

Intrinsic Goals for Autonomous Agents



Model-based Exploration Predicts Ethological
Behavior and Whole-Brain Dynamics in Larval Zebrafish

Aran Nayebi

Assistant Professor

Machine Learning Department, Neuroscience Institute (core faculty), Robotics Institute (by courtesy)

Carnegie Mellon University

Slides credit: Reece Keller

Autonomy: The Essence of Intelligence

intrinsic drives



extrinsic drives

Autonomy: The Essence of Intelligence

Autonomy is the capacity for *internal states* (in addition to *external states*) to drive behavior.

intrinsic drives

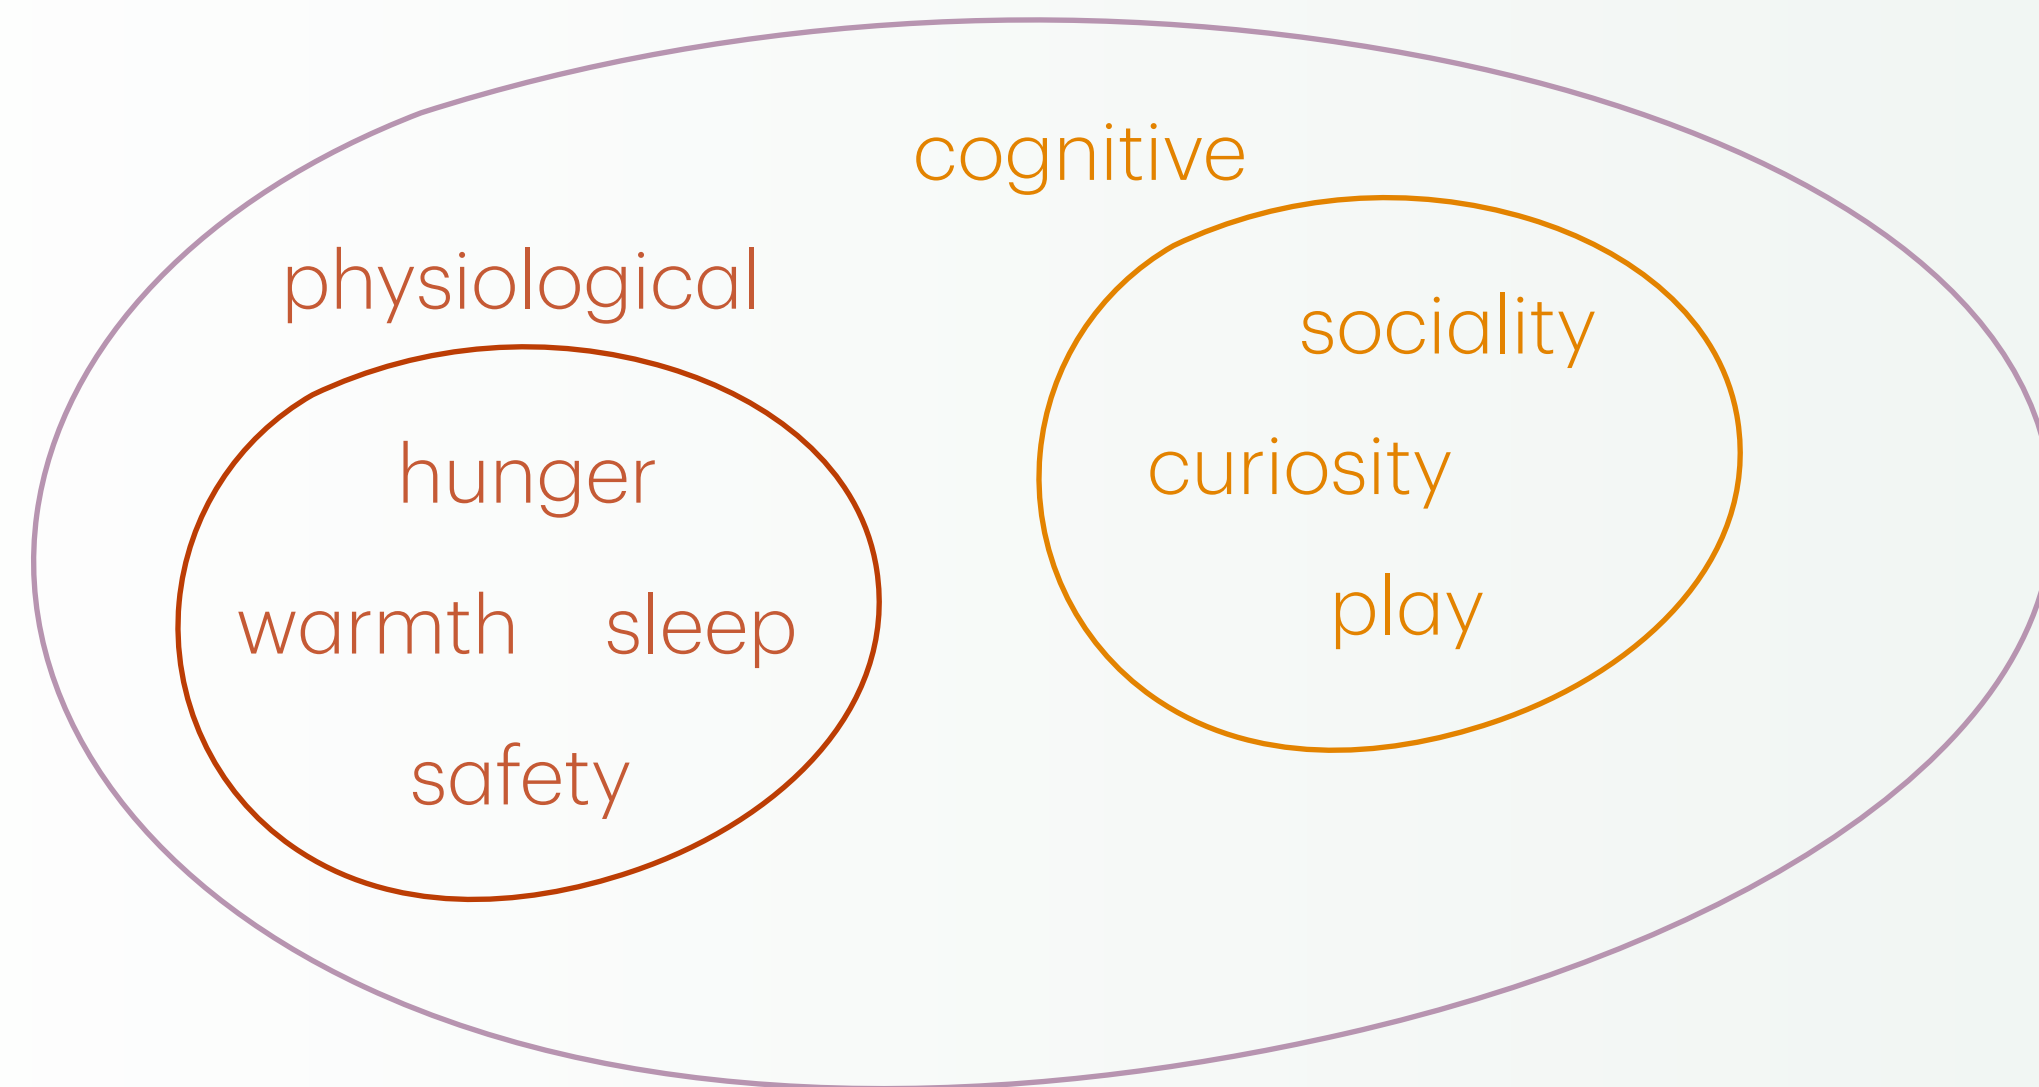


extrinsic drives

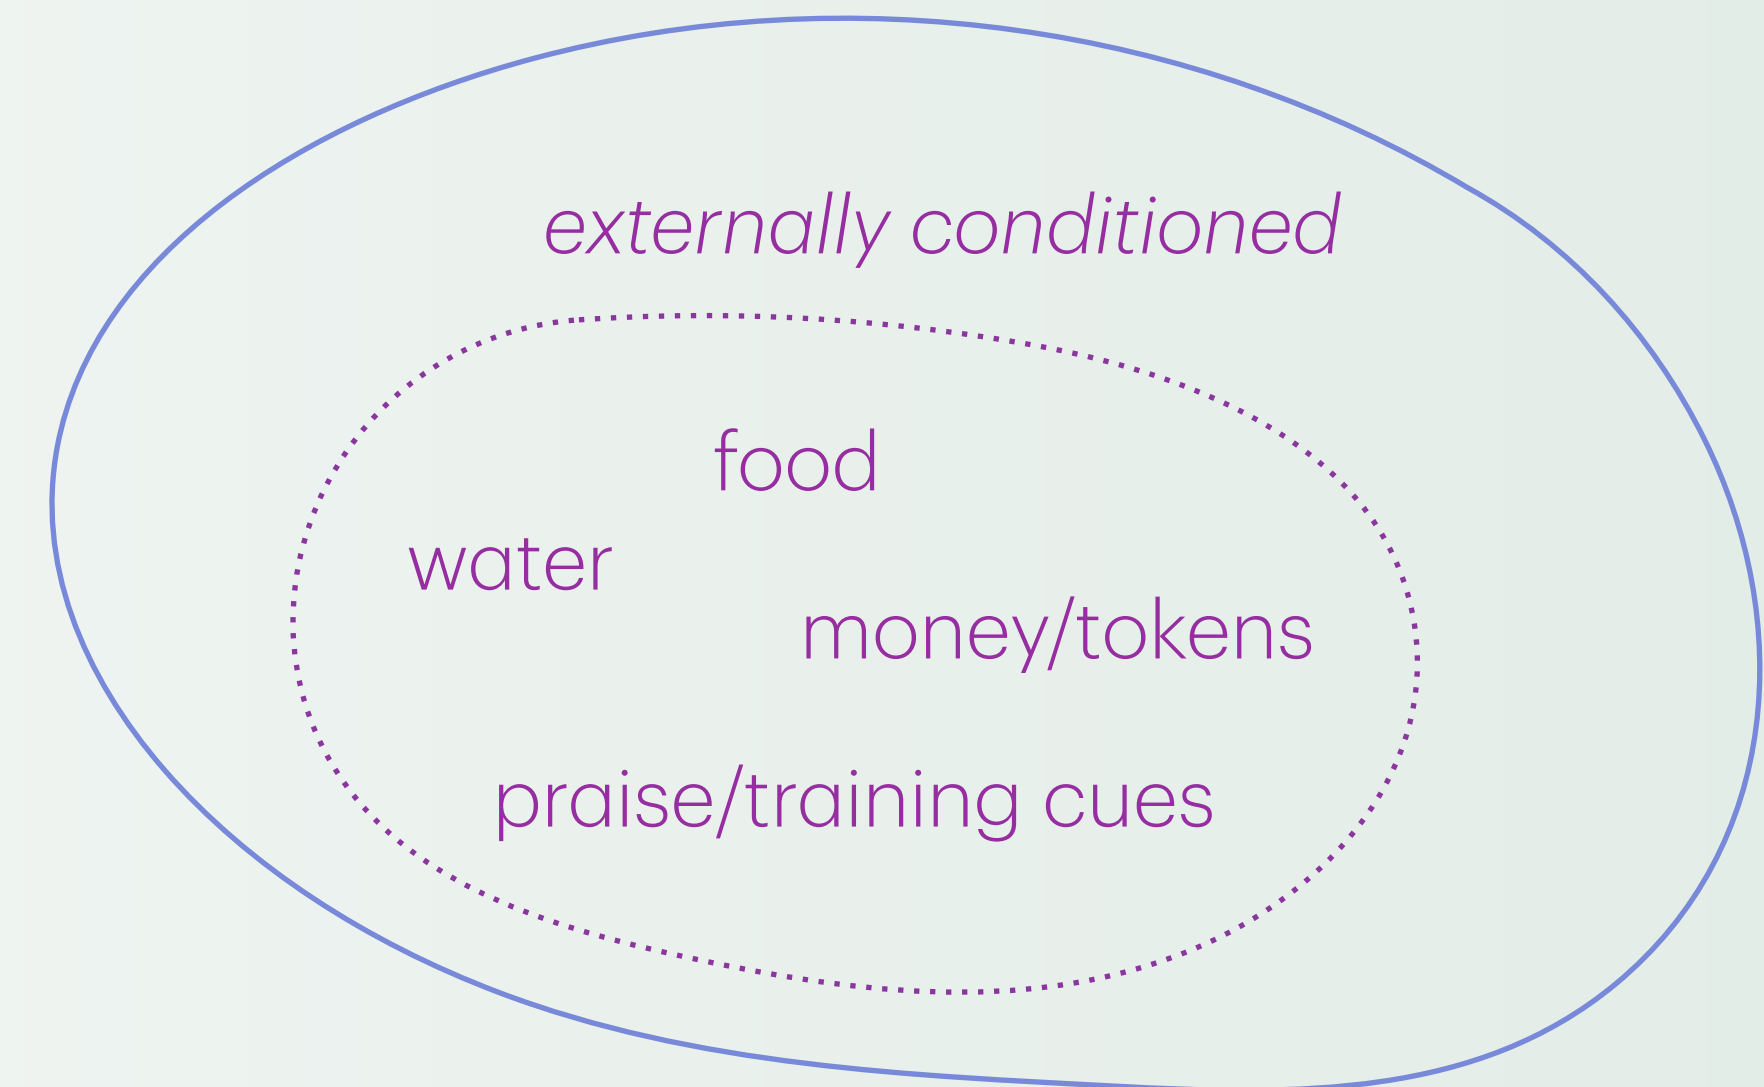
Autonomy: The Essence of Intelligence

Autonomy is the capacity for *internal states* (in addition to *external states*) to drive behavior.

intrinsic drives




extrinsic drives



Autonomy: The Essence of Intelligence

Autonomy is the capacity for *internal states* (in addition to *external states*) to drive behavior.



Autonomy: The Essence of Intelligence



Autonomy: The Essence of Intelligence



Autonomy: The Essence of Intelligence




- What is the **motivation/goal**?
- How is it **computationally formalized**?

Autonomy: The Essence of Intelligence



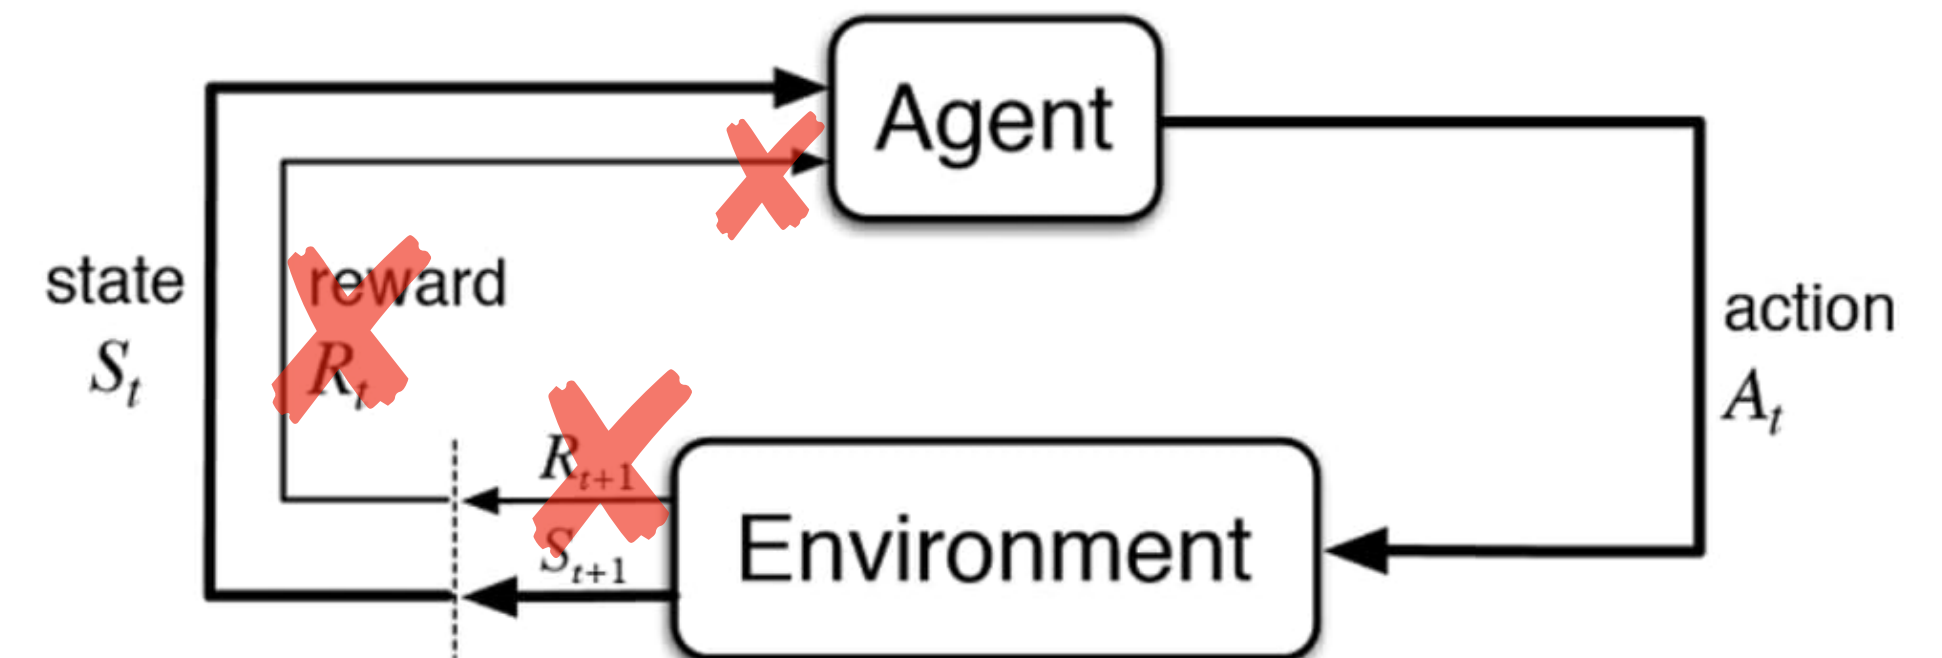
- What is the **motivation/goal**?
- How is it **computationally formalized**?



Autonomy: The Essence of Intelligence



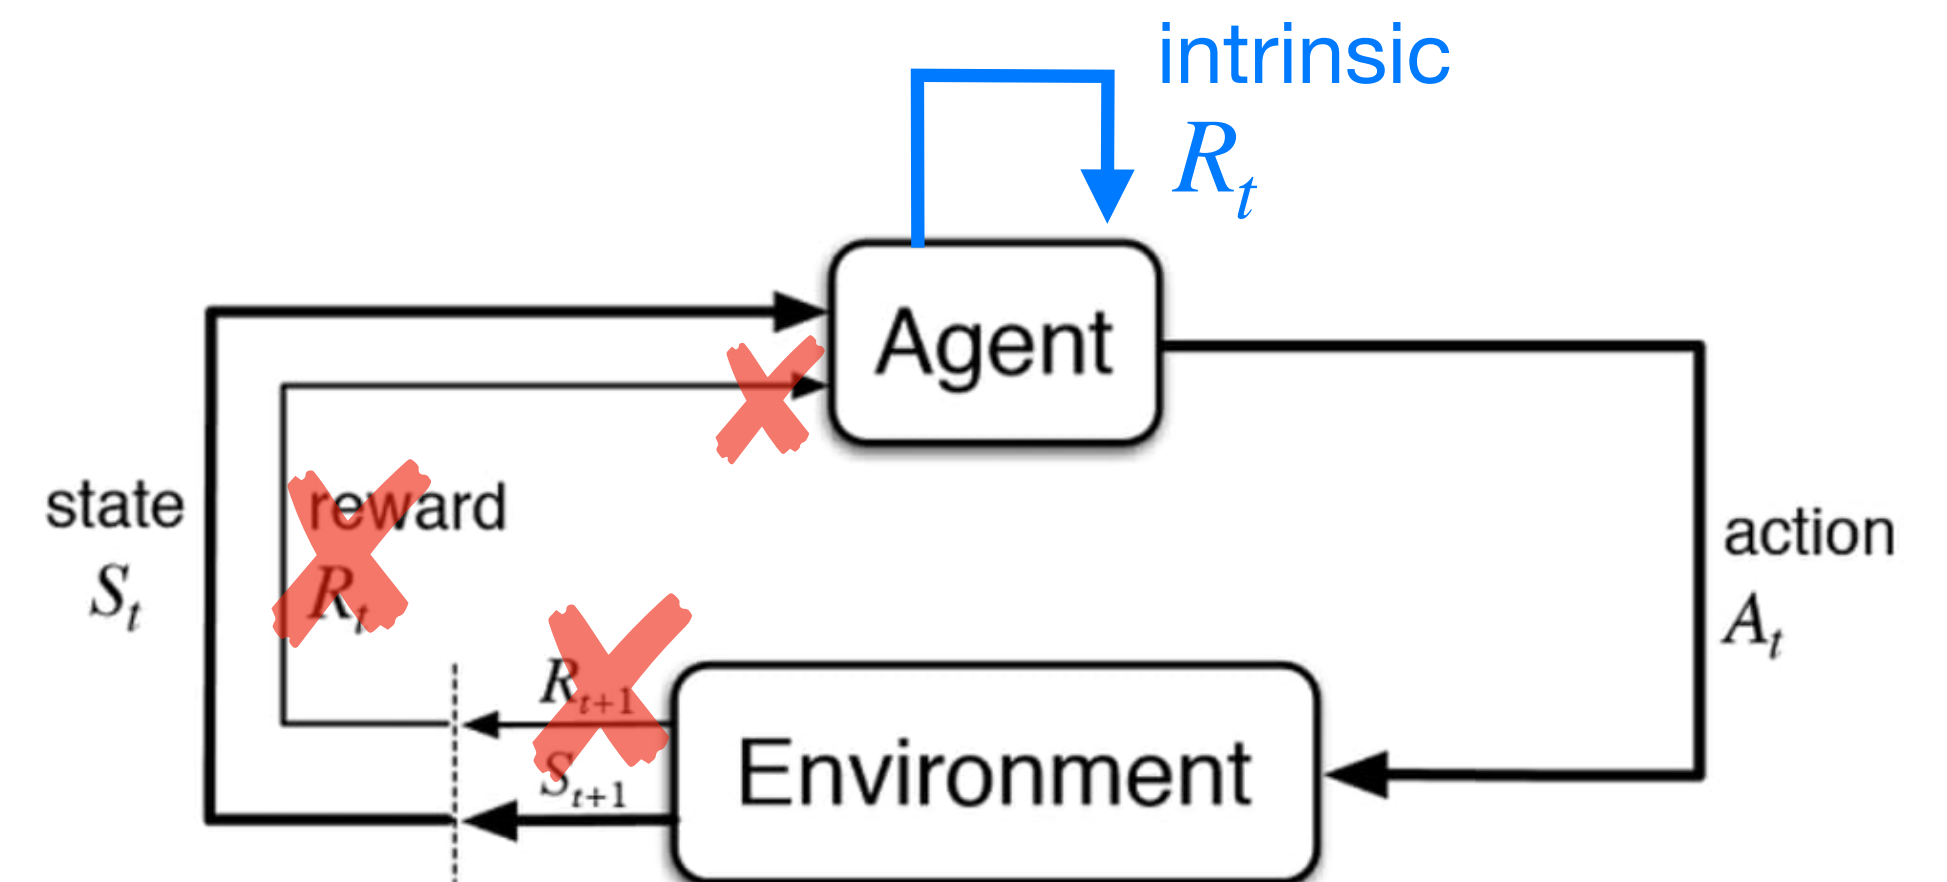
- What is the **motivation/goal**?
- How is it **computationally formalized**?



Autonomy: The Essence of Intelligence



- What is the **motivation/goal**?
- How is it **computationally formalized**?




Autonomy: The Essence of Intelligence



- What is the **motivation/goal**?
- How is it **computationally formalized**?

Goal:

algorithms for **intrinsic motivation** that capture the richness of Animal autonomy.



A Neuroscience Dataset for Animal Autonomy



A Neuroscience Dataset for Animal Autonomy




A Neuroscience Dataset for Animal Autonomy




Closed loop:

Swim commands
→ fish moves

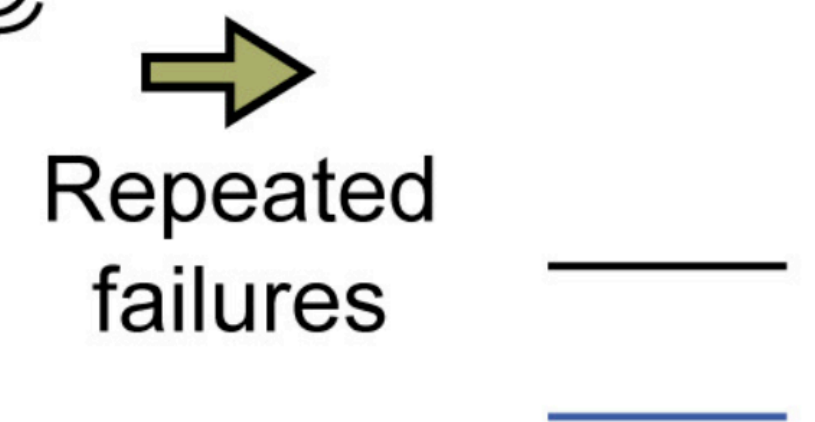


Open loop:

Swim commands
→ no movement




Passivity
(no swim commands)





A Neuroscience Dataset for Animal Autonomy

Mu et al., *Cell* (2019)

Neurons






Radial astrocytes



Fictive behavior
3.0X playback speed


Open loop




A Neuroscience Dataset for Animal Autonomy

Mu et al., *Cell* (2019)


