EPFL

Neuroscience and Machine Learning

Alexander Mathis
alexander.mathis@epfl.ch

Phys-754 Lecture series on scientific machine learning
September 26, 2024

Part 2

ML models for neuroscience

Biological Intelligence ⟷ Artificial Intelligence

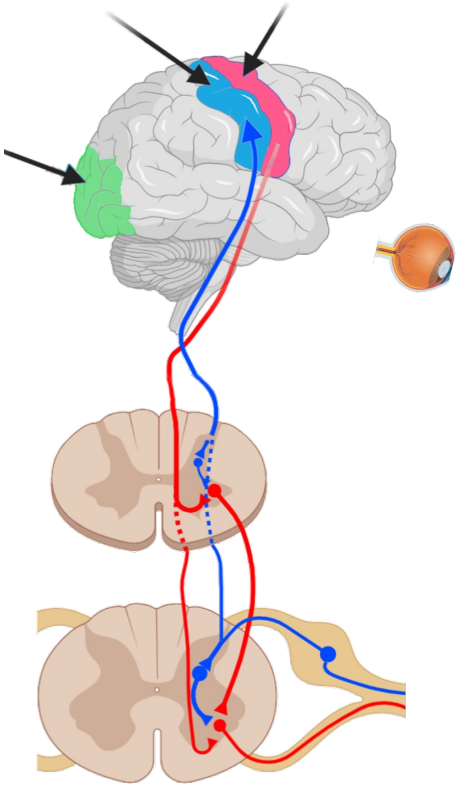

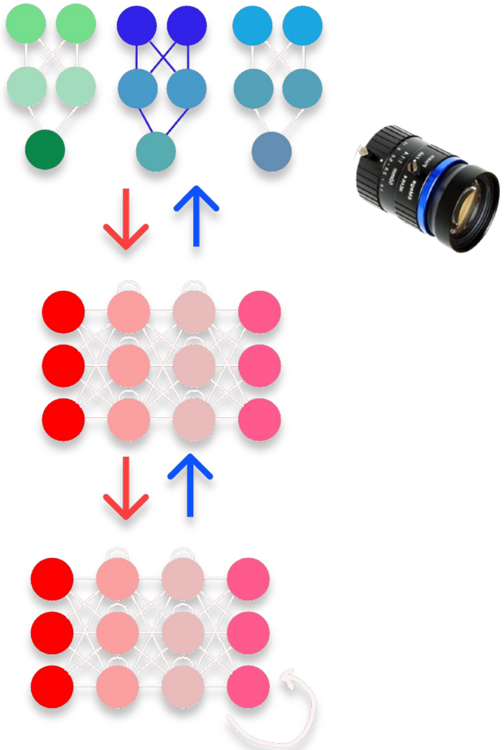

Hausmann*, Marin-Vargas*, Mathis**, Mathis**, Curr op. in Neuro., 2021

Temporal and spatial scales in neuroscience

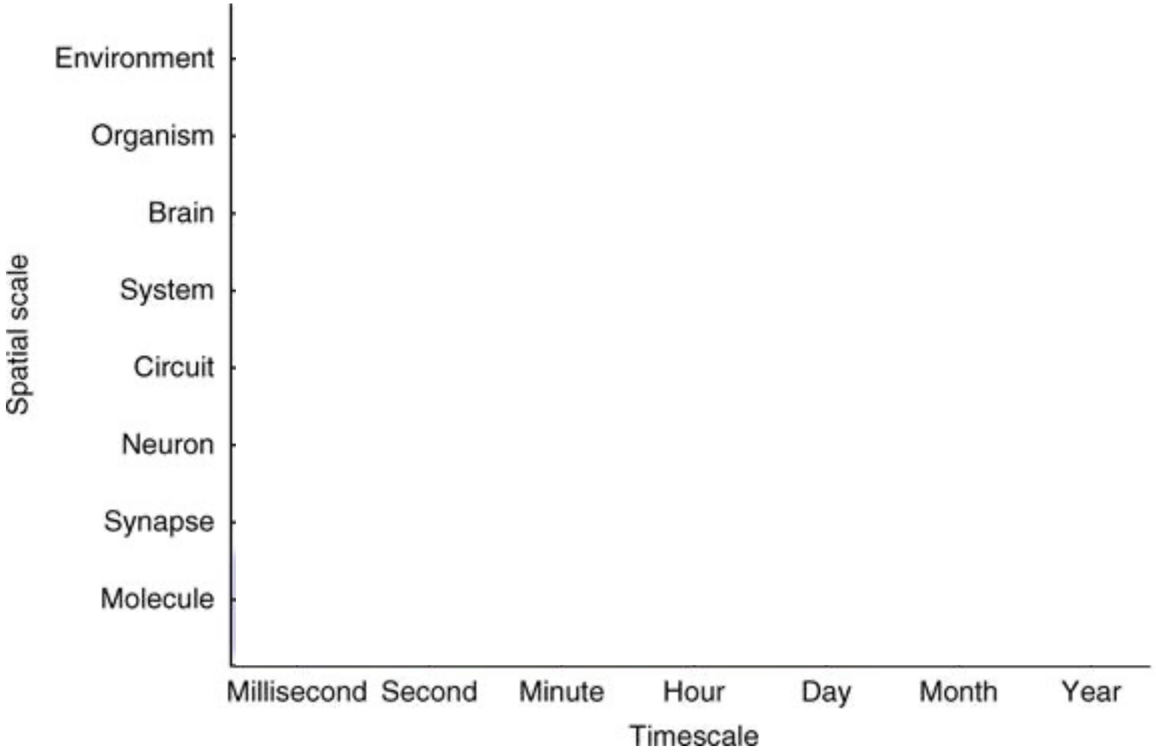

Basset, Nat Neuro 2017

Temporal and spatial scales in neuroscience

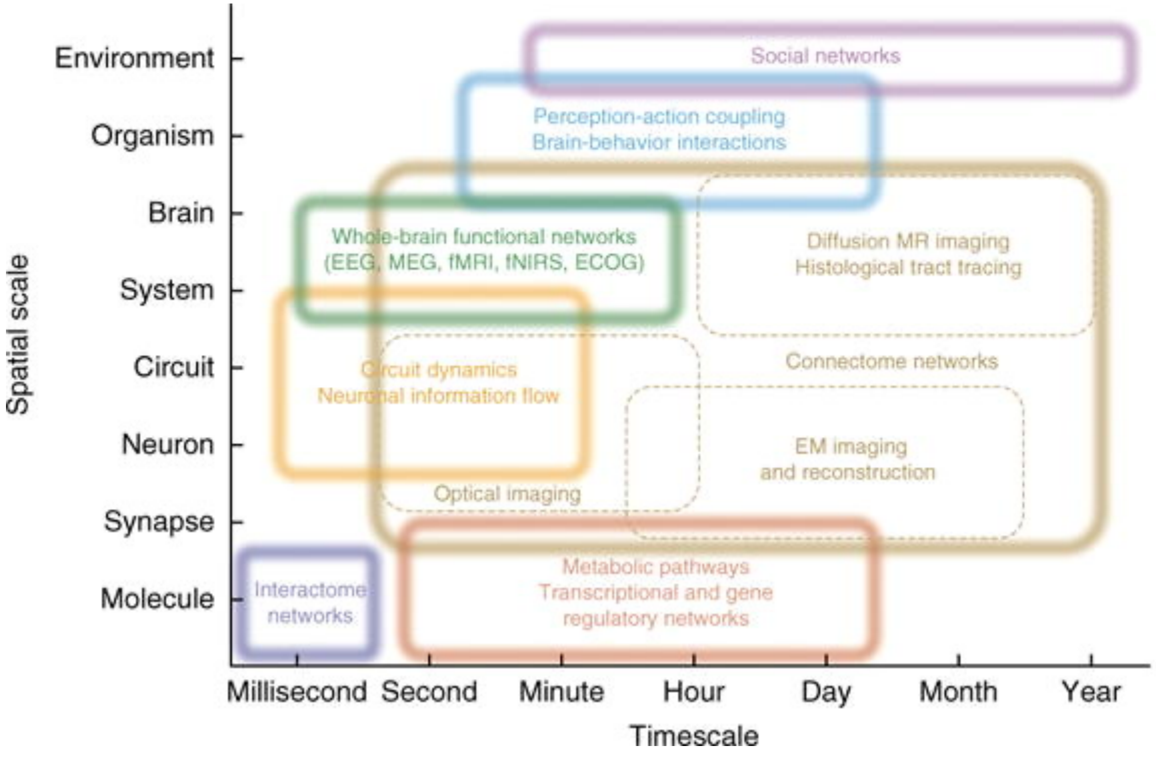

Basset, Nat Neuro 2017

Task-driven modeling: linking behavior to circuits

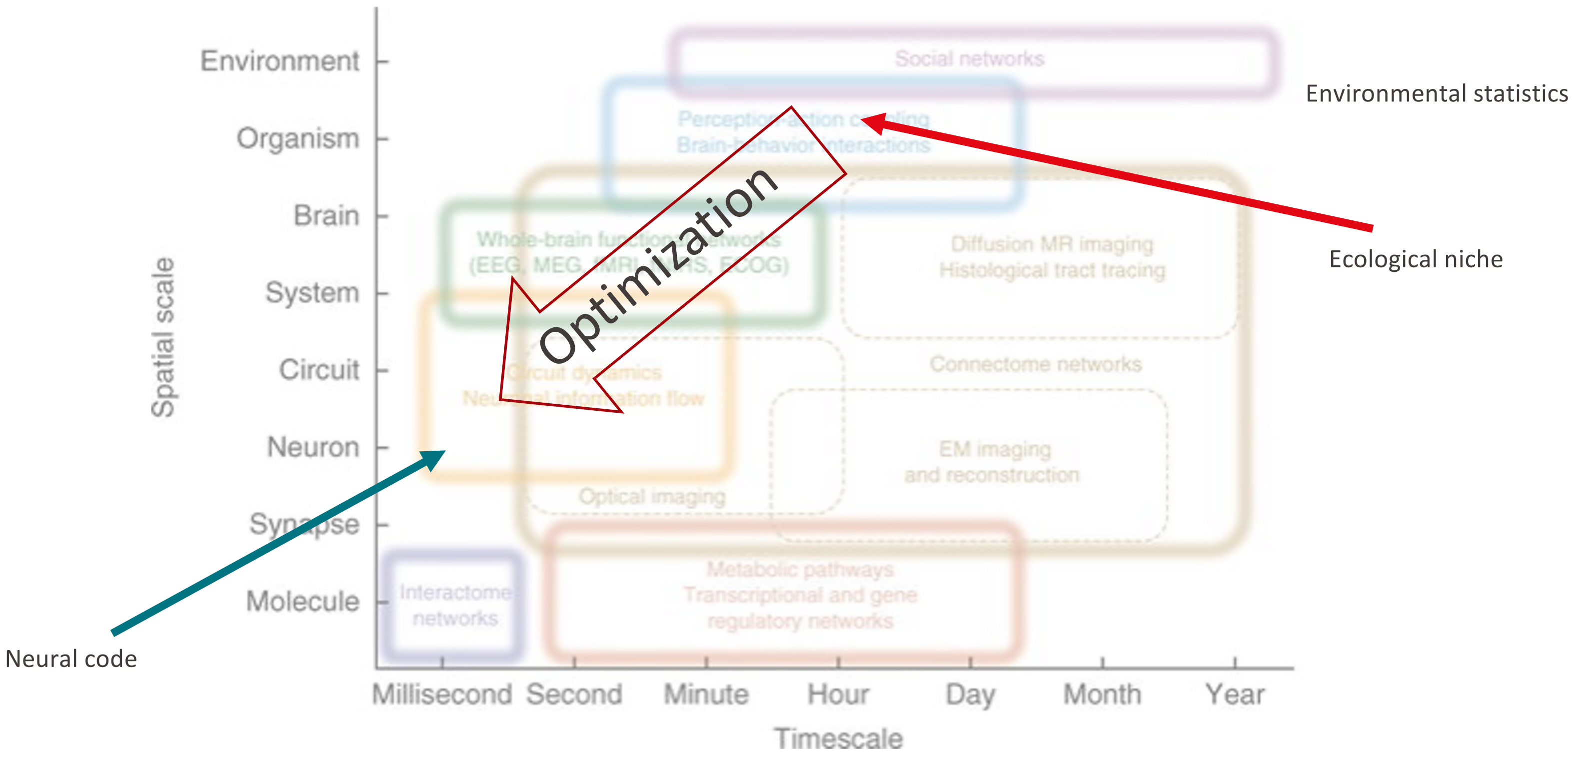

Machine learning

Architecture

Task / objective

Dataset

Optimization method (learning rule)

ML model

Machine learning	Neuroscience
Architecture	Circuits
Task / objective	Ecological niche
Dataset	Environment
Optimization method (learning rule)	Natural selection + synaptic plasticity

ML model

Using deep neural networks as goal-driven models of a system

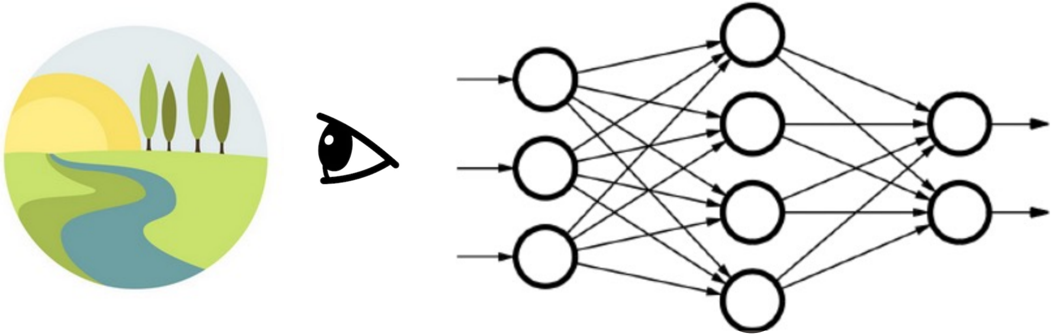

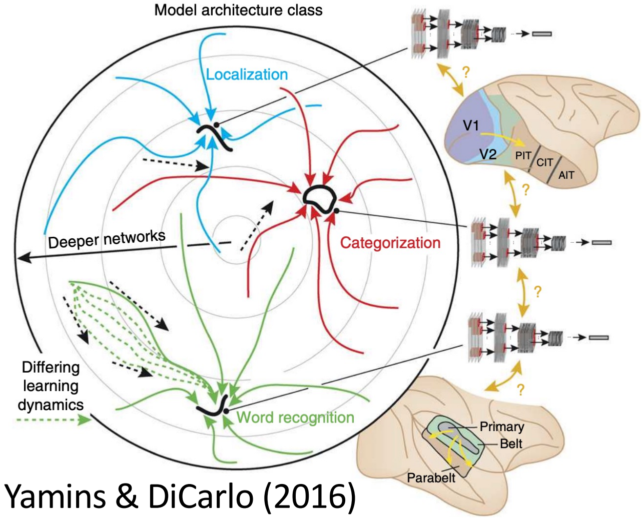

Yamins & DiCarlo (2016)

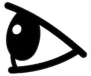

Vision: Yamins et al. (2014) -> object recog.

