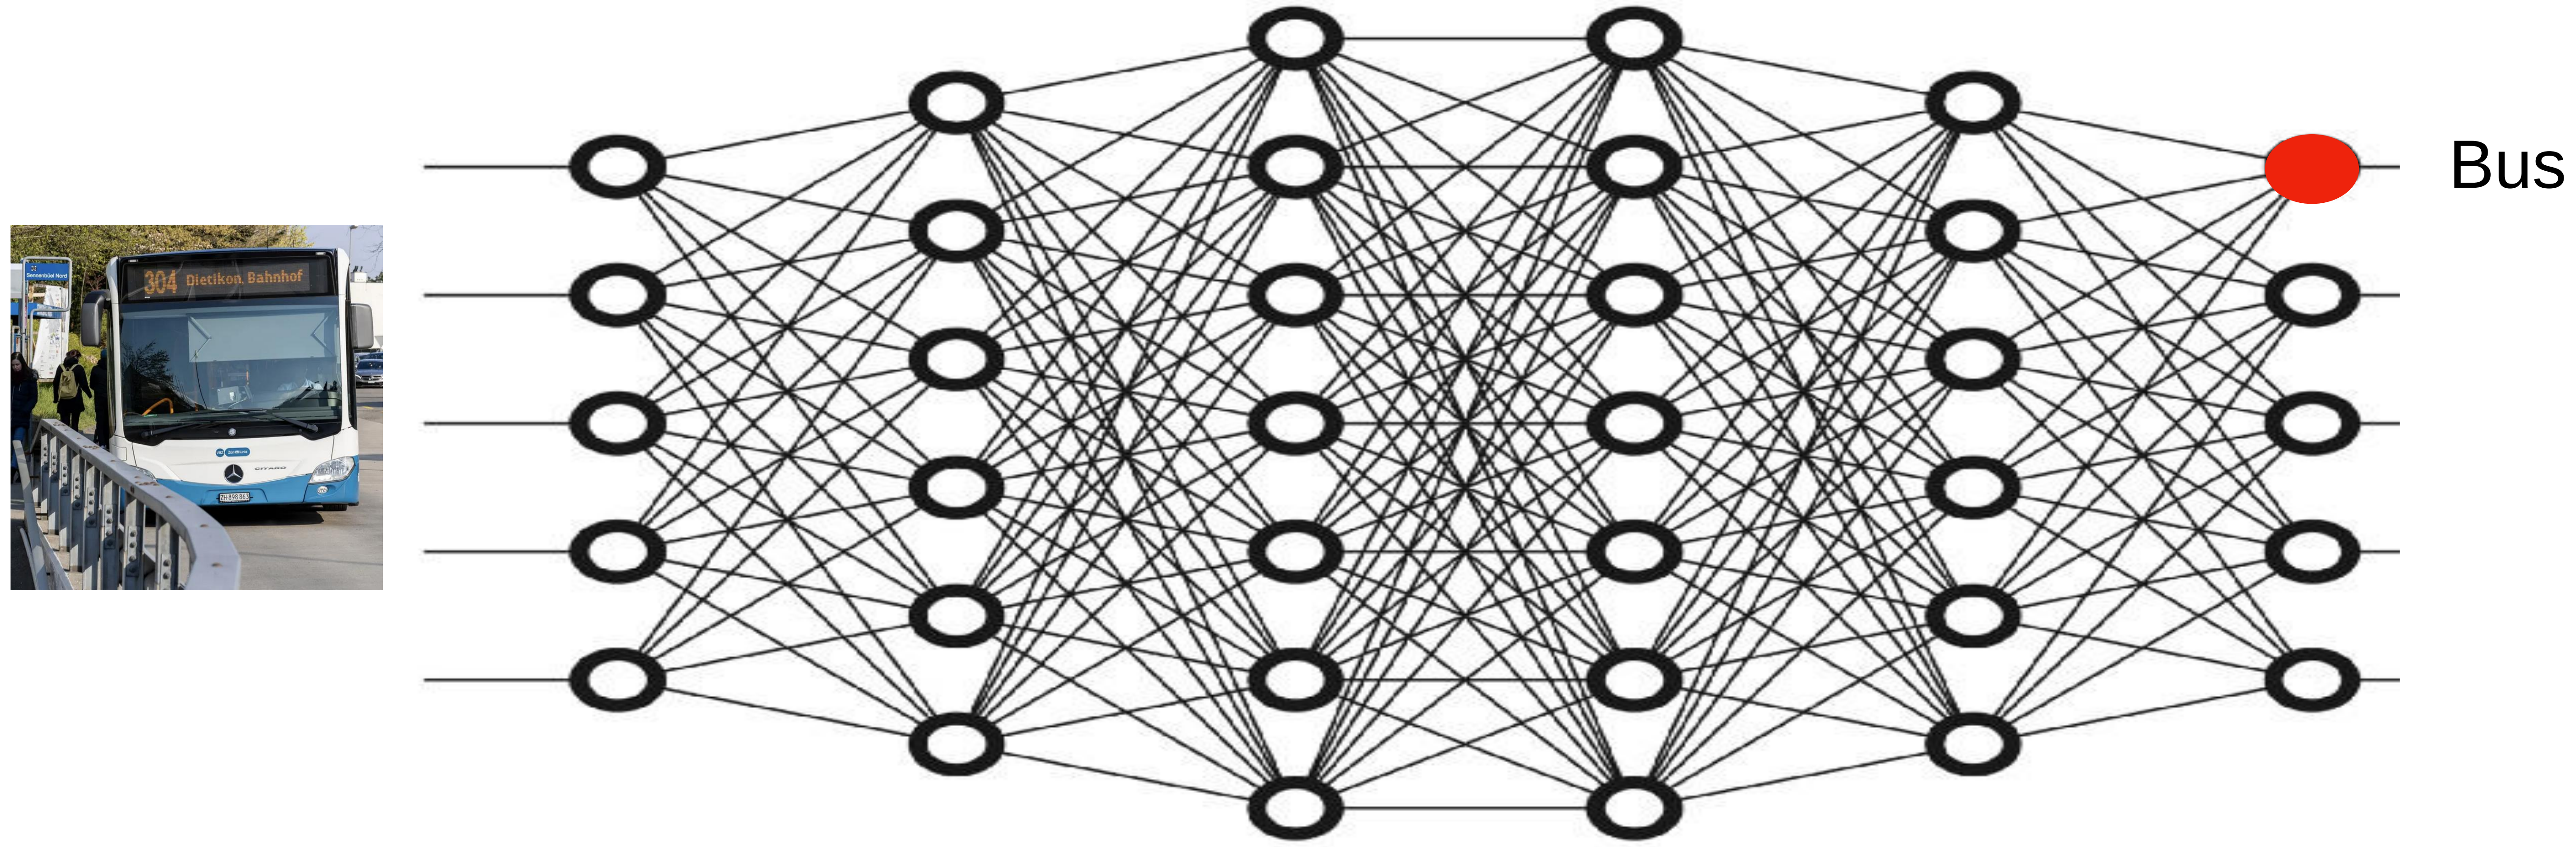

Ehret et al. final revisions @ Nature Neurosc.