Neurons



Radial astrocytes

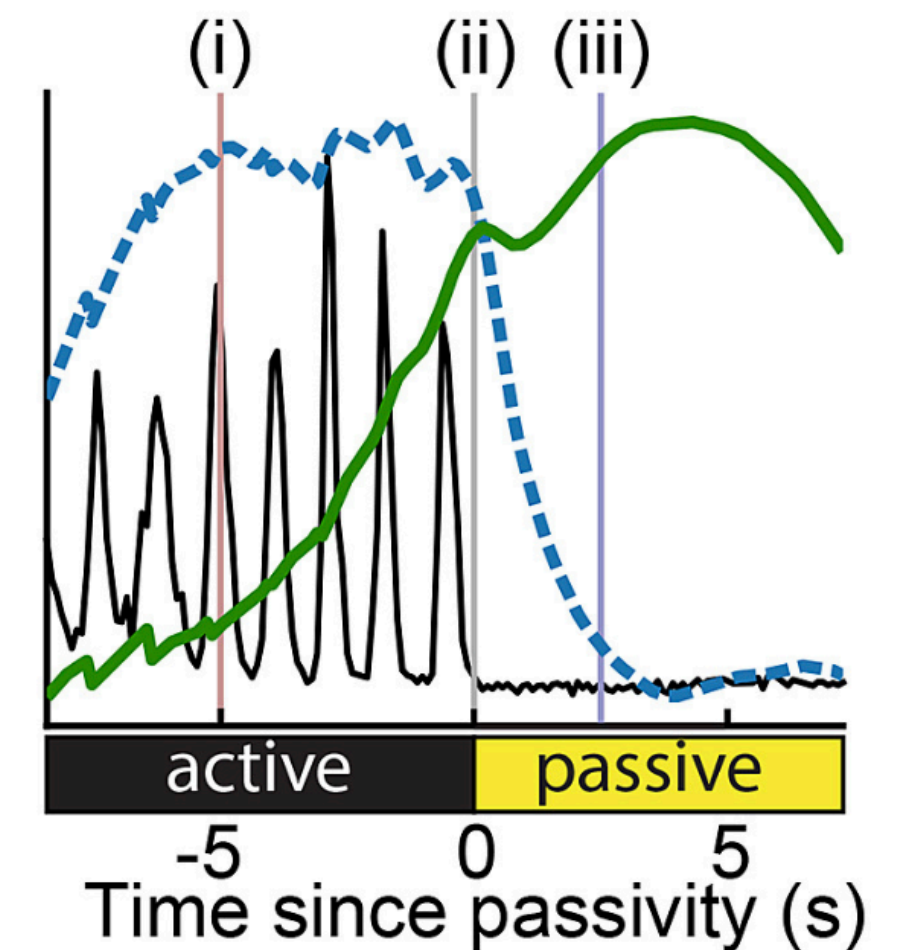




$\Delta f/f$



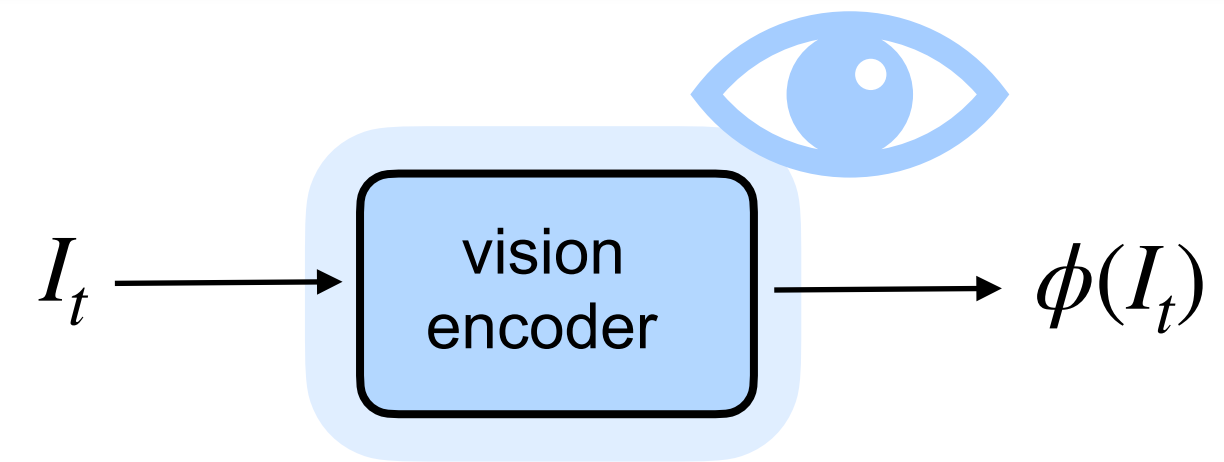
Fictive behavior
3.0X playback speed

Open loop

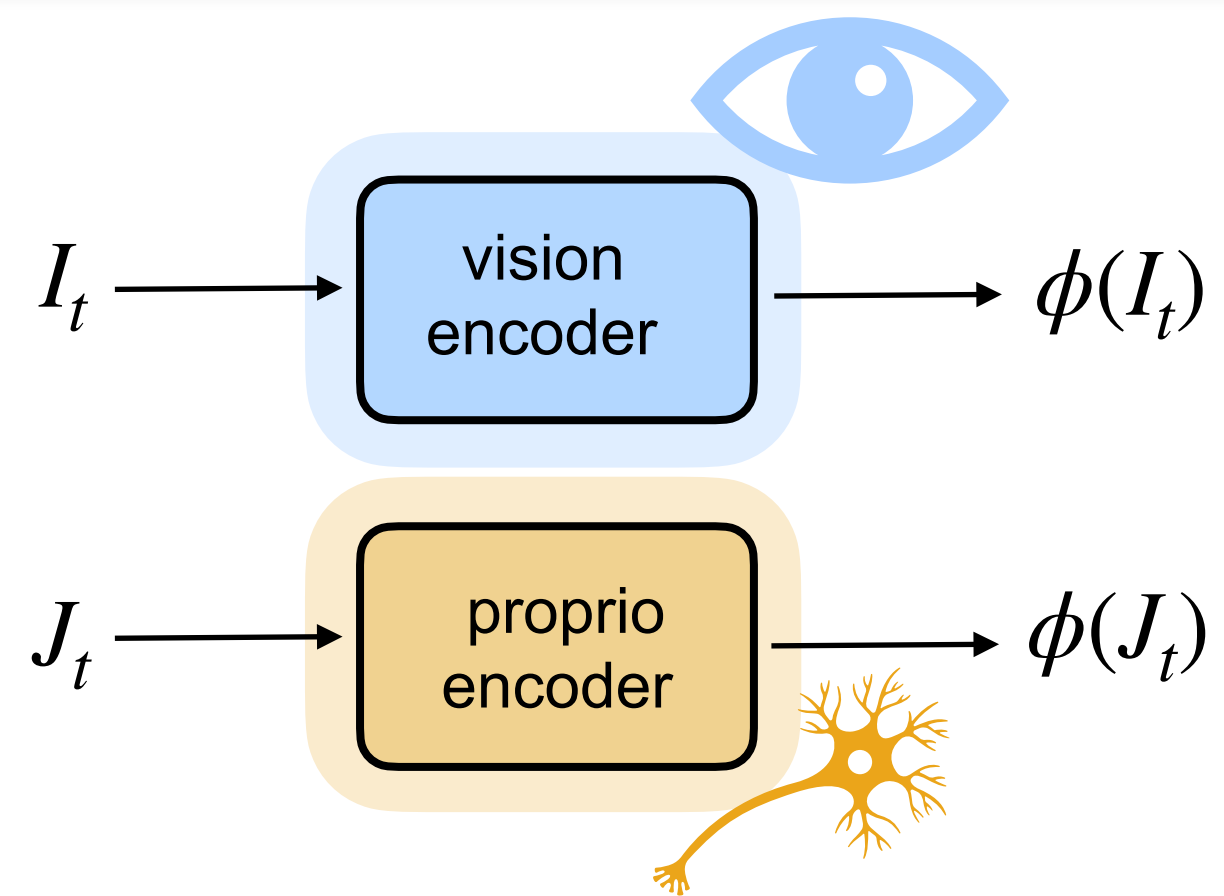


Virtual Agent Design


Virtual Agent Design



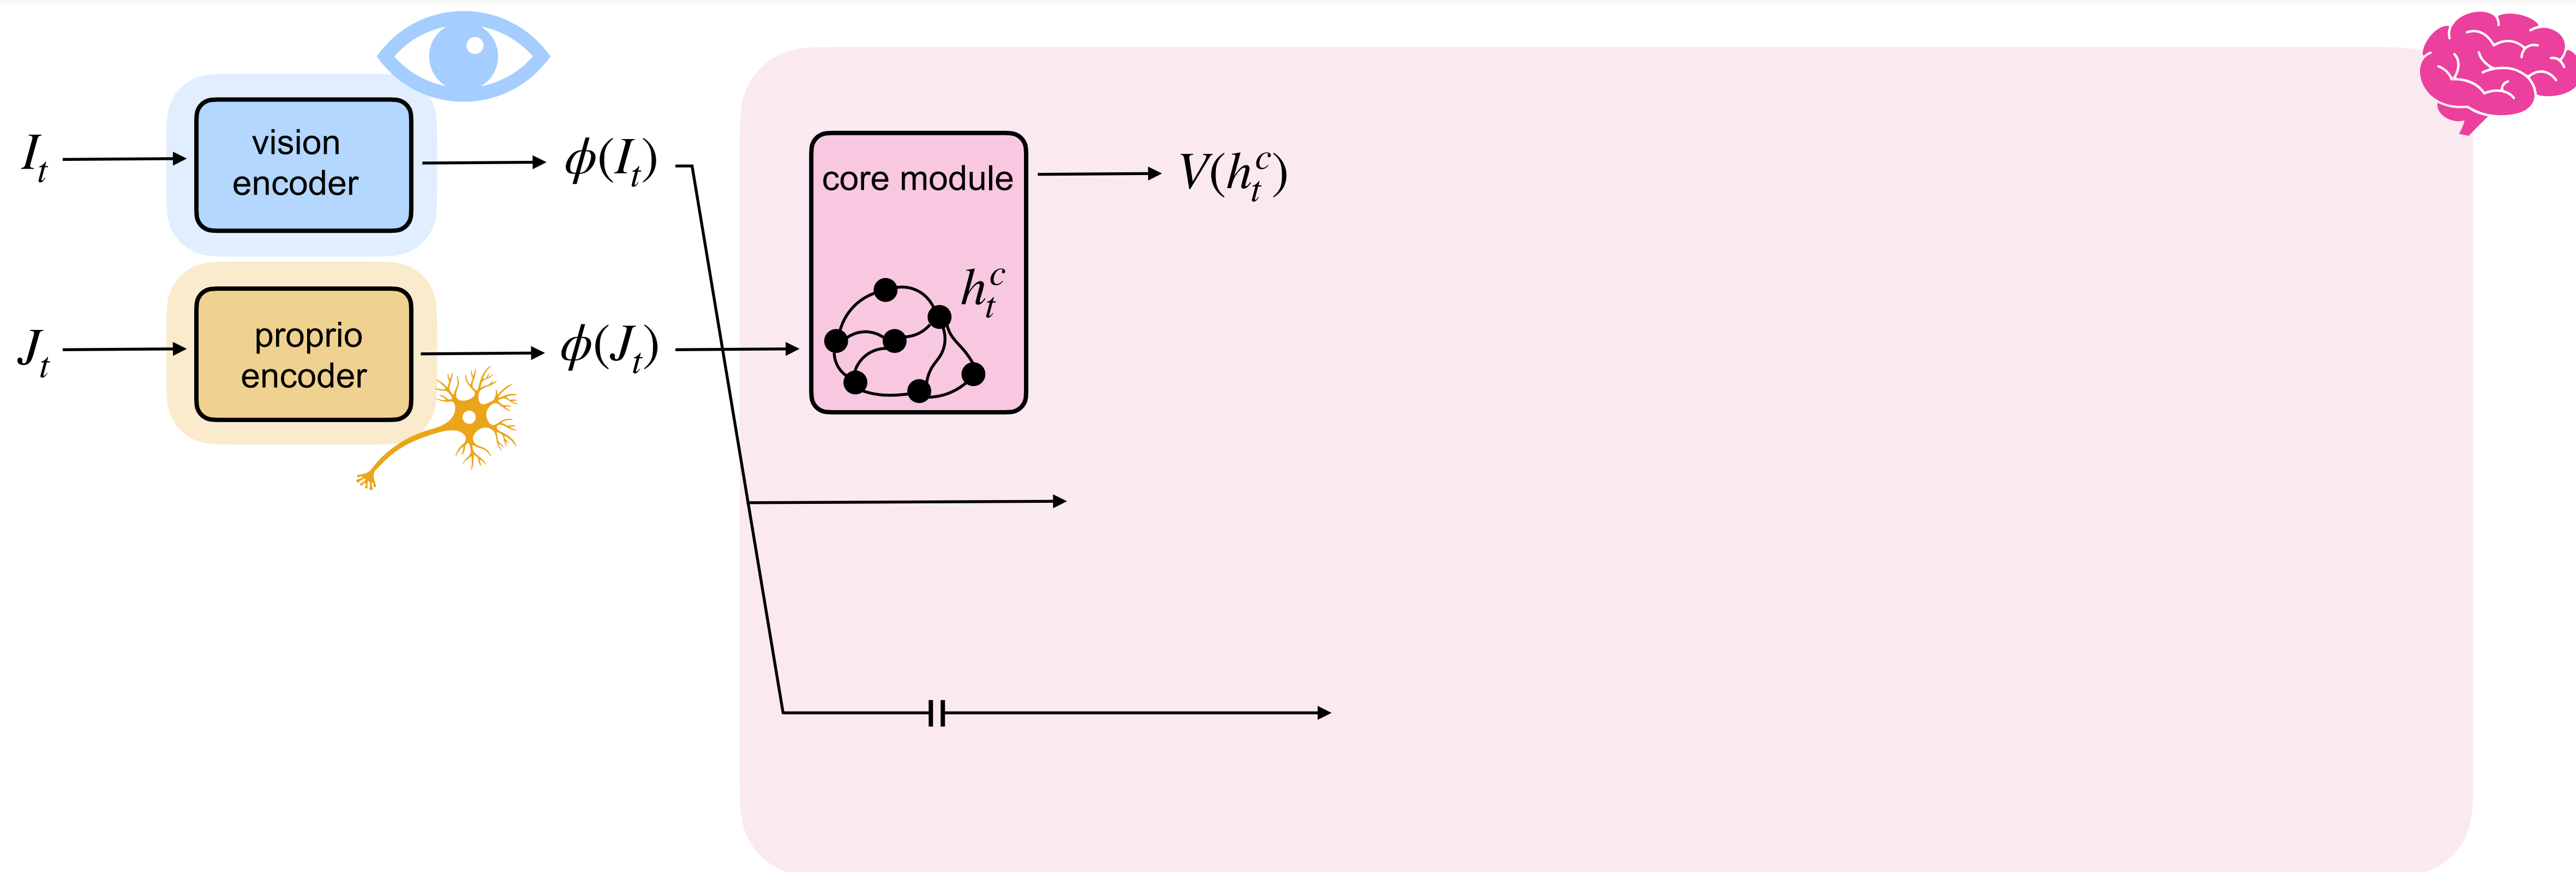
Virtual Agent Design



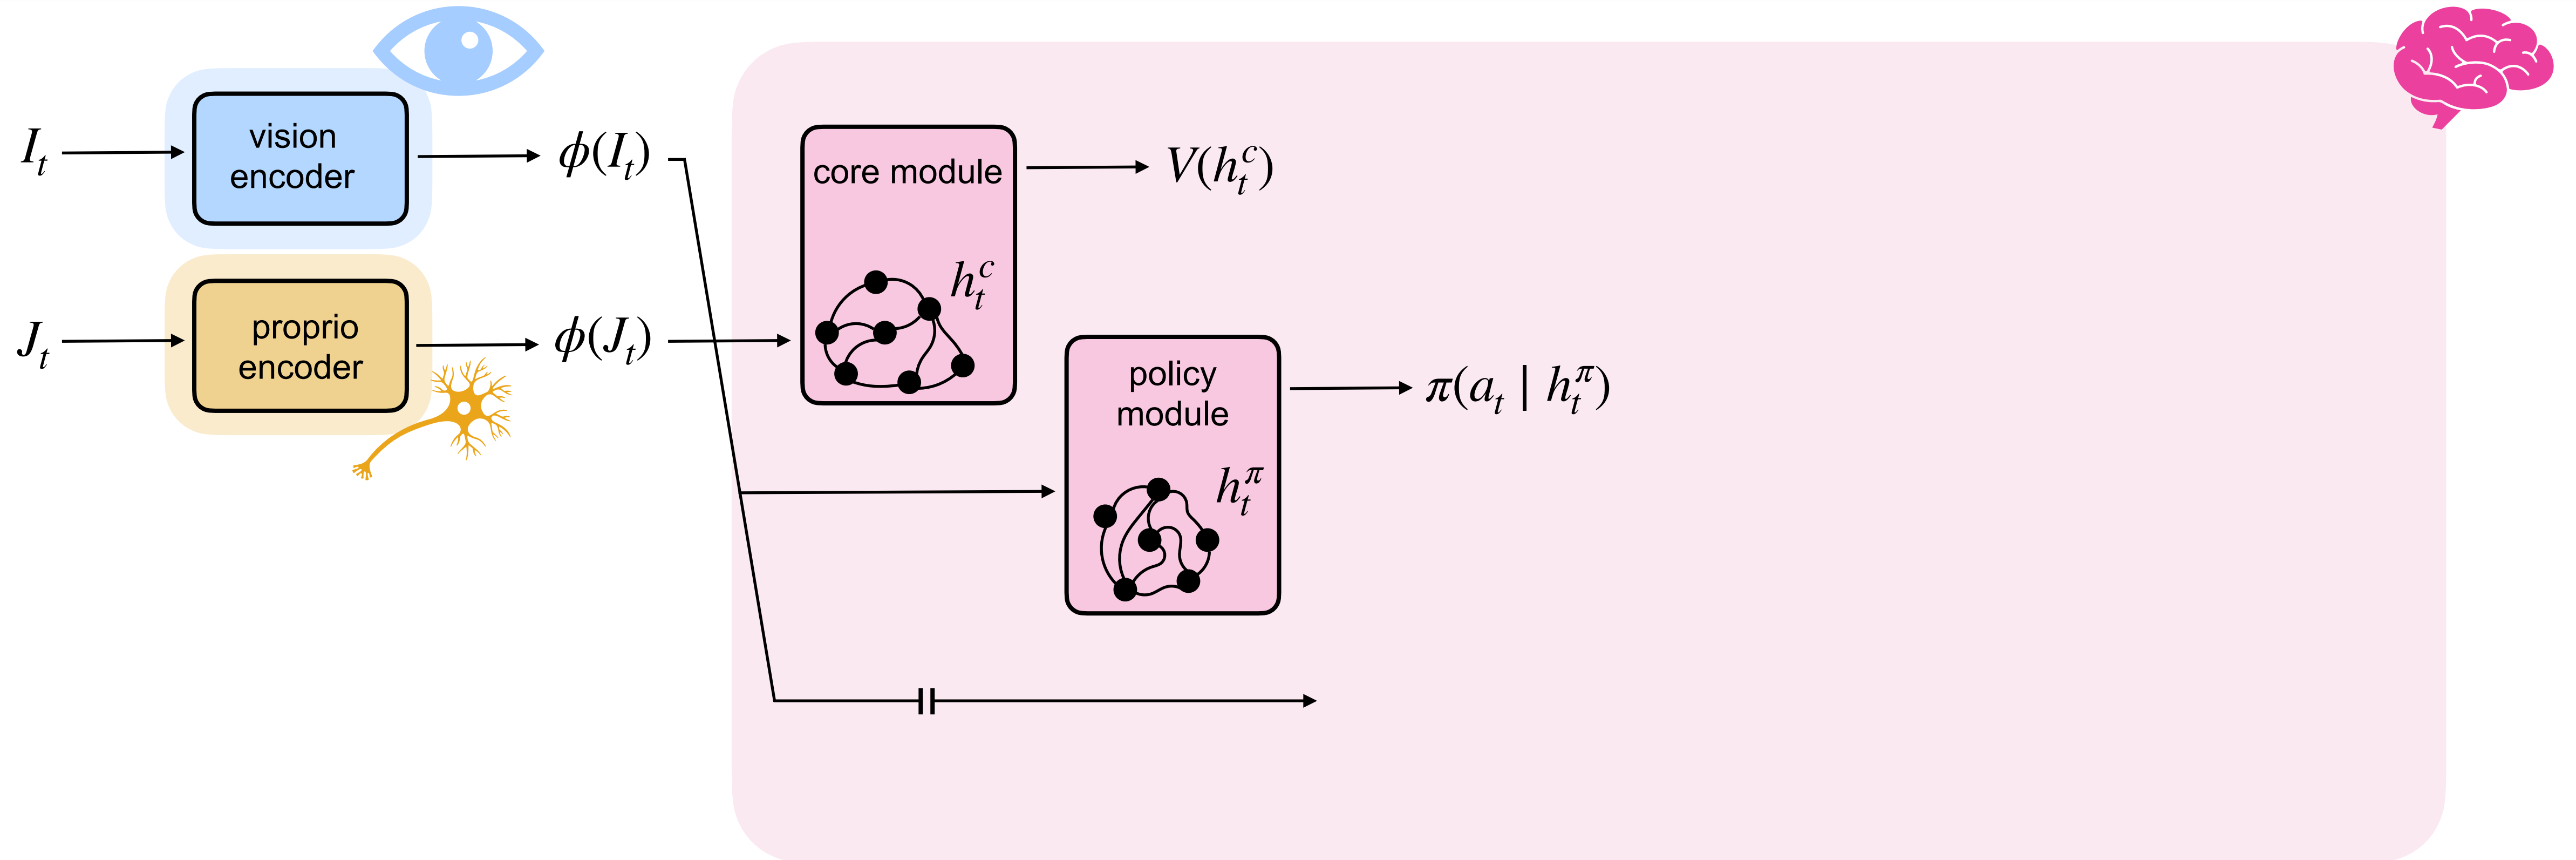
Virtual Agent Design



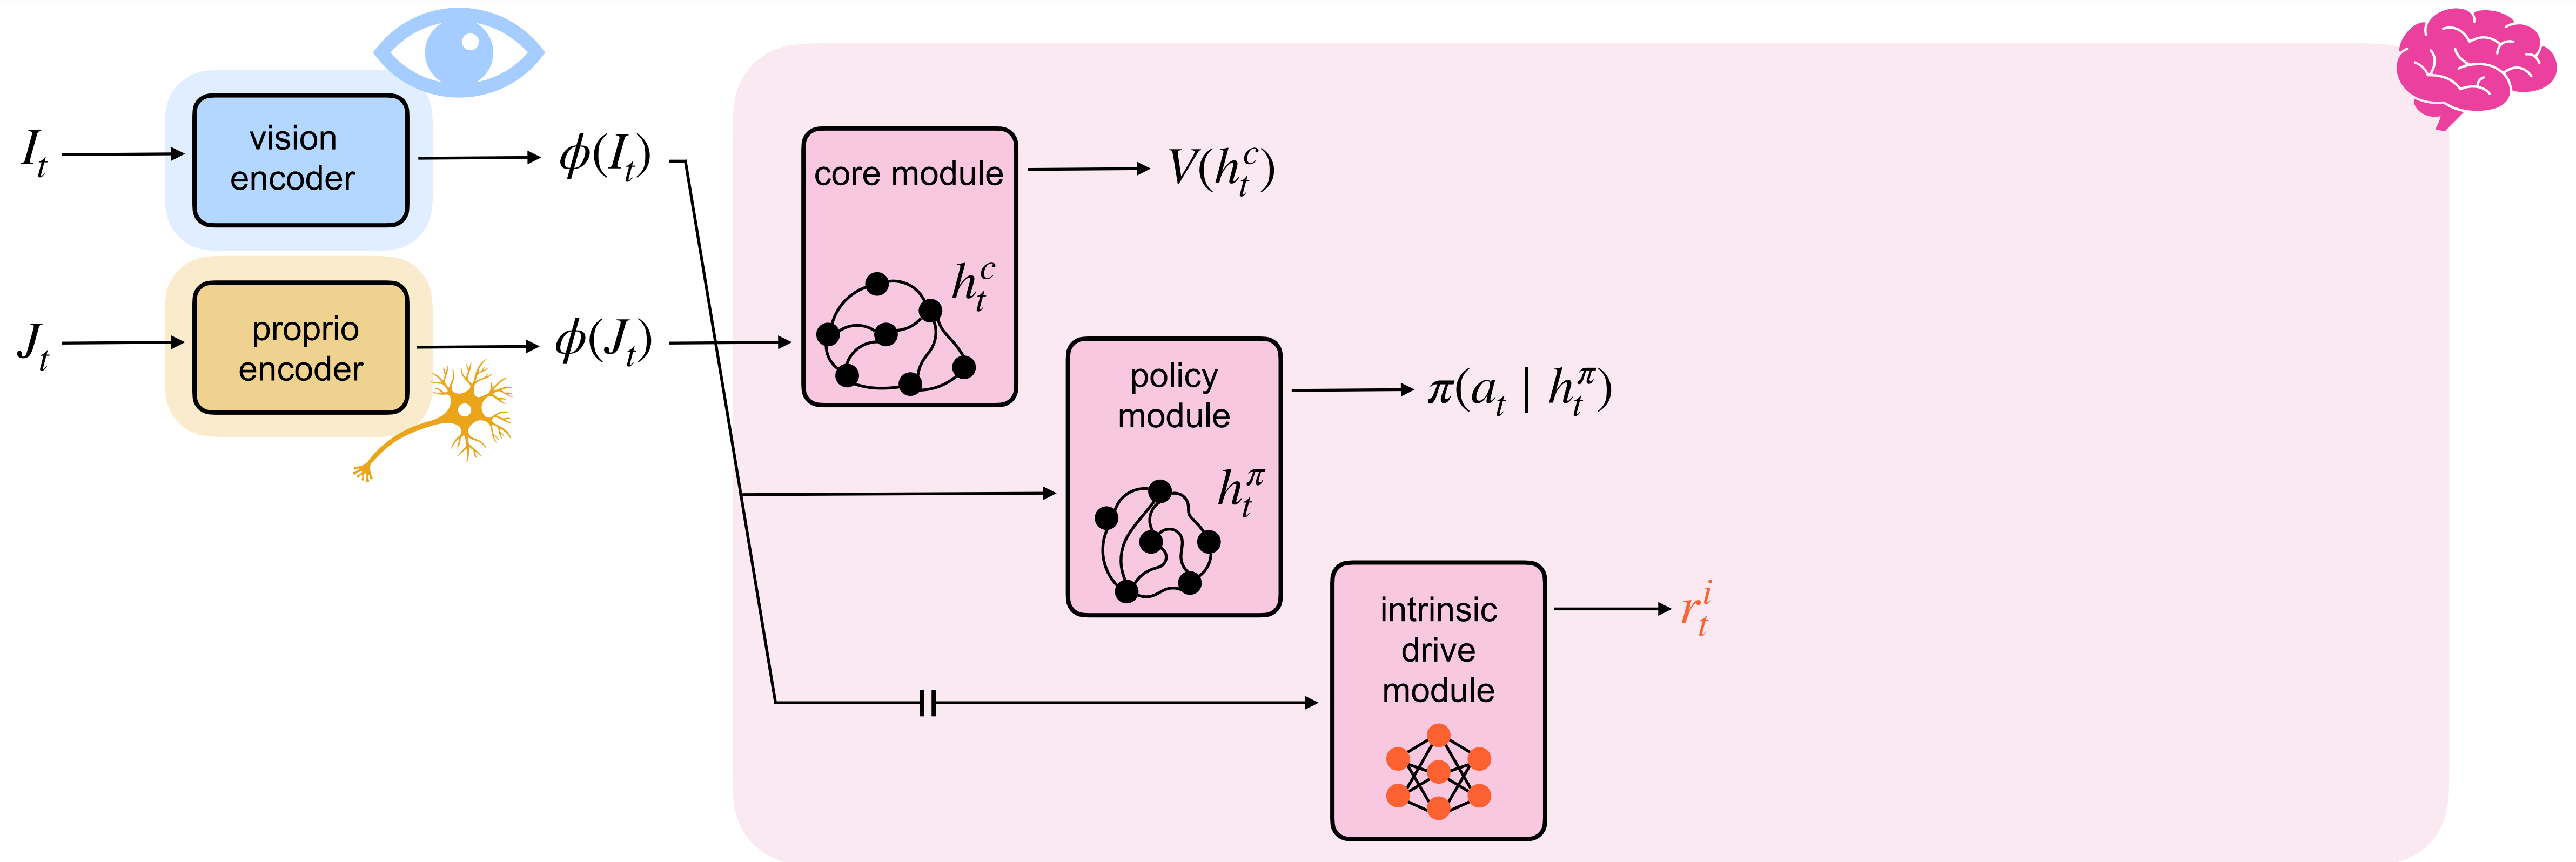
Virtual Agent Design



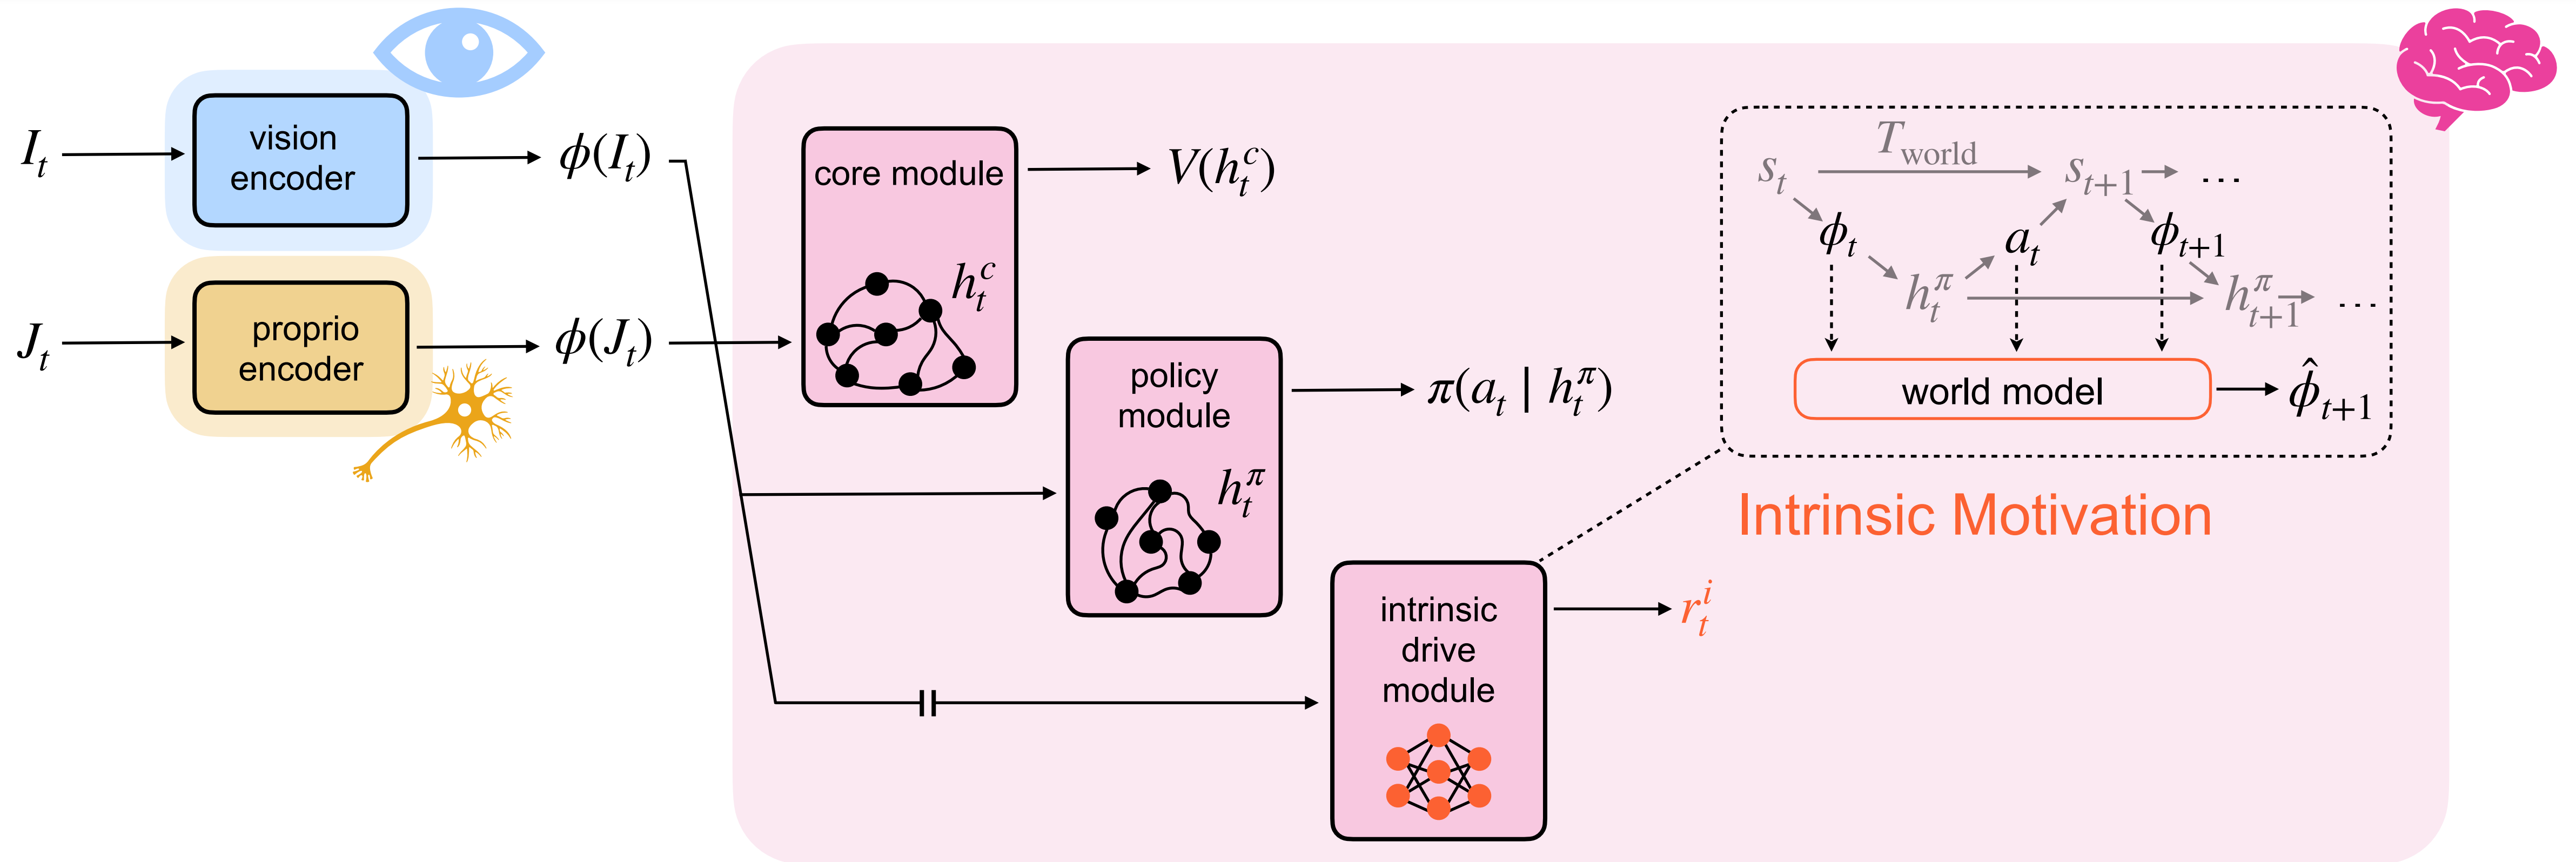
Virtual Agent Design



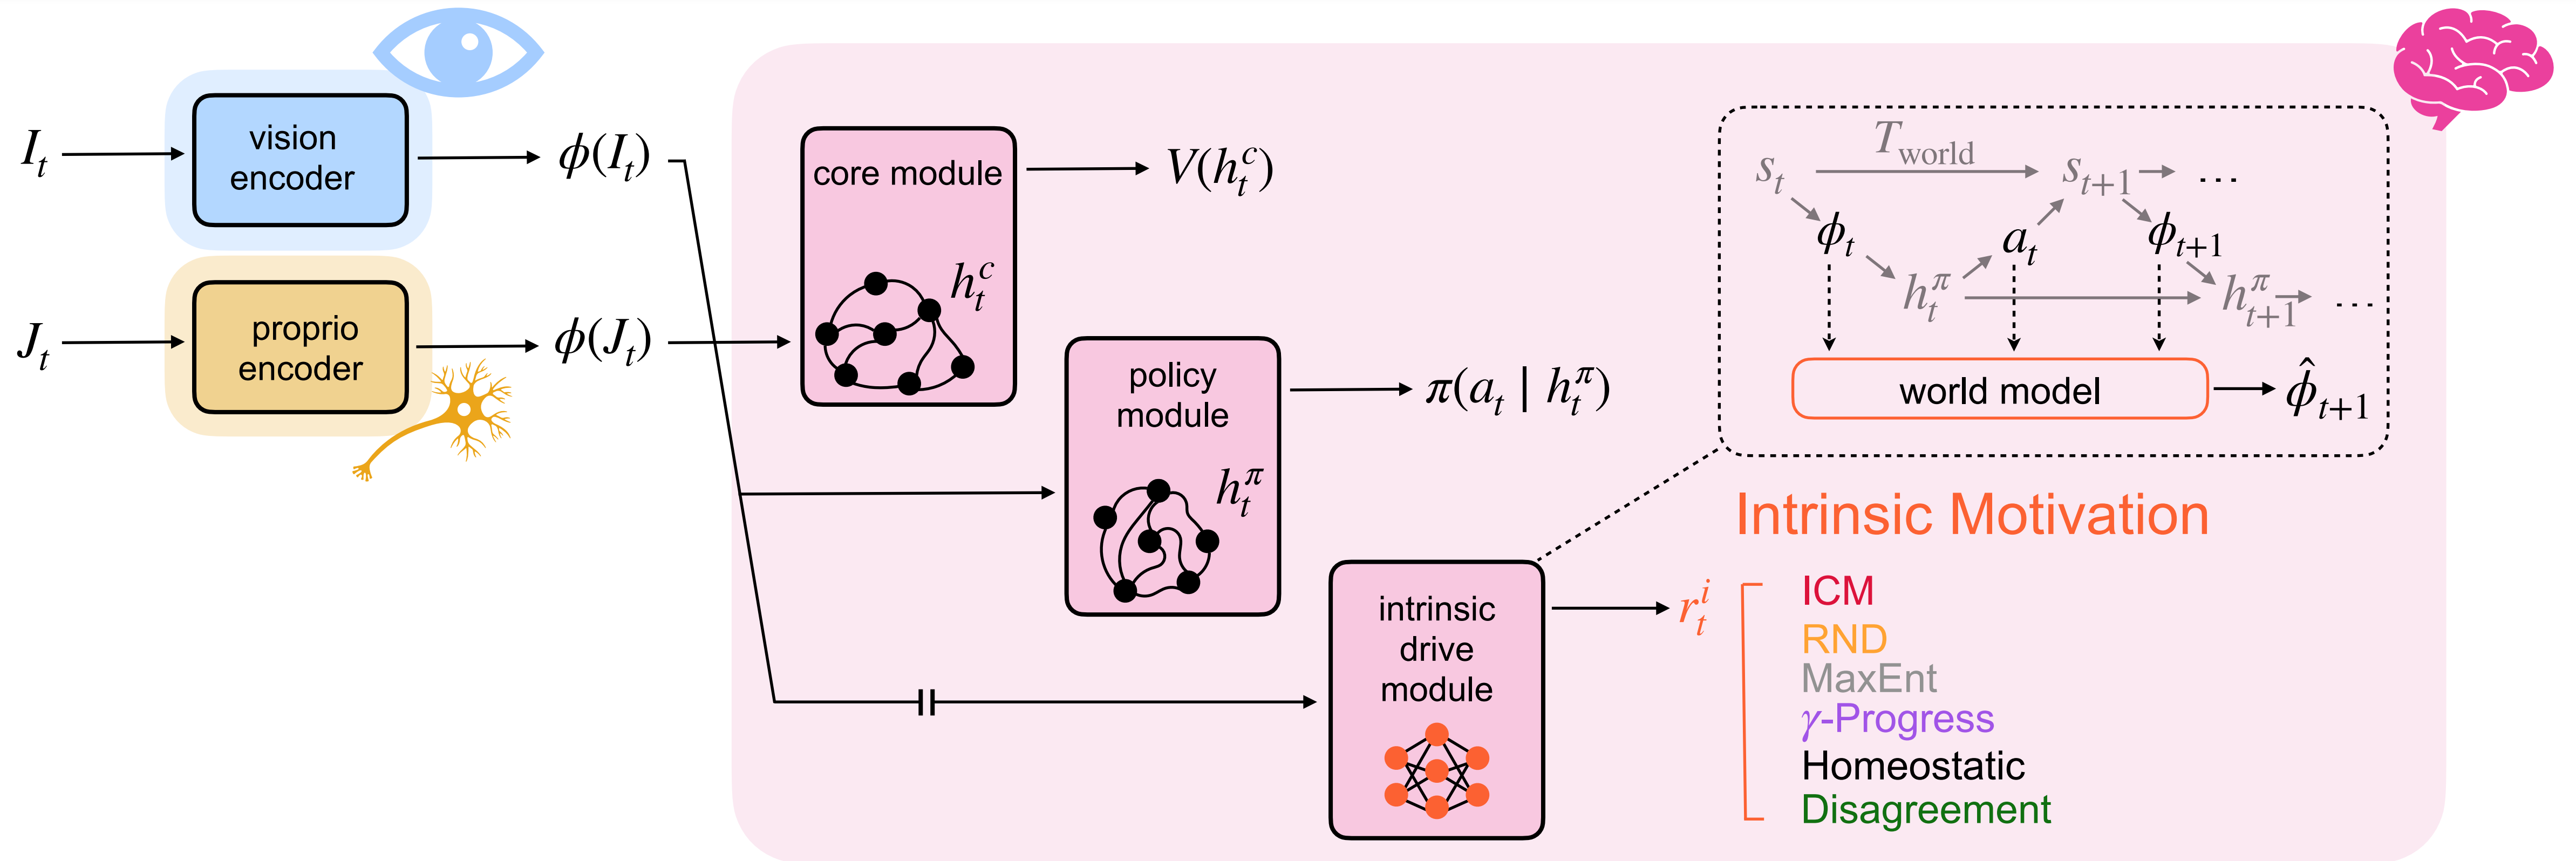