Audition: Kell et al. (2018)- speech recognition, speaker identification, natural sound identification

Barrel Cortex: Zhuang et al. (2017)

Cognition: Mante et al. (2013)

Proprioception: Sandbrink et al. (2023)

Ventral visual pathway

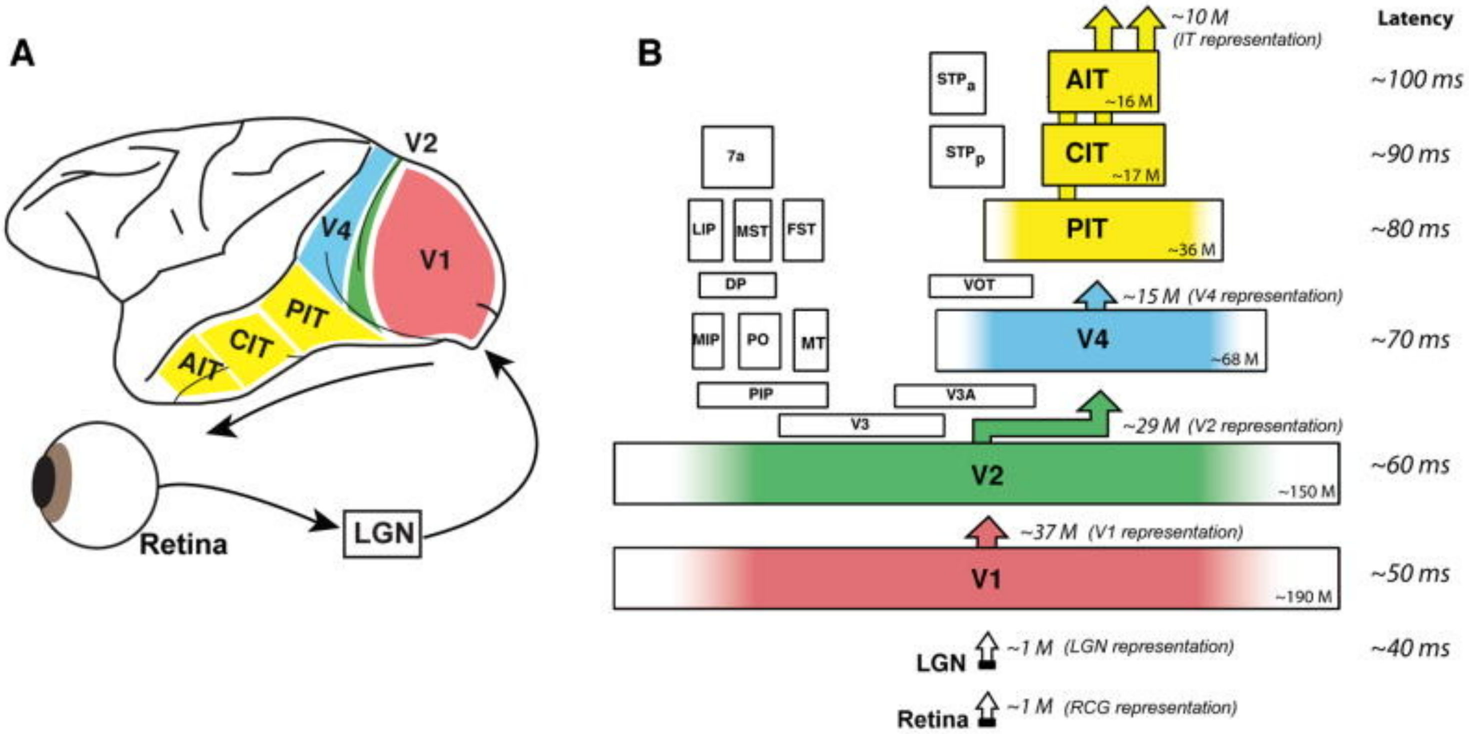

DiCarlo, Neuron 2012

IT neurons are nonlinear

Example neuron

ı)