Conclusions Part II: The Nature of Semantic Representations

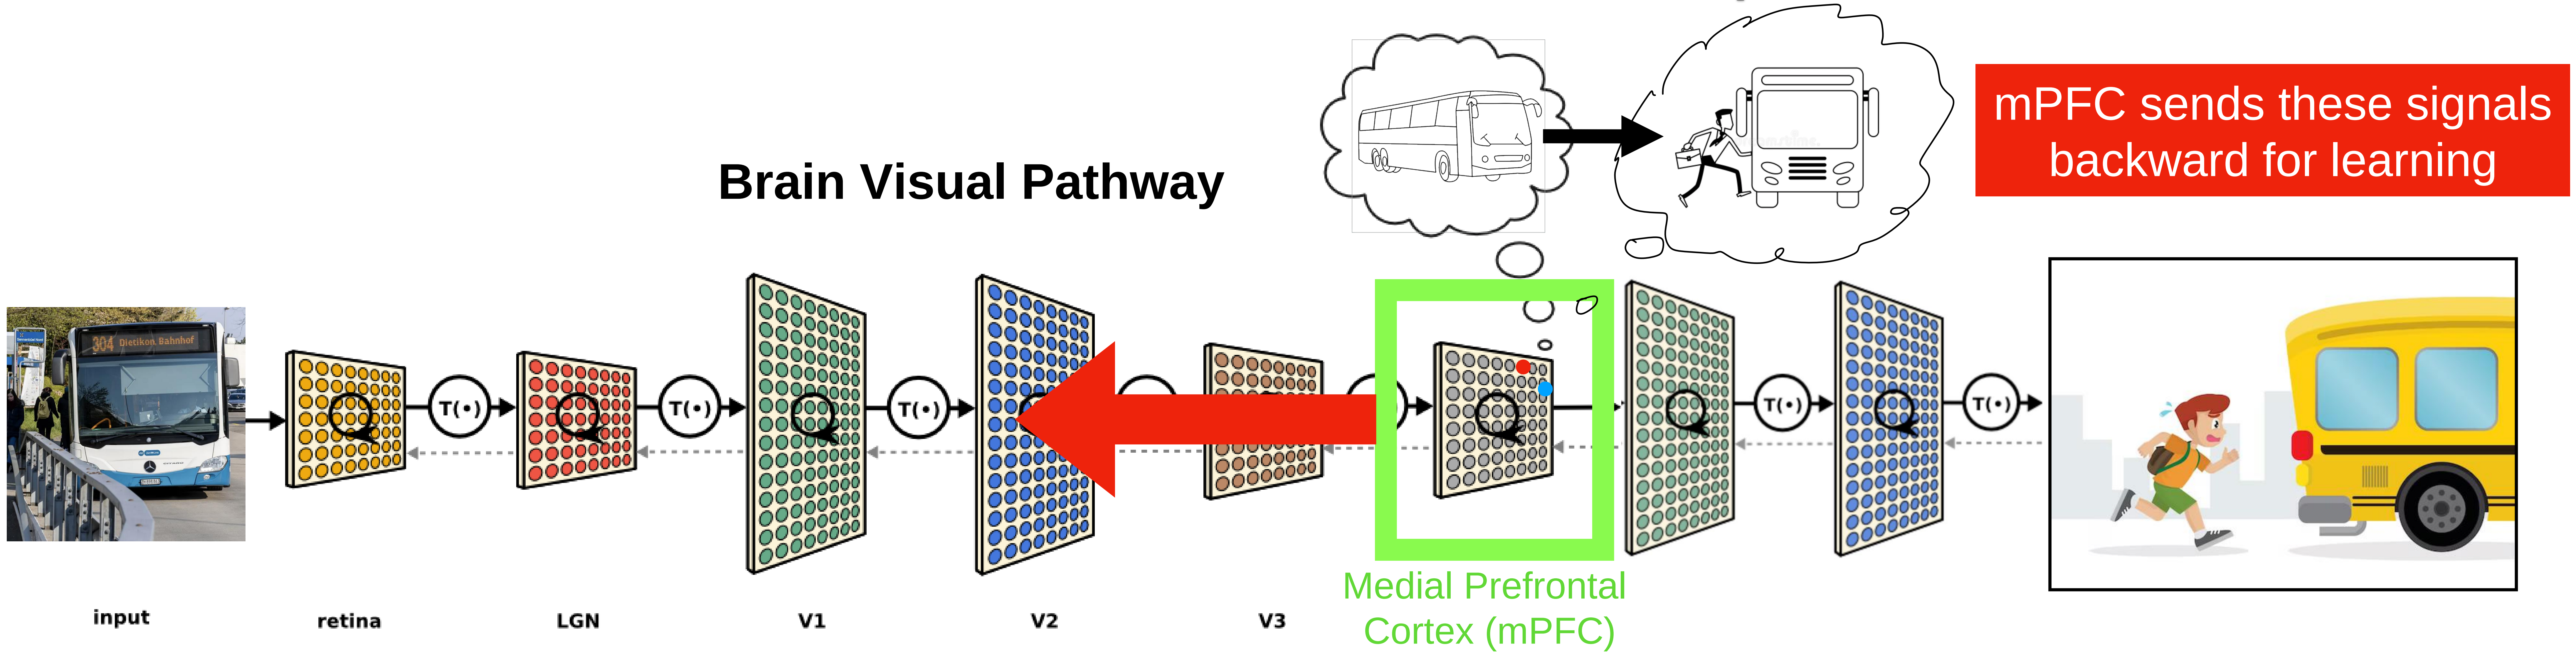

Conclusions Part II: The Nature of Semantic Representations

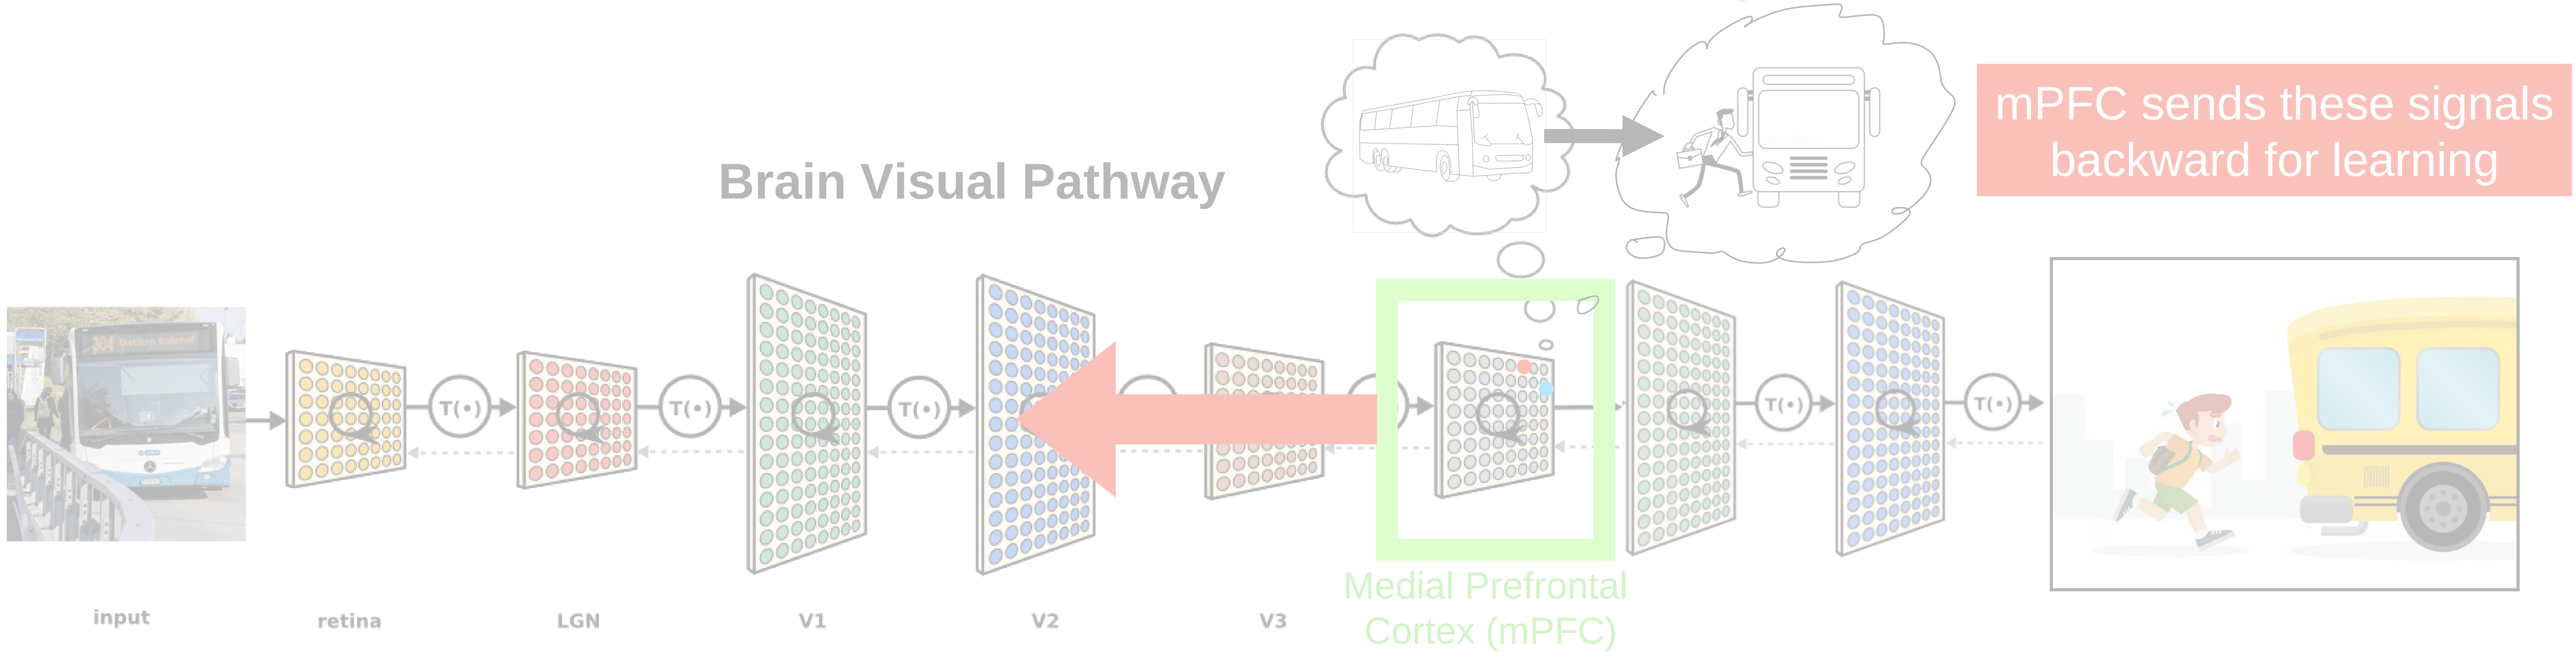

We are using motor signals to supervise the hierarchical learning of objects.

ZUKUNFTSBLOG · DIGITALISIERUNG

KI muss lernen wie ein Kind

Neue KI-Systeme verblüffen, doch Menschen reichen sie nicht das Wasser. Benjamin Grewe plädiert deshalb dafür, dass intelligente Maschinen von morgen so lernen wie kleine Kinder.