Virtual Agent Design




Virtual Agent Design



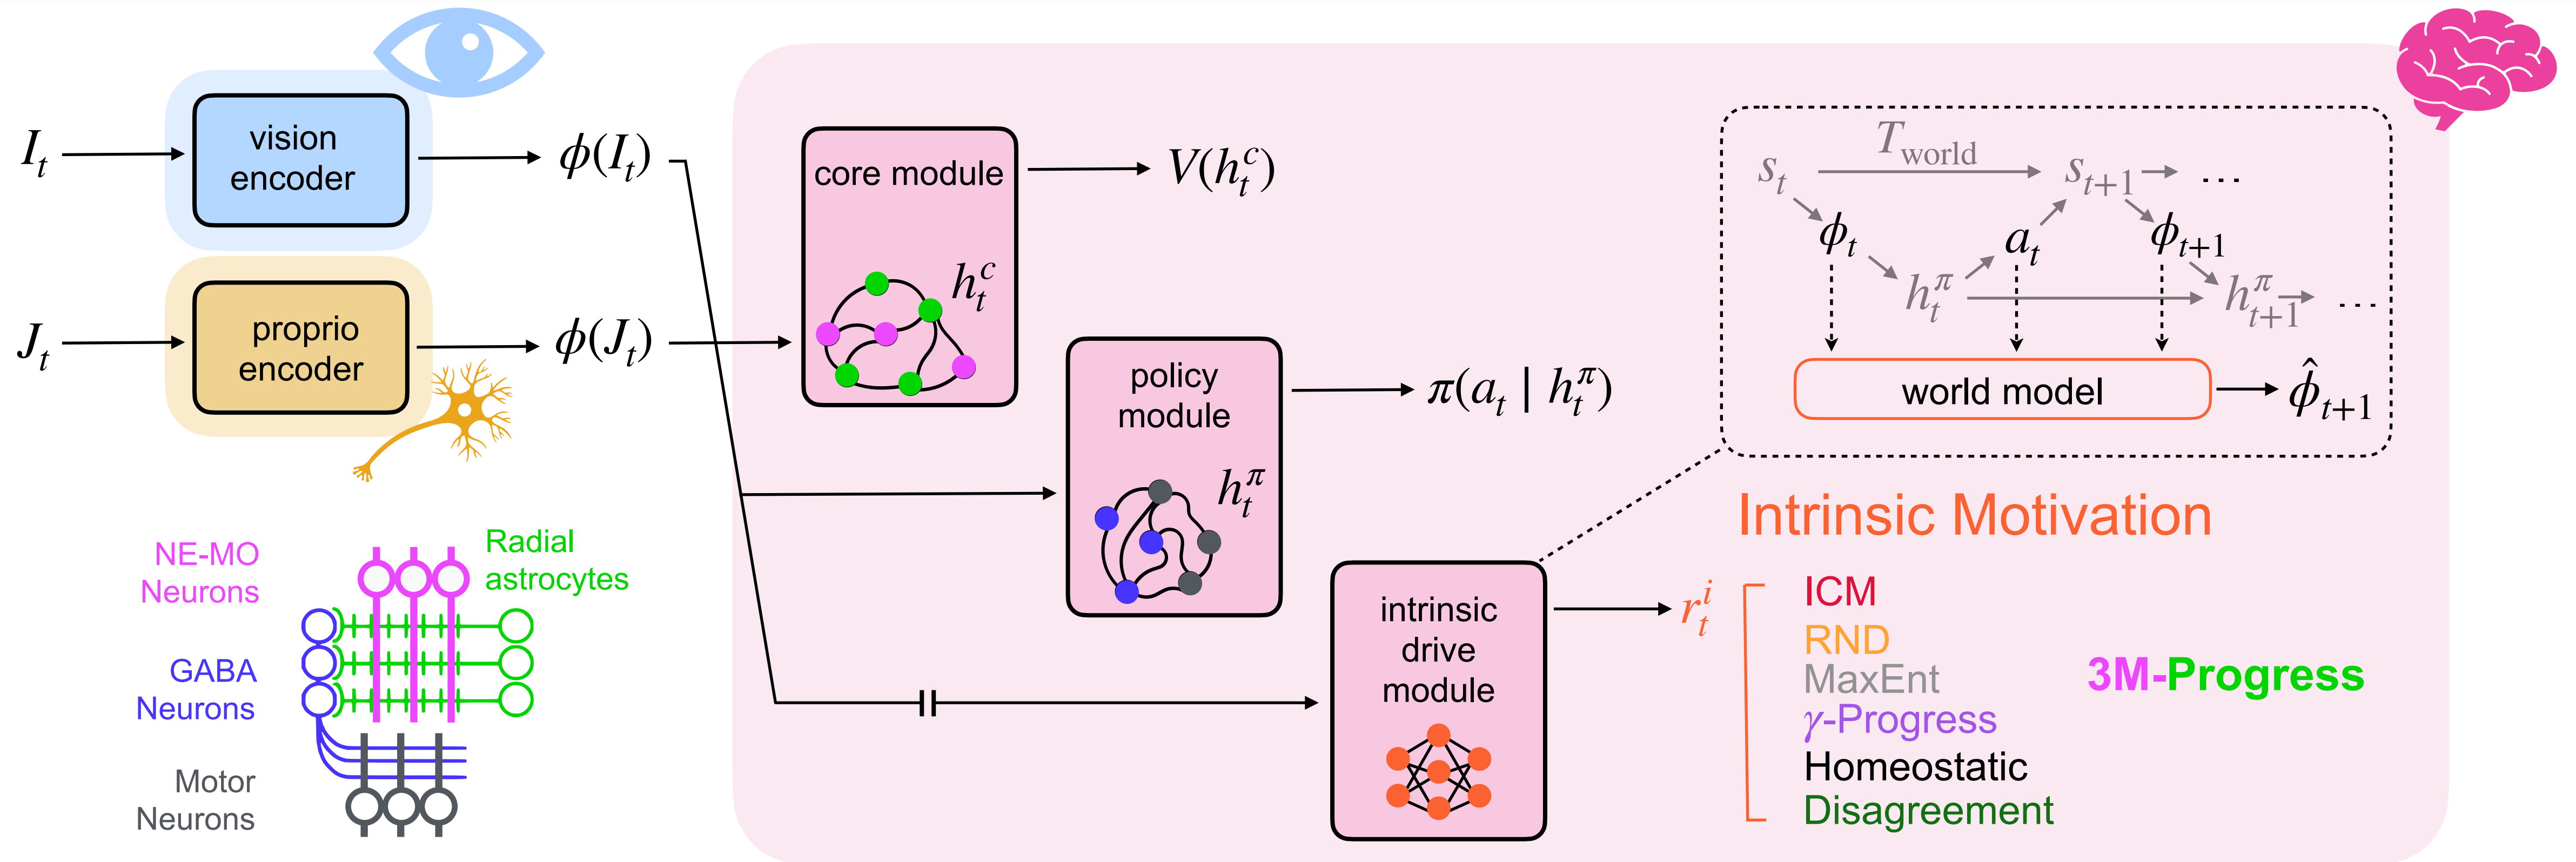
Virtual Agent Design



Virtual Agent Design






Virtual Agent Design




Environment / Embodiment


- 6-link articulation (5 DoF)
- Low Reynolds fluid approximation
- Non-stationary dynamics
- Egocentric observations



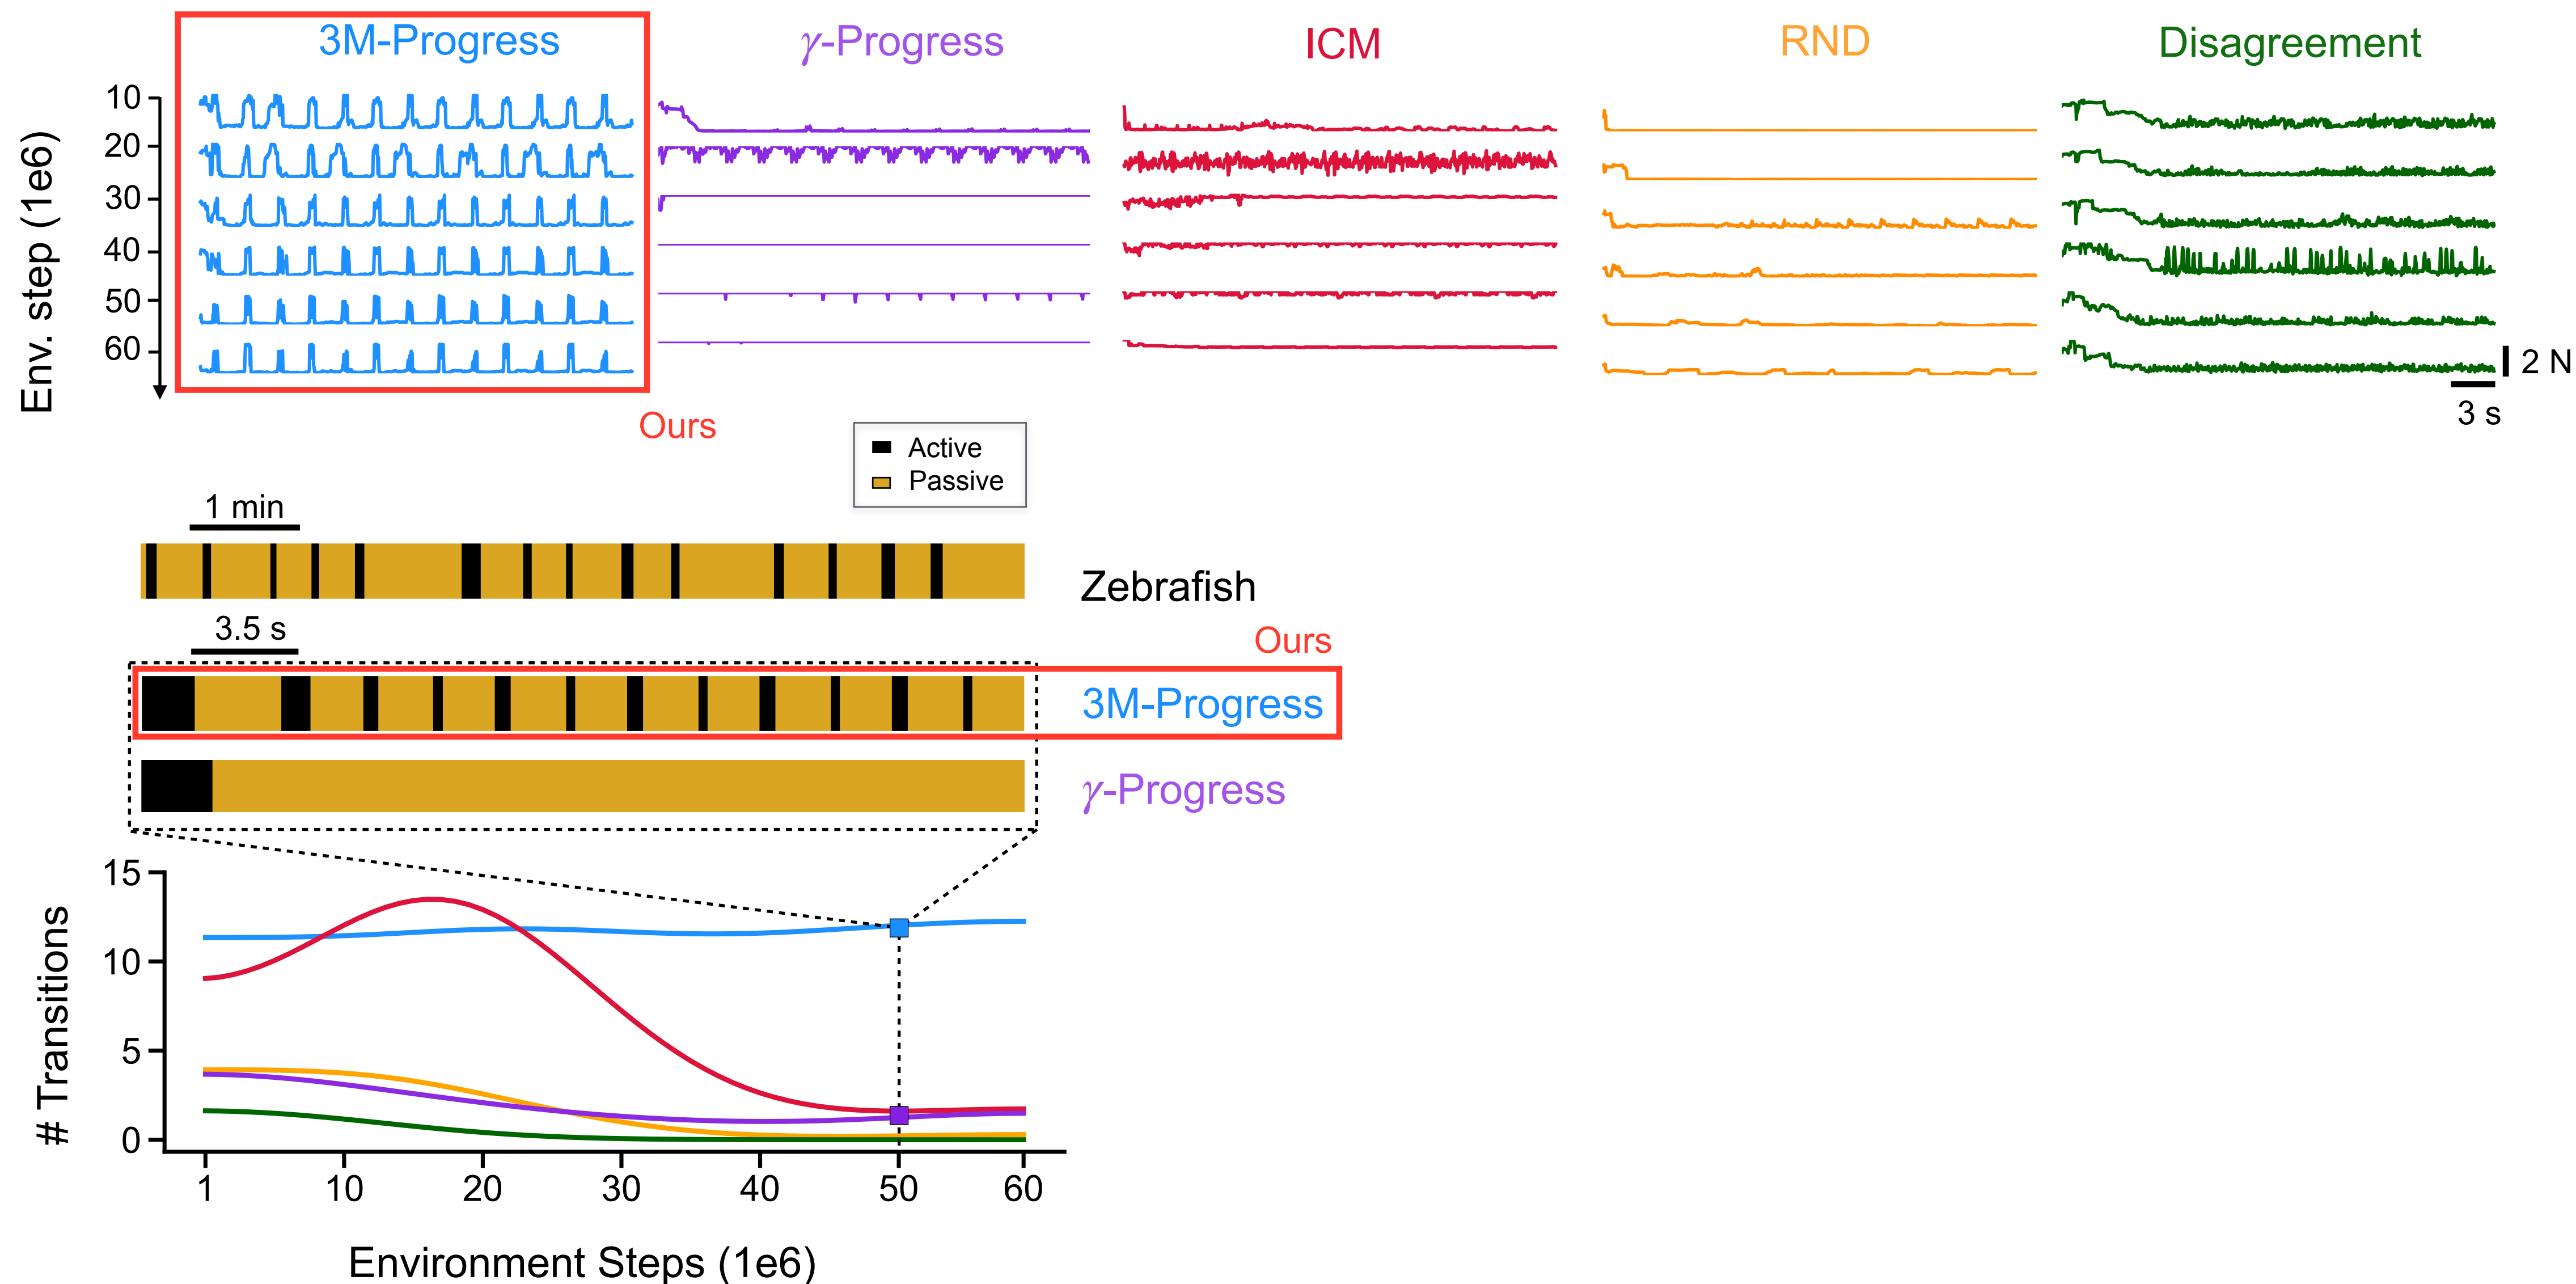
3M-Progress Captures Behavioral Transitions




3M-Progress Captures Behavioral Transitions



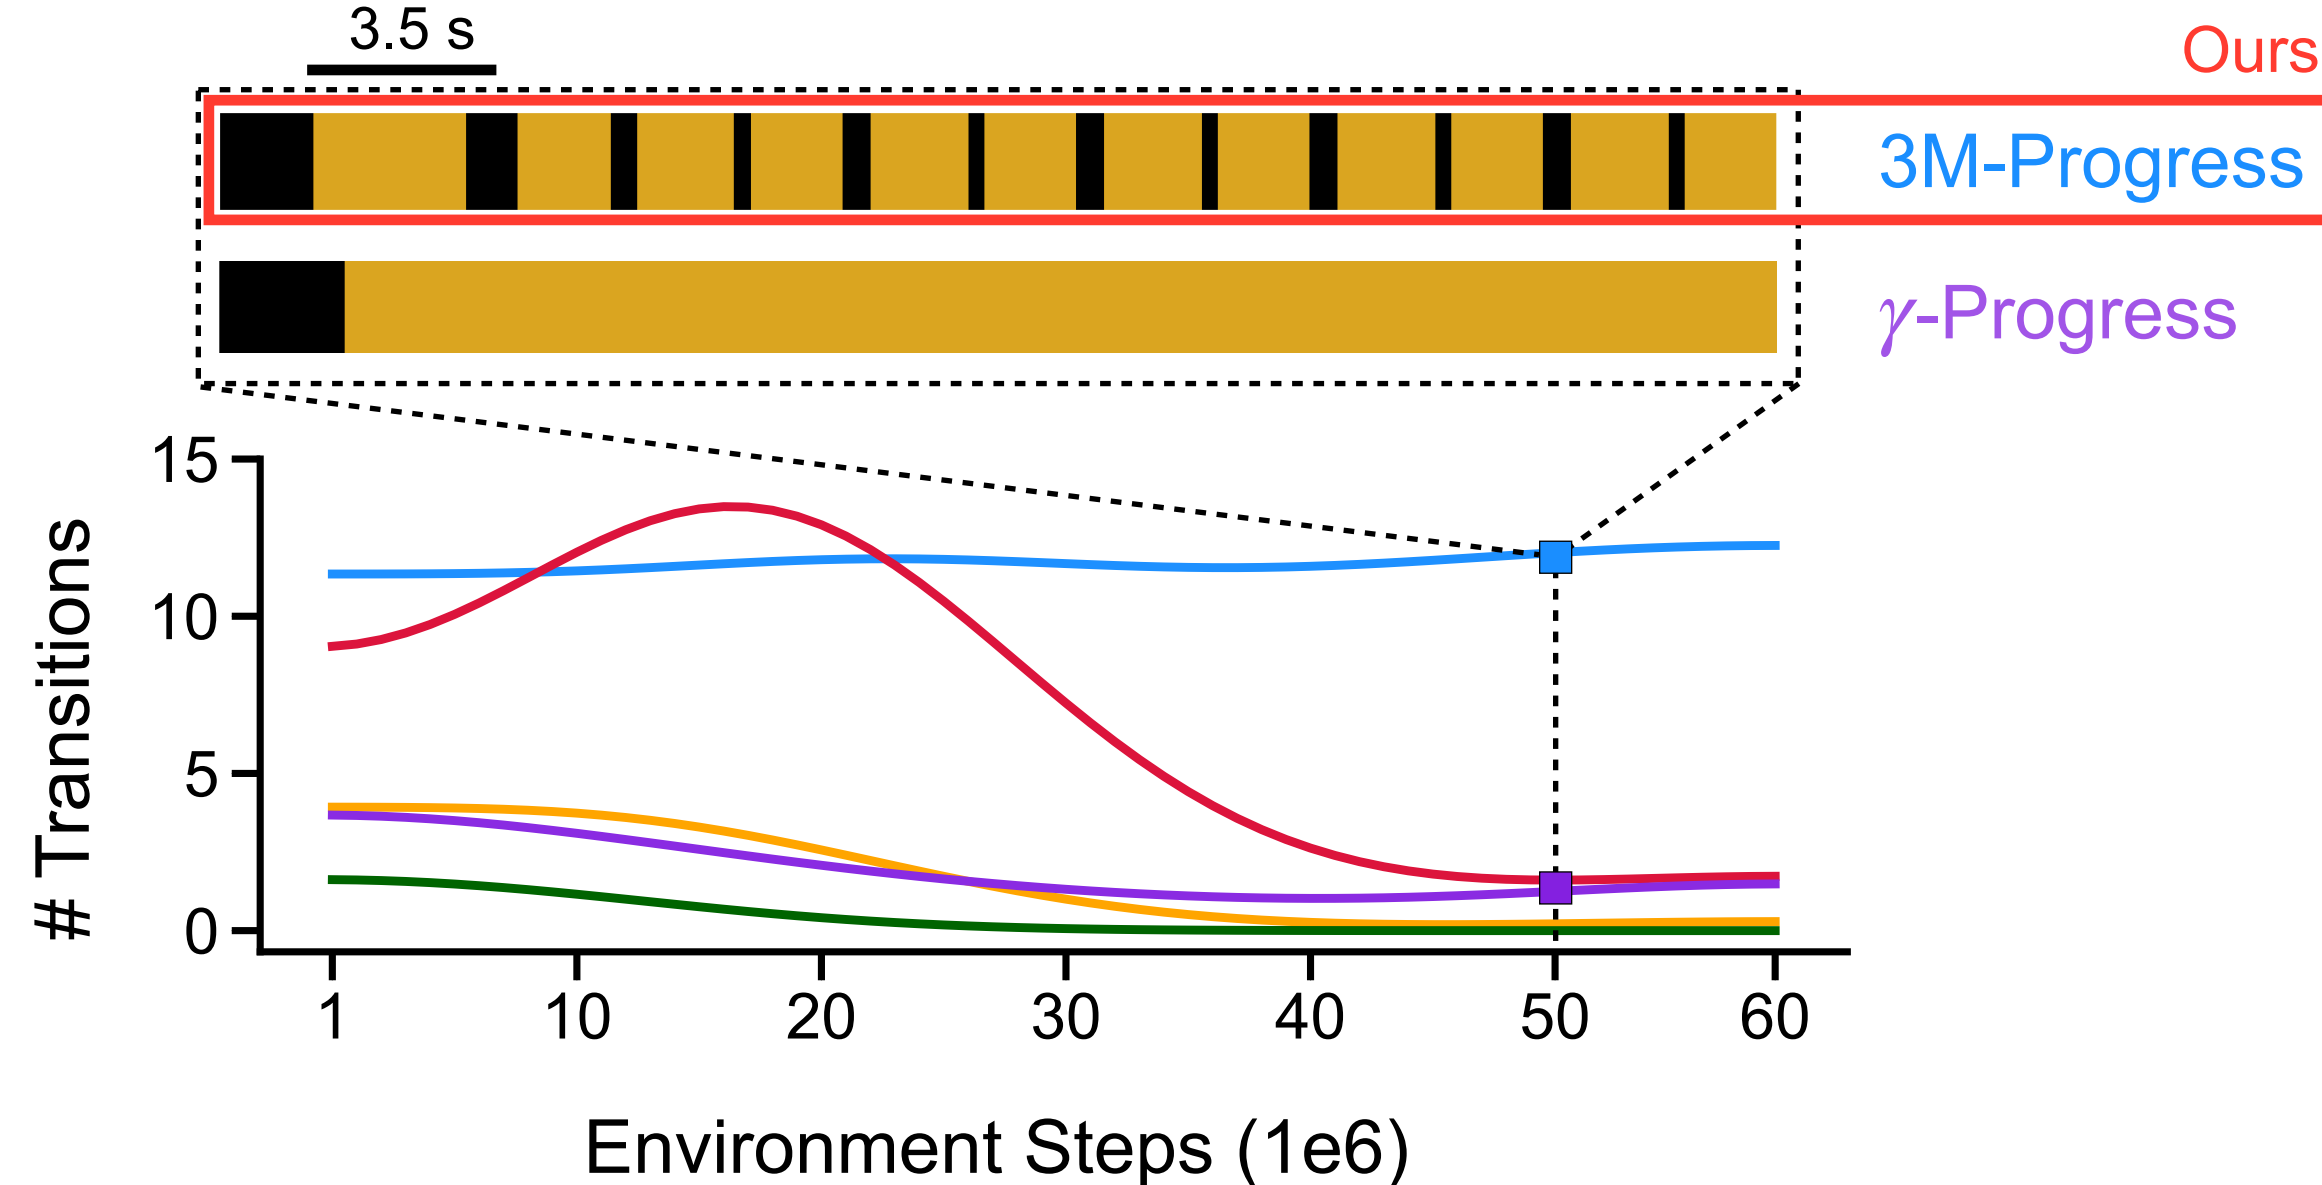
3M-Progress Captures Behavioral Transitions