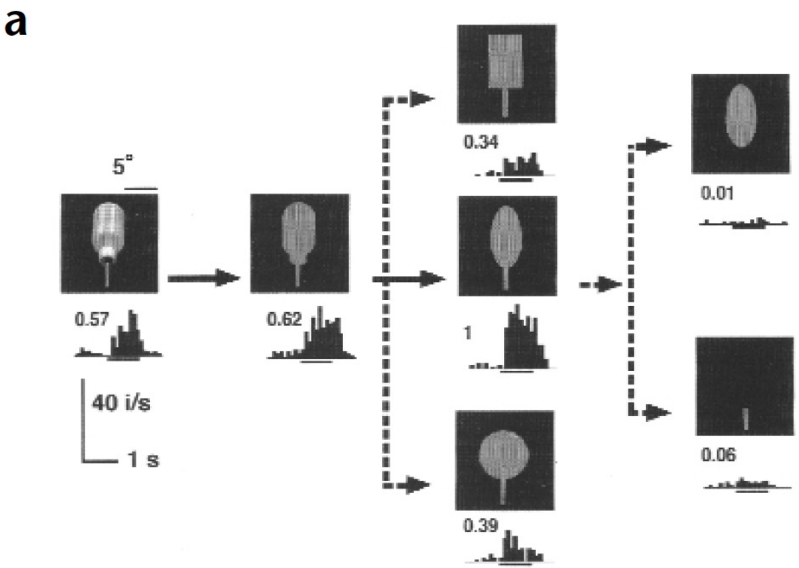

Riesenhuber & Poggio, Nat Neuro 1999

Core-object recognition paradigm

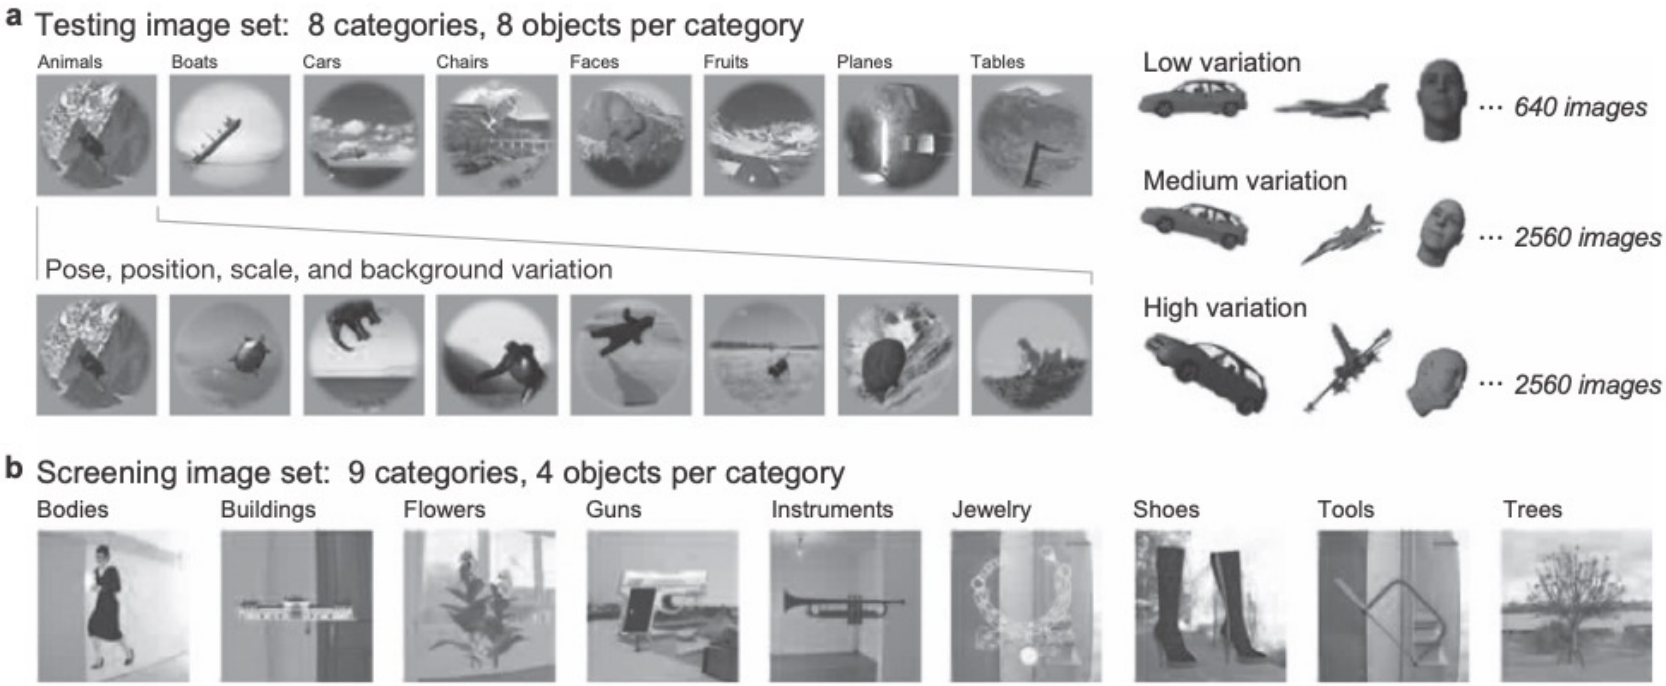

Yamins et al., PNAS 2014

Example higher-order visual cortex response

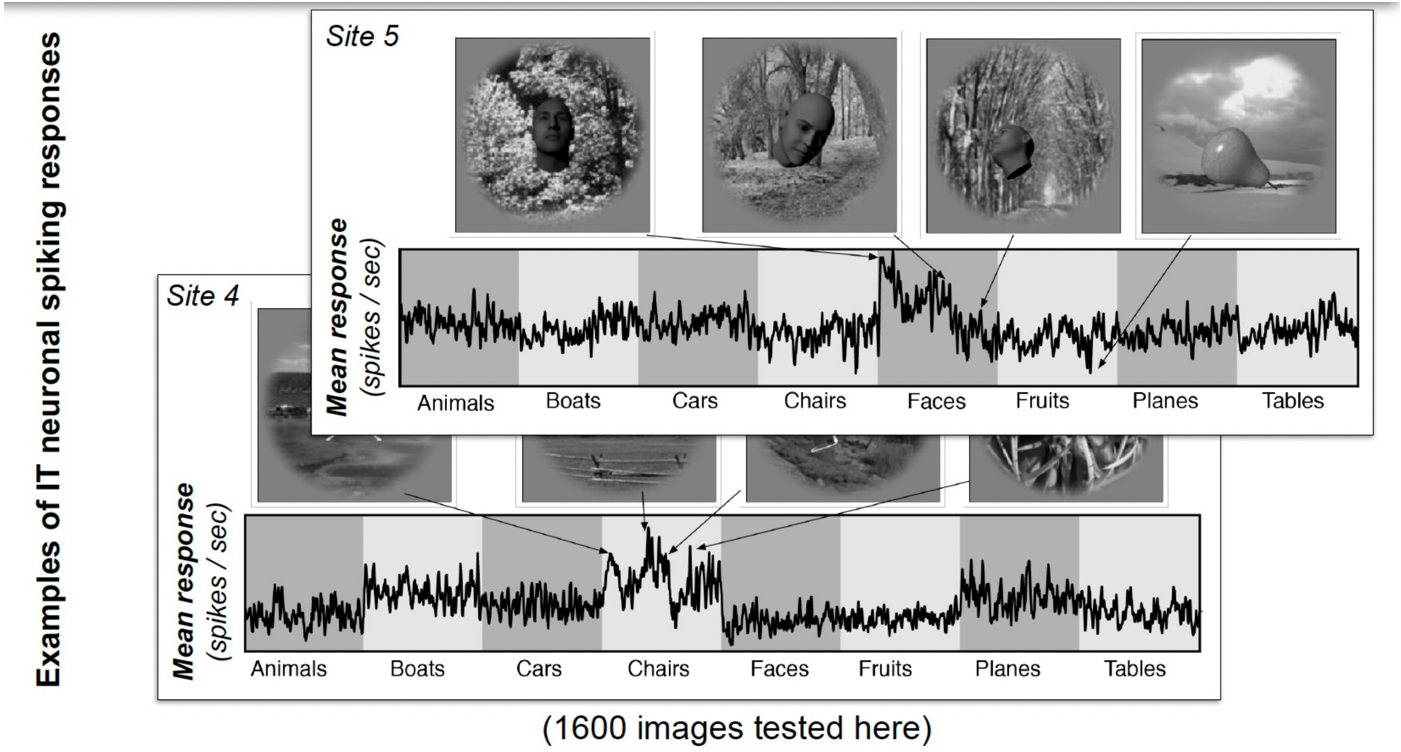

Slide from Jim DiCarlo, MIT

Model building

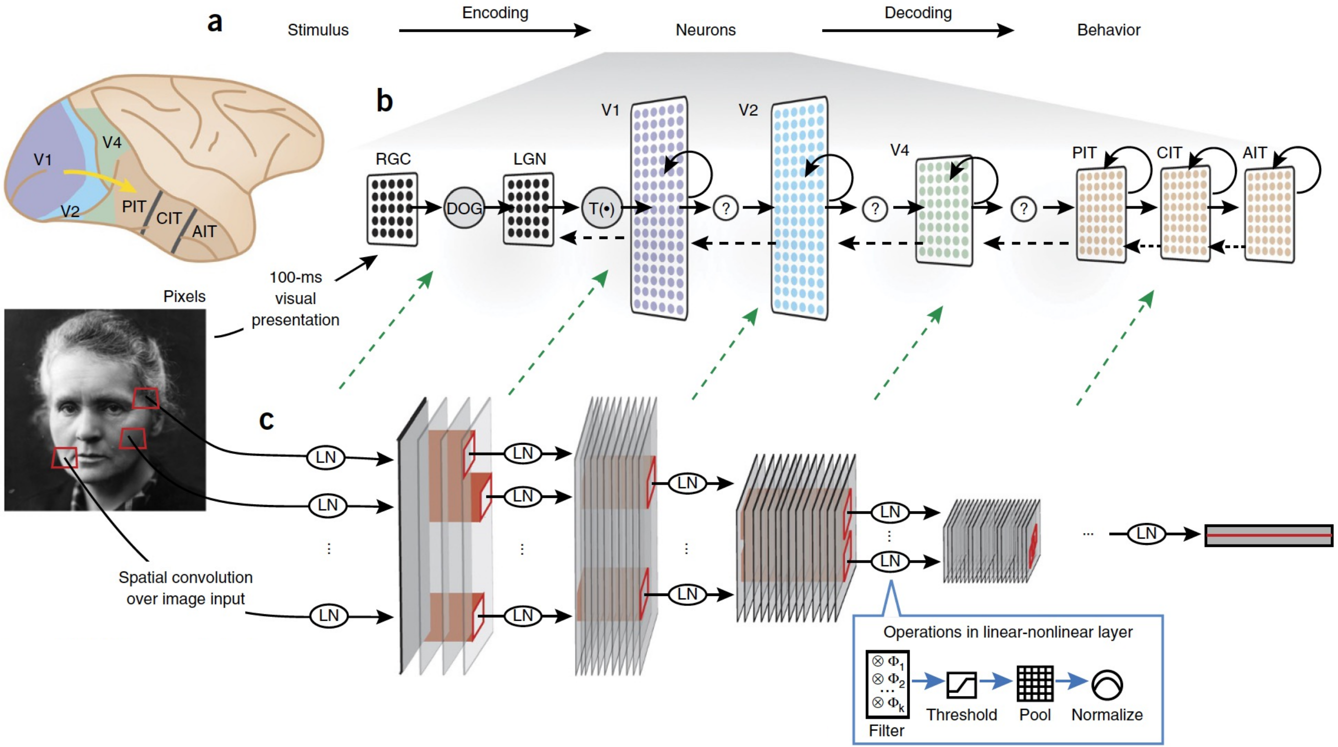

Yamins and Di Carlo, Nat Neuro 2016

Task-trained models, matches brain's "hidden neurons"

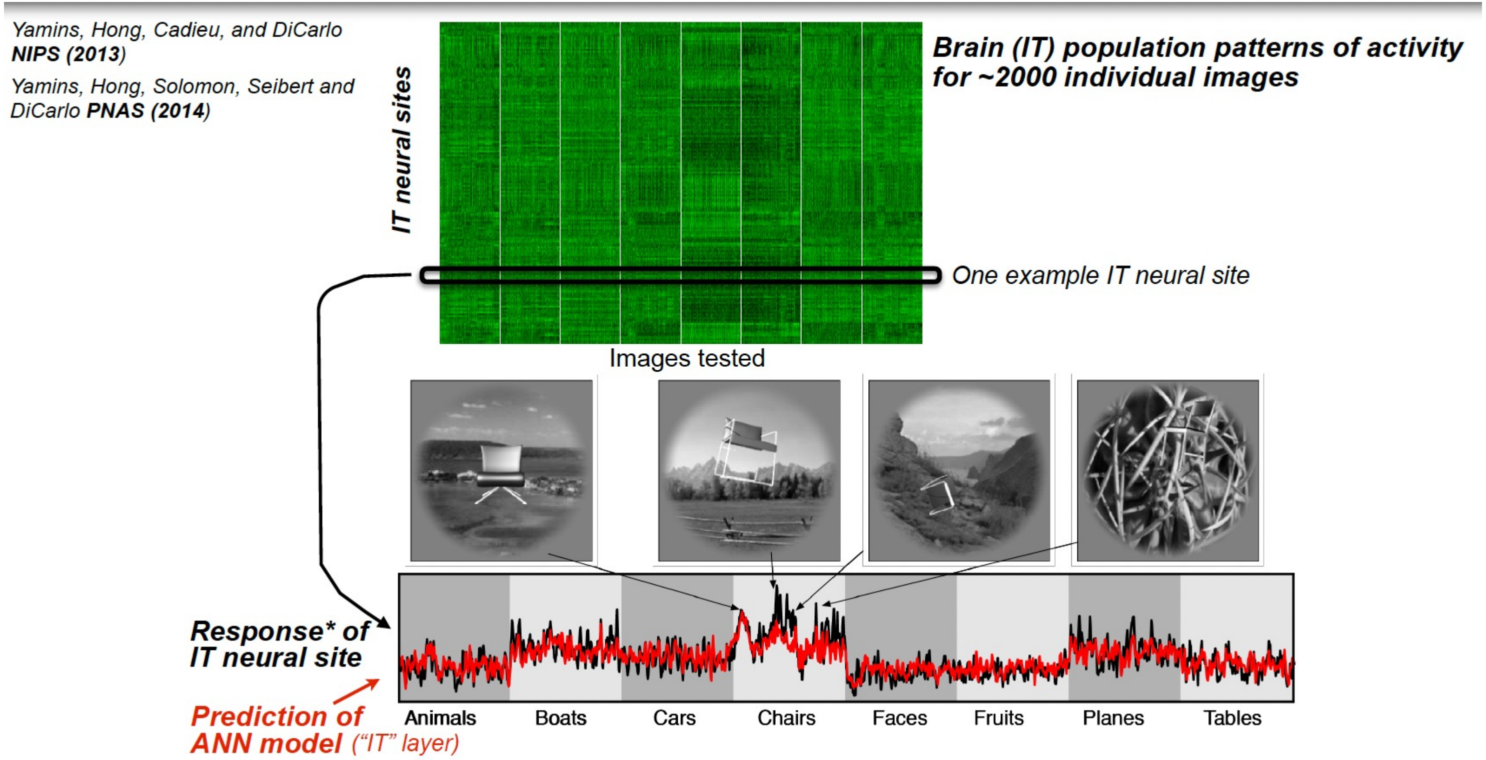

Slide from Jim DiCarlo, MIT

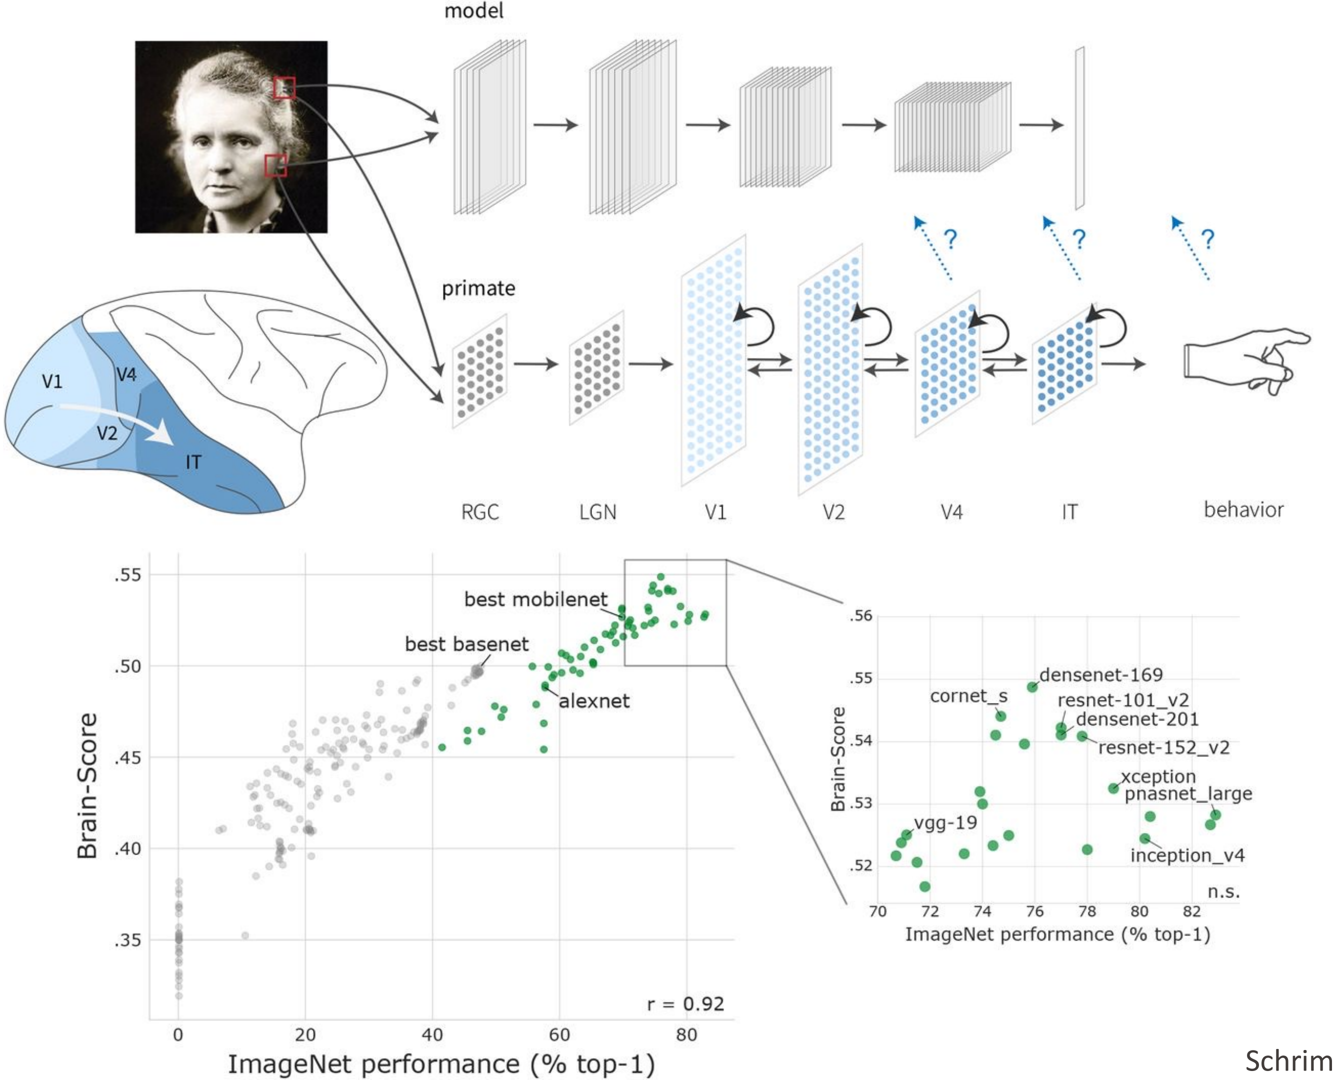

Schrimpf et al. Neuron 2020

Where do the “labels” come from?

Training **CNNs on object recognition** gives rise to neural represenations that resemble the representations in the brain (for various sensory systems)

However:

- Learning was based on many *labeled* examples, which is unlikely the case for the brain
- Could one not exploit the images, themselves without the labels?