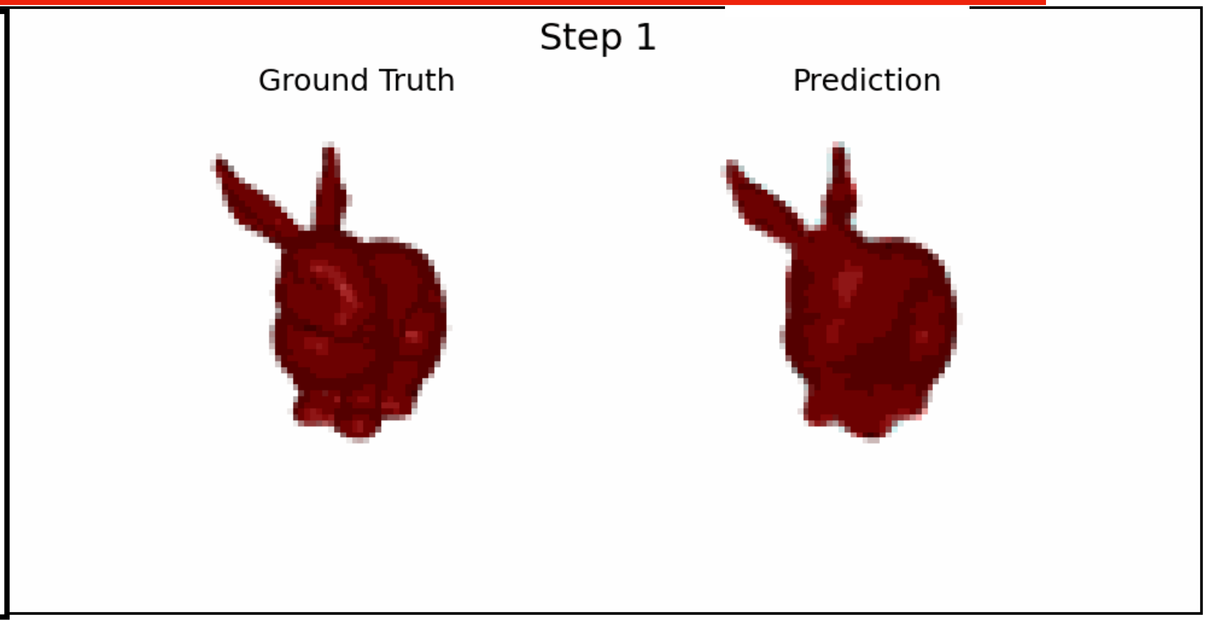

Hamza Keurti

Keurti et al., 2023, ICLM

ETH zürich

Developing Bio-Inspired AI Technologies that generate Behaviour

institute of neuroinformatics

Prompt:

- Wipe the table with thin cloth.
- Pick cube and place in bowl.
- Pick lying bottle from bowl.
- Pick coffee cup from pedestal
- Push bowl around the table.
- Pick standing bottle from pedestal.
- Pick cube and place in bowl, with distractors.

Camera Image
Table Scence

Robotic Action Transformer (RAT)