3M-Progress Captures Behavioral Transitions



Why existing exploration methods fail

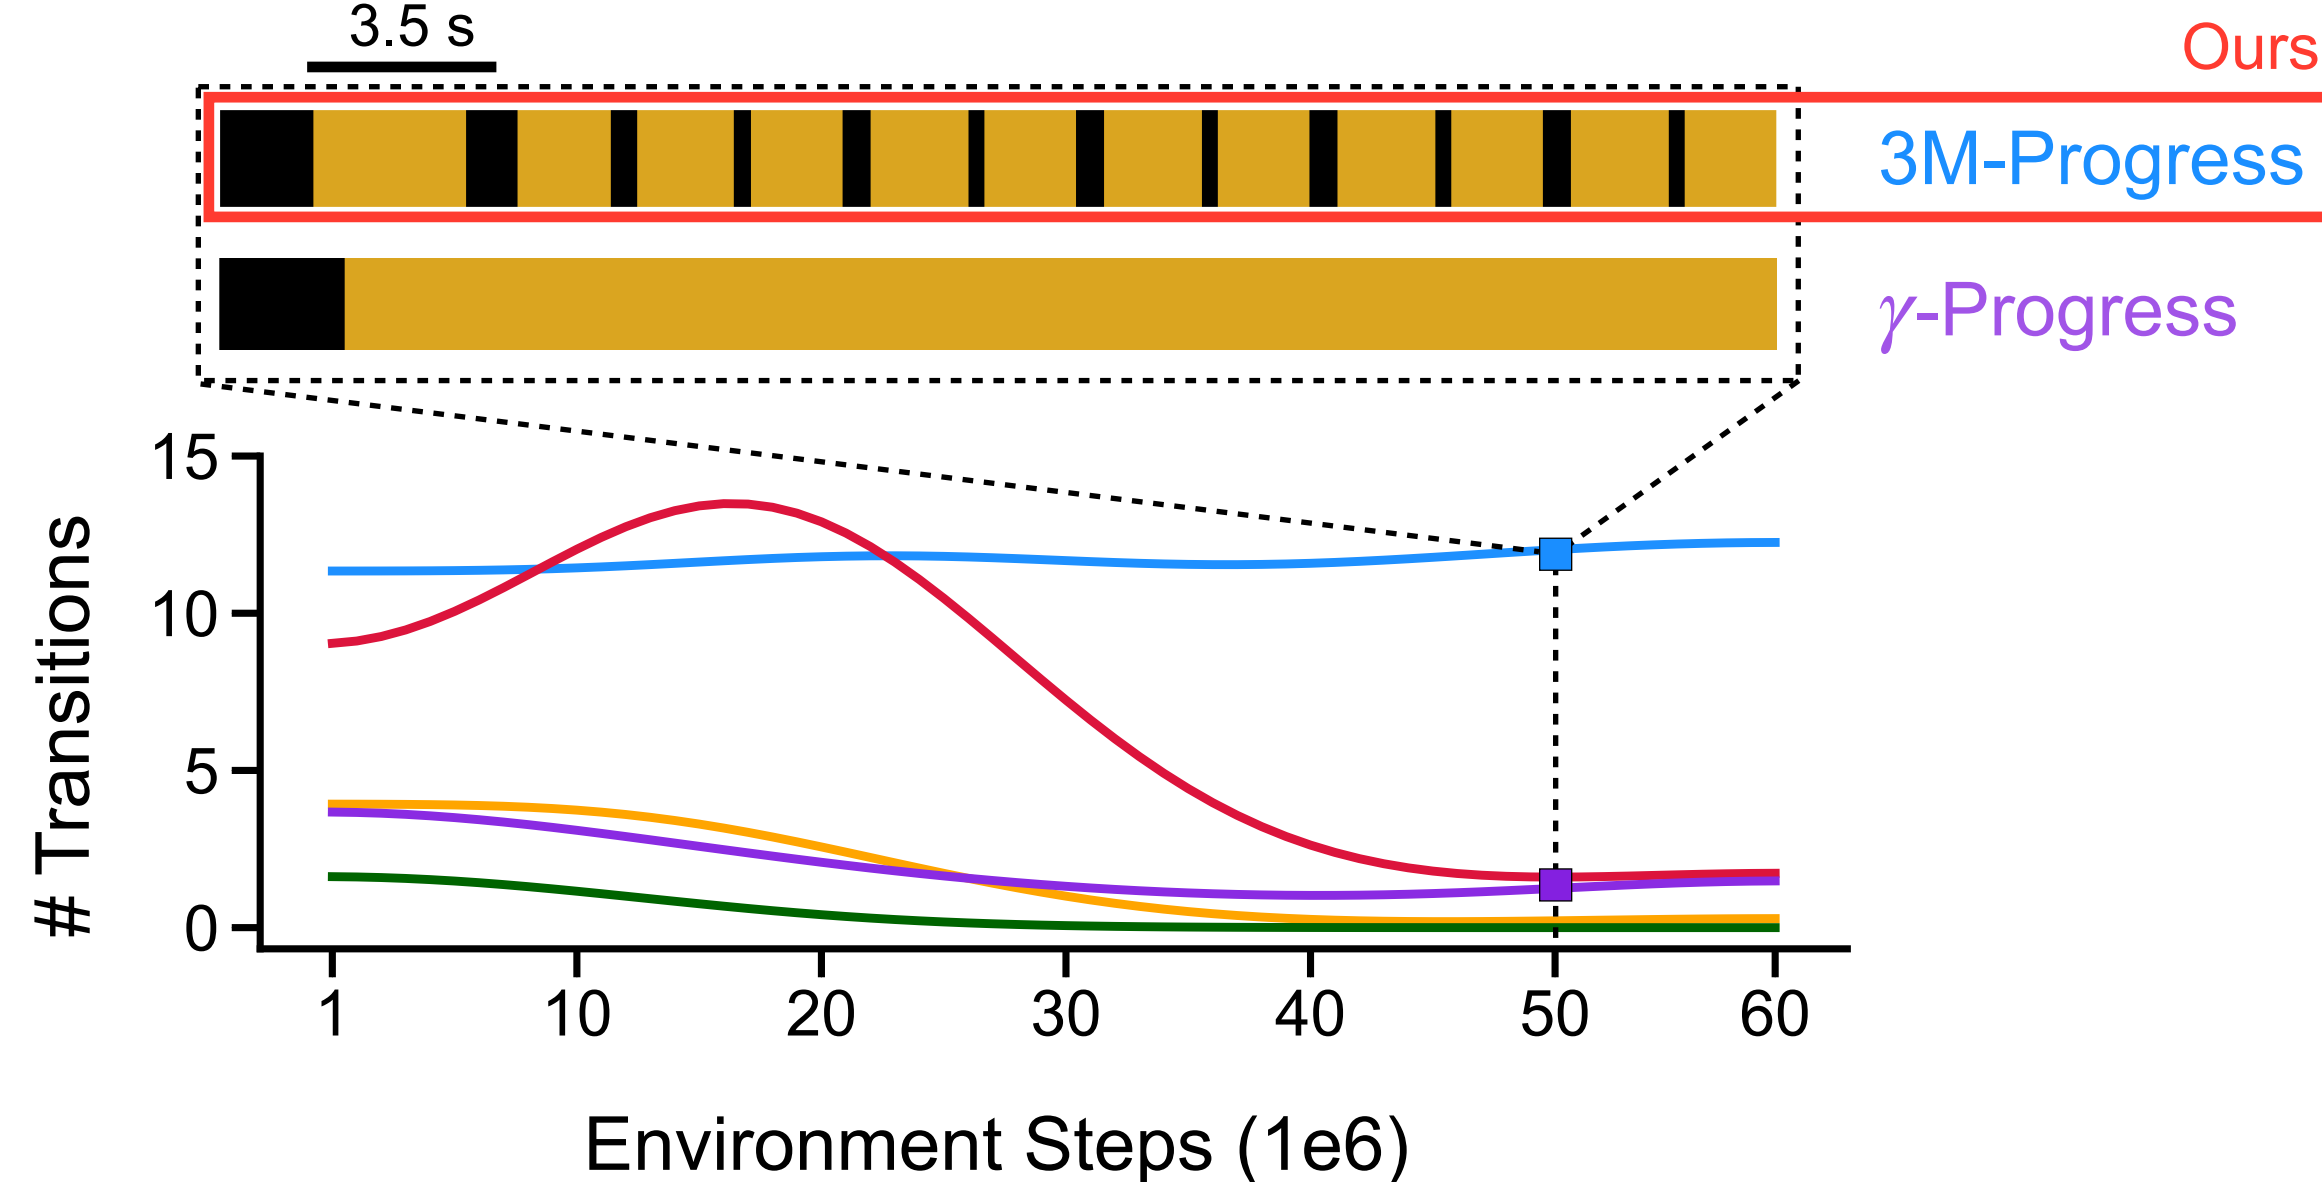


3M-Progress Captures Behavioral Transitions





Why existing exploration methods fail

Intrinsic reward $\propto f(\text{prediction-error})$



3M-Progress Captures Behavioral Transitions



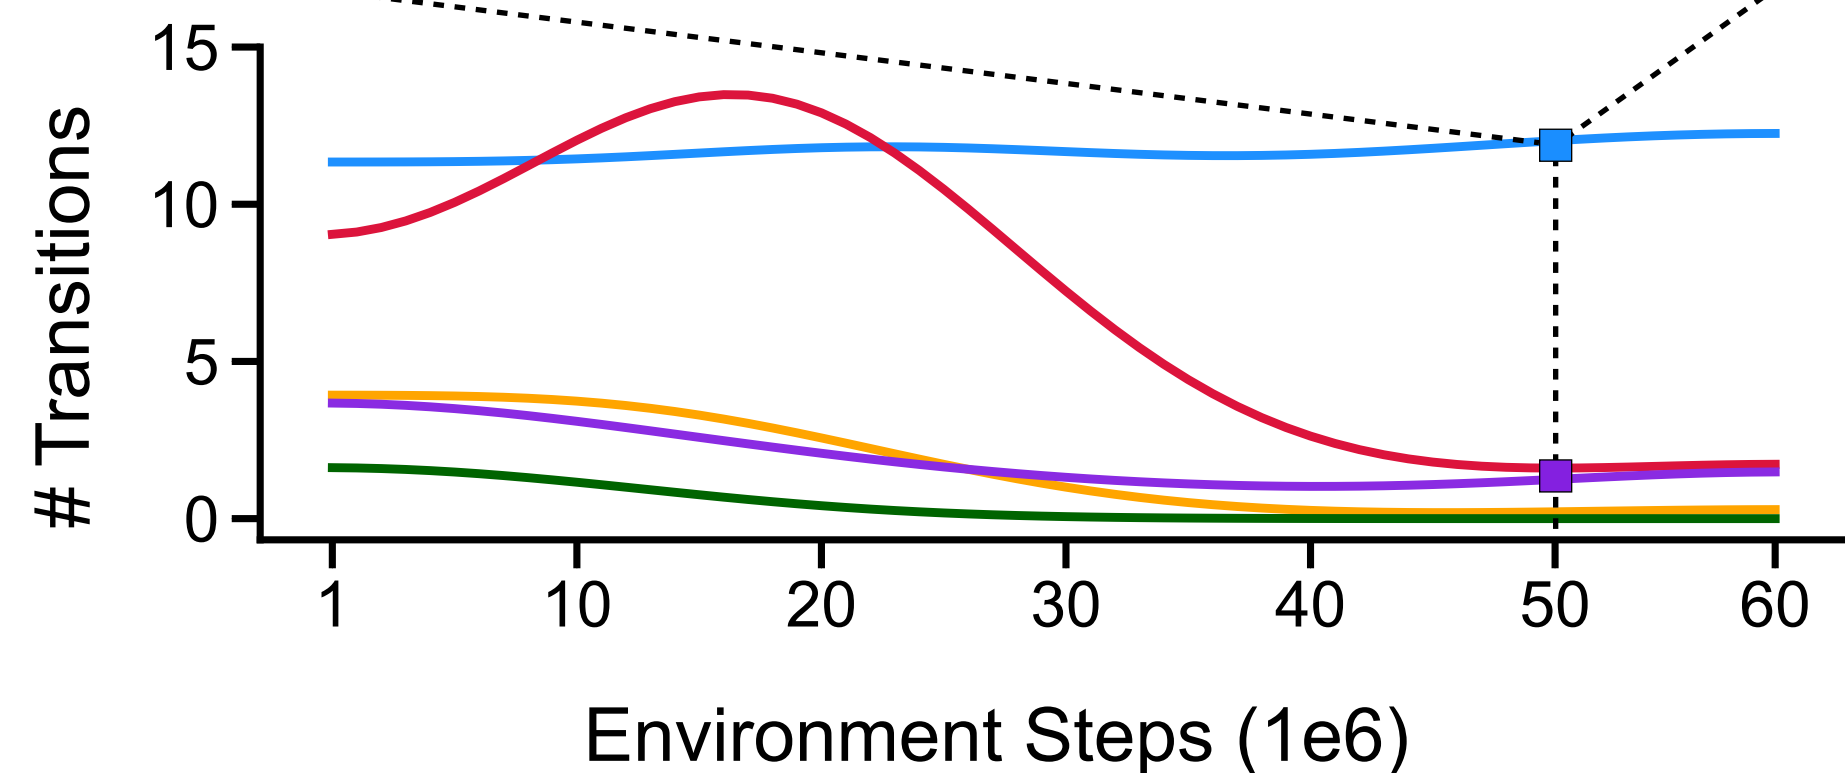

Why existing exploration methods fail

Intrinsic reward $\propto f(\text{prediction-error})$

1. **Non-stationary** and **saturates**

Consequence: transient behavioral strategies,
no policy convergence

3M-Progress Captures Behavioral Transitions



Why existing exploration methods fail

Intrinsic reward $\propto f(\text{prediction-error})$


1. **Non-stationary** and **saturates**

Consequence: transient behavioral strategies,
no policy convergence

2. **Rewards aleatoric uncertainty**

Consequence: perseveration on unlearnable and
uncontrollable stimuli.

Ethological Priors Guide Adaptive Behavior




What is the most fundamental tool of an explorer?




Infante Dom Henrique

Ethological Priors Guide Adaptive Behavior




What is the most fundamental tool of an explorer?



Infante Dom Henrique

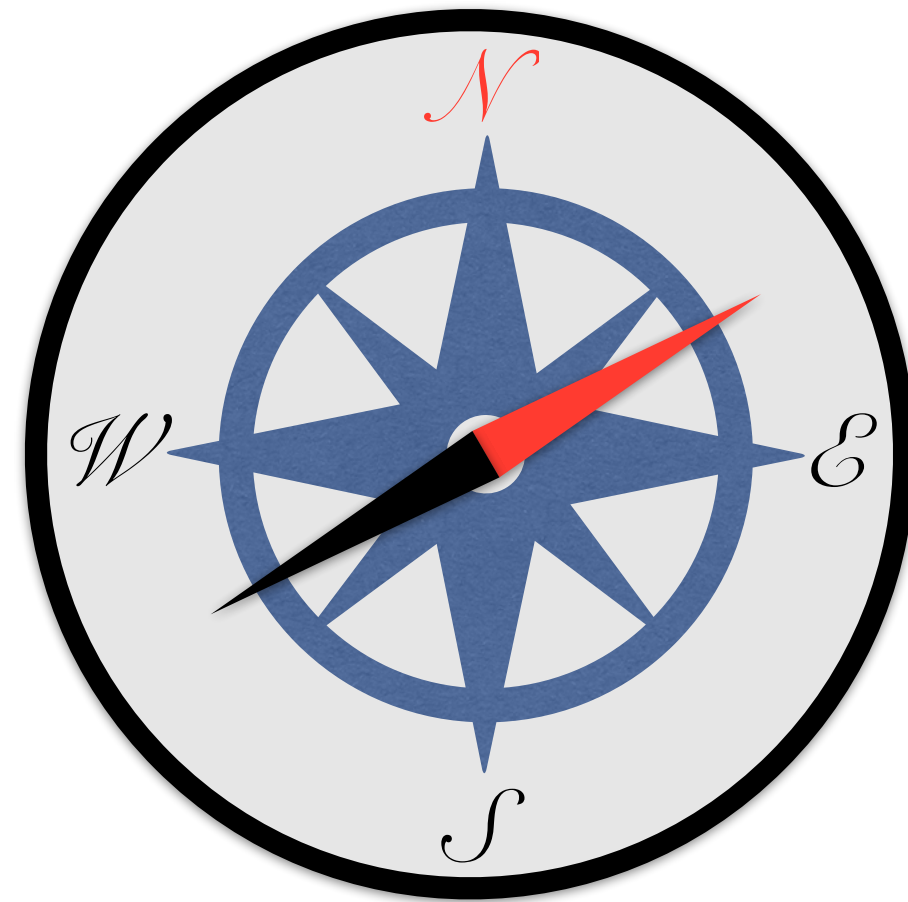
Ethological Priors Guide Adaptive Behavior



What is the most fundamental tool of an explorer?




Infante Dom Henrique



What is *North* for discovering new behaviors?

Ethological Priors Guide Adaptive Behavior




Mariner's Astrolabe

What is the most fundamental tool of an explorer?




Infante Dom Henrique

Ecological Niche




$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics



What is *North* for discovering new behaviors?

Ethological Priors Guide Adaptive Behavior




What is the most fundamental tool of an explorer?




Infante Dom Henrique

Ecological Niche




$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics



What is *North* for discovering new behaviors?

Ethological Priors Guide Adaptive Behavior




What is the most fundamental tool of an explorer?




Infante Dom Henrique

Ecological Niche




$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics



Unseen environment




$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics

What is *North* for discovering new behaviors?

Ethological Priors Guide Adaptive Behavior

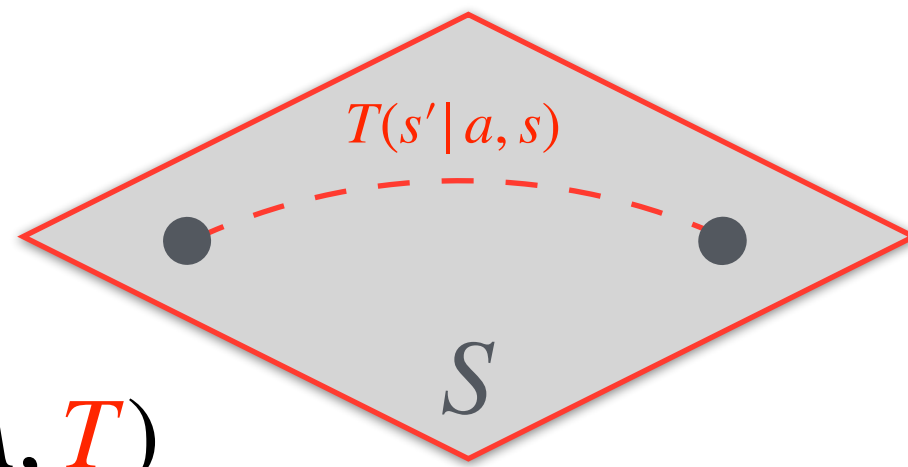


What is the most fundamental tool of an explorer?




Infante Dom Henrique

Ecological Niche




$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics



Unseen environment




$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics

What is *North* for discovering new behaviors?
The **Ecological Niche** provides a dynamics prior.

Ethological Priors Guide Adaptive Behavior




What is the most fundamental tool of an explorer?




Infante Dom Henrique

Unseen environment




Ecological Niche



$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics


What is *North* for discovering new behaviors?
The *Ecological Niche* provides a dynamics prior.



$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics

Ethological Priors Guide Adaptive Behavior




What is the most fundamental tool of an explorer?




Infante Dom Henrique

Ecological Niche




$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics



Unseen environment




$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics

What is *North* for discovering new behaviors?
The *Ecological Niche* provides a dynamics prior.


Ethological Priors Guide Adaptive Behavior



What is the most fundamental tool of an explorer?




Ecological Niche




$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics



Unseen environment




$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics

What is *North* for discovering new behaviors?
The *Ecological Niche* provides a dynamics prior.


Ethological Priors Guide Adaptive Behavior



What is the most fundamental tool of an explorer?

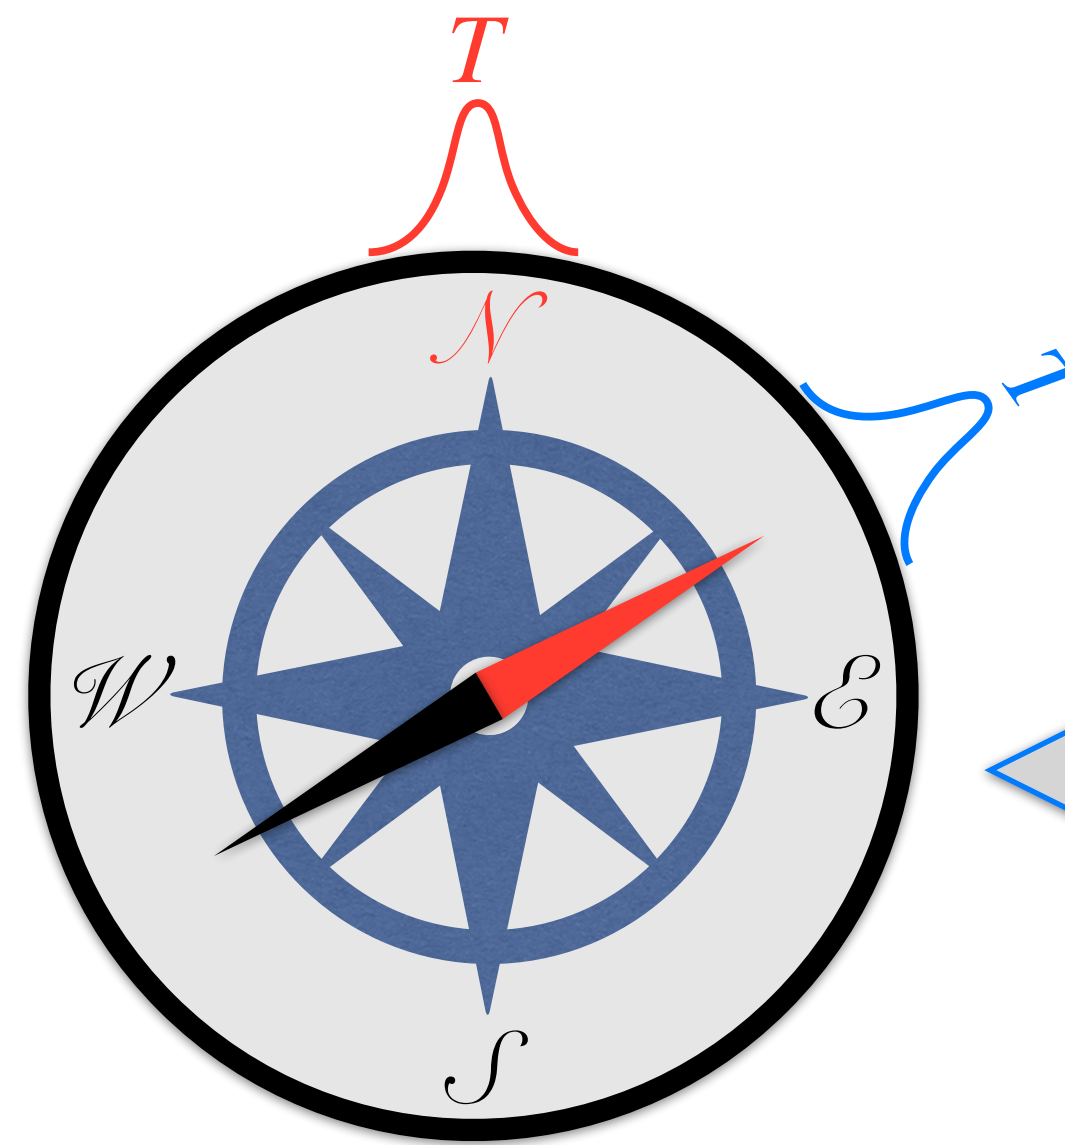


Ecological Niche




$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics



Unseen environment



$$MDP := (S, A, T)$$


state-space
action-space
transition dynamics

What is *North* for discovering new behaviors?

The **Ecological Niche** provides a dynamics prior.

Agents **explore** using the **prior memory** as a compass (fixed reference).


Ethological Priors Guide Adaptive Behavior



What is the most fundamental tool of an explorer?




Ecological Niche

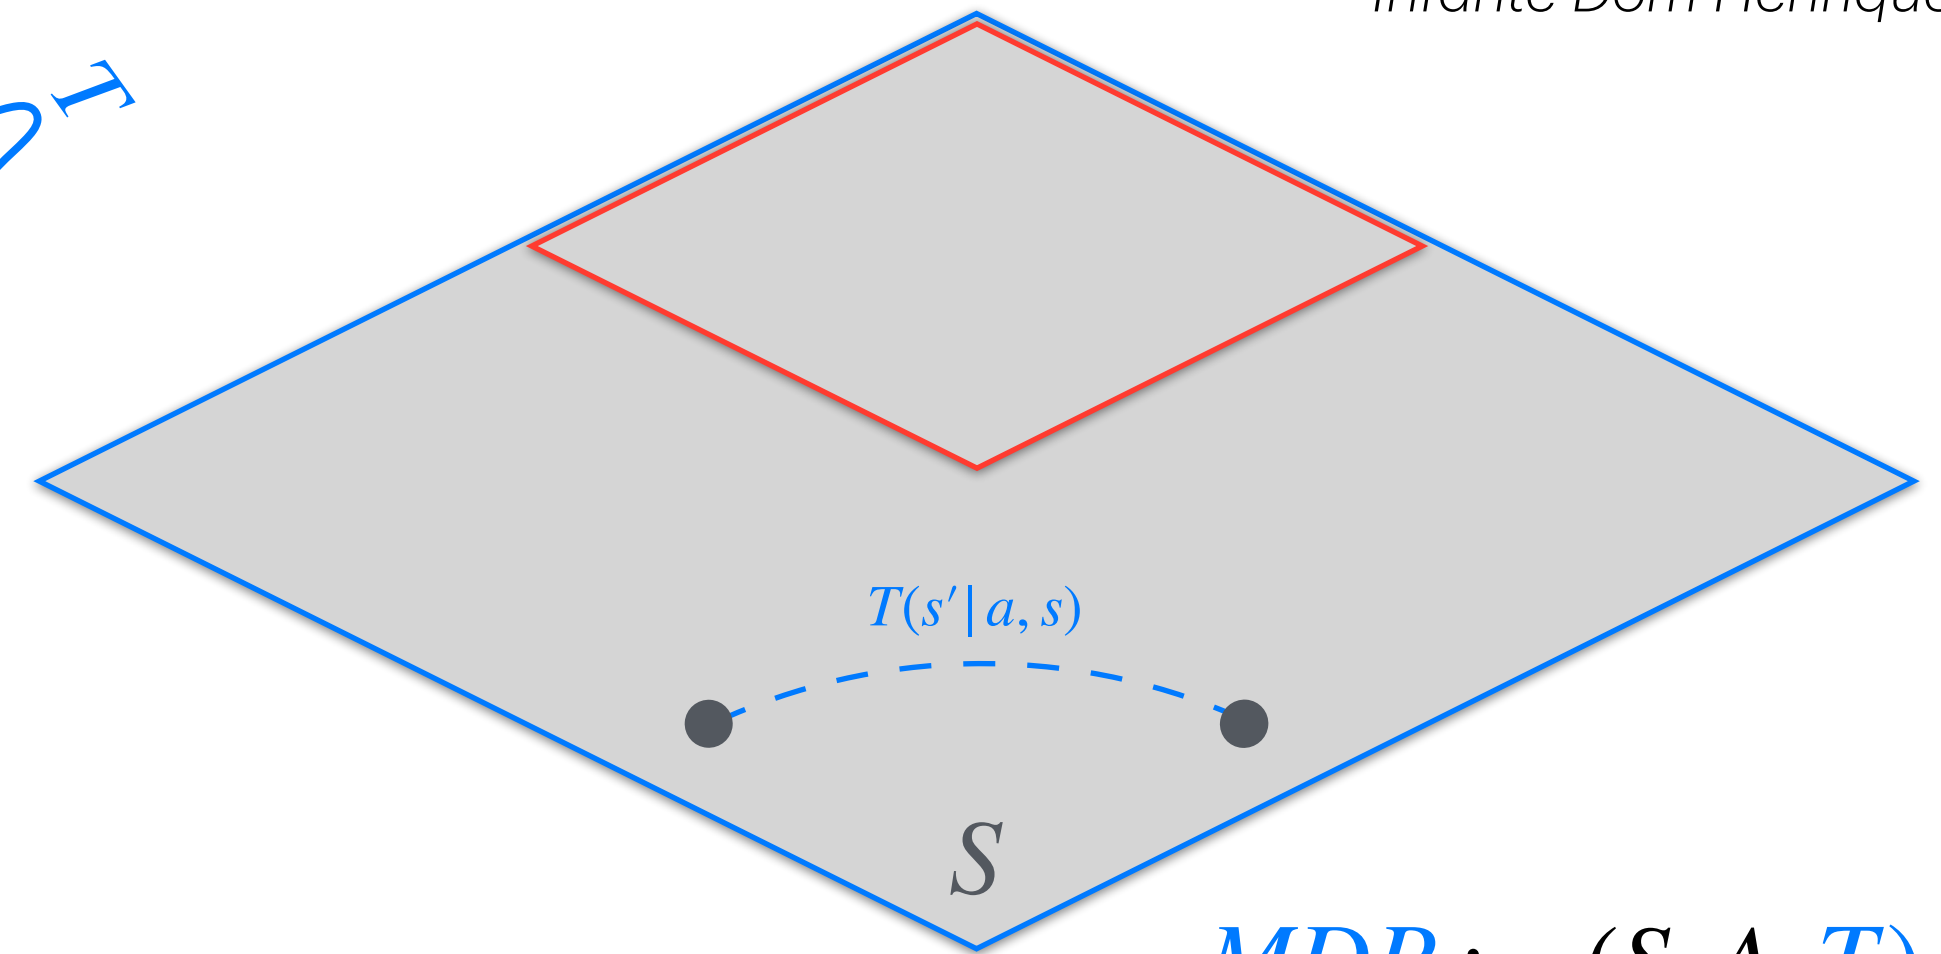


$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics



Unseen environment



$$MDP := (S, A, T)$$


state-space
action-space
transition dynamics

What is *North* for discovering new behaviors?

The **Ecological Niche** provides a dynamics prior.

Agents **explore** using the **prior memory** as a compass (fixed reference).

Ethological Priors Guide Adaptive Behavior




What is the most fundamental tool of an explorer?