Self-supervised learning

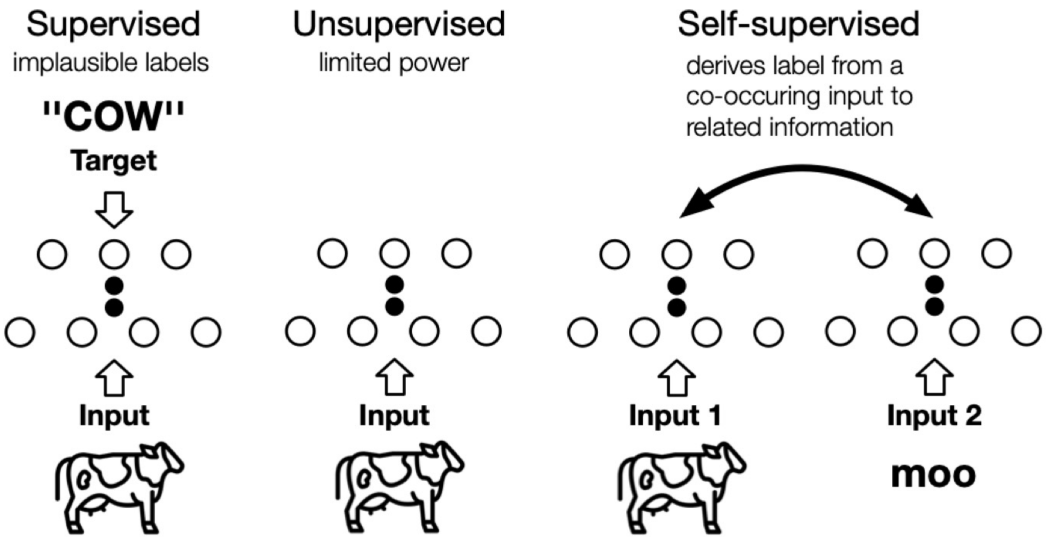

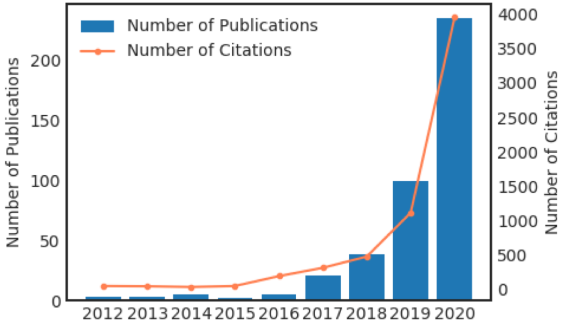

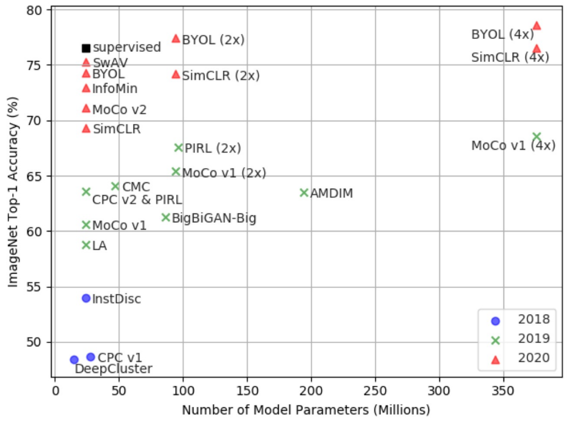

Liu et al., IEEE Transactions on Knowledge and data engineering 2023

Comparing unsupervised tasks

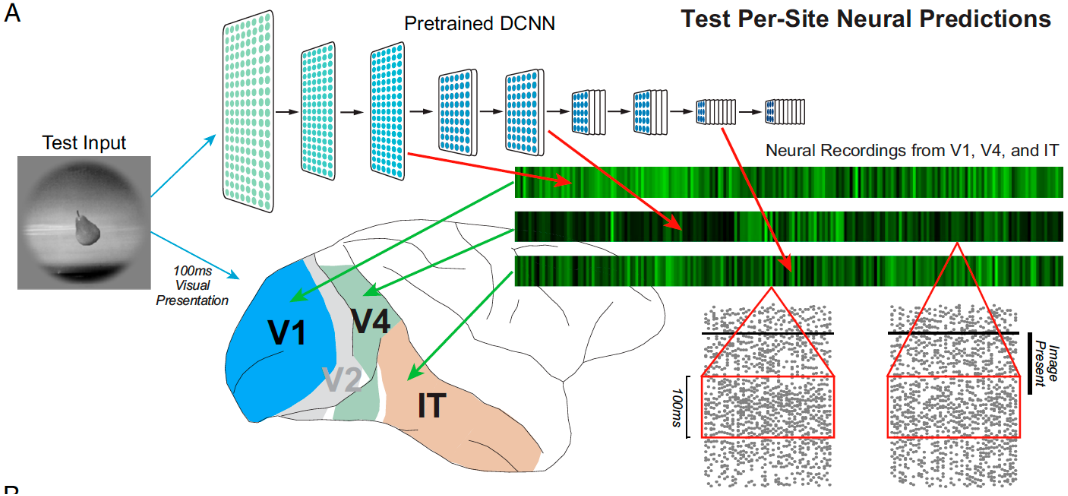

Zhuang et al., PNAS 2020

Comparing unsupervised tasks

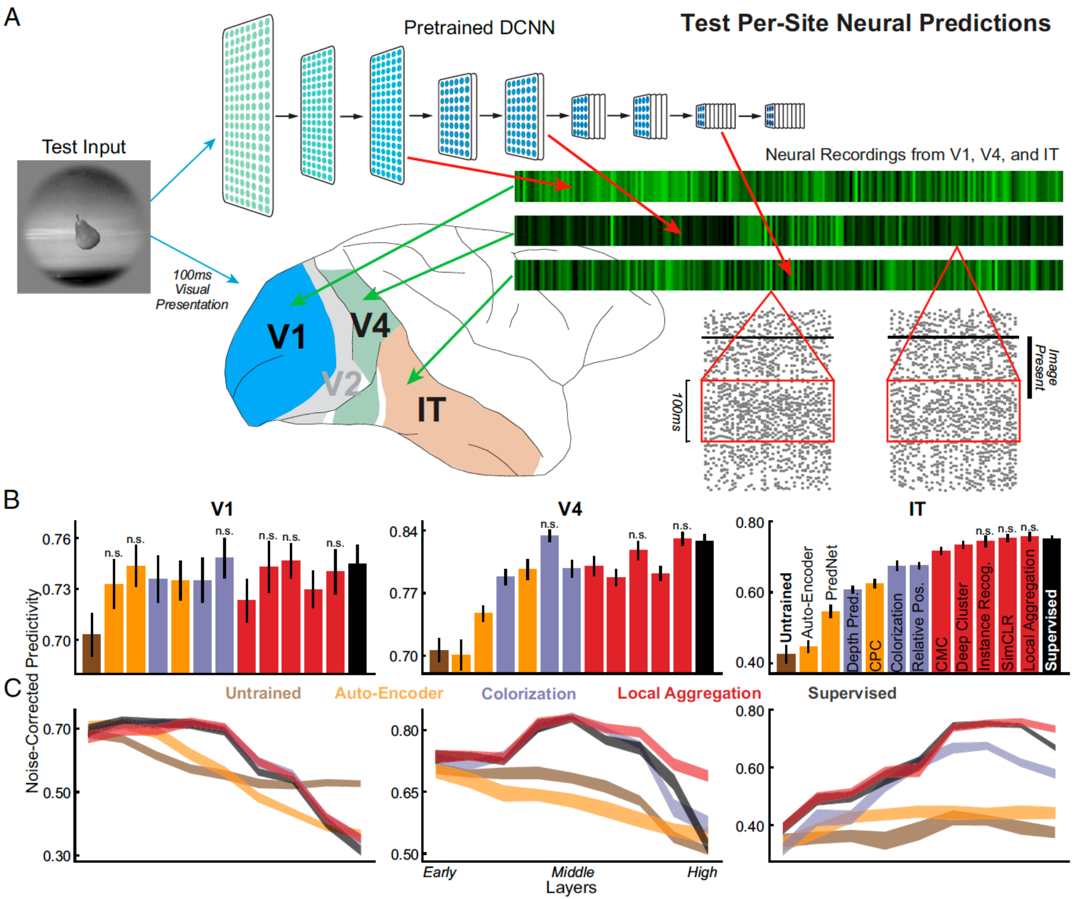

Zhuang et al., PNAS 2020

Task-driven modeling of sensory systems

This approach has been *used* in vision, audition, touch and thermoception, proprioception.

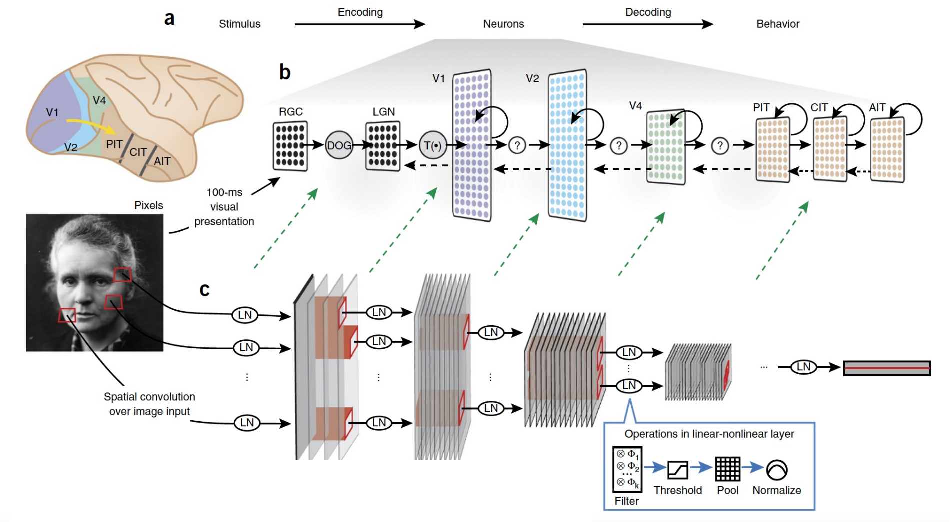

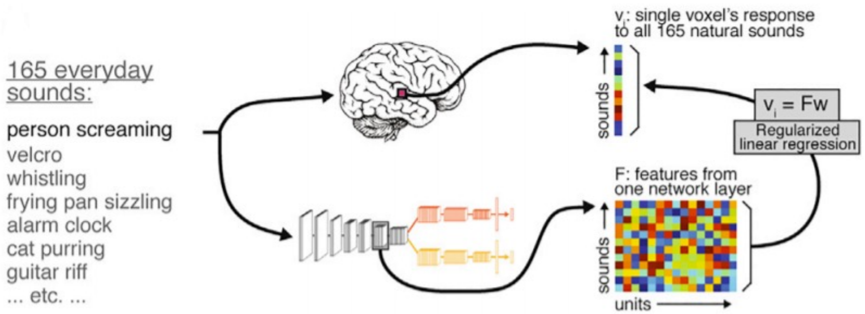

Yamins and Di Carlo, Nat Neuro 2016

Kell et al., Neuron 2018

Using deep neural networks as goal-driven models of a system


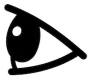
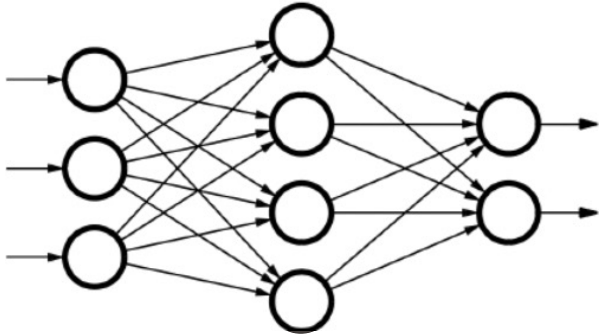

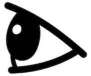
Vision: Yamins et al. (2014) -> object recog.


Audition: Kell et al. (2018)- speech recognition, speaker identification, natural sound identification


Barrel Cortex: Zhuang et al. (2017)


Cognition: Mante et al. (2013)


Proprioception: Sandbrink et al. (2023)

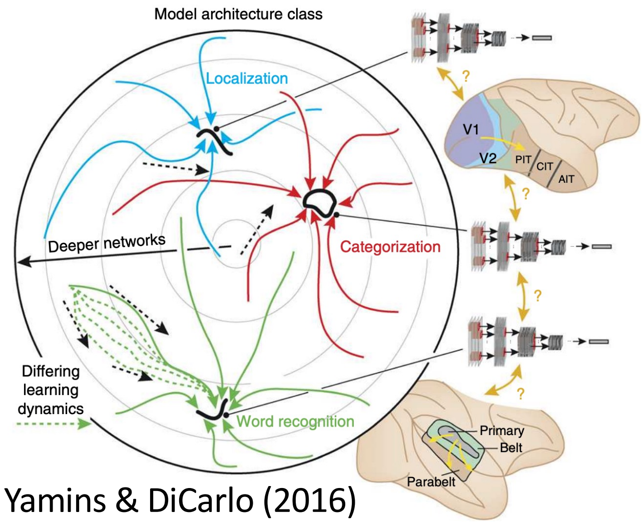

Yamins & DiCarlo (2016)

Revisiting Marr's levels

Bottom-up:

Are "principles of implementation/algorithms" good at solving computational problems?

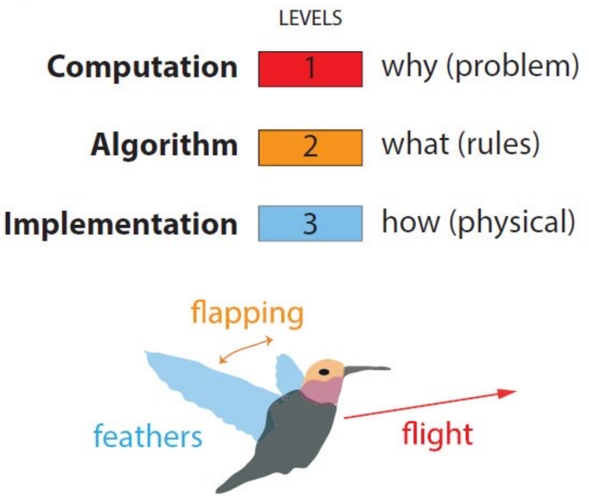

Illustration from Krakauer et al. 2017, Neuron

Top-down (Normative):

Which task-driven models provide good implementations?

EPFL

Adaptability of the sensorimotor system

Neuroscience ⟷ Artificial Intelligence

How do ants estimate the distance?