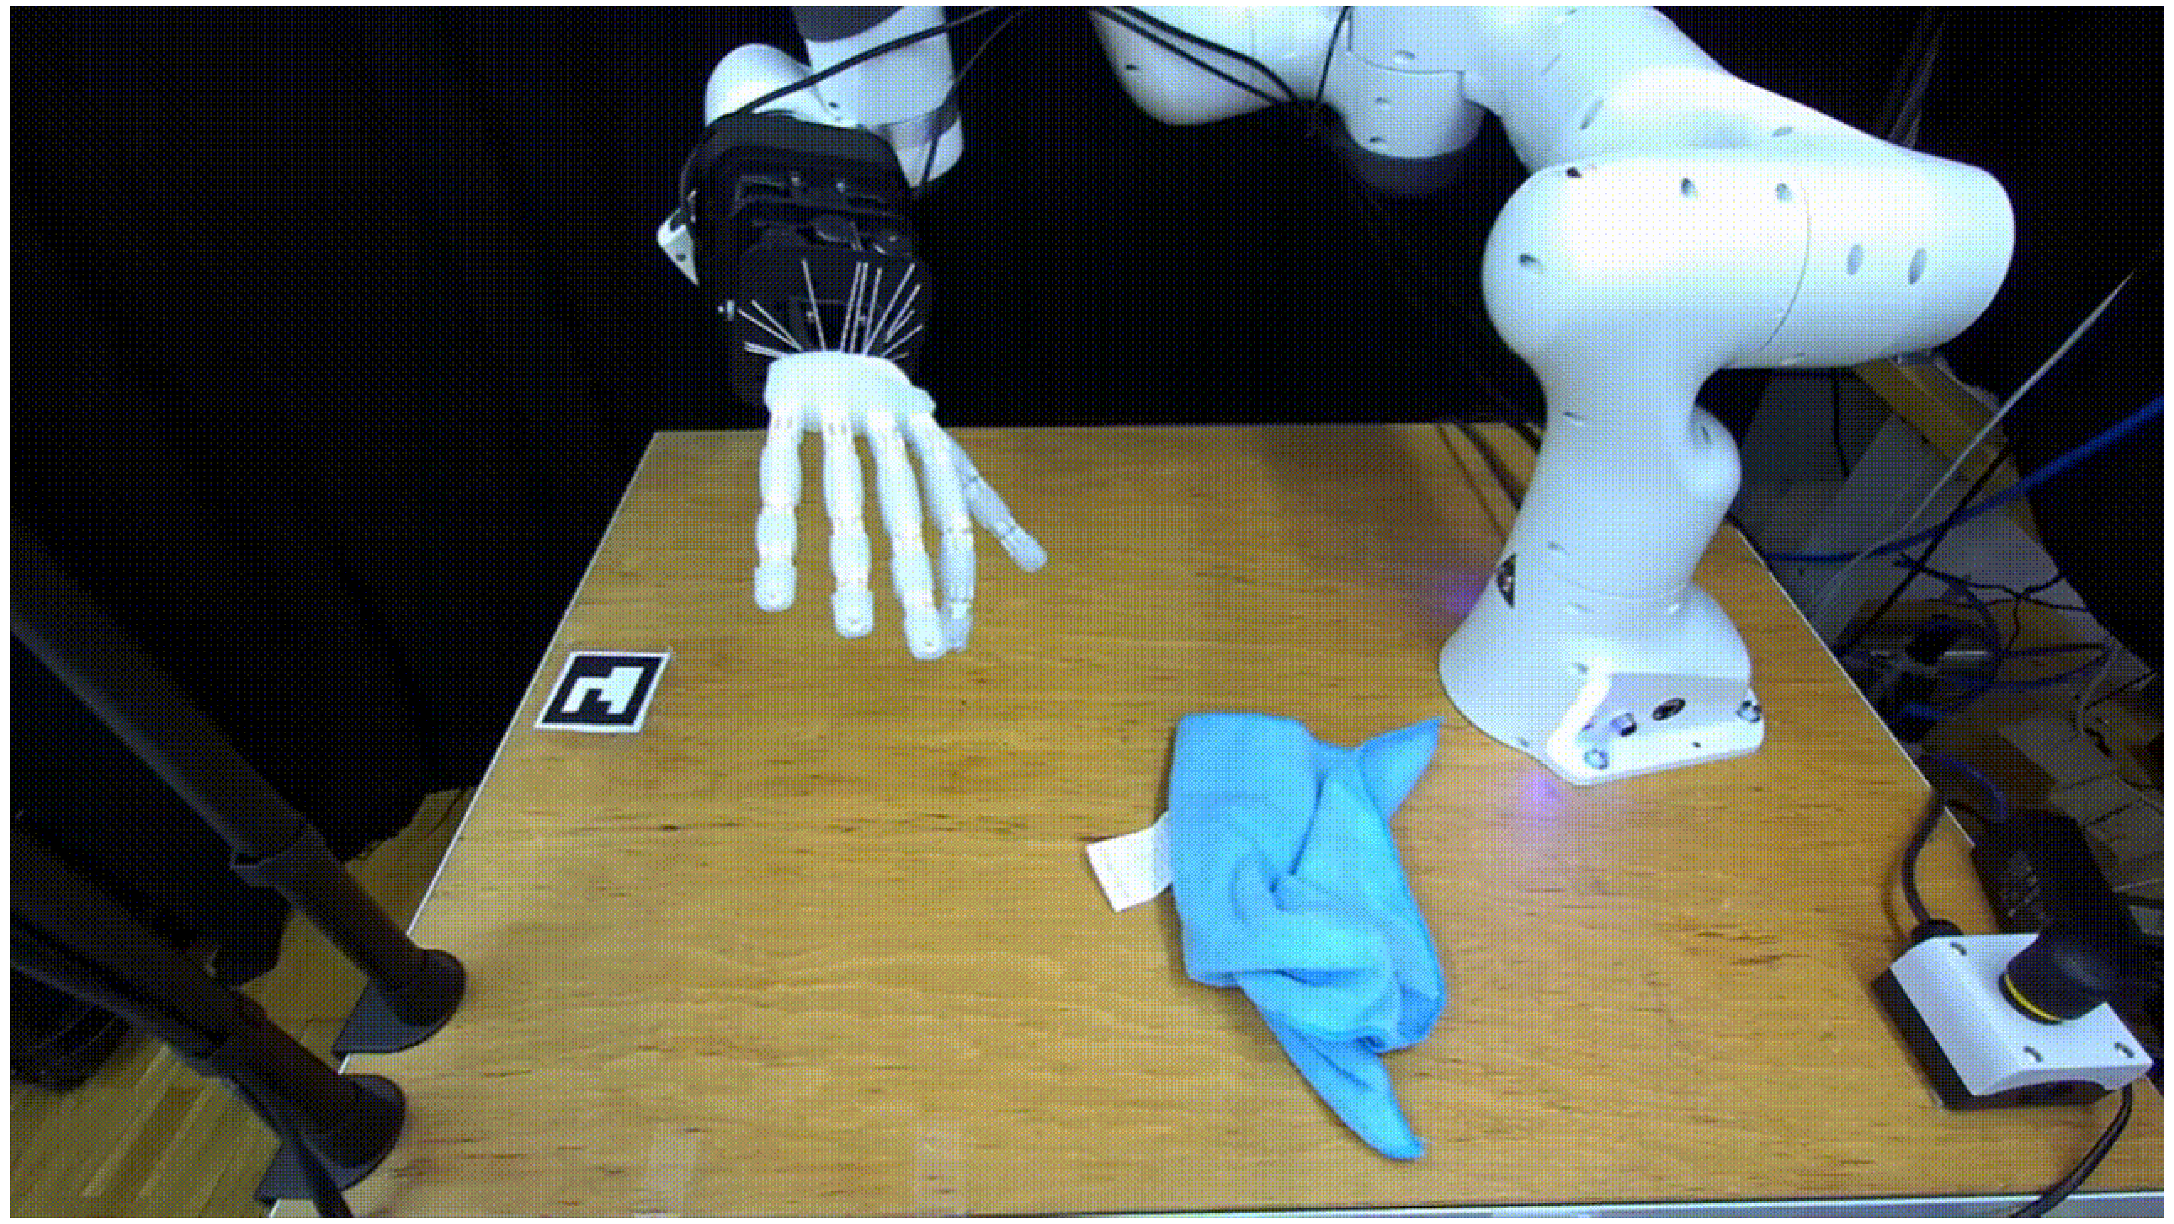

Movie Credit Elvis Nava

Developed @ ETH AI CENTER

ChatGPT:	Prompt to Prompt
Dalle:	Prompt to Image
RAT:	Prompt to Action