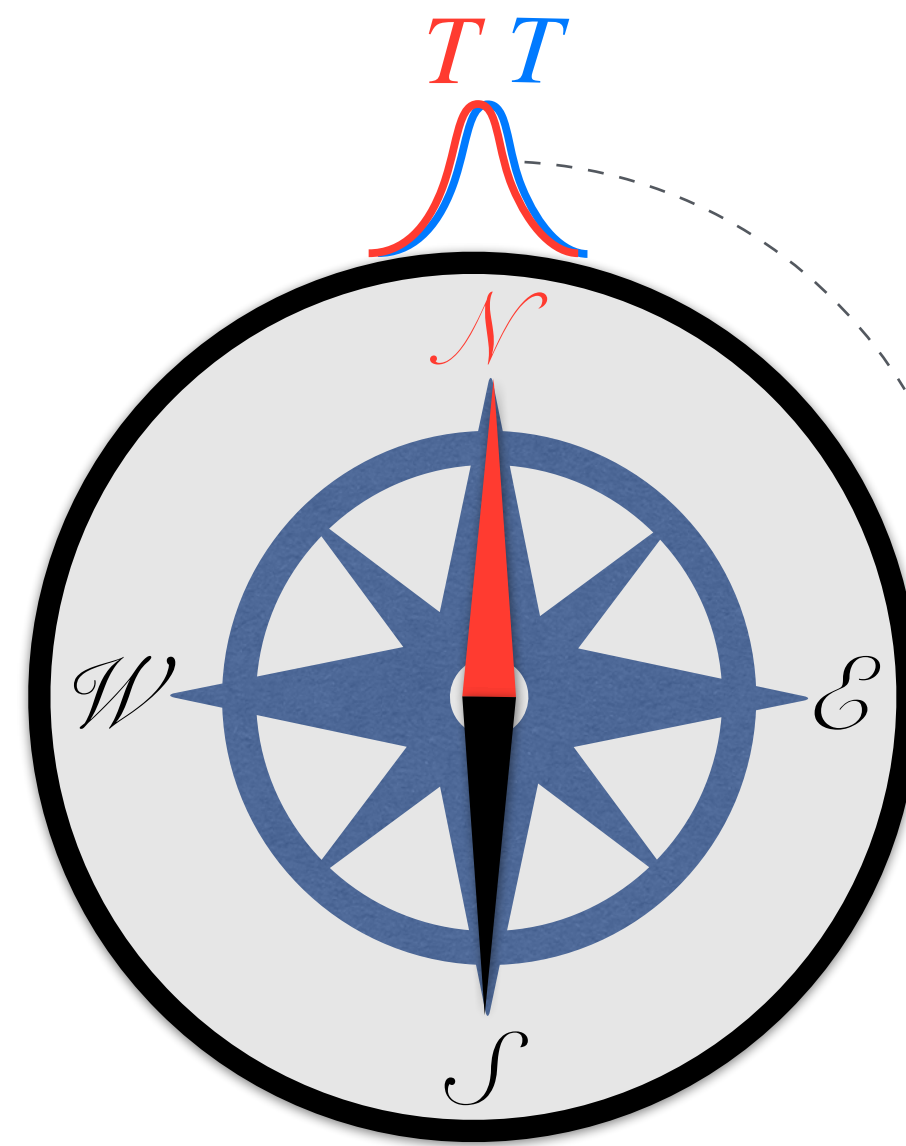
Infante Dom Henrique

Ecological Niche

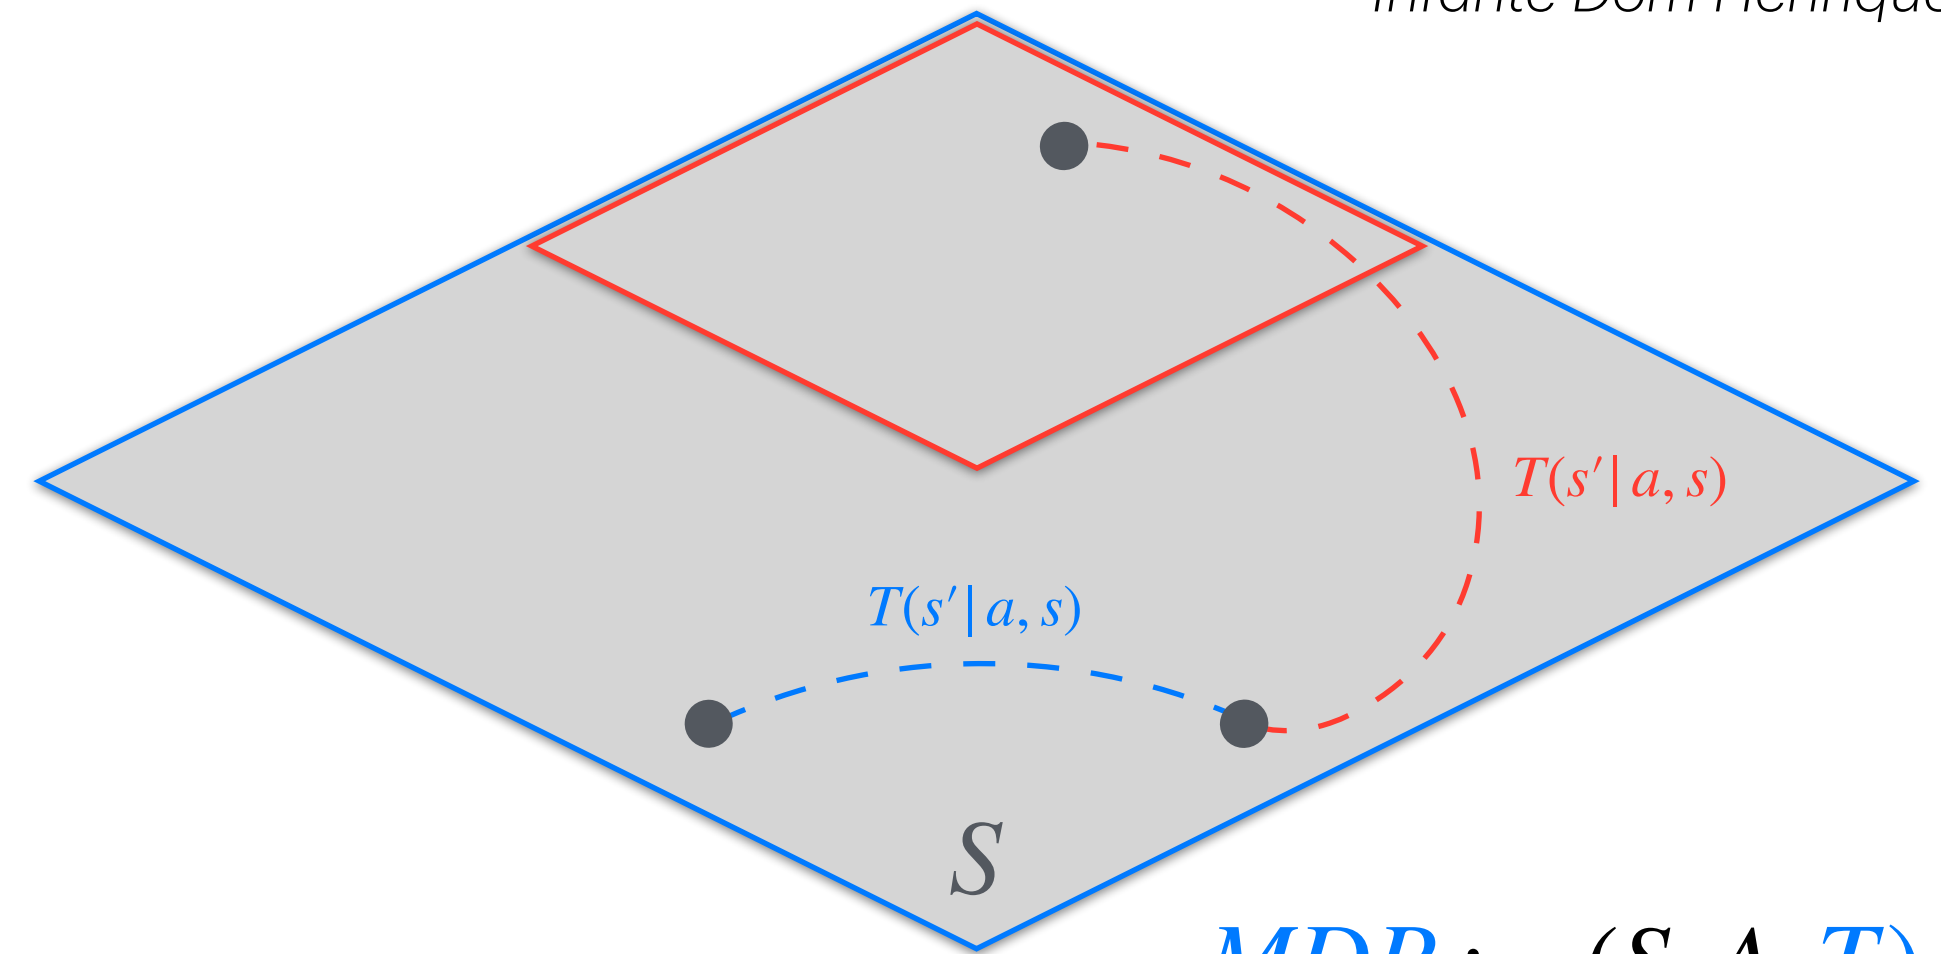


$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics



Unseen environment



$$MDP := (S, A, T)$$


state-space
action-space
transition dynamics

What is *North* for discovering new behaviors?

The **Ecological Niche** provides a dynamics prior.

Agents **explore** using the **prior memory** as a compass (fixed reference).

Ethological Priors Guide Adaptive Behavior




What is the most fundamental tool of an explorer?




Infante Dom Henrique

Ecological Niche

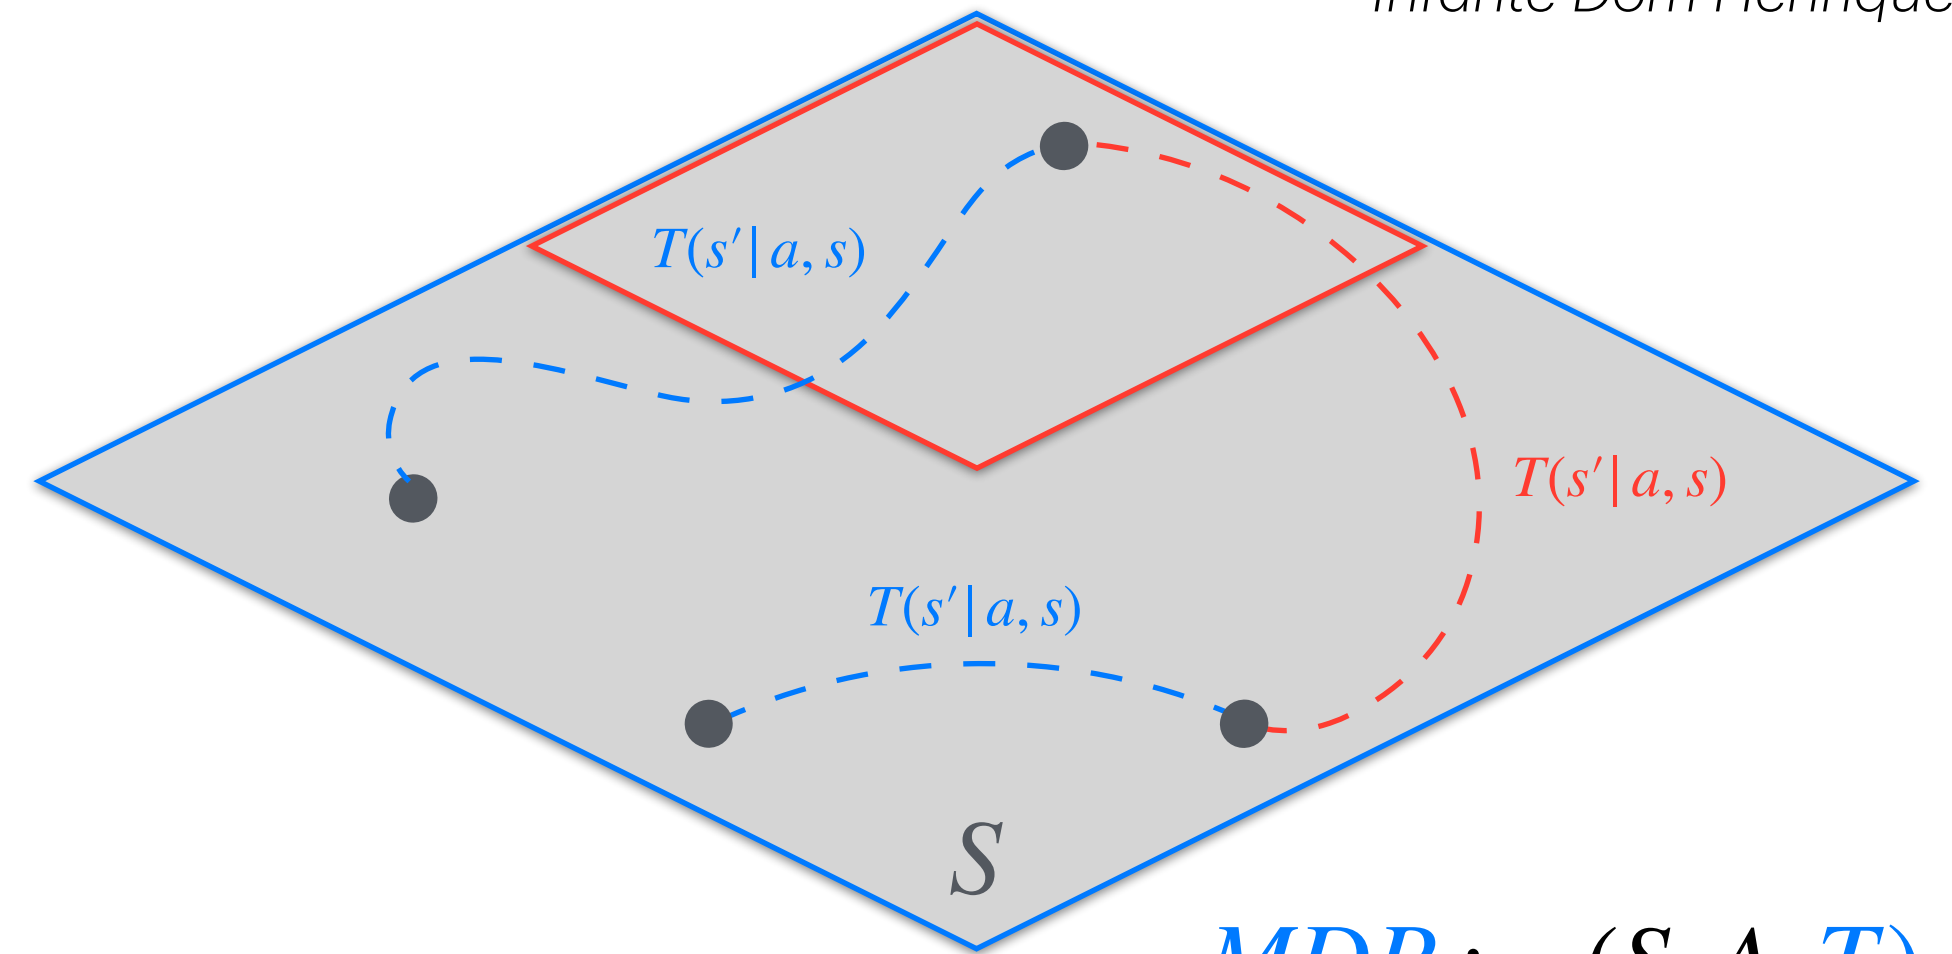


$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics



Unseen environment



$$MDP := (S, A, T)$$


state-space
action-space
transition dynamics

What is *North* for discovering new behaviors?

The **Ecological Niche** provides a dynamics prior.

Agents **explore** using the **prior memory** as a compass (fixed reference).

Ethological Priors Guide Adaptive Behavior




What is the most fundamental tool of an explorer?




Infante Dom Henrique

Ecological Niche

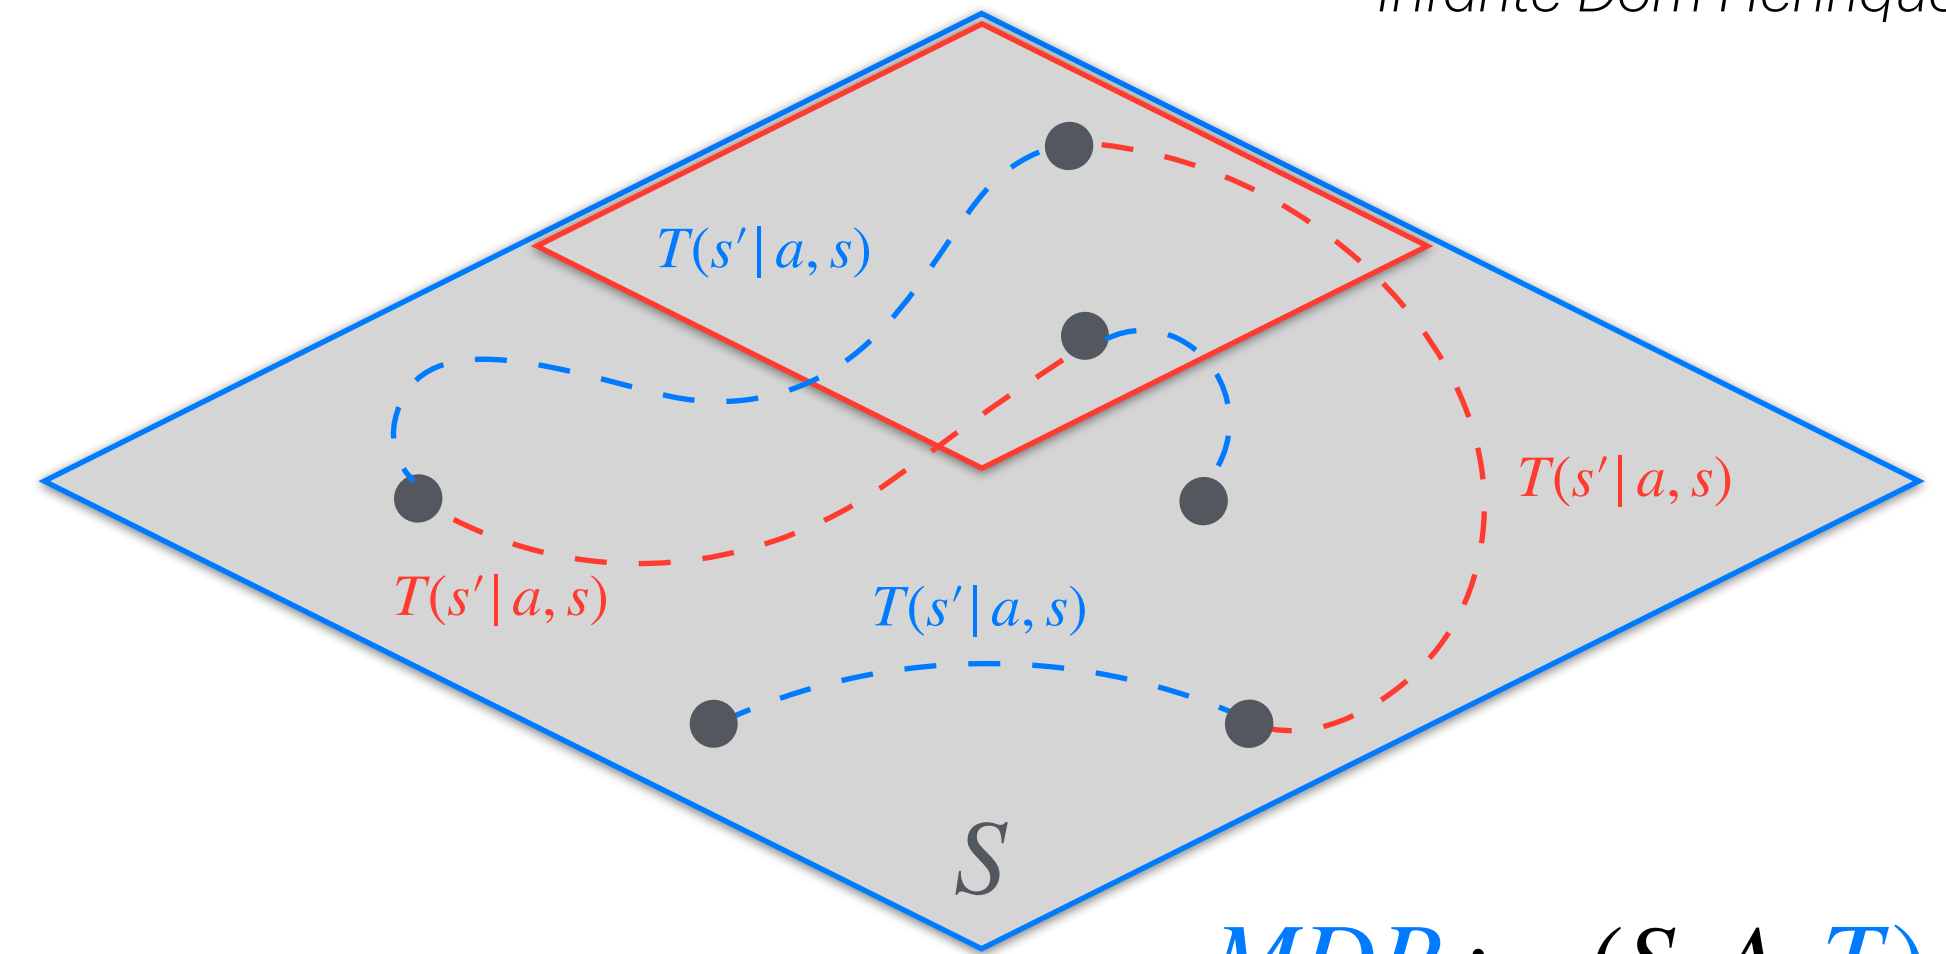


$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics



Unseen environment



$$MDP := (S, A, T)$$

state-space
action-space
transition dynamics

What is *North* for discovering new behaviors?

The **Ecological Niche** provides a dynamics prior.

Agents **explore** using the **prior memory** as a compass (fixed reference).

Ethological Priors Guide Adaptive Behavior

Ethological Priors Guide Adaptive Behavior

Active Pretraining

Ethological Priors Guide Adaptive Behavior

Active Pretraining

Online Exploration
with **Ethological Memory**

Ethological Priors Guide Adaptive Behavior

■ ethological



Active Pretraining

Online Exploration
with **Ethological Memory**

Ethological Priors Guide Adaptive Behavior

■ ethological

$T_1(s' | s, a)$



Active Pretraining

Online Exploration
with **Ethological Memory**

Ethological Priors Guide Adaptive Behavior


■ ethological



$$T_1(s' | s, a)$$

distill via experience

$$\omega_\theta(s' | s, a)$$




Active Pretraining

Online Exploration
with **Ethological Memory**

Ethological Priors Guide Adaptive Behavior

■ ethological



$$T_1(s' | s, a)$$

distill via experience

$$\omega_\theta(s' | s, a)$$



Active Pretraining

Online Exploration
with **Ethological Memory**

Ethological Priors Guide Adaptive Behavior


■ ethological



$$T_1(s' | s, a)$$


distill via experience

$$\omega_\theta(s' | s, a)$$



Active Pretraining


■ unethological



Online Exploration
with **Ethological Memory**

Ethological Priors Guide Adaptive Behavior


■ ethological



$T_1(s' | s, a)$



distill via experience

$\omega_\theta(s' | s, a)$



Active Pretraining

■ unethological




$T_2(s' | s, a)$

Online Exploration
with **Ethological Memory**

Ethological Priors Guide Adaptive Behavior


■ ethological



$$T_1(s' | s, a)$$



distill via experience

$$\omega_\theta(s' | s, a)$$



Active Pretraining

■ unethological



$$T_2(s' | s, a)$$

distill via experience

$$\omega_\theta$$


$$\omega_\beta(s' | s, a)$$



Online Exploration
with **Ethological Memory**

Ethological Priors Guide Adaptive Behavior


■ ethological



$T_1(s' | s, a)$



distill via experience

$\omega_\theta(s' | s, a)$



Active Pretraining

■ unethological



$T_2(s' | s, a)$

distill via experience

ω_θ

ϵ_t


$\omega_\beta(s' | s, a)$



Online Exploration
with **Ethological Memory**

Ethological Priors Guide Adaptive Behavior


■ ethological



$T_1(s' | s, a)$



distill via experience

$\omega_\theta(s' | s, a)$



Active Pretraining

■ unethological



$T_2(s' | s, a)$

distill via experience

ω_θ


ϵ_t

$\omega_\beta(s' | s, a)$



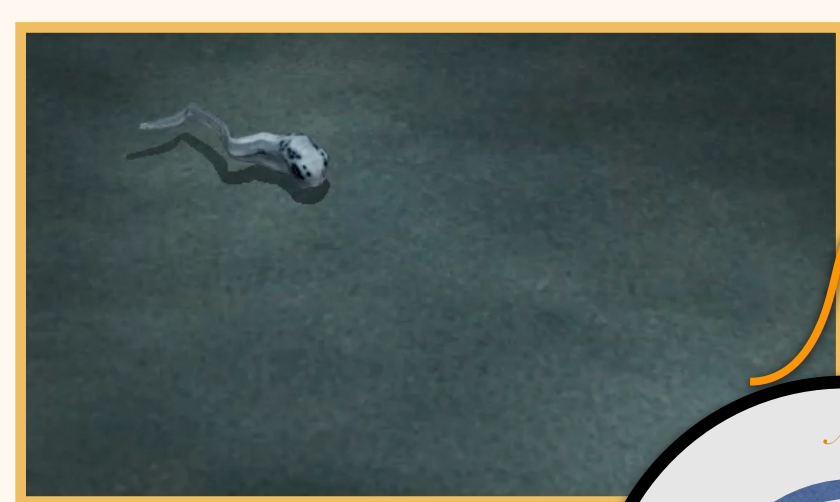
Online Exploration
with **Ethological Memory**

3M: Model-Memory-Mismatch



Ethological Priors Guide Adaptive Behavior


■ ethological



$T_1(s' | s, a)$



distill via experience

$\omega_\theta(s' | s, a)$



Active Pretraining

■ unethological



$T_2(s' | s, a)$

distill via experience

ω_θ


ϵ_t

$\omega_\beta(s' | s, a)$

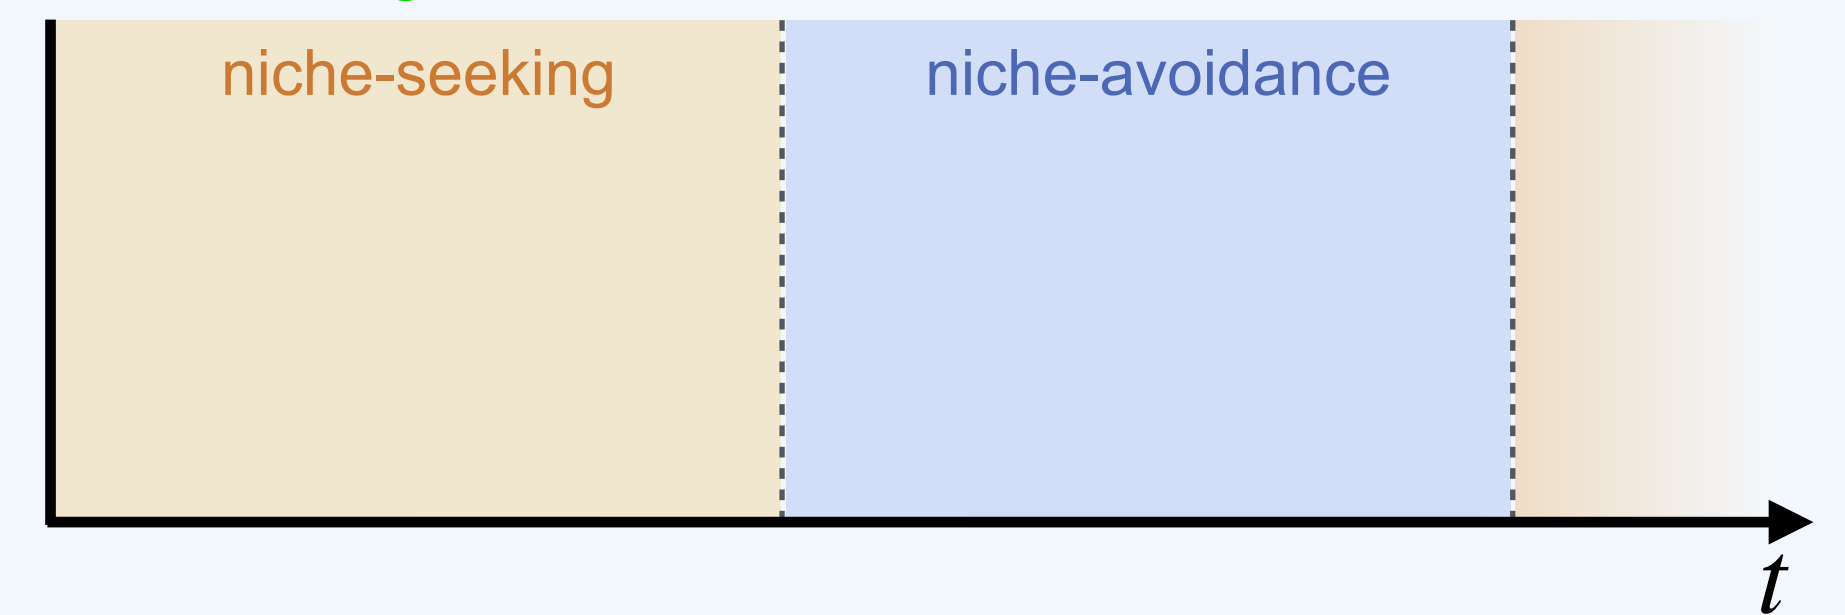


Online Exploration
with **Ethological Memory**

3M: Model-Memory-Mismatch



3M-Progress



Ethological Priors Guide Adaptive Behavior

■ ethological

$T_1(s' | s, a)$ distill via experience

$\omega_\theta(s' | s, a)$

Active Pretraining

The diagram illustrates the active pretraining phase. It features two images of a fish: one in a natural, dark environment and one in a simulated environment with a similar background. A compass rose is positioned below the images, with cardinal directions labeled: North (N), South (S), East (E), and West (W). An orange curve represents the probability distribution $\omega_\theta(s' | s, a)$. An arrow labeled "distill via experience" points from the simulated environment towards the distribution curve.

■ unethological

$T_2(s' | s, a)$ distill via experience

ω_θ

ϵ_t

$\omega_\beta(s' | s, a)$

Online Exploration with Ethological Memory

The diagram illustrates the online exploration phase. It features two images of a fish: one in a simulated environment with blue and black diagonal stripes, and one in a natural environment. A compass rose is positioned below the images, with cardinal directions labeled: North (N), South (S), East (E), and West (W). Two probability distribution curves are shown: an orange curve ω_θ and a blue curve $\omega_\beta(s' | s, a)$. A dashed purple arrow labeled ϵ_t points from the orange curve to the blue curve. An arrow labeled "distill via experience" points from the simulated environment towards the blue curve.