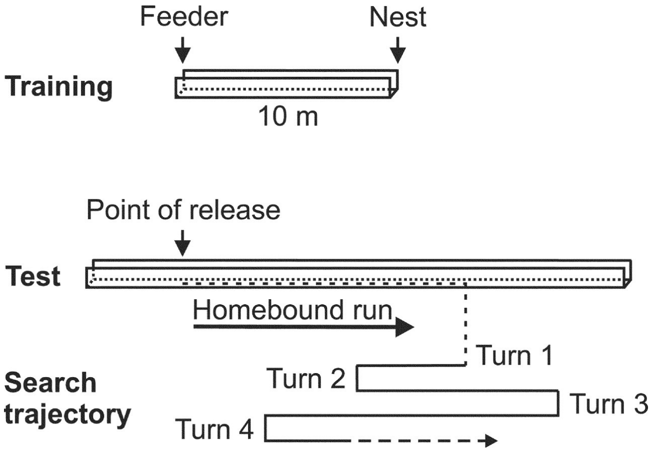

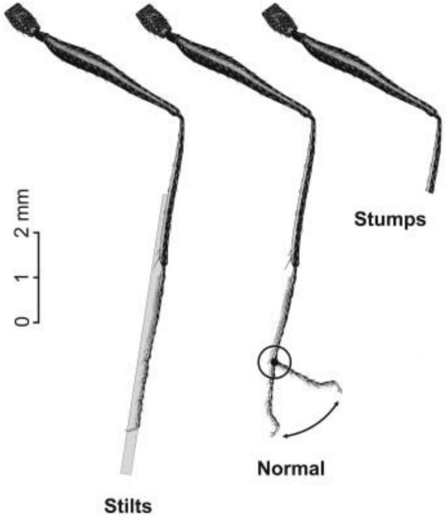

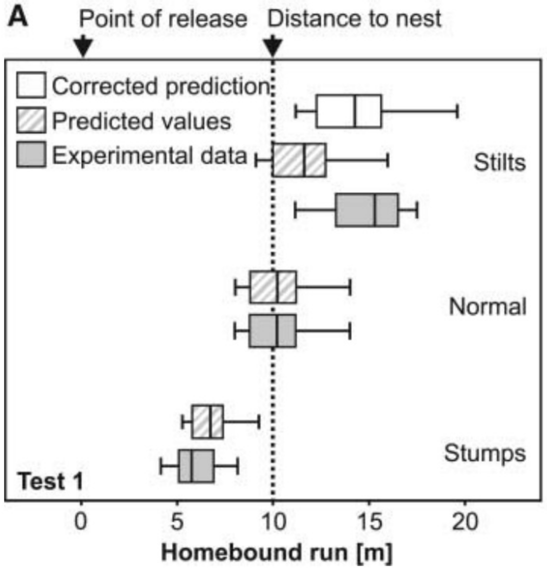

Wittlinger, Wehner & Wolf, Science 2006

Learning to walk with a changing body

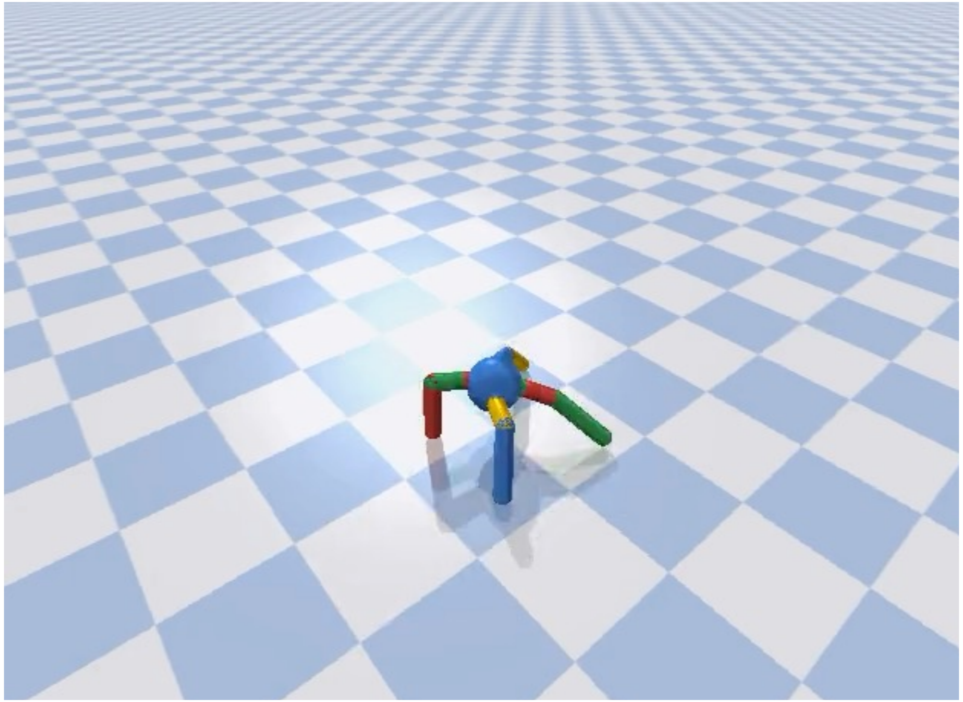

Ant (OpenAI Gym)

- Quadruped robot
- Action size: 8
- State size: 28

- Objective: run as fast as possible

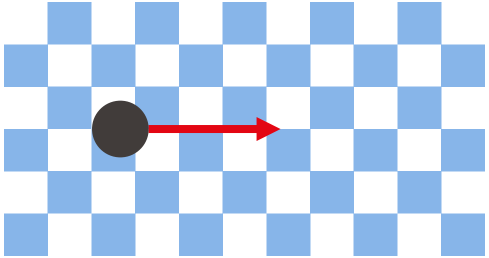

Chiappa, Marin Vargas, Mathis, Neurips 2022

Baseline: MLP policy

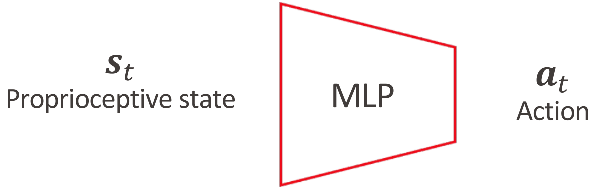

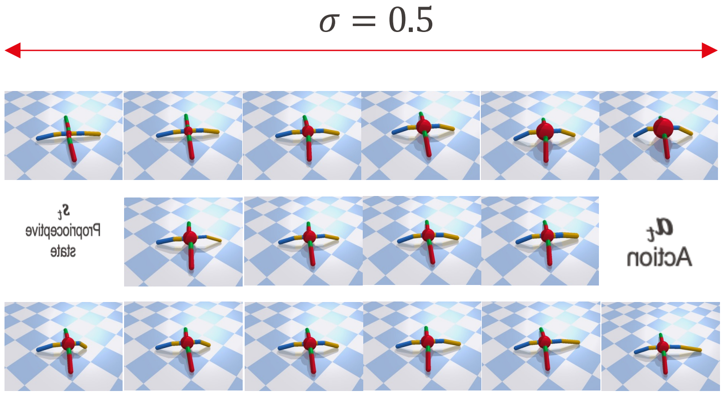

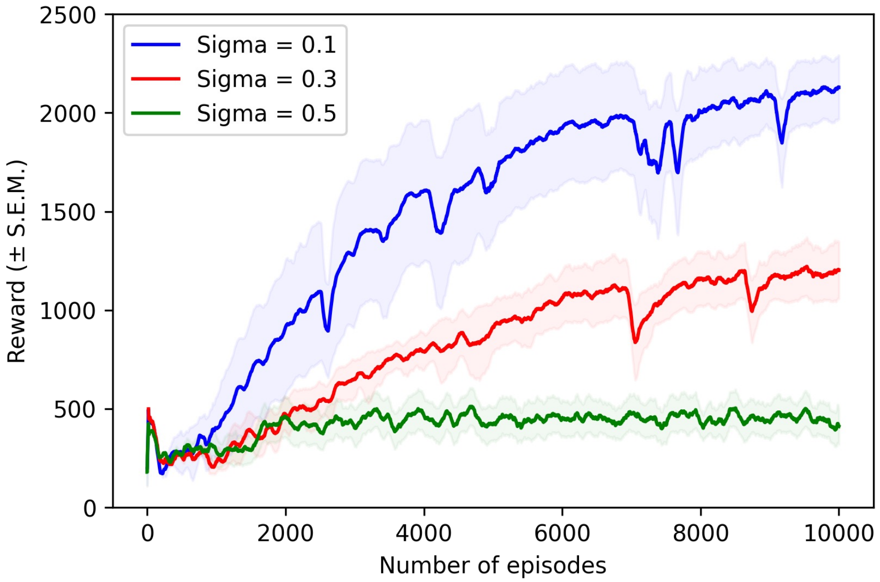

Chiappa, Marin Vargas, Mathis, Neurips 2022

Baseline: MLP policy

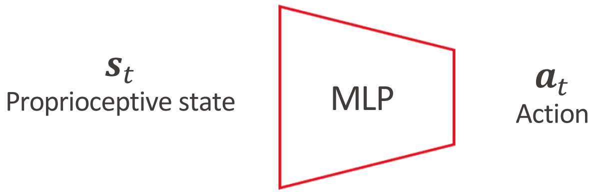

Morphology encoding policy (aka oracle)