Nava et al., 2023, TMLR

ETH zürich

Developing Bio-Inspired AI Technologies that generate (virtual) Behaviour

ALPINEAI

Industry Partner of

ETH AI CENTER

ETH zürich

Developing Bio-Inspired AI Technologies that generate (Virtual) Behaviour

ALPINEAI

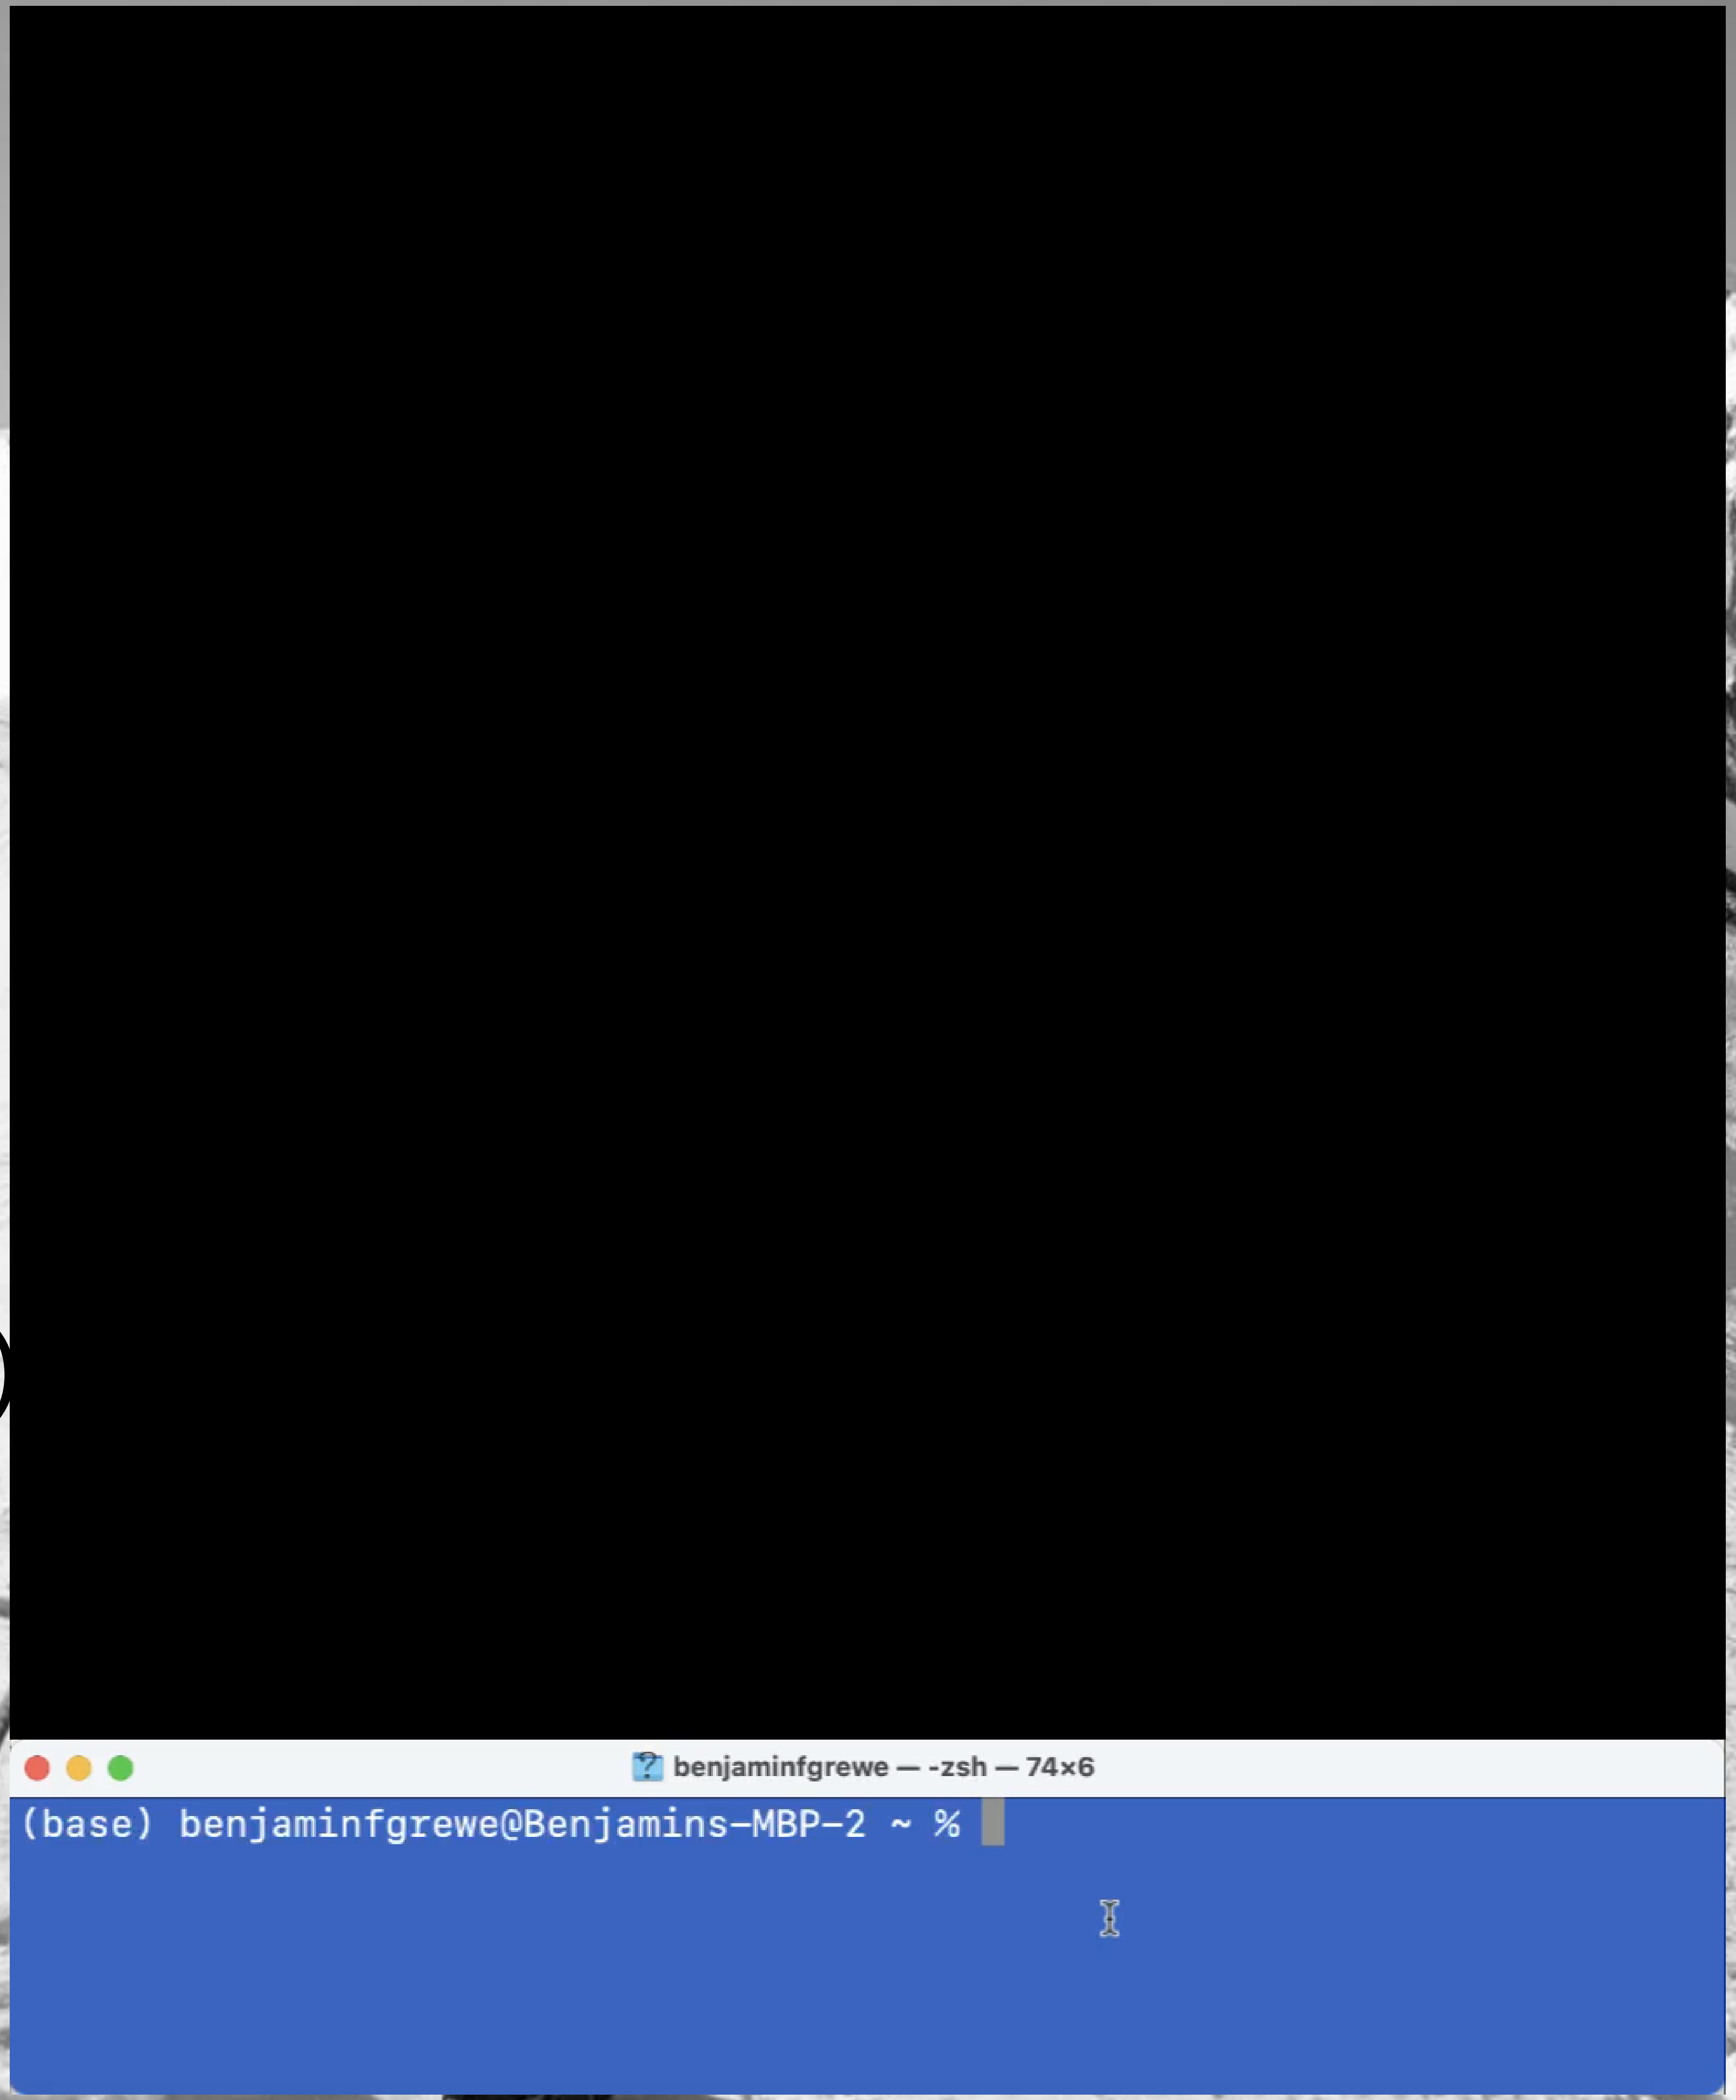

Virtual Action Transformer (VAT)