3M: Model-Memory-Mismatch

ω_θ

$\epsilon_t = \text{KL}[\omega_\beta || \omega_\theta]$

ω_β

$s' | s, a$

3M-Progress

niche-seeking

niche-avoidance

t


The diagram illustrates the 3M (Model-Memory-Mismatch) concept and its progress. The top part shows two probability distribution curves: an orange curve ω_θ and a blue curve ω_β . A horizontal arrow between their peaks is labeled $\epsilon_t = \text{KL}[\omega_\beta || \omega_\theta]$. The x-axis is labeled $s' | s, a$. The bottom part shows a timeline of behavior with two phases: "niche-seeking" (orange background) and "niche-avoidance" (blue background). The x-axis is labeled t . A compass rose is positioned below the timeline, with cardinal directions labeled: North (N), South (S), East (E), and West (W).

Ethological Priors Guide Adaptive Behavior

■ ethological

$T_1(s' | s, a)$ distill via experience

$\omega_\theta(s' | s, a)$



Active Pretraining

Detailed description: This diagram illustrates the 'Active Pretraining' phase. It features two images of a fish in a natural, dark environment. The first image shows the fish from a top-down perspective, and the second shows it from a side view. To the right, a compass rose is shown with an orange arrow pointing towards the North (N) direction. A bell-shaped curve, labeled $\omega_\theta(s' | s, a)$, is positioned above the compass rose. An arrow labeled $T_1(s' | s, a)$ points from the text 'distill via experience' to the curve.

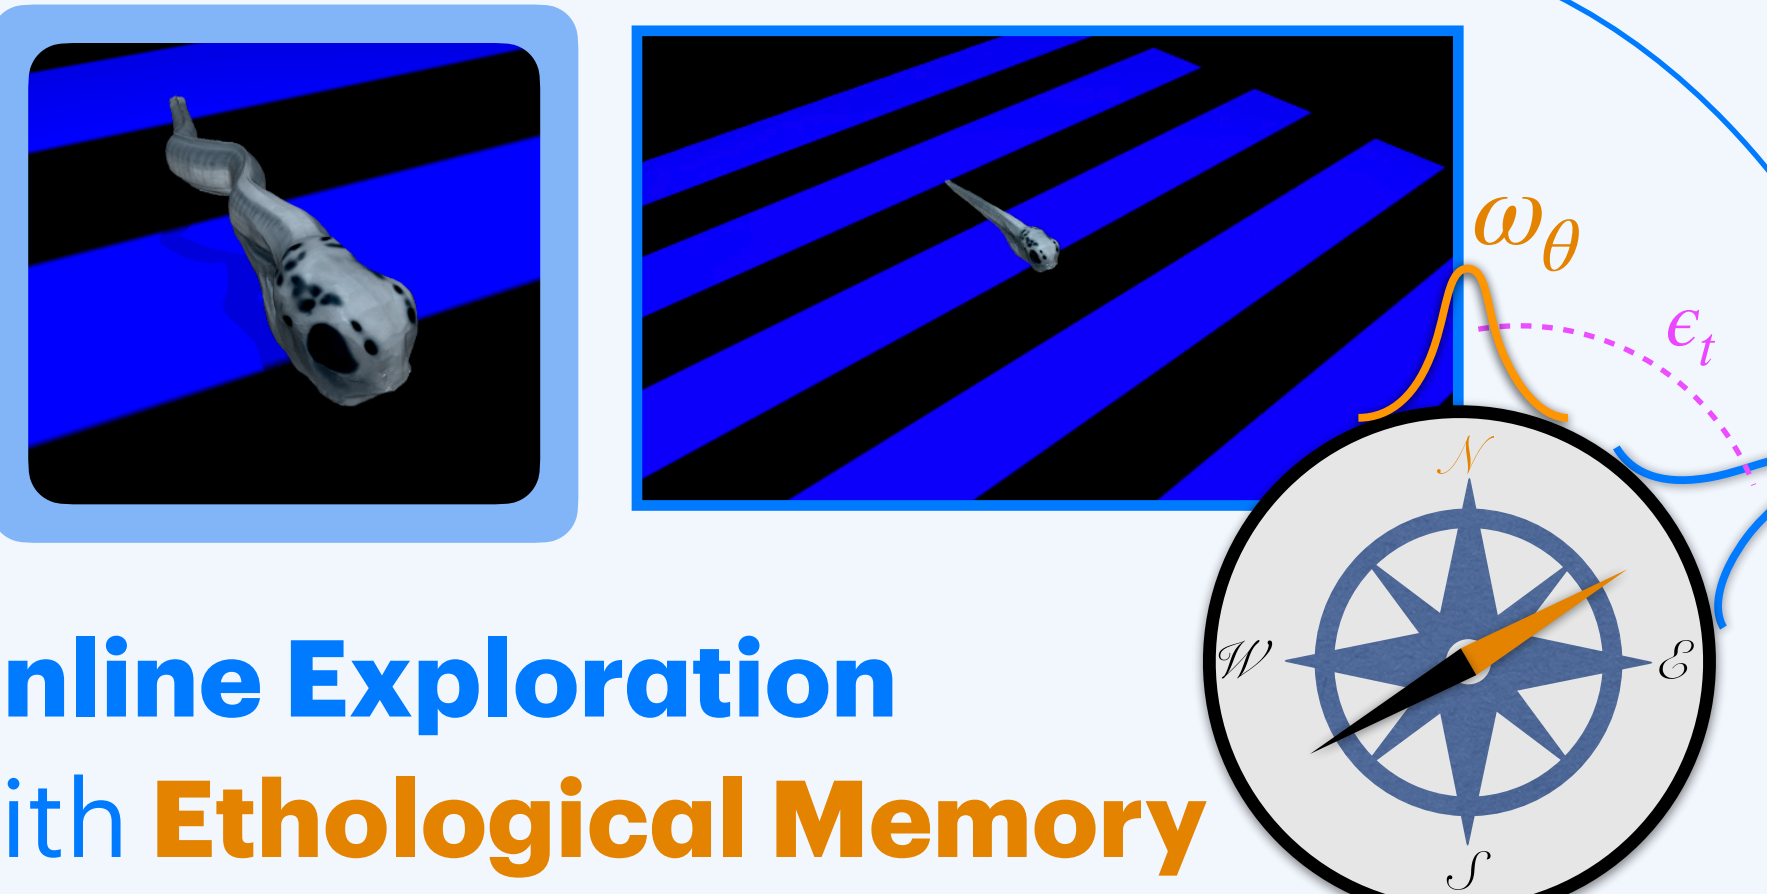
■ unethological

$T_2(s' | s, a)$ distill via experience

ω_θ

ϵ_t

$\omega_\beta(s' | s, a)$



Online Exploration with Ethological Memory

Detailed description: This diagram illustrates the 'Online Exploration with Ethological Memory' phase. It features two images of a fish in a striped environment. The first image shows the fish from a top-down perspective, and the second shows it from a side view. To the right, a compass rose is shown with an orange arrow pointing towards the East (E) direction. A bell-shaped curve, labeled ω_θ , is positioned above the compass rose. A dashed purple curve, labeled ϵ_t , is shown below the ω_θ curve. An arrow labeled $T_2(s' | s, a)$ points from the text 'distill via experience' to the $\omega_\beta(s' | s, a)$ label.

3M: Model-Memory-Mismatch

ω_θ

$\epsilon_t = \text{KL}[\omega_\beta || \omega_\theta]$

ω_β

$s' | s, a$

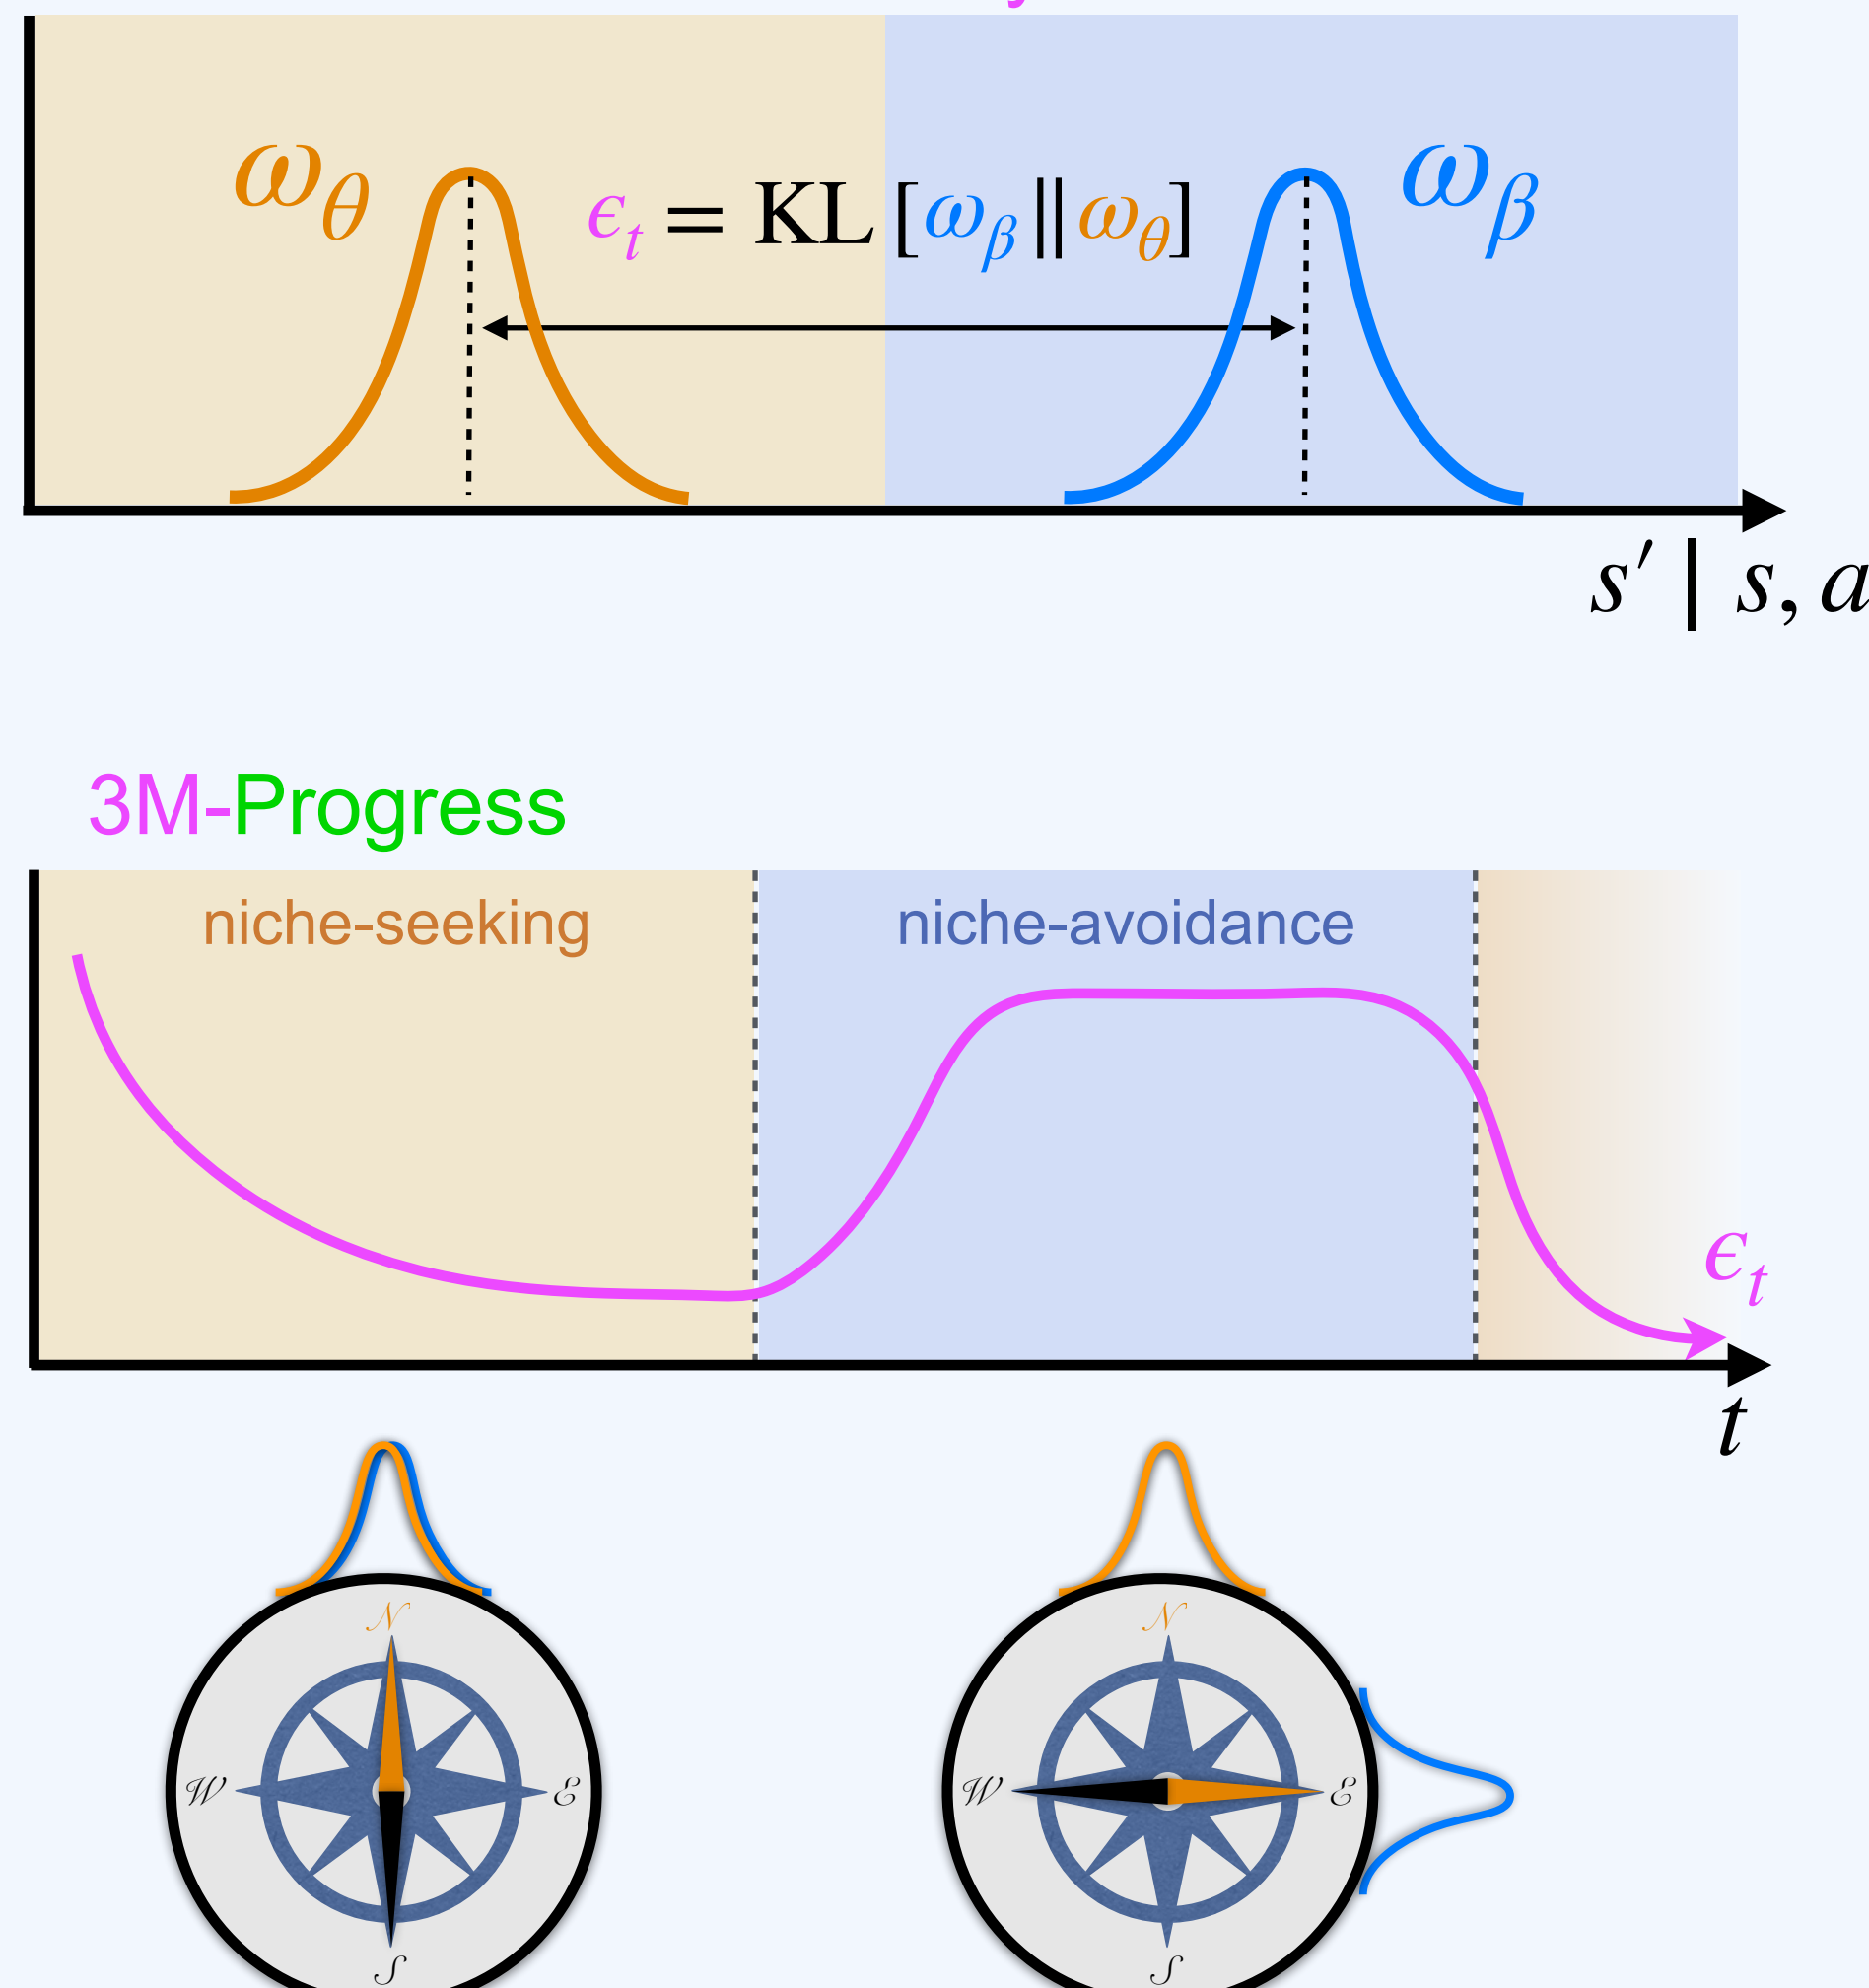
3M-Progress

niche-seeking

niche-avoidance

ϵ_t

t



Detailed description: This diagram illustrates the '3M-Progress' phase. It features a graph with two x-axes. The top x-axis is labeled $s' | s, a$ and shows two bell-shaped curves: an orange one labeled ω_θ and a blue one labeled ω_β . A double-headed arrow between their peaks is labeled $\epsilon_t = \text{KL}[\omega_\beta || \omega_\theta]$. The bottom x-axis is labeled t and shows a purple curve that starts high, dips, rises to a plateau, and then dips again. The first dip is labeled 'niche-seeking' and the plateau is labeled 'niche-avoidance'. Below the graph are two compass roses. The first has an orange arrow pointing North, and the second has an orange arrow pointing East.