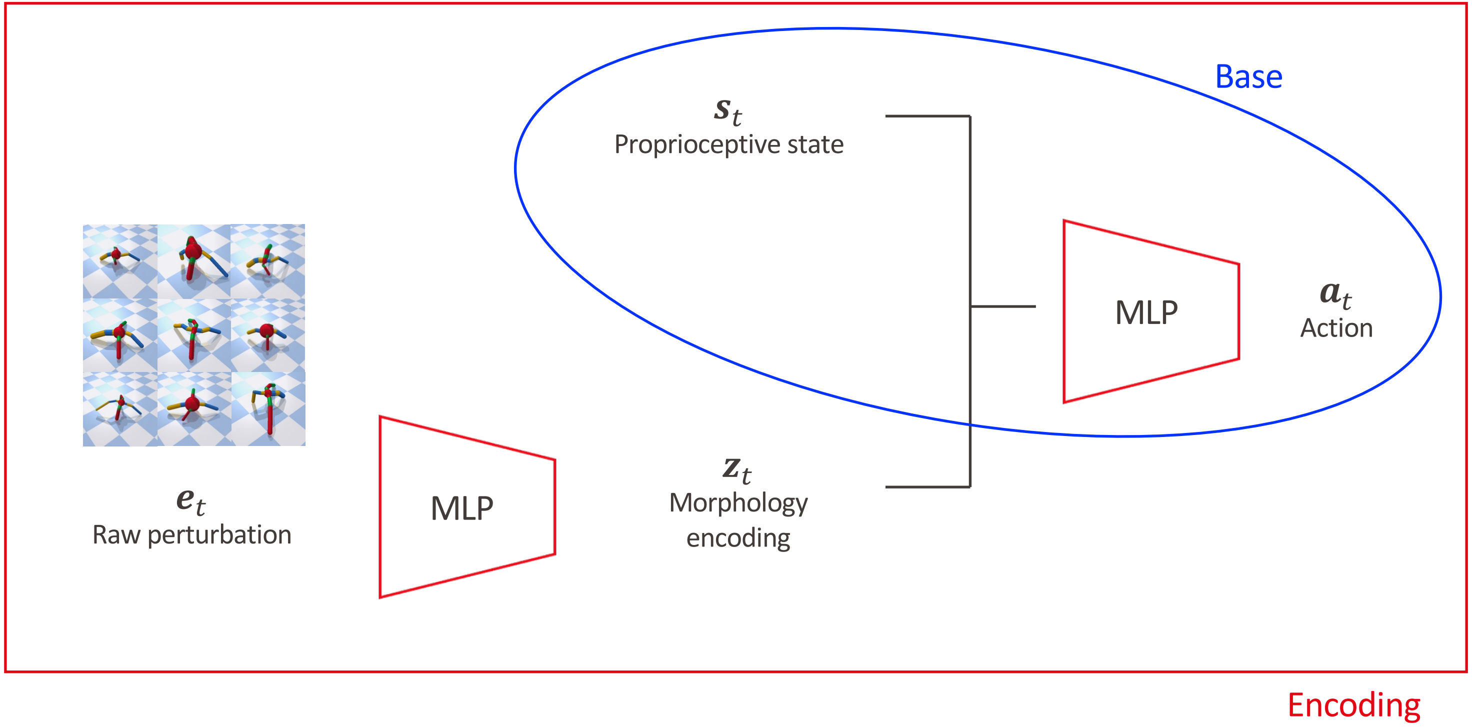

Morphology encoding policy

$\sigma = 0.1$

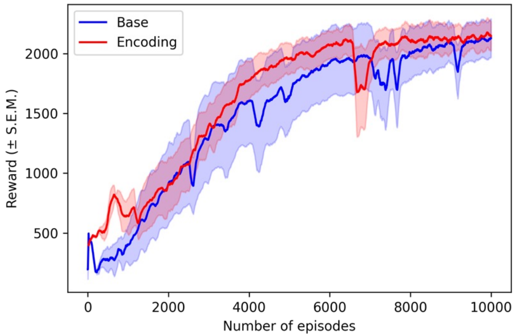

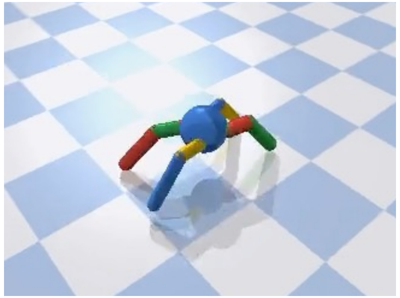

$\sigma = 0.3$

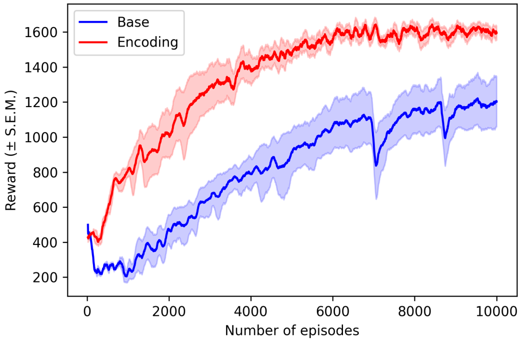

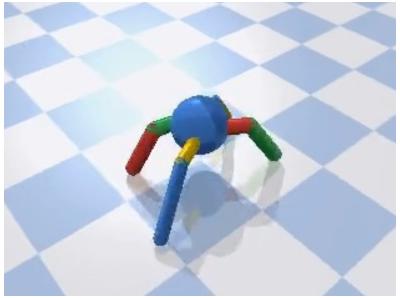

$\sigma = 0.5$

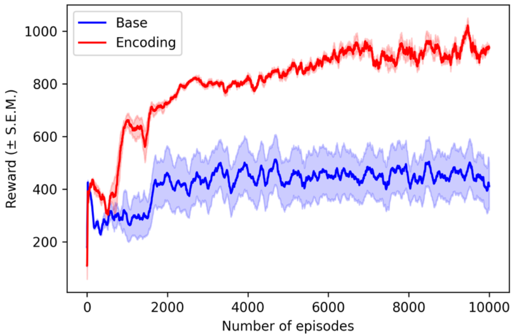

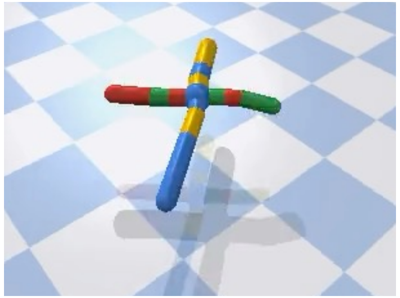

Chiappa, Marin Vargas, Mathis, Neurips 2022

Morphology encoding from experience

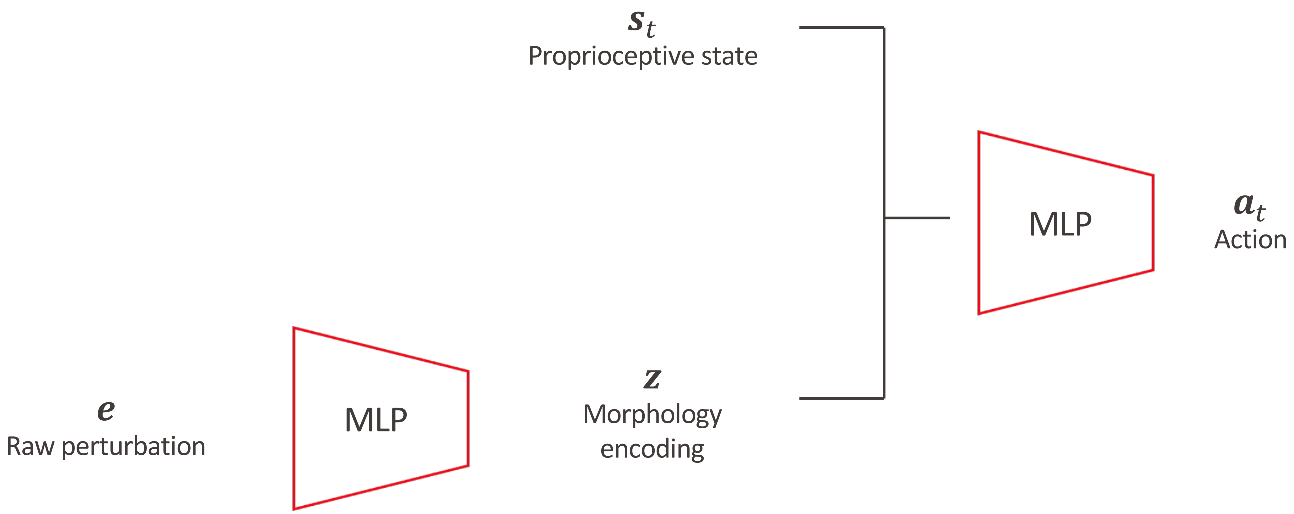

- Kumar, A., Fu, Z., Pathak, D., & Malik, J, “RMA: Rapid Motor Adaptation for Legged Robots, 2021

Morphology encoding from experience

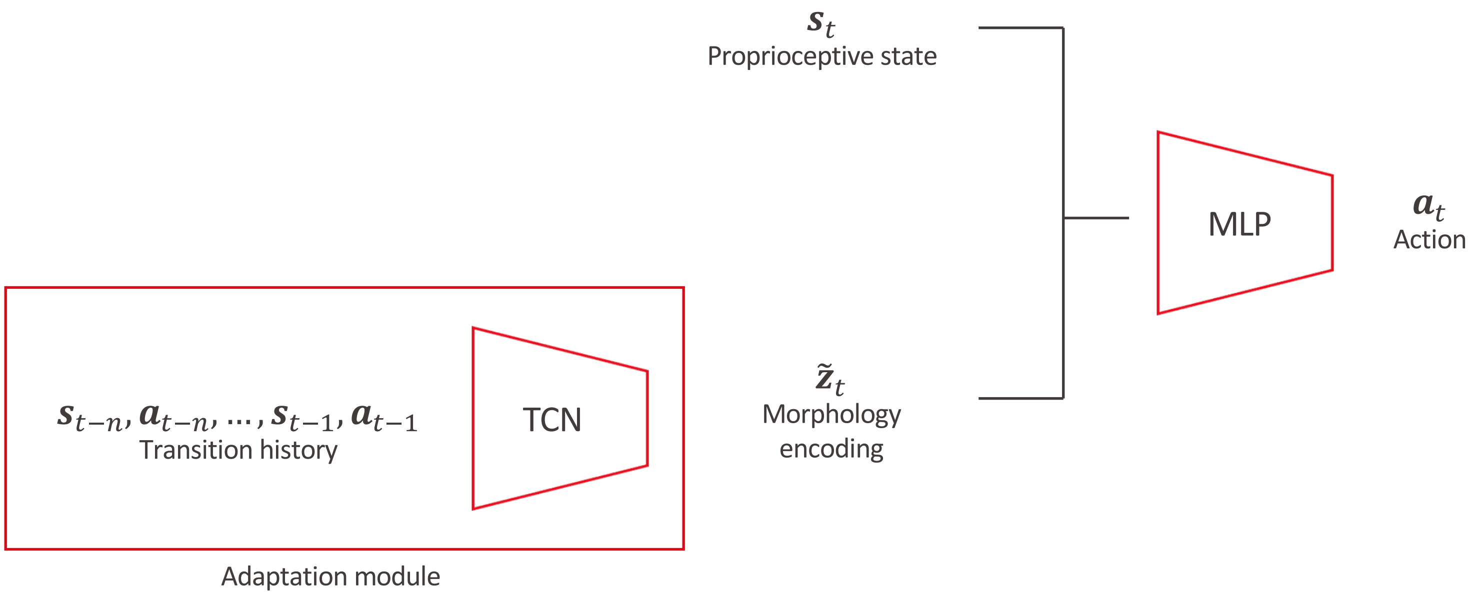

- Kumar, A., Fu, Z., Pathak, D., & Malik, J, "RMA: Rapid Motor Adaptation for Legged Robots, 2021

Morphology encoding from experience

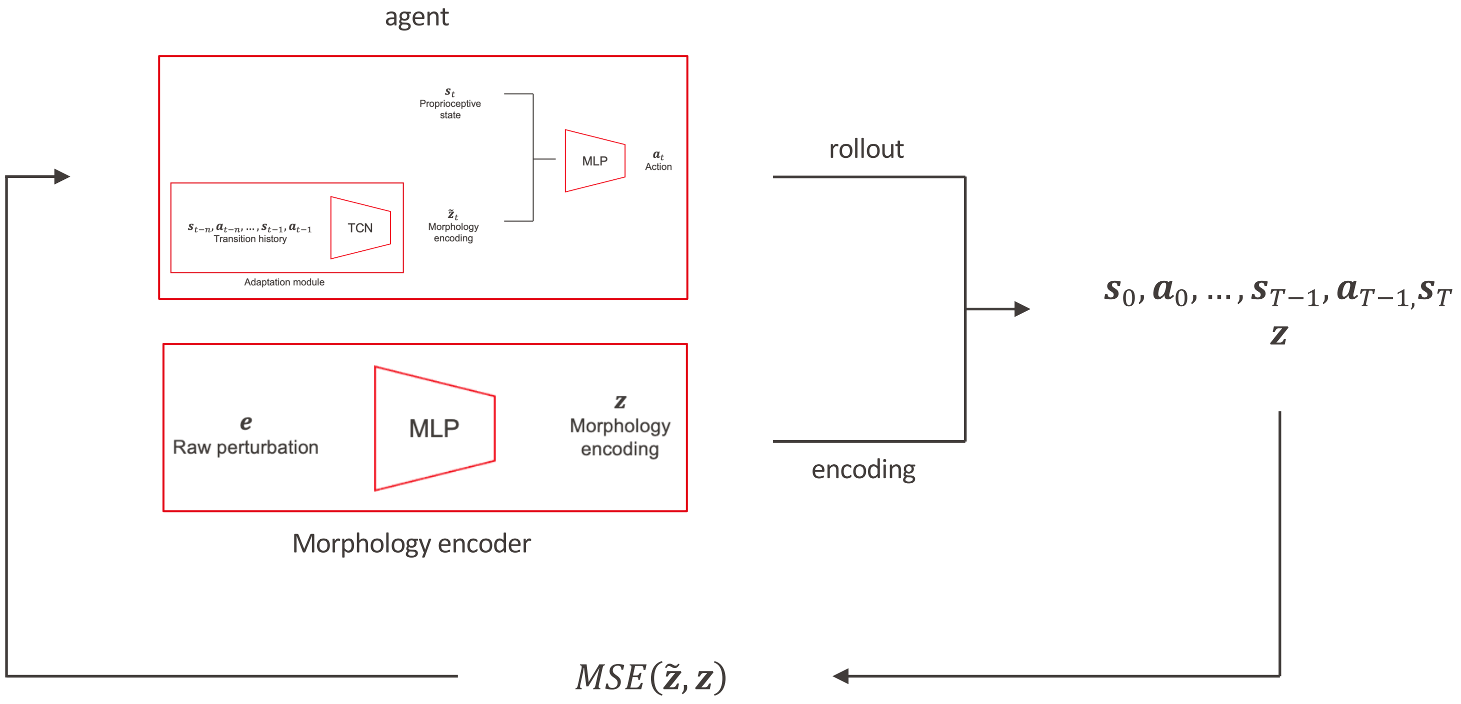

- Kumar, A., Fu, Z., Pathak, D., & Malik, J, "RMA: Rapid Motor Adaptation for Legged Robots, 2021

Learning with a perturbed body

100 perturbations per intensity level

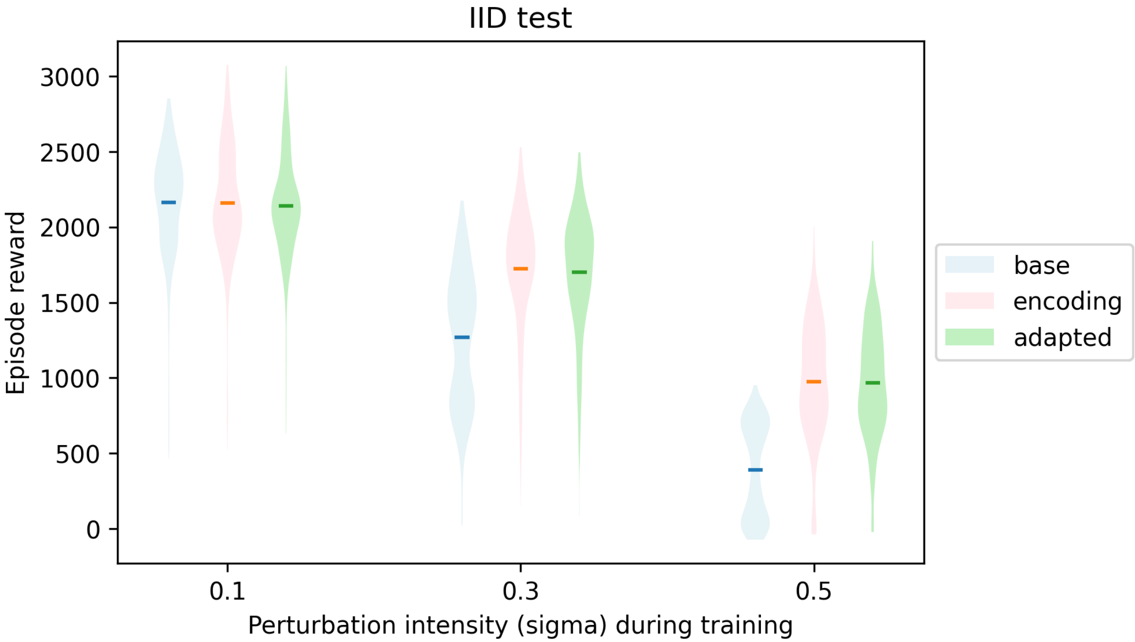

Is the 2-step training necessary?