Industry Partner of

ETH AI CENTER

Summary: Advancing Neuroscience and AI

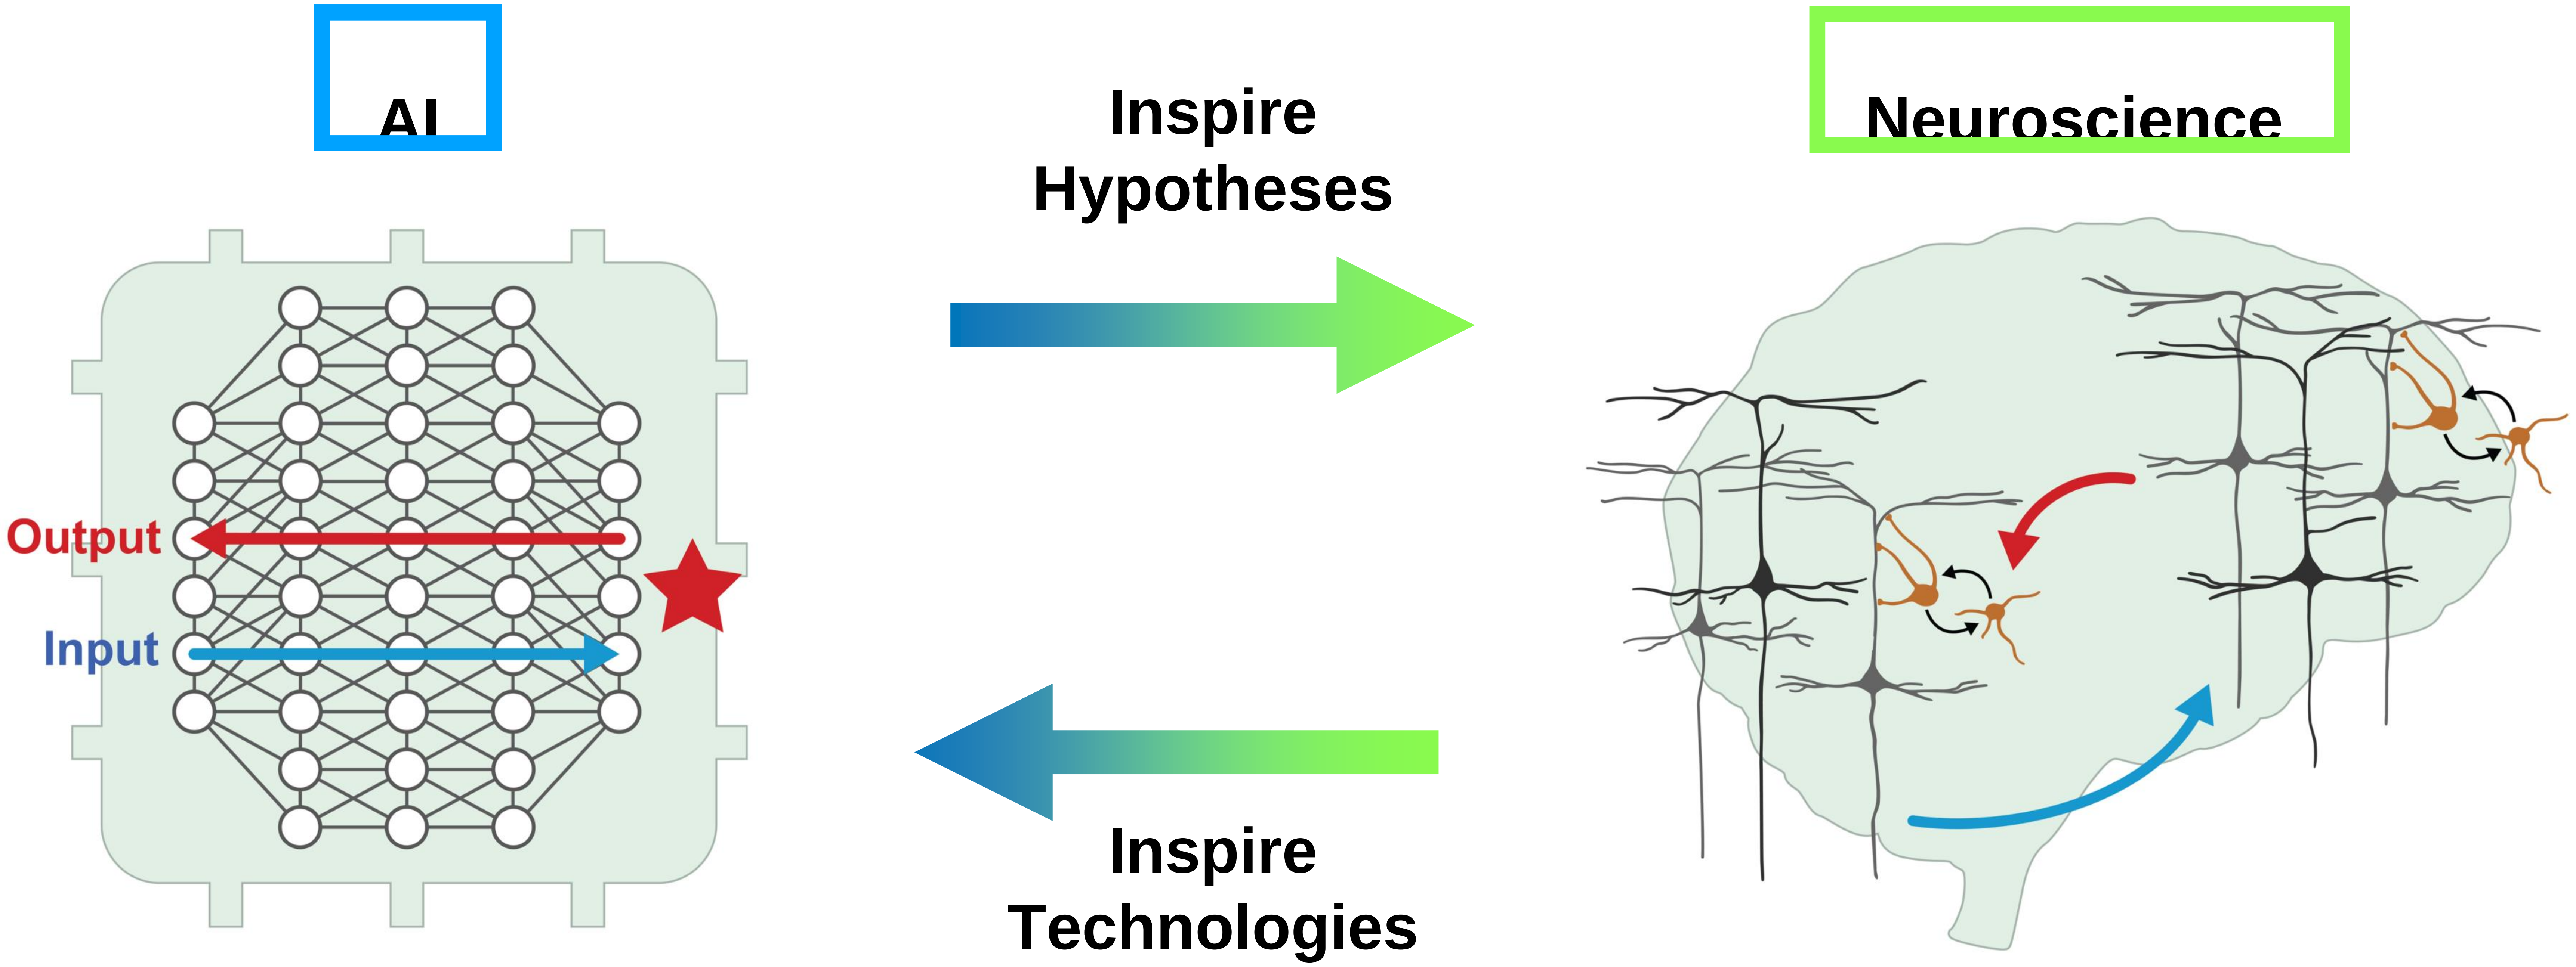

Artificial Intelligence

Natural Intelligence

goal driven behaviour (like a human).

representations that allow the generation of actions.

ETH zürich

Acknowledgements

institute of neuroinformatics

Part I: Learning in Hierarchical (Deep) Cortical Networks

AI → **Neuroscience**


Pau Aceituno


Matilde T. Farinha


Alexander Meulemans


Francesco Laessig

Part II: Understanding Hierarchical Neuronal Representations in Brain

AI ← **Neuroscience**


Benjamin Ehret


Liz Ann Amadei


Roman Boehringer


Valerio Mante

ETH zürich

Acknowledgements

institute of neuroinformatics

benjamin@alpineai.ch

bgrewe@ethz.ch