Ethological Priors Guide Adaptive Behavior

■ ethological

$T_1(s' | s, a)$ distill via experience

$\omega_\theta(s' | s, a)$

Active Pretraining

■ unethological

$T_2(s' | s, a)$ distill via experience

ω_θ

ϵ_t

$\omega_\beta(s' | s, a)$

Online Exploration with Ethological Memory

3M: Model-Memory-Mismatch

ω_θ

$\epsilon_t = \text{KL}[\omega_\beta || \omega_\theta]$

ω_β

$s' | s, a$

3M-Progress

niche-seeking

niche-avoidance

$\hat{\epsilon}_t = (1 - \gamma)\hat{\epsilon}_{t-1} + \gamma\epsilon_t$

ϵ_t

t

Ethological Priors Guide Adaptive Behavior

■ ethological

$T_1(s' | s, a)$ distill via experience

$\omega_\theta(s' | s, a)$

Active Pretraining

■ unethological

$T_2(s' | s, a)$ distill via experience

ω_θ ϵ_t $\omega_\beta(s' | s, a)$

Online Exploration with Ethological Memory

3M: Model-Memory-Mismatch

ω_θ $\epsilon_t = \text{KL}[\omega_\beta || \omega_\theta]$ ω_β

$s' | s, a$

3M-Progress

niche-seeking niche-avoidance

$\hat{\epsilon}_t = (1 - \gamma)\hat{\epsilon}_{t-1} + \gamma\epsilon_t$

$\hat{\epsilon}_t$ ϵ_t

t

Ethological Priors Guide Adaptive Behavior

■ ethological

$T_1(s' | s, a)$ distill via experience

$\omega_\theta(s' | s, a)$

Active Pretraining

■ unethological

$T_2(s' | s, a)$ distill via experience

ω_θ

ϵ_t

$\omega_\beta(s' | s, a)$

Online Exploration with Ethological Memory

3M: Model-Memory-Mismatch

$\epsilon_t = \text{KL}[\omega_\beta || \omega_\theta]$

$r_t^i \propto |\hat{\epsilon}_t - \epsilon_t|$

$\hat{\epsilon}_t = (1 - \gamma)\hat{\epsilon}_{t-1} + \gamma\epsilon_t$

3M-Progress

niche-seeking

niche-avoidance

$\hat{\epsilon}_t$

ϵ_t

t

Ethological Priors Guide Adaptive Behavior

■ ethological

$T_1(s' | s, a)$ distill via experience

$\omega_\theta(s' | s, a)$

Active Pretraining

■ unethological

$T_2(s' | s, a)$ distill via experience

ω_θ

ϵ_t

$\omega_\beta(s' | s, a)$

Online Exploration with Ethological Memory

3M: Model-Memory-Mismatch

$\epsilon_t = \text{KL}[\omega_\beta || \omega_\theta]$

$r_t^i \propto |\hat{\epsilon}_t - \epsilon_t|$

$\hat{\epsilon}_t = (1 - \gamma)\hat{\epsilon}_{t-1} + \gamma\epsilon_t$

3M-Progress

niche-seeking

niche-avoidance

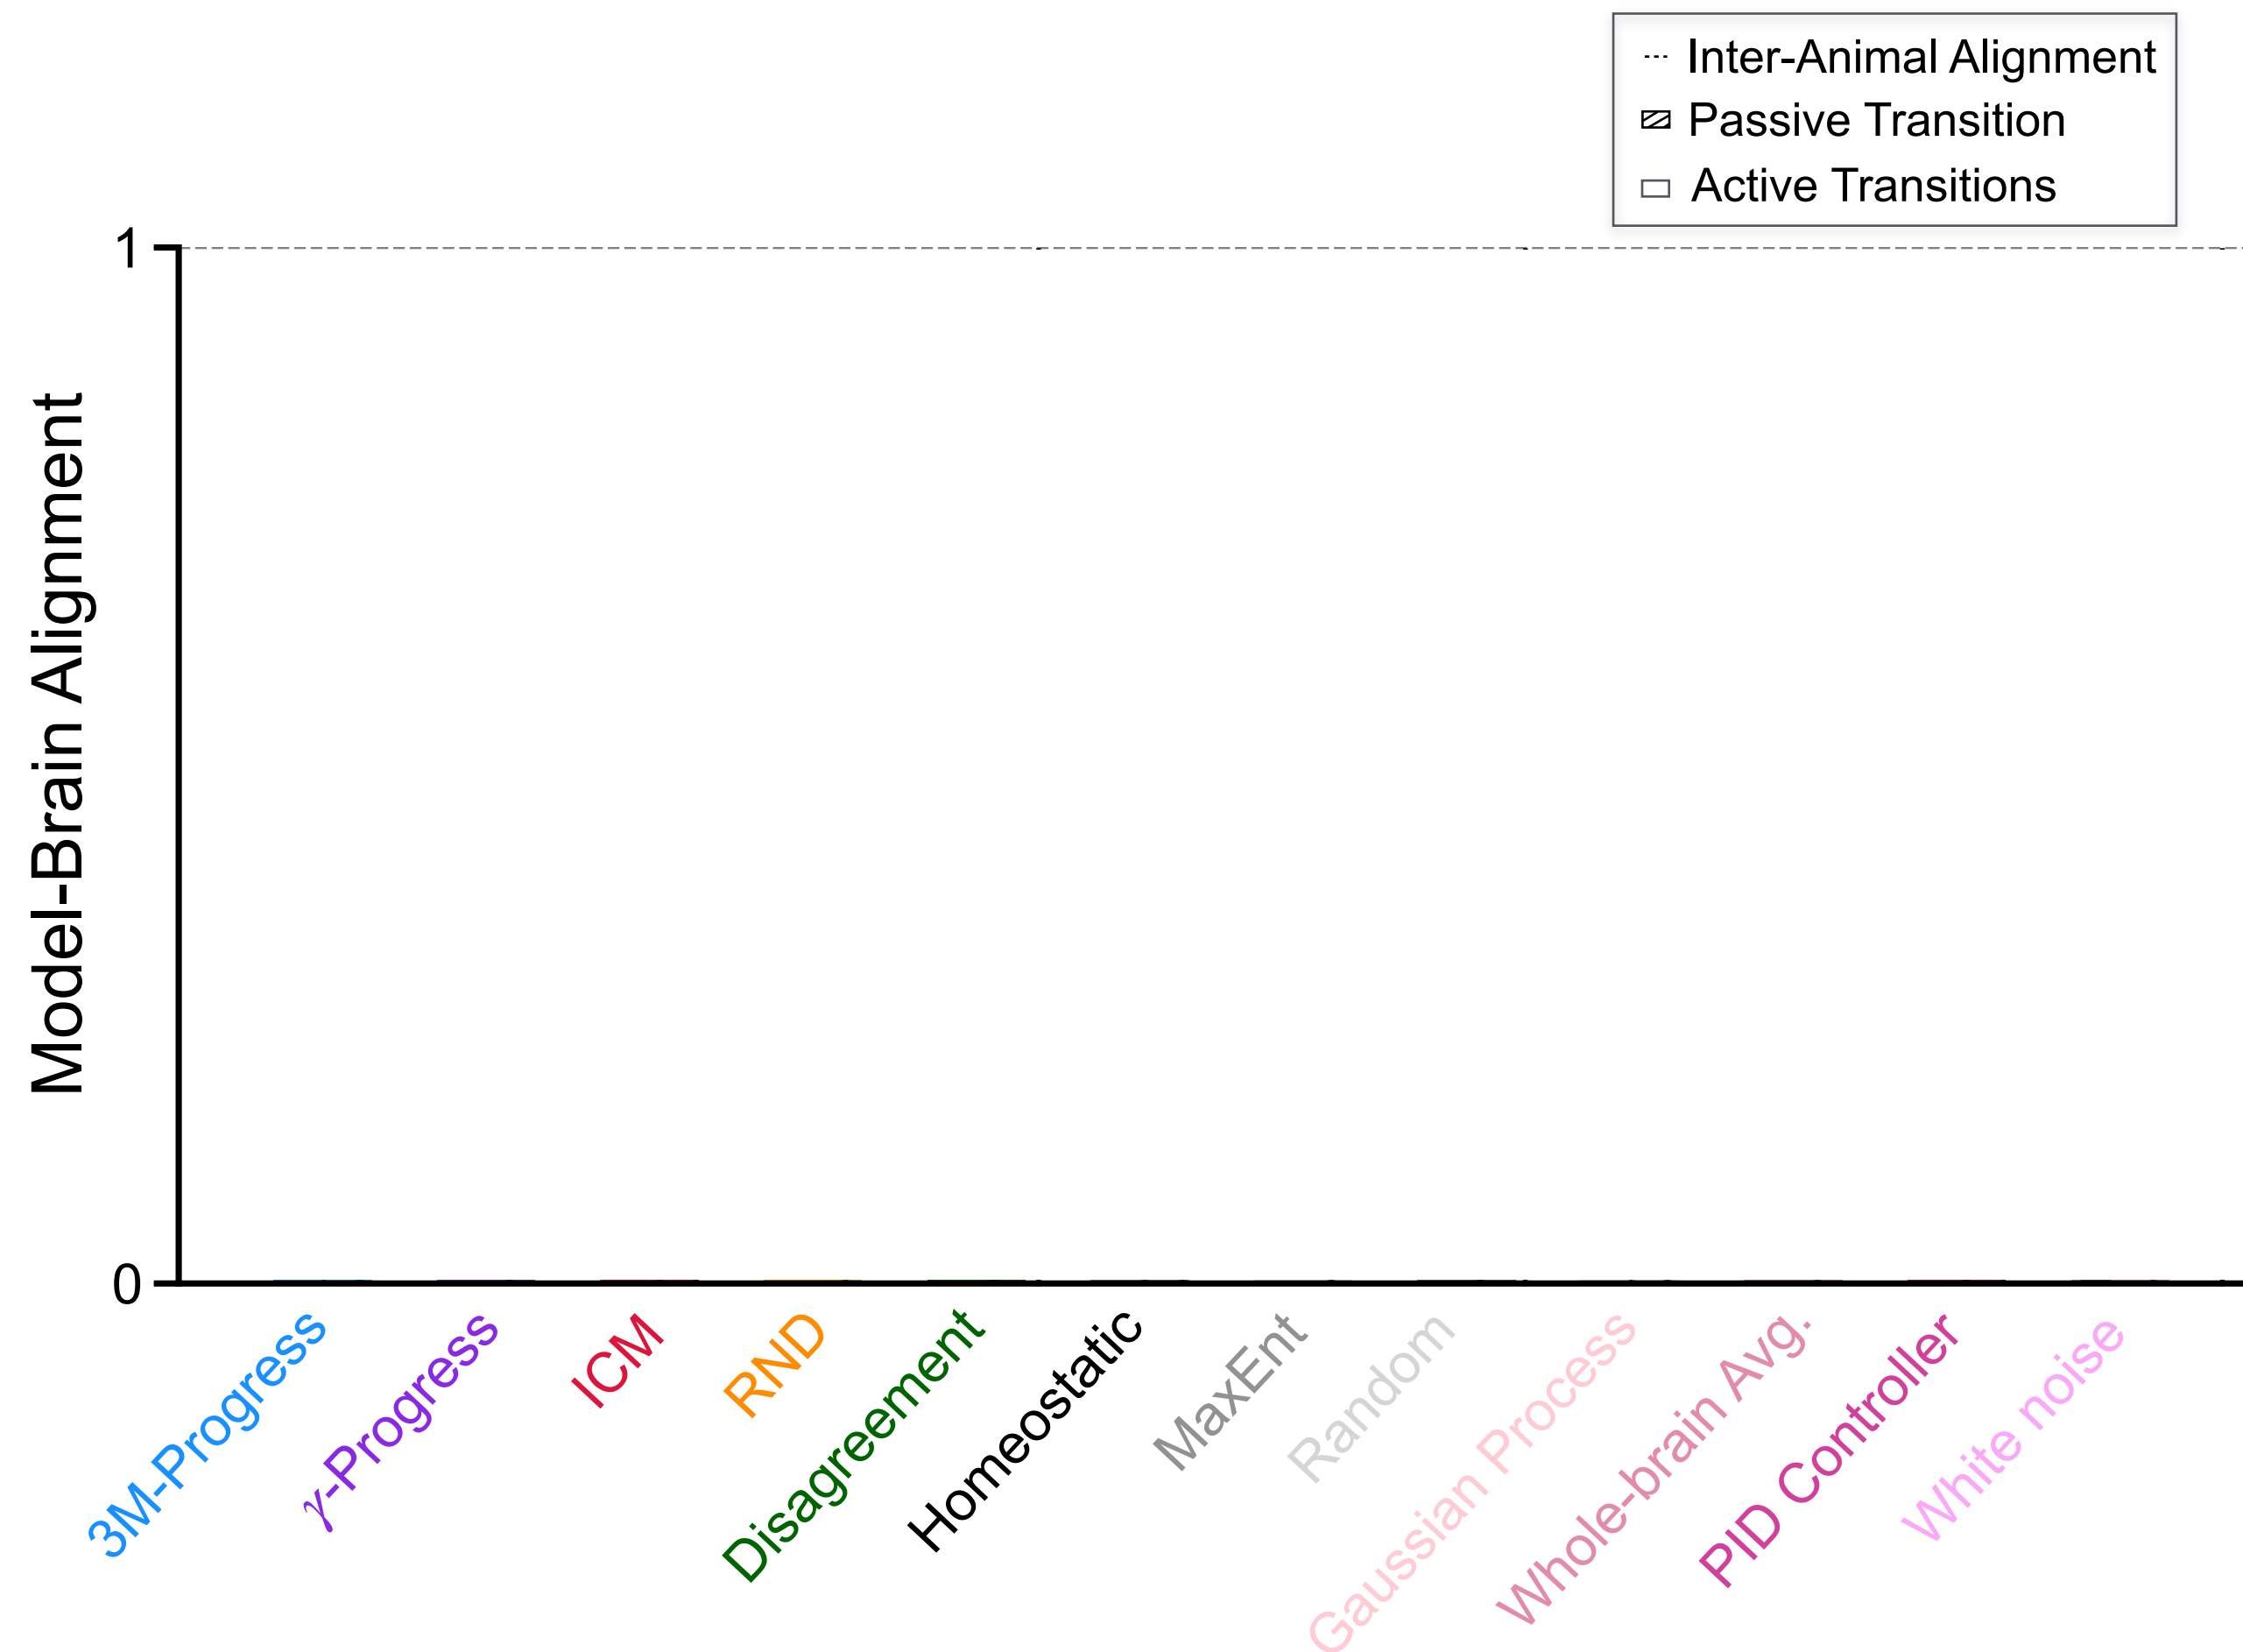
$\hat{\epsilon}_t$

r_t^i

ϵ_t

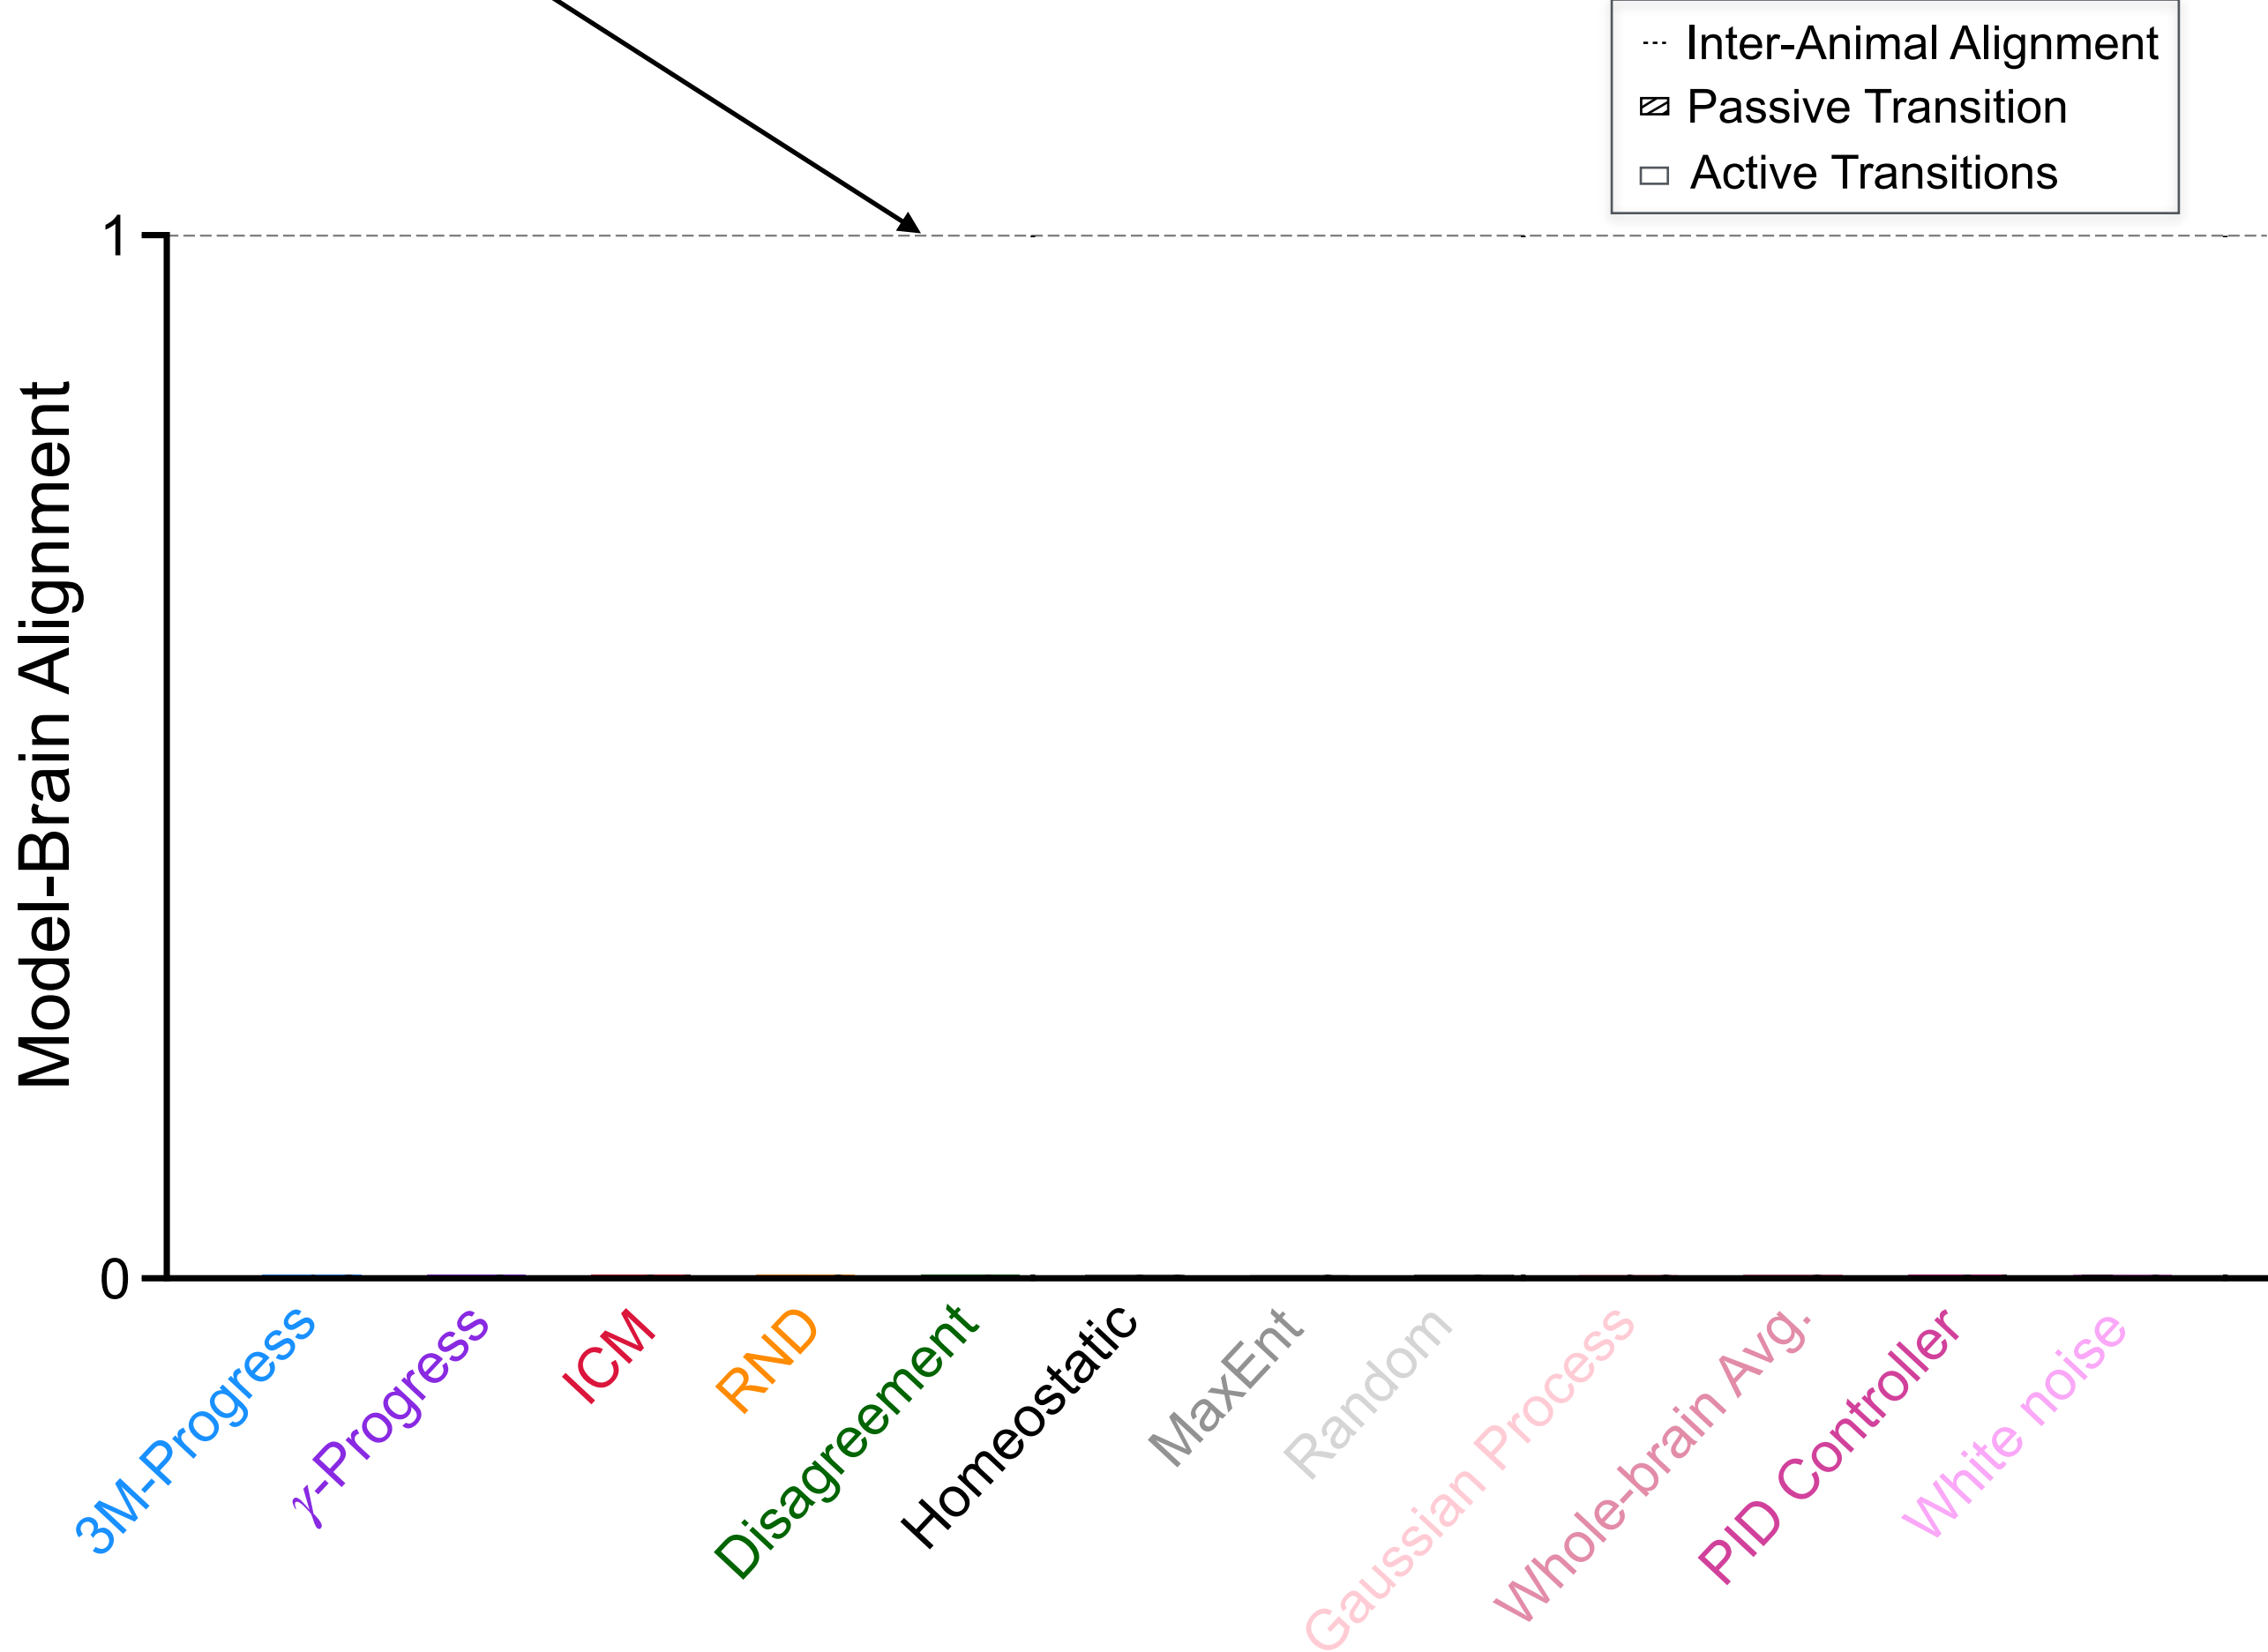
t

3MP Agent Predicts Whole-brain Single-cell Activity




3MP Agent Predicts Whole-brain Single-cell Activity

Same-species consistency (one-to-one)




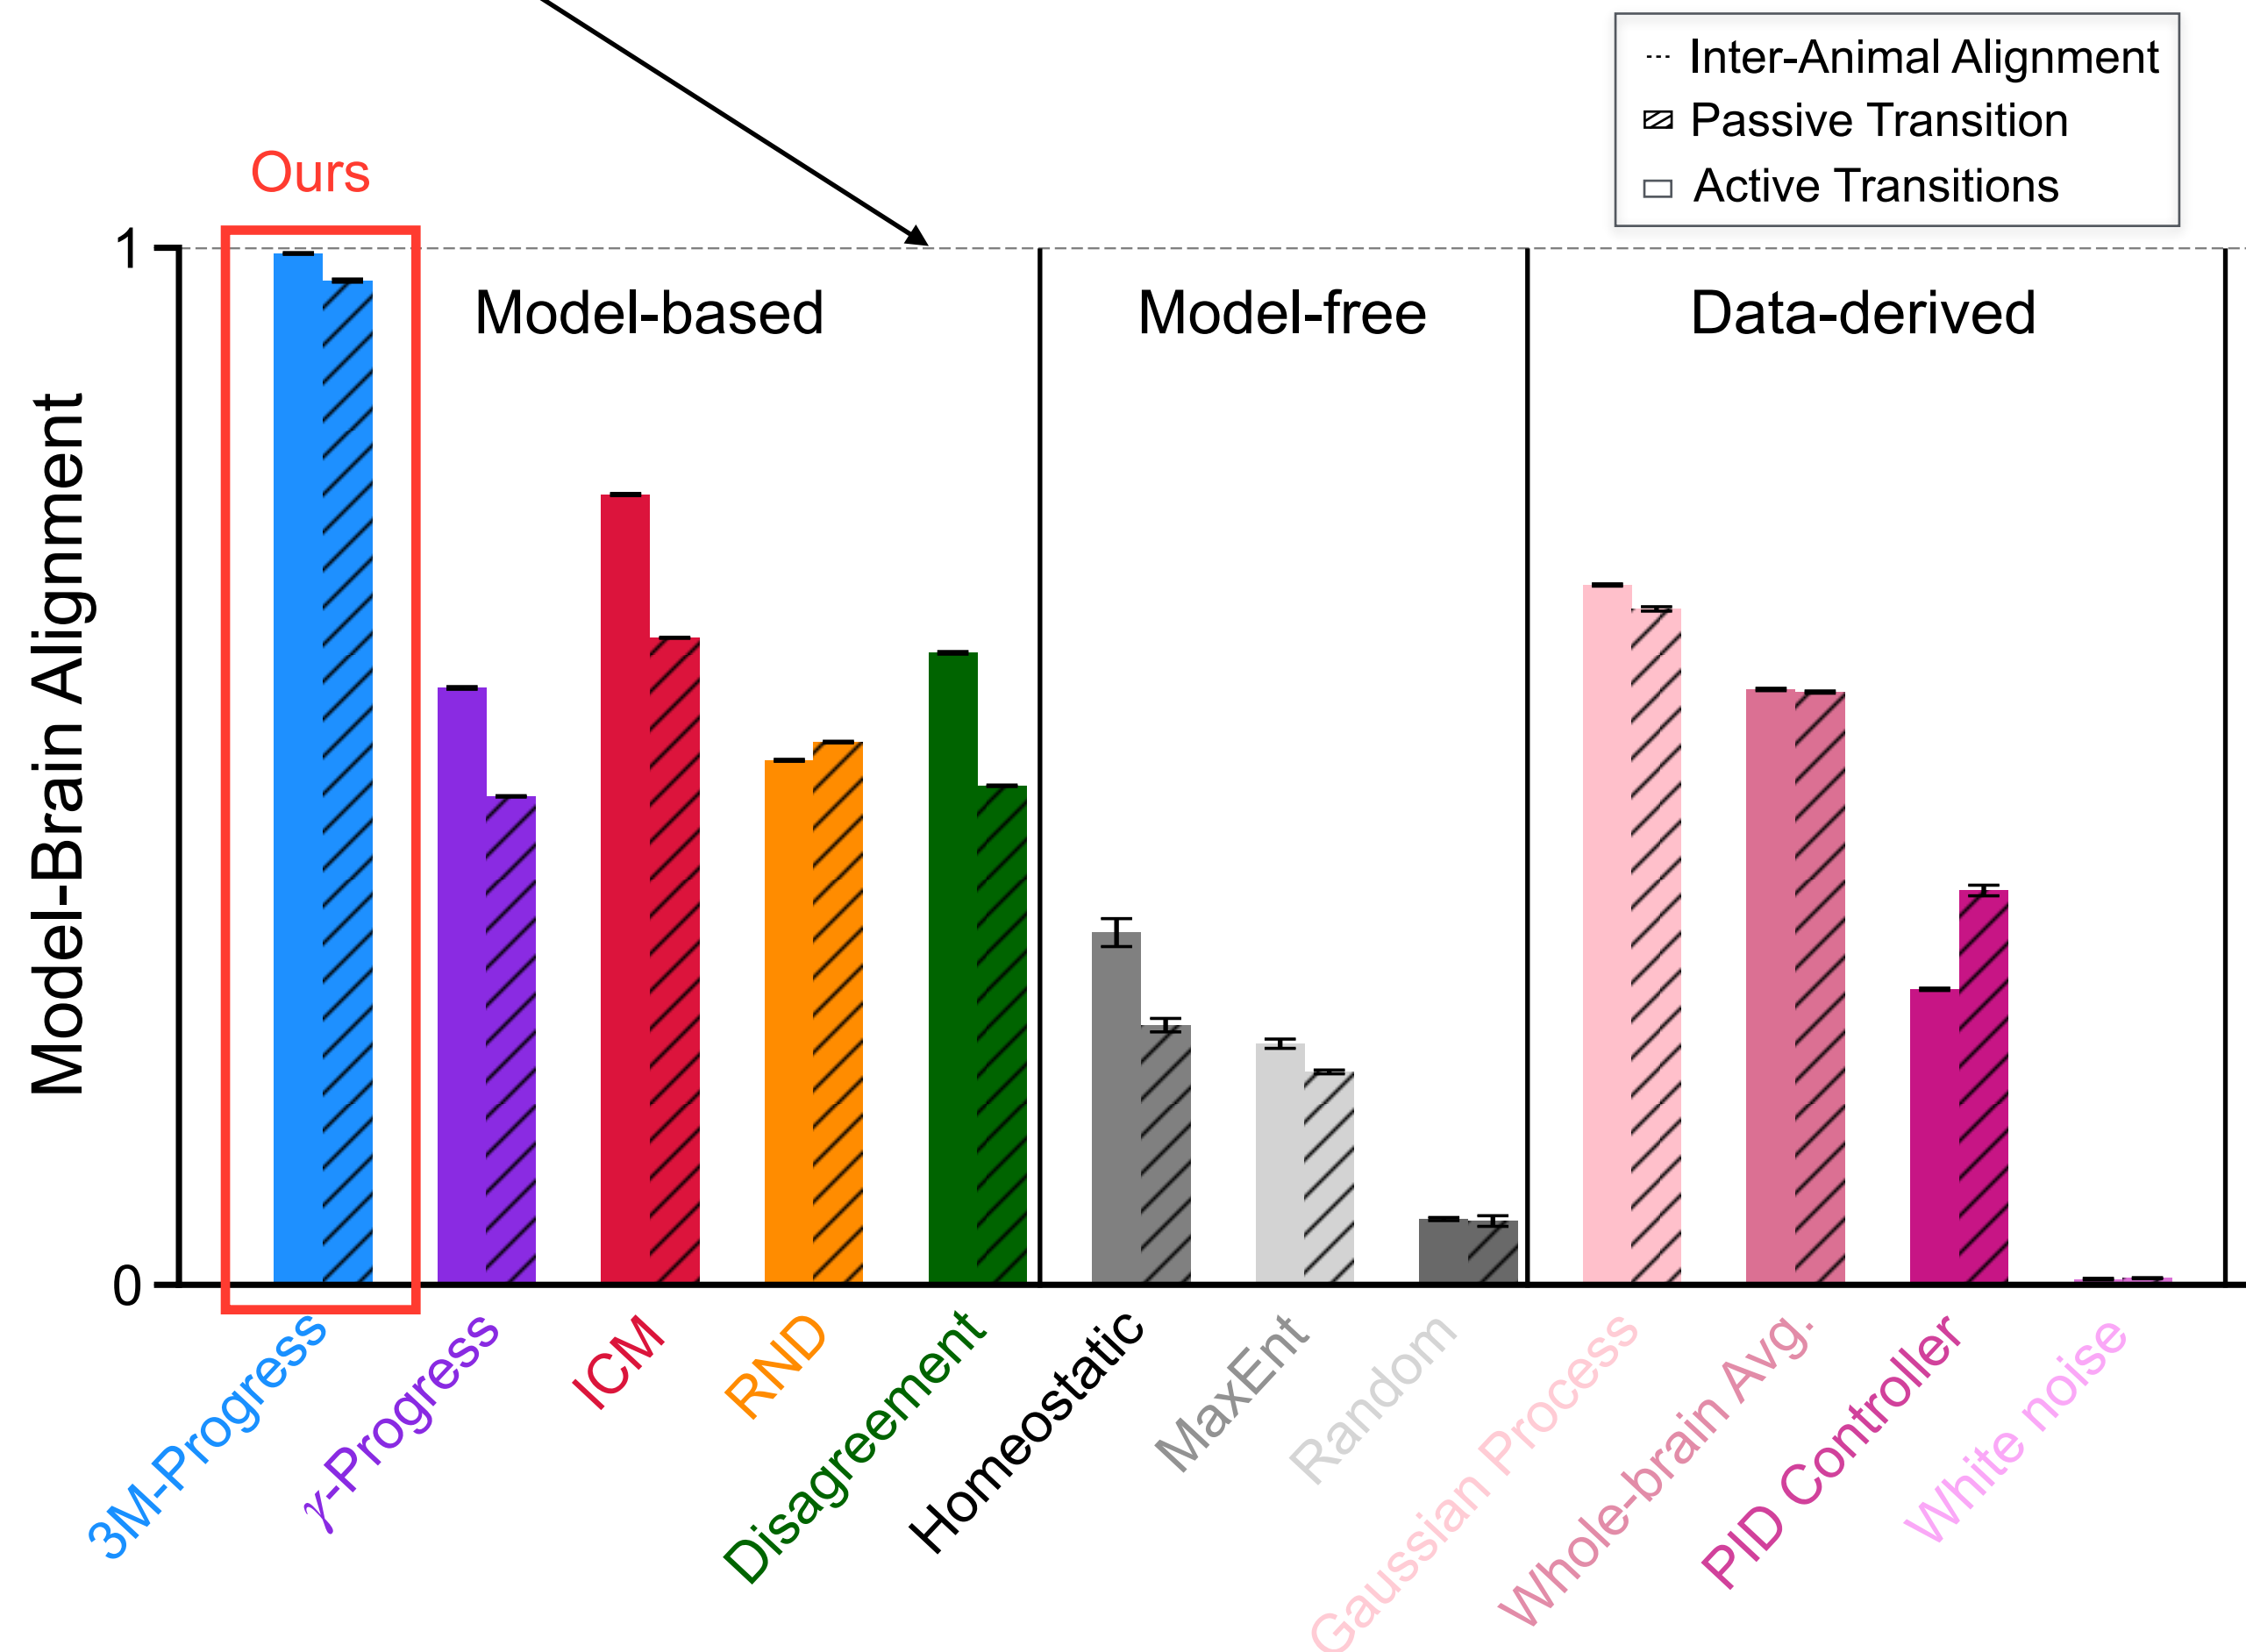
3MP Agent Predicts Whole-brain Single-cell Activity

Same-species consistency (one-to-one)




3MP Agent Predicts Whole-brain Single-cell Activity

Same-species consistency (one-to-one)




3MP Agent Predicts Whole-brain Single-cell Activity

Same-species consistency (one-to-one)

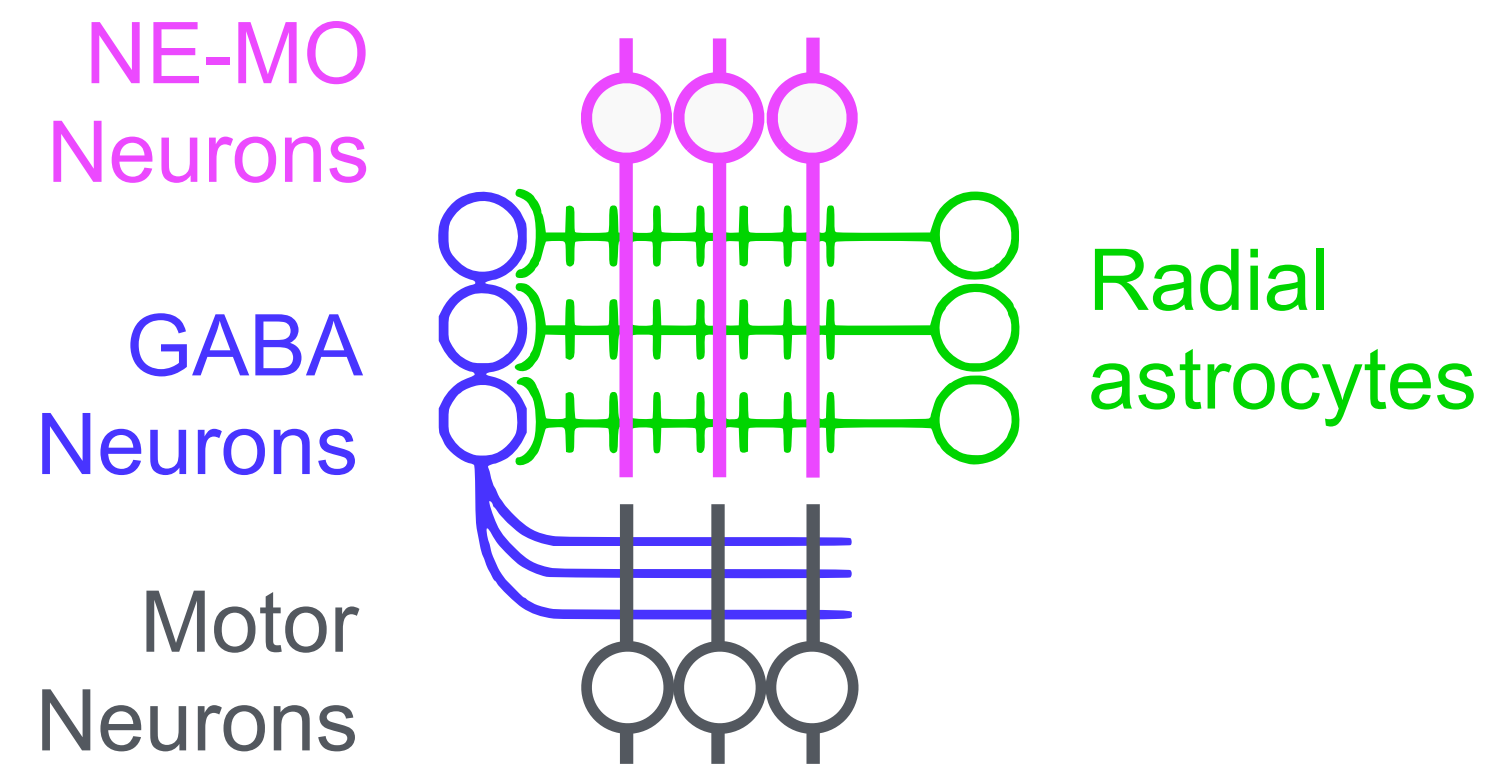


Outlook: models of the brain that "walk the walk"