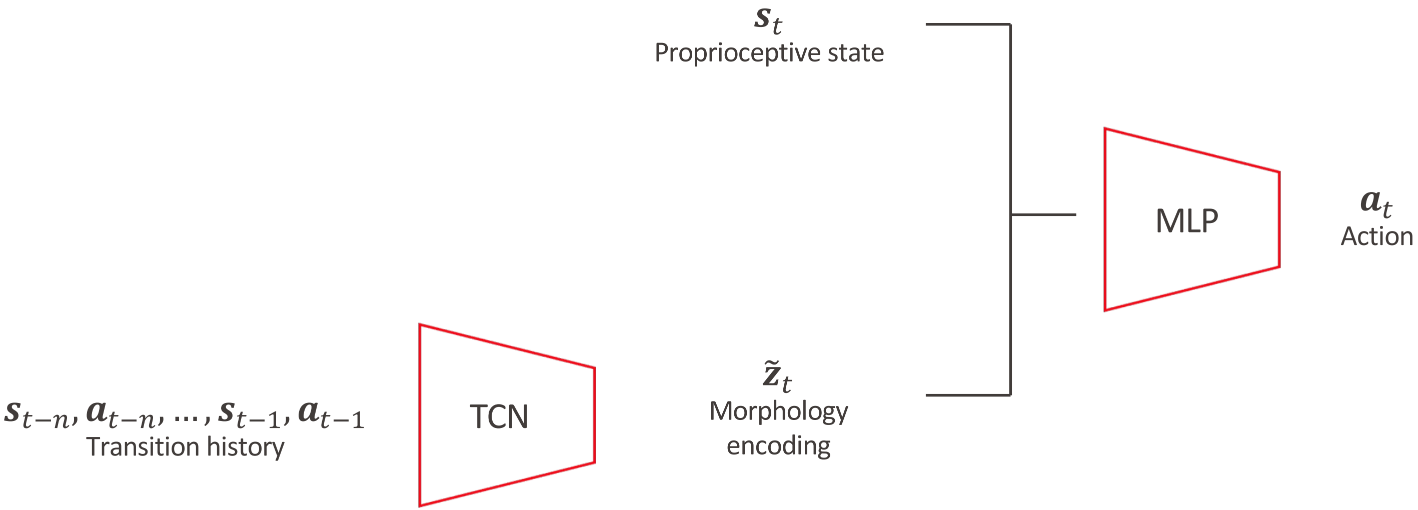

Training the CNN encoder end-to-end

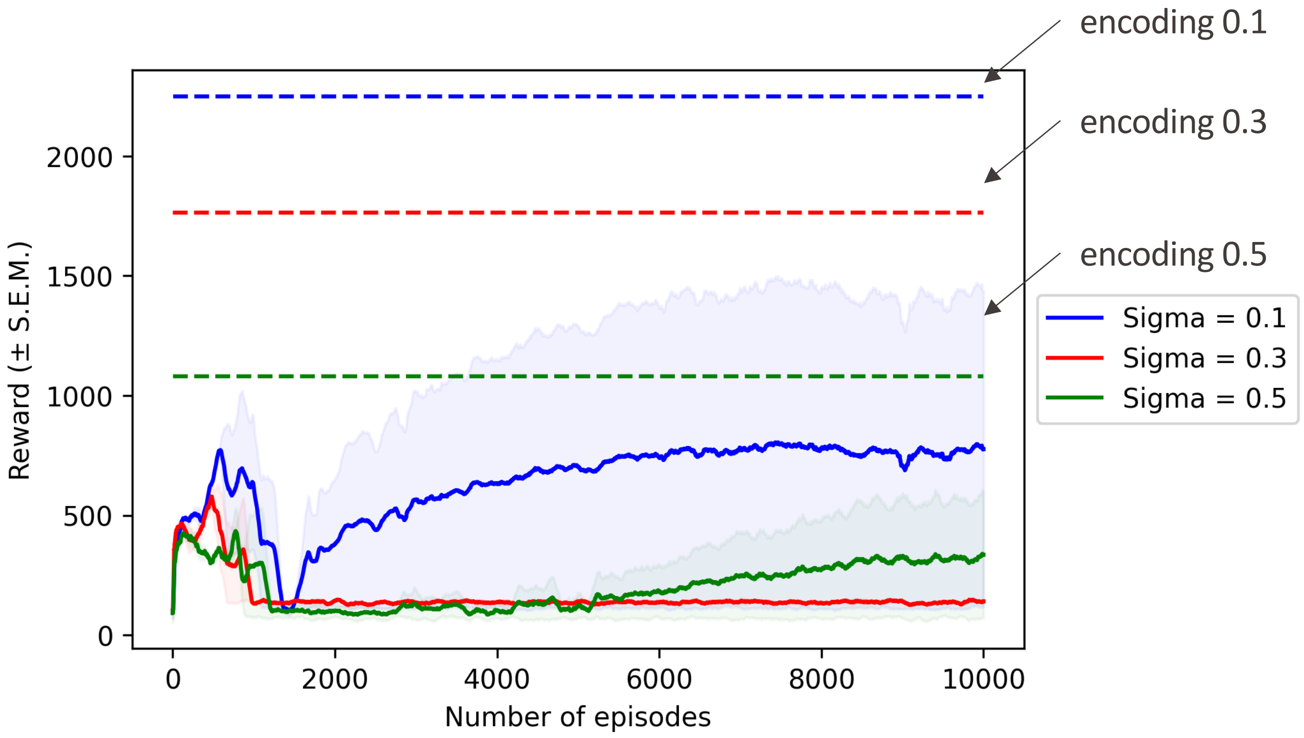

Principles from the brain: Distributed sensing and control

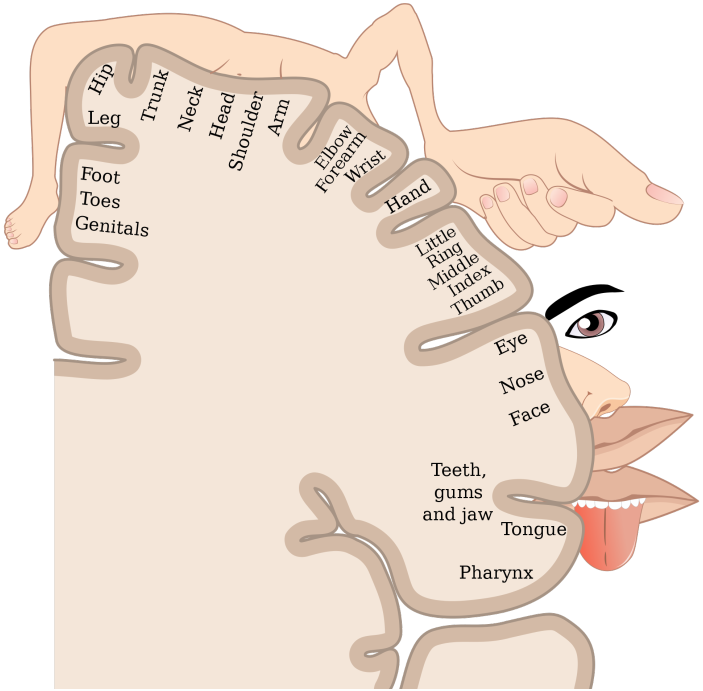

- Independent low-level processing
- High-level proprioceptive input integration
- Distributed control