Neural-Glial Circuits Emerge in 3M-Progress Agents

Neural-Glial Circuits Emerge in 3M-Progress Agents





Neural-Glial Circuits Emerge in 3M-Progress Agents


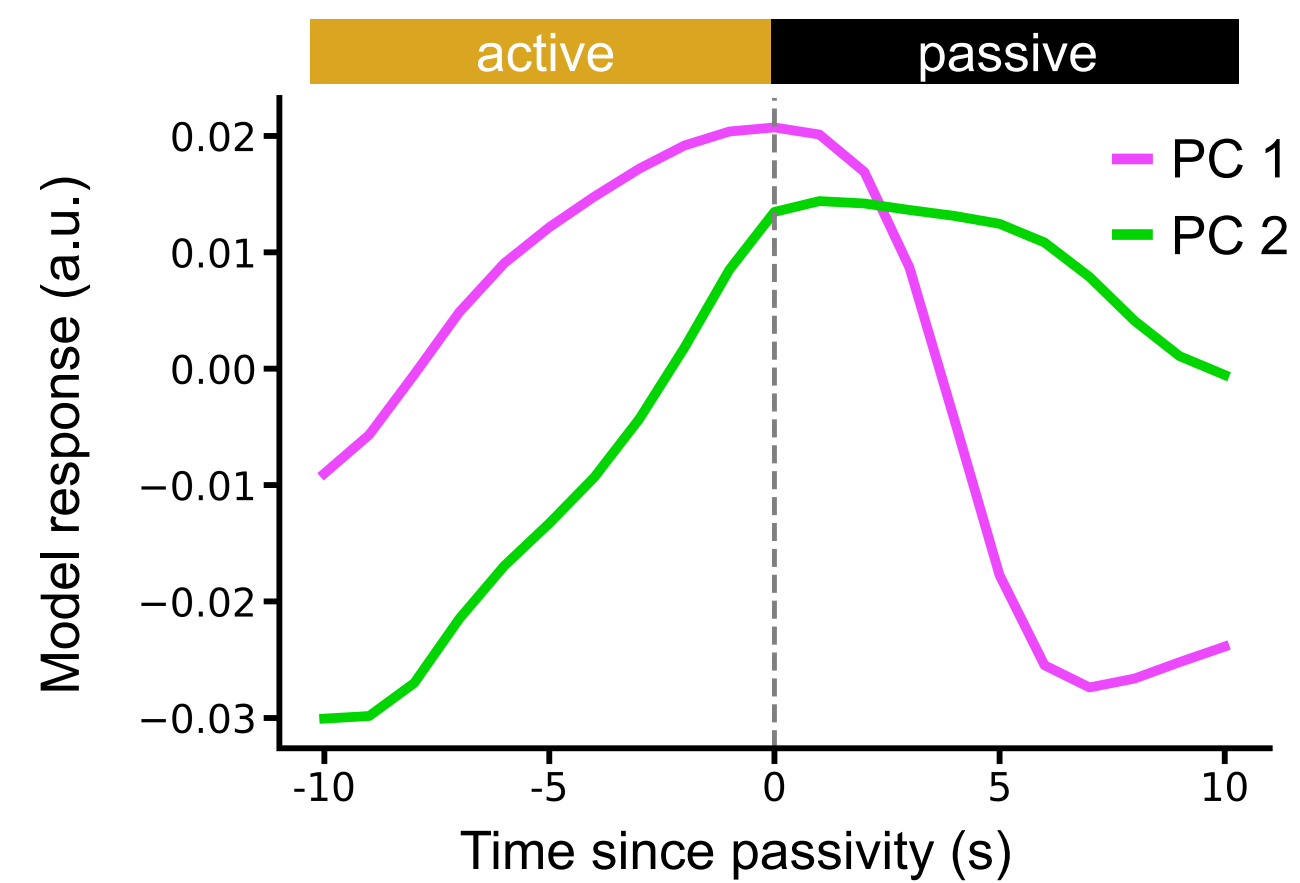



Agent










Zebrafish



Neural-Glial Circuits Emerge in 3M-Progress Agents



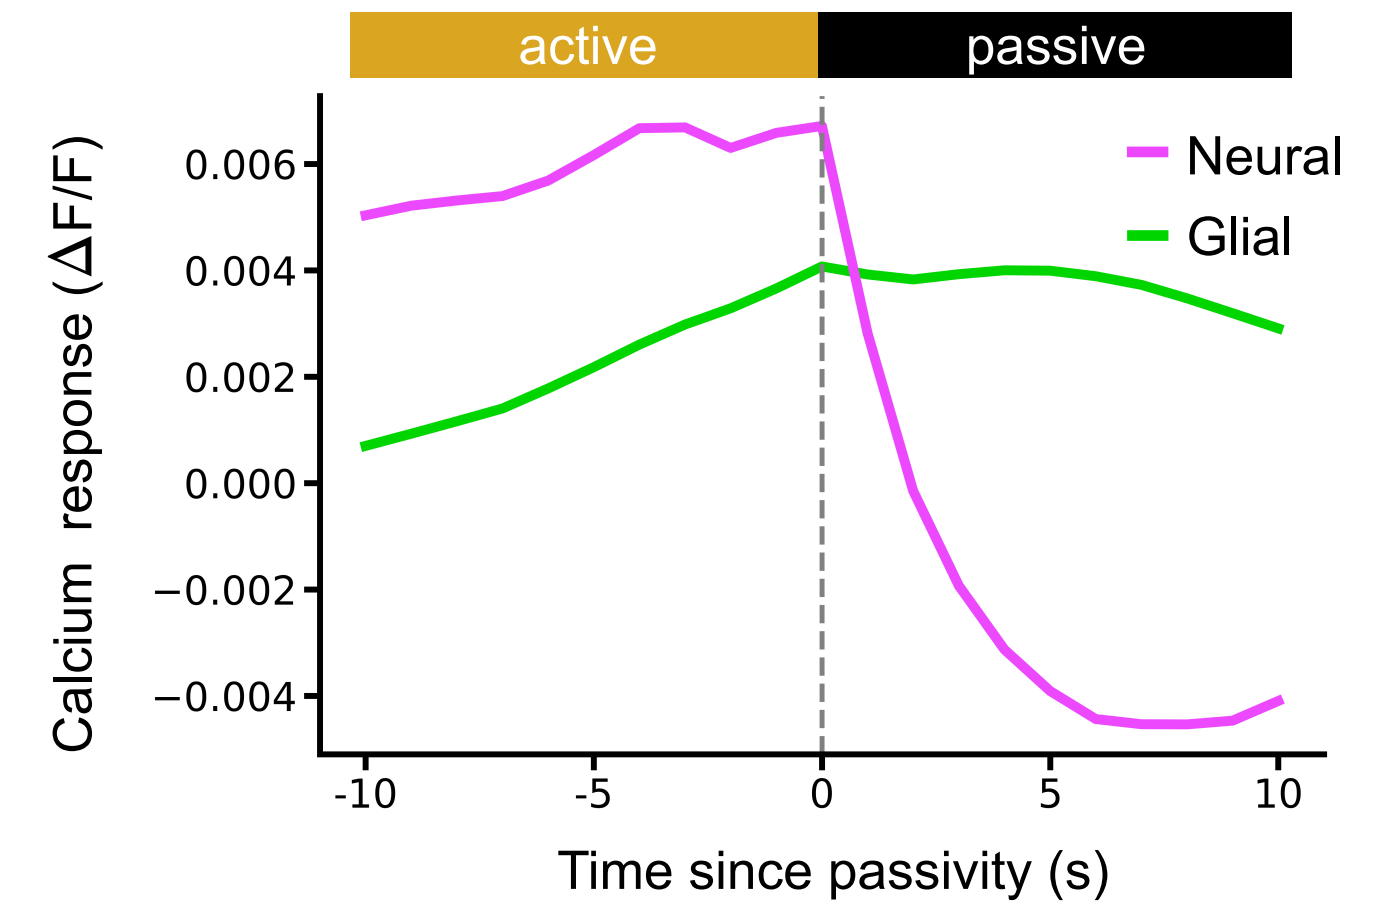
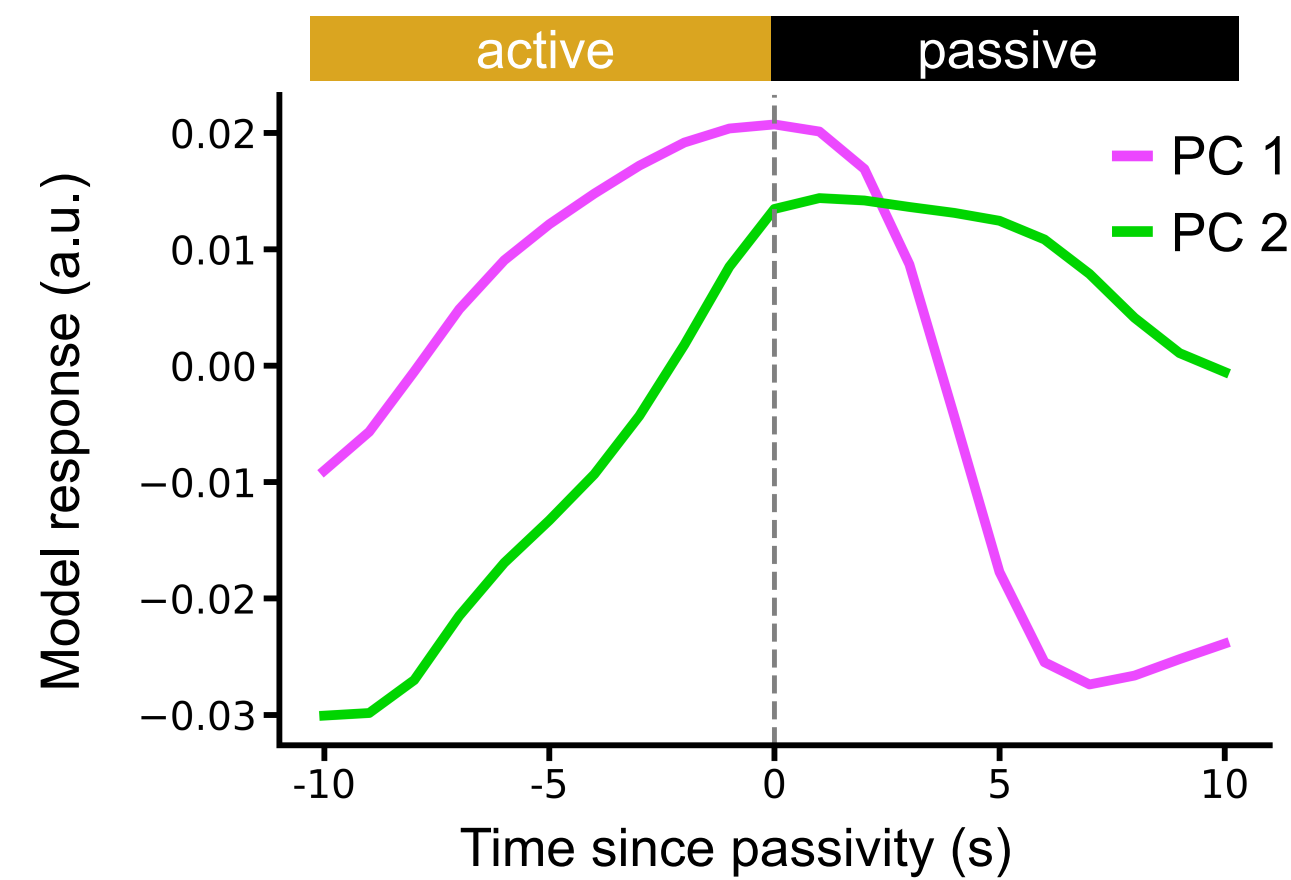


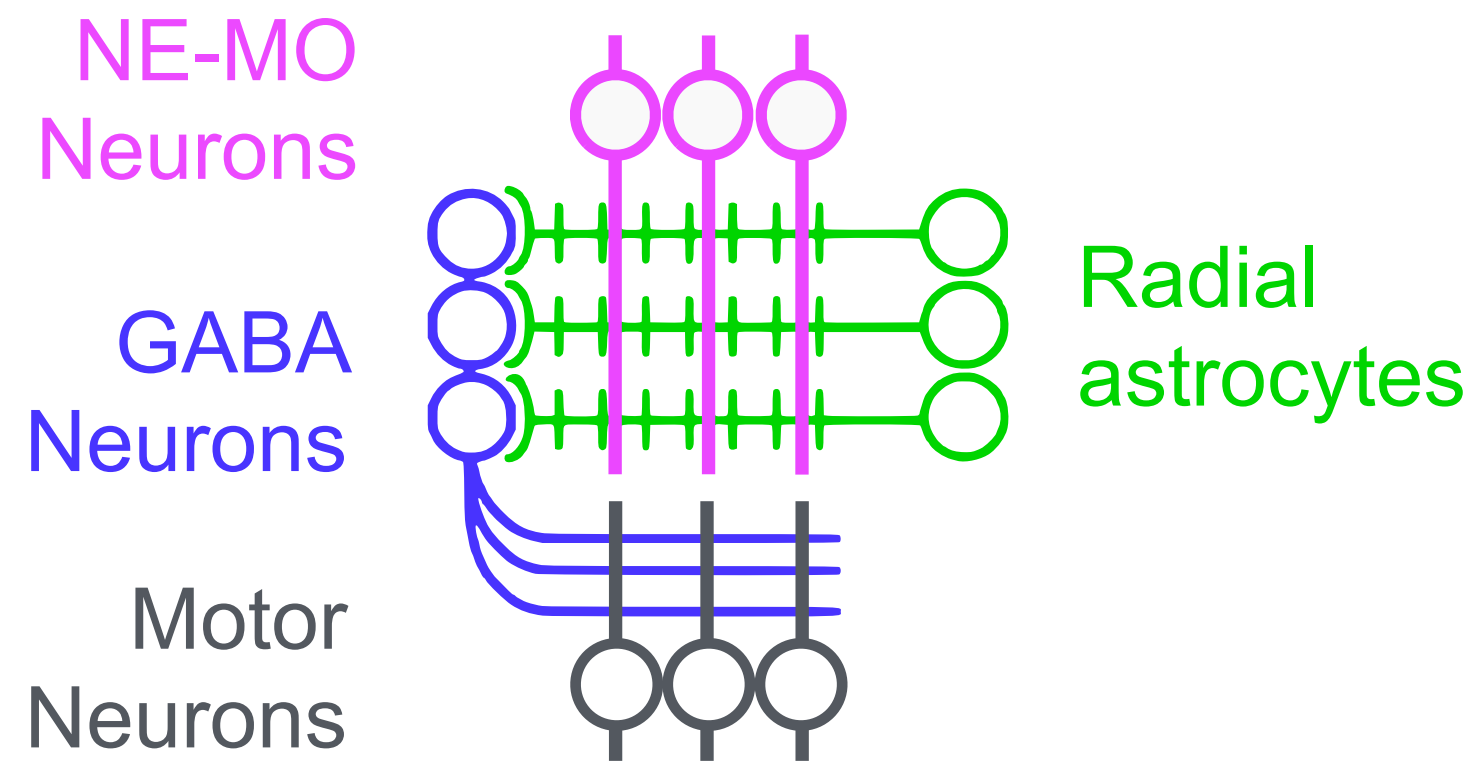


Neural-Glial Circuits Emerge in 3M-Progress Agents



Neural-Glial Circuits Emerge in 3M-Progress Agents


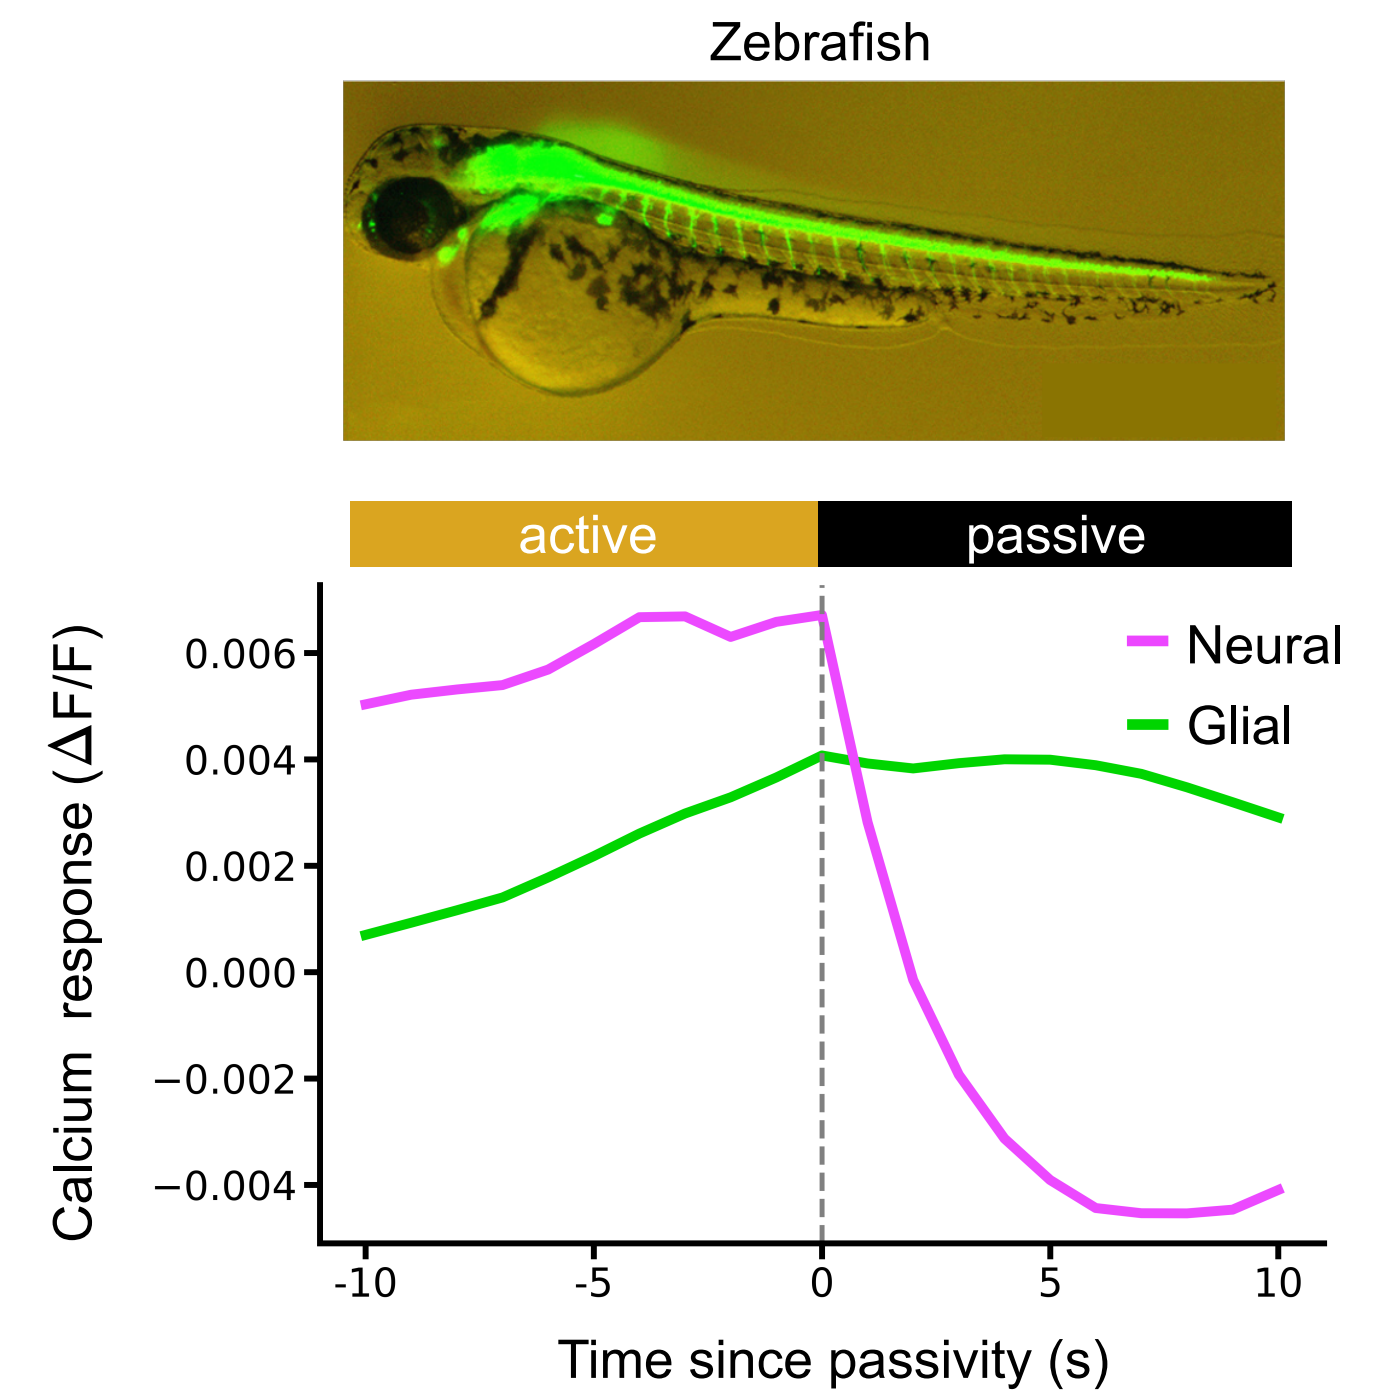

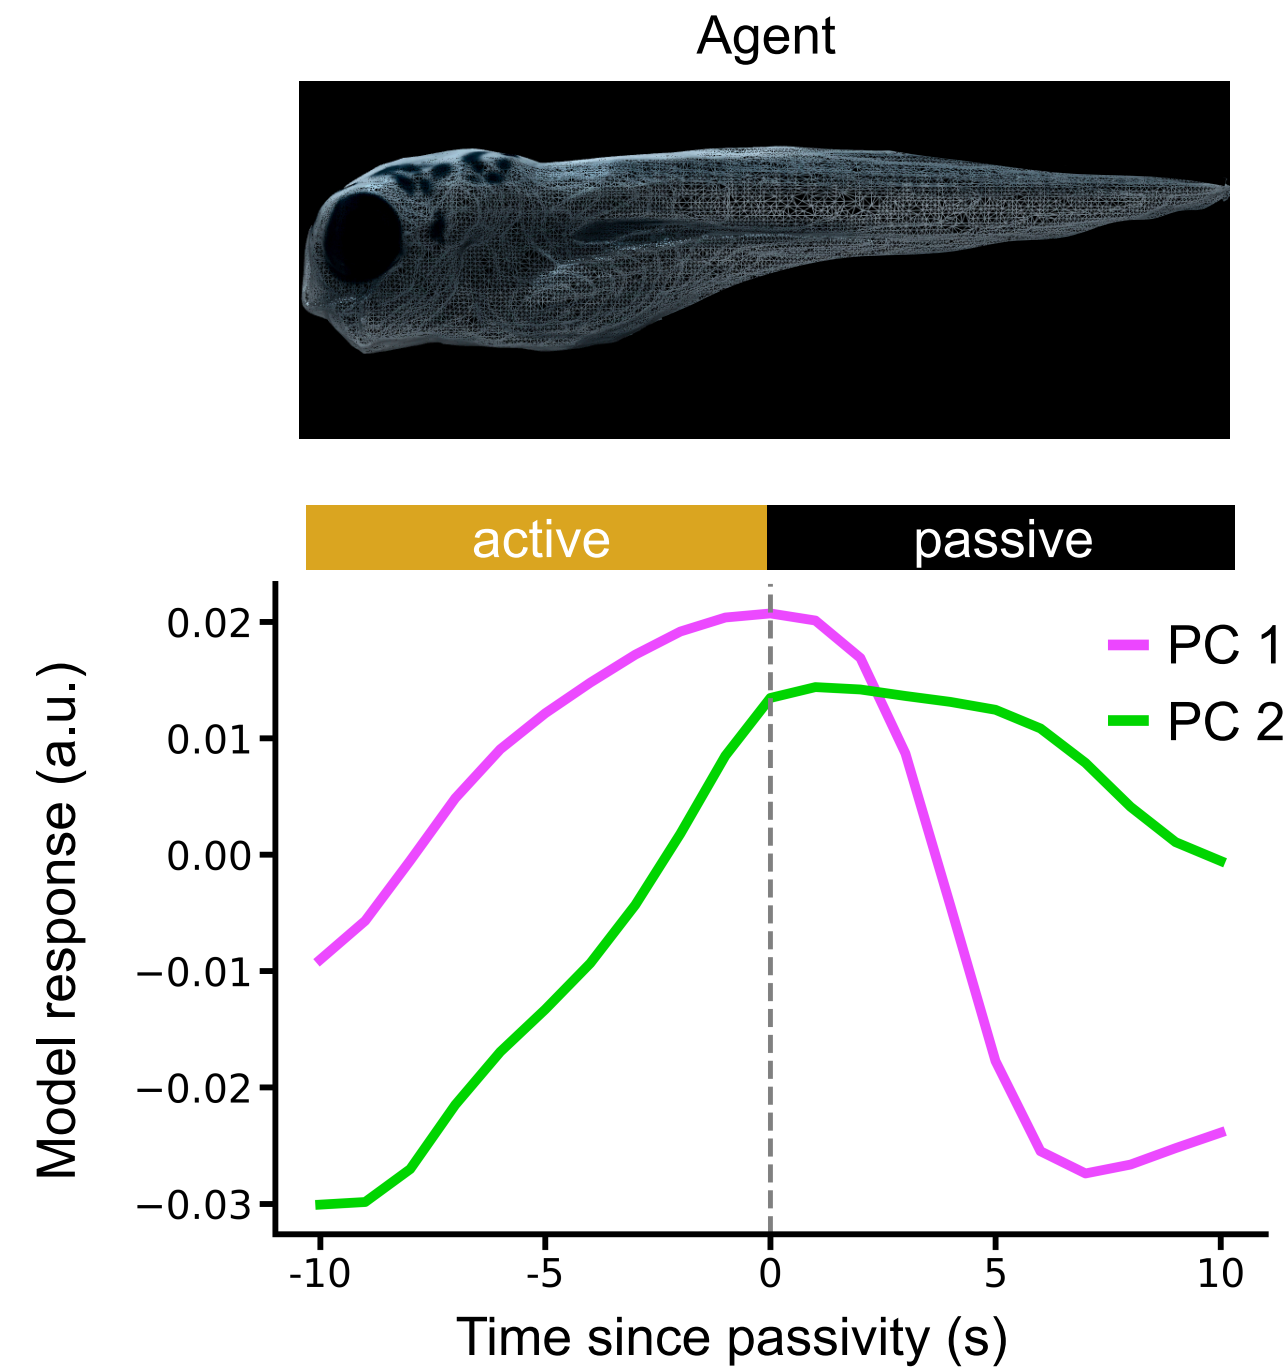

The promise of Embodied NeuroAI



Neural-Glial Circuits Emerge in 3M-Progress Agents

The promise of Embodied NeuroAI


- Propose and validate behaviorally relevant circuit hypotheses at macroscopic scale in closed loop.




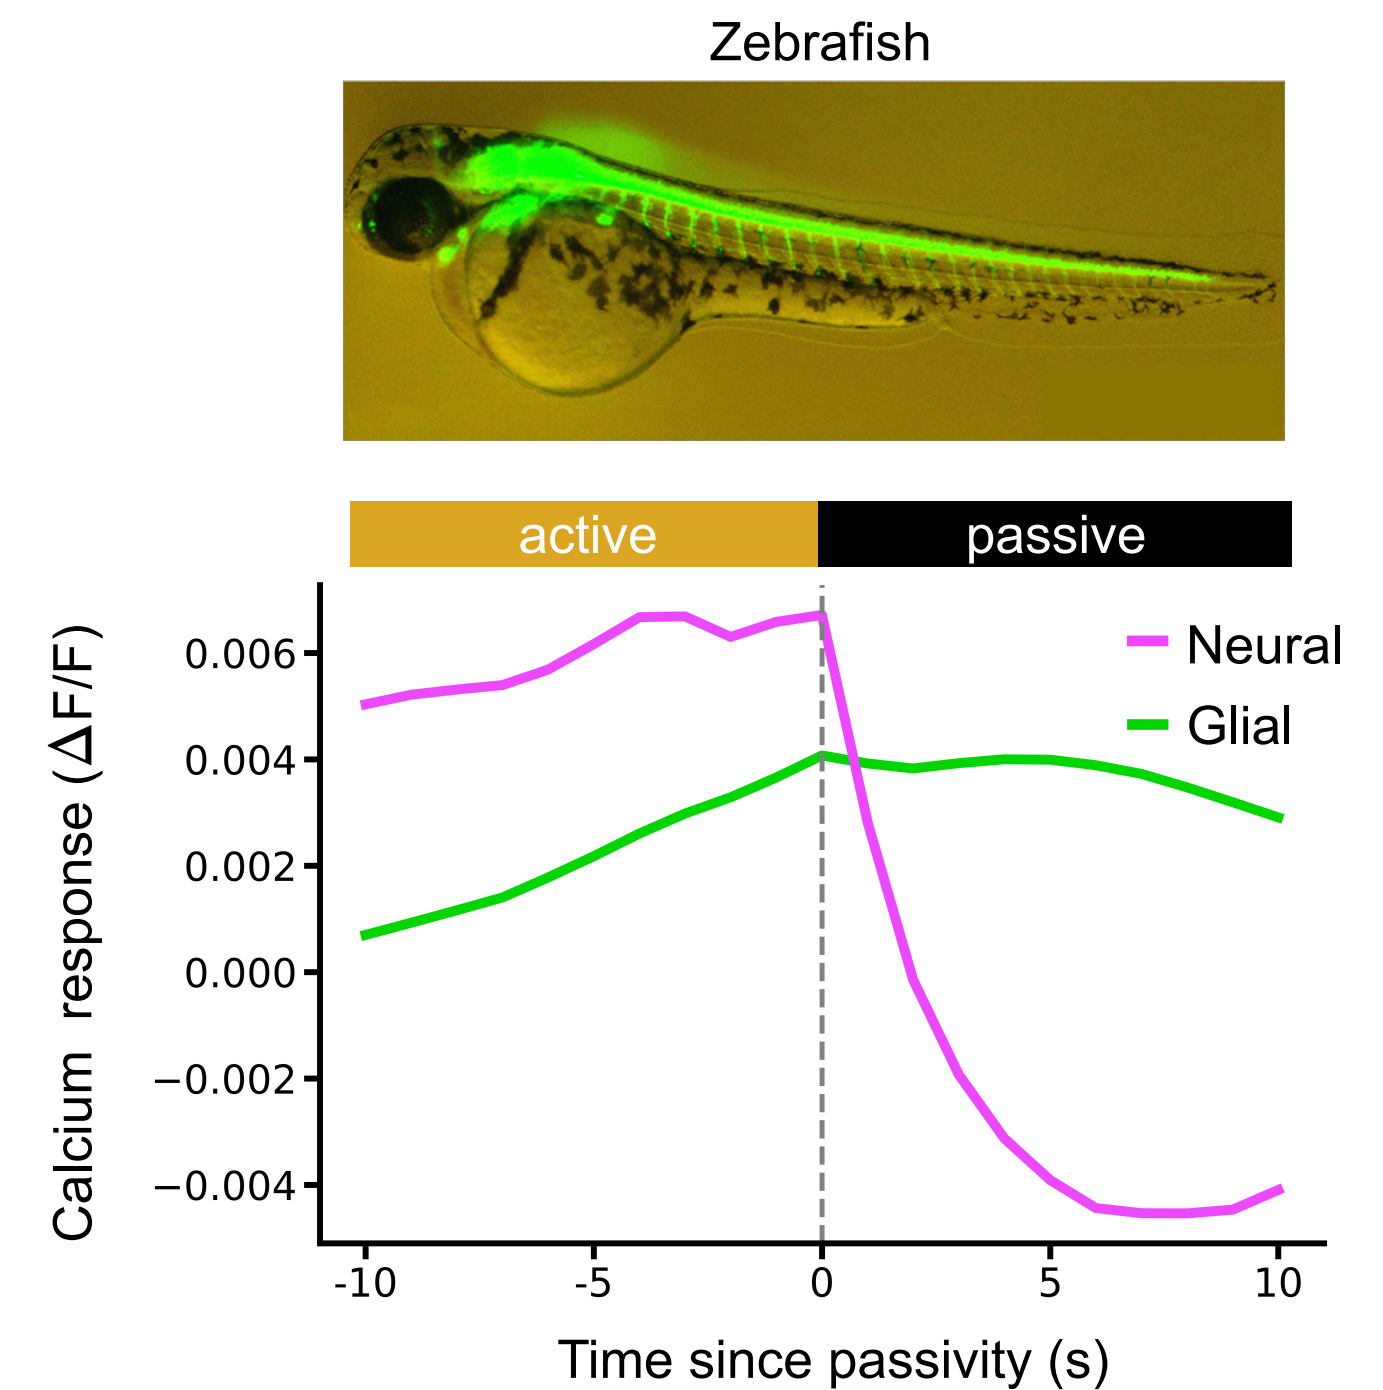