DMAP's brain inspired architecture

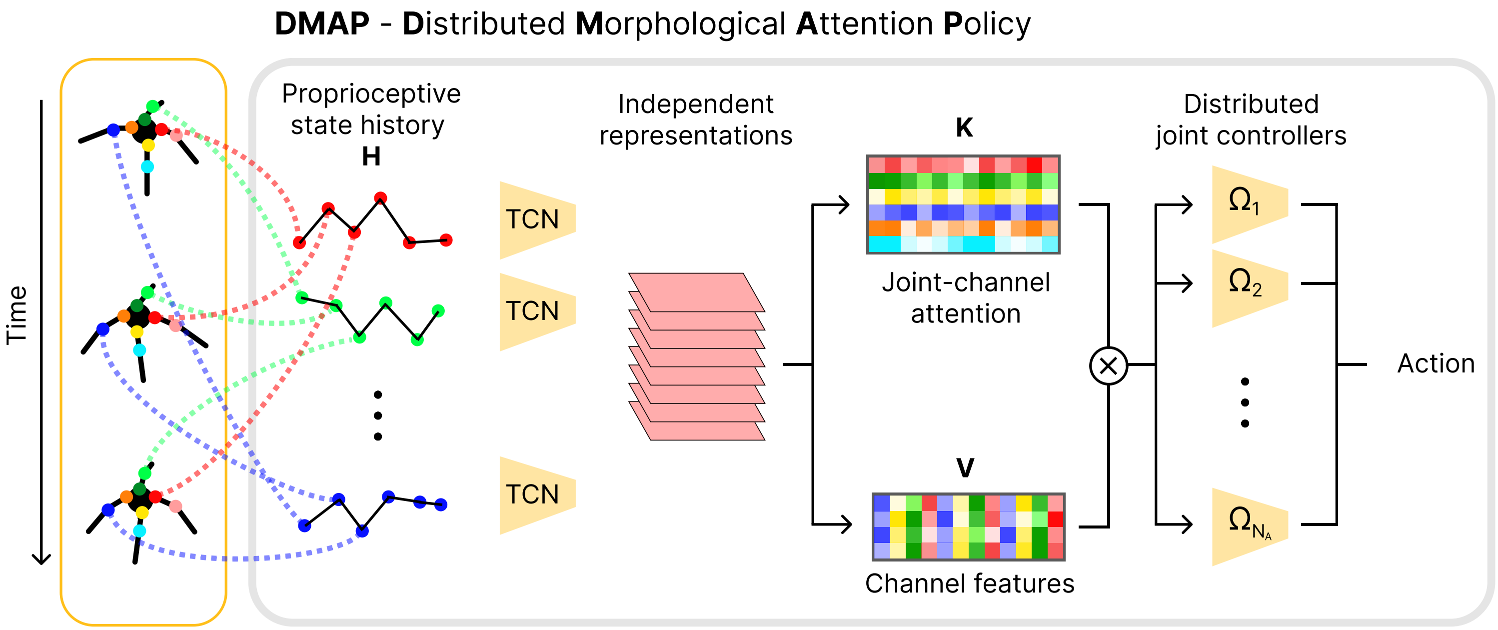

Chiappa, Marin Vargas, Mathis, Neurips 2022

DMAP performance comparison

Sigma = 0.1

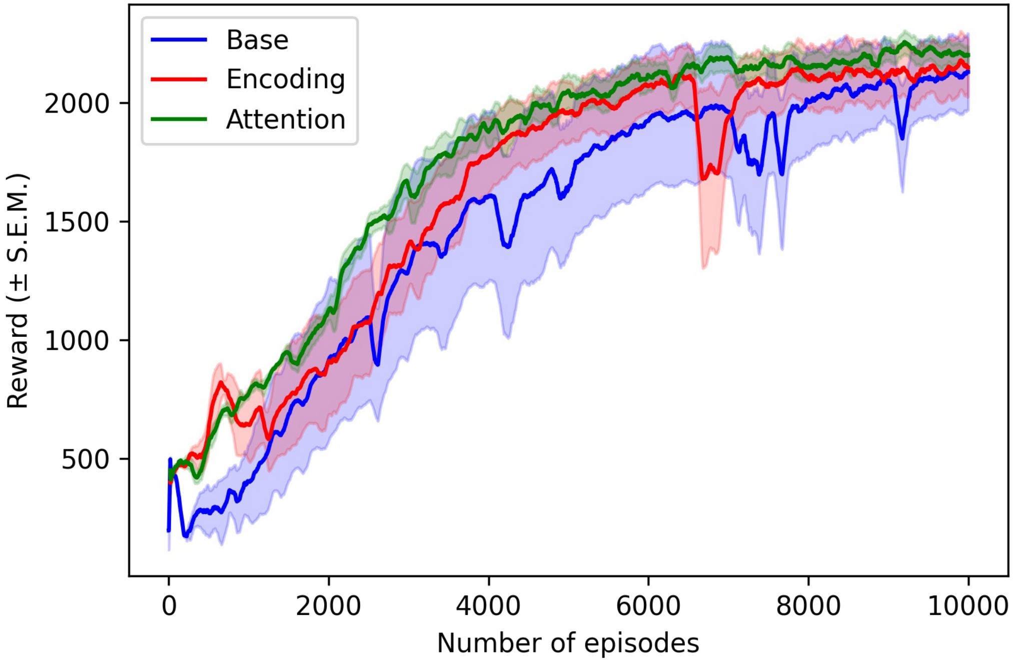

Chiappa, Marin Vargas, Mathis, Neurips 2022

A. Mathis

DMAP performance comparison

Sigma = 0.3

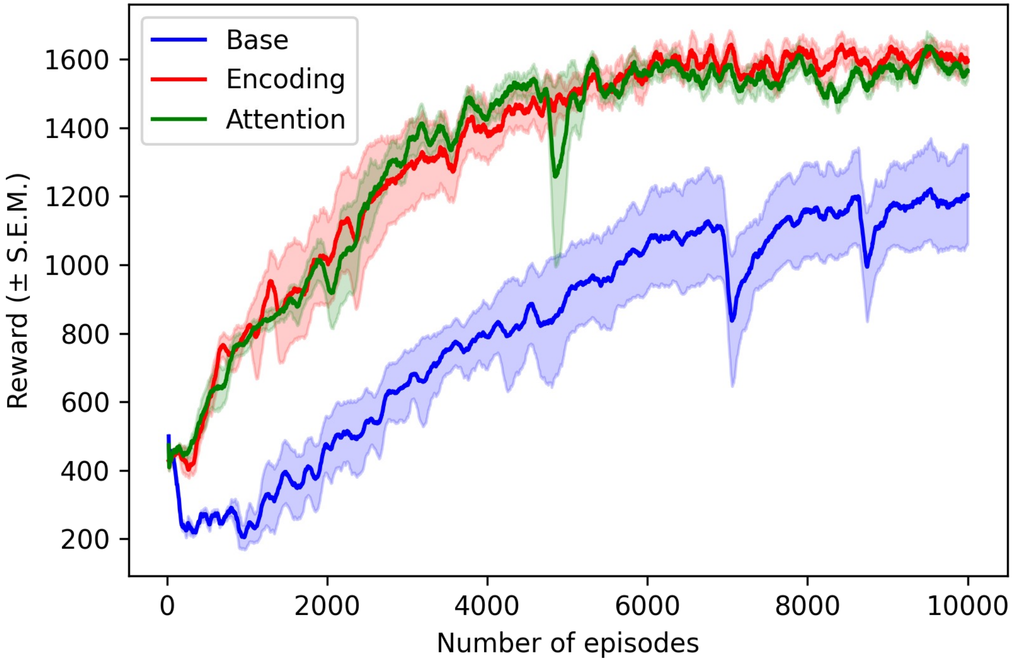

DMAP performance comparison

Sigma = 0.5

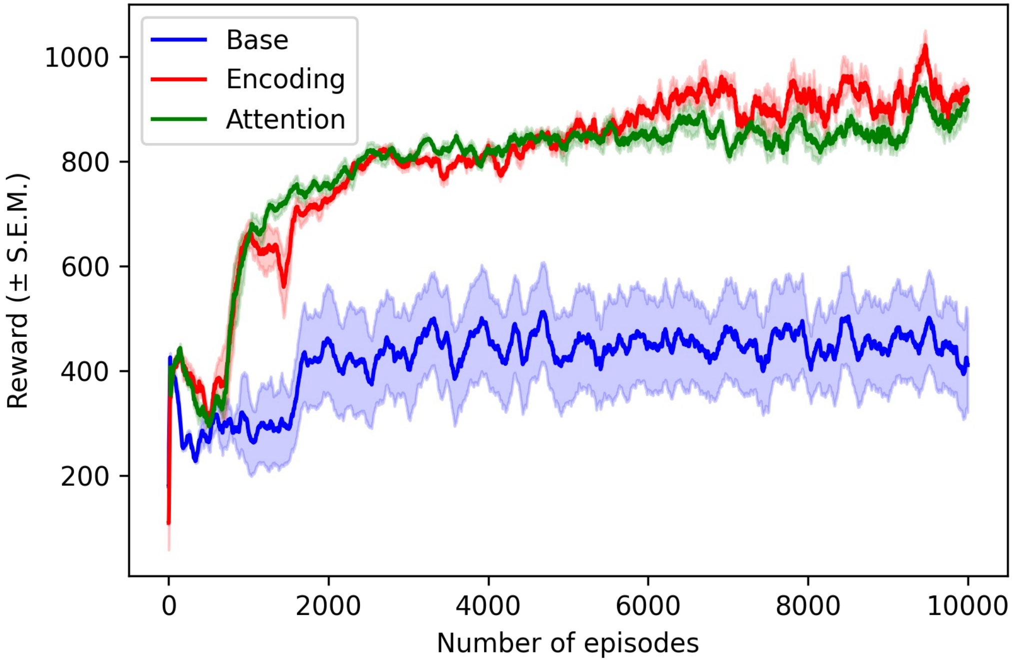

EPFL

Results also hold across different morphologies

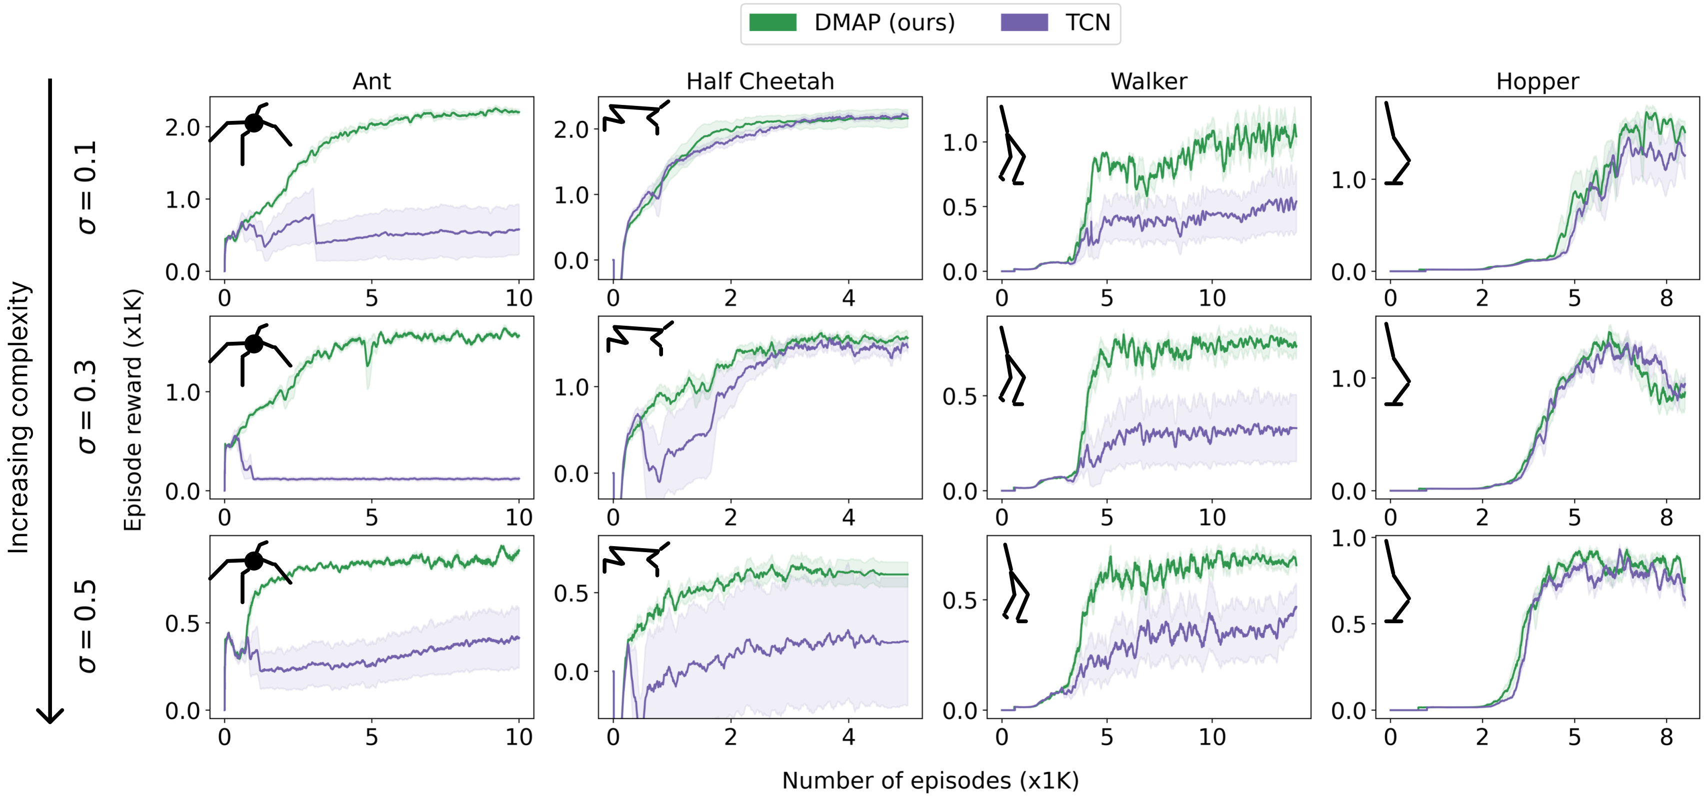

Chiappa, Marin Vargas, Mathis, Neurips 2022

Analysis of the learned attention weights

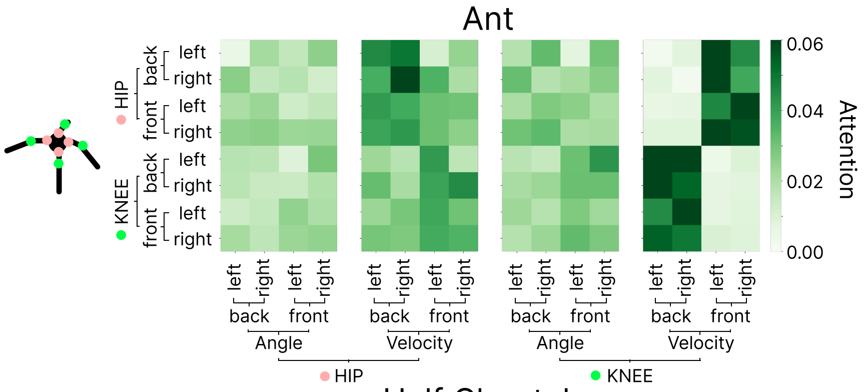

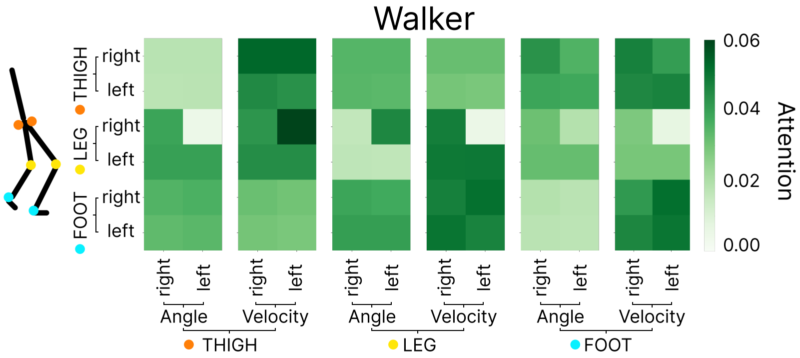

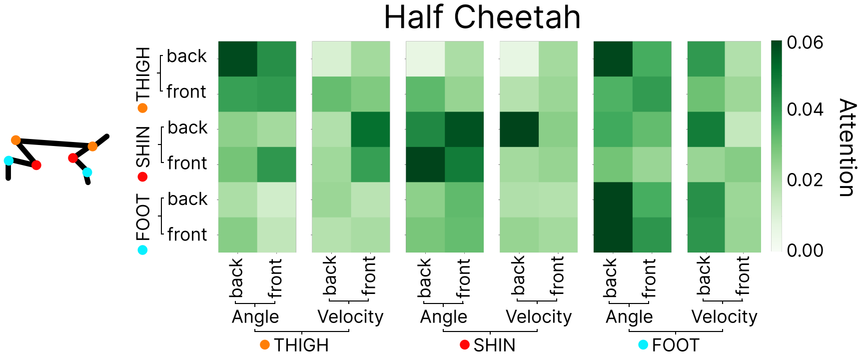

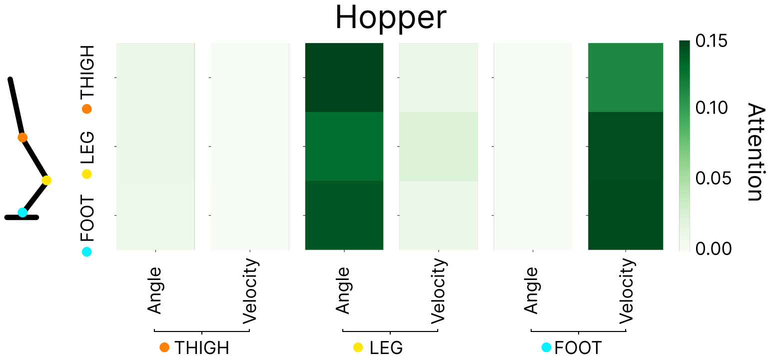

EPFL

Attention is gait-cycle specific

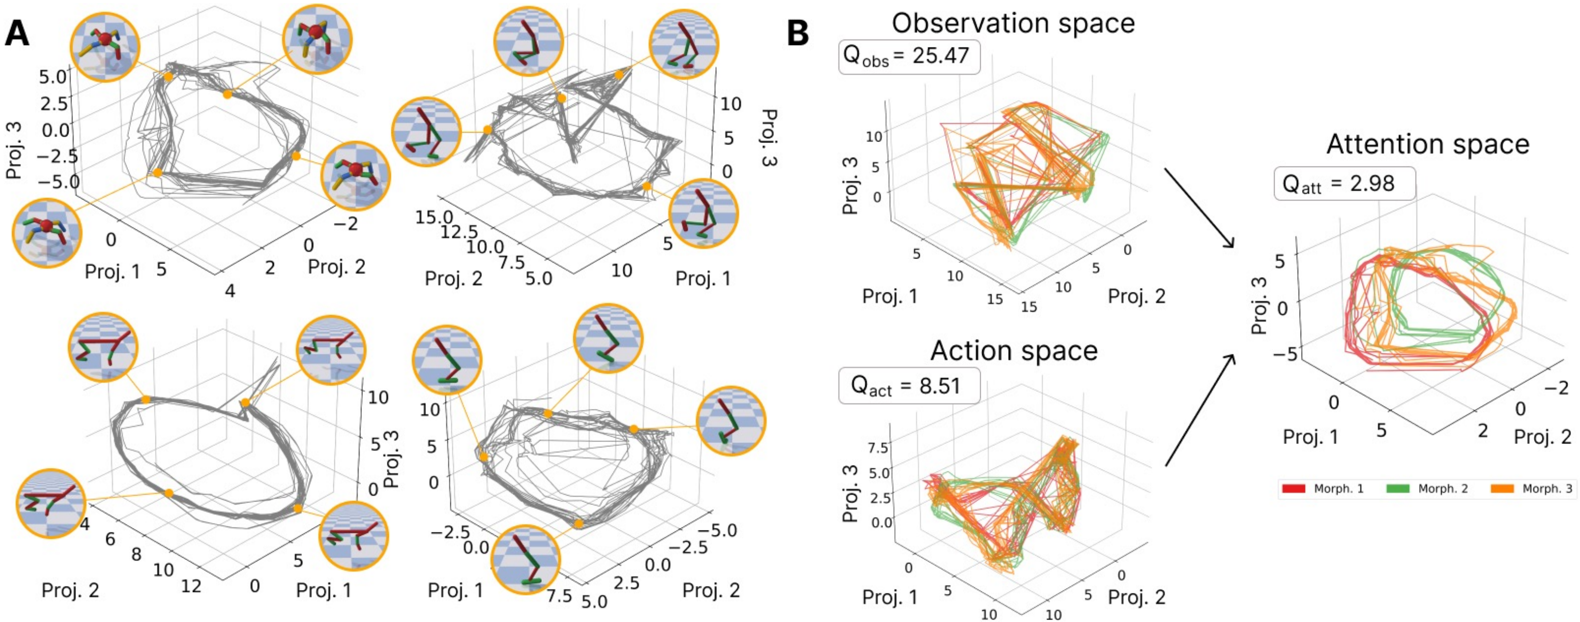

Inductive biases

- We discussed a brain-inspired inductive bias (DMAP) that implicitly can deal with changing bodies (better than other policies)
- You'll discuss Tony Zador's Perspective in Nat. Comm 2019

A critique of pure learning and what artificial neural networks can learn from animal brains

Anthony M. Zador[1]

Artificial neural networks (ANNs) have undergone a revolution, catalyzed by better supervised learning algorithms. However, in stark contrast to young animals (including humans), training such networks requires enormous numbers of labeled examples, leading to the belief that animals must rely instead mainly on unsupervised learning. Here we argue that most animal behavior is not the result of clever learning algorithms—supervised or unsupervised—but is encoded in the genome. Specifically, animals are born with highly structured brain connectivity, which enables them to learn very rapidly. Because the wiring diagram is far too complex to be specified explicitly in the genome, it must be compressed through a "genomic bottleneck". The genomic bottleneck suggests a path toward ANNs capable of rapid learning.

Conclusions

Modeling biological systems:

Task-driven modeling with ML provides insights into brain function (vison, audition, language…)

Vargas, A. M., Bisi, A., Chiappa, A. S., Versteeg, C., Miller, L. E., & Mathis, A. (2024). Task-driven neural network models predict neural dynamics of proprioception. *Cell*

Conversely:

Principles from Neuro/Psychology may provide better ML solutions (CNNs, RL, …)

Hassabis et al (2017): Neuroscience-Inspired Artificial Intelligence, Neuron

Zador (2019): A critique of pure learning and what artificial neural networks can learn from animal brains, Nature Communication

Chiappa, A. S., Tano, P., Patel, N., Ingster, A., Pouget, A., & Mathis, A. (2024). Acquiring musculoskeletal skills with curriculum-based reinforcement learning. Neuron (in press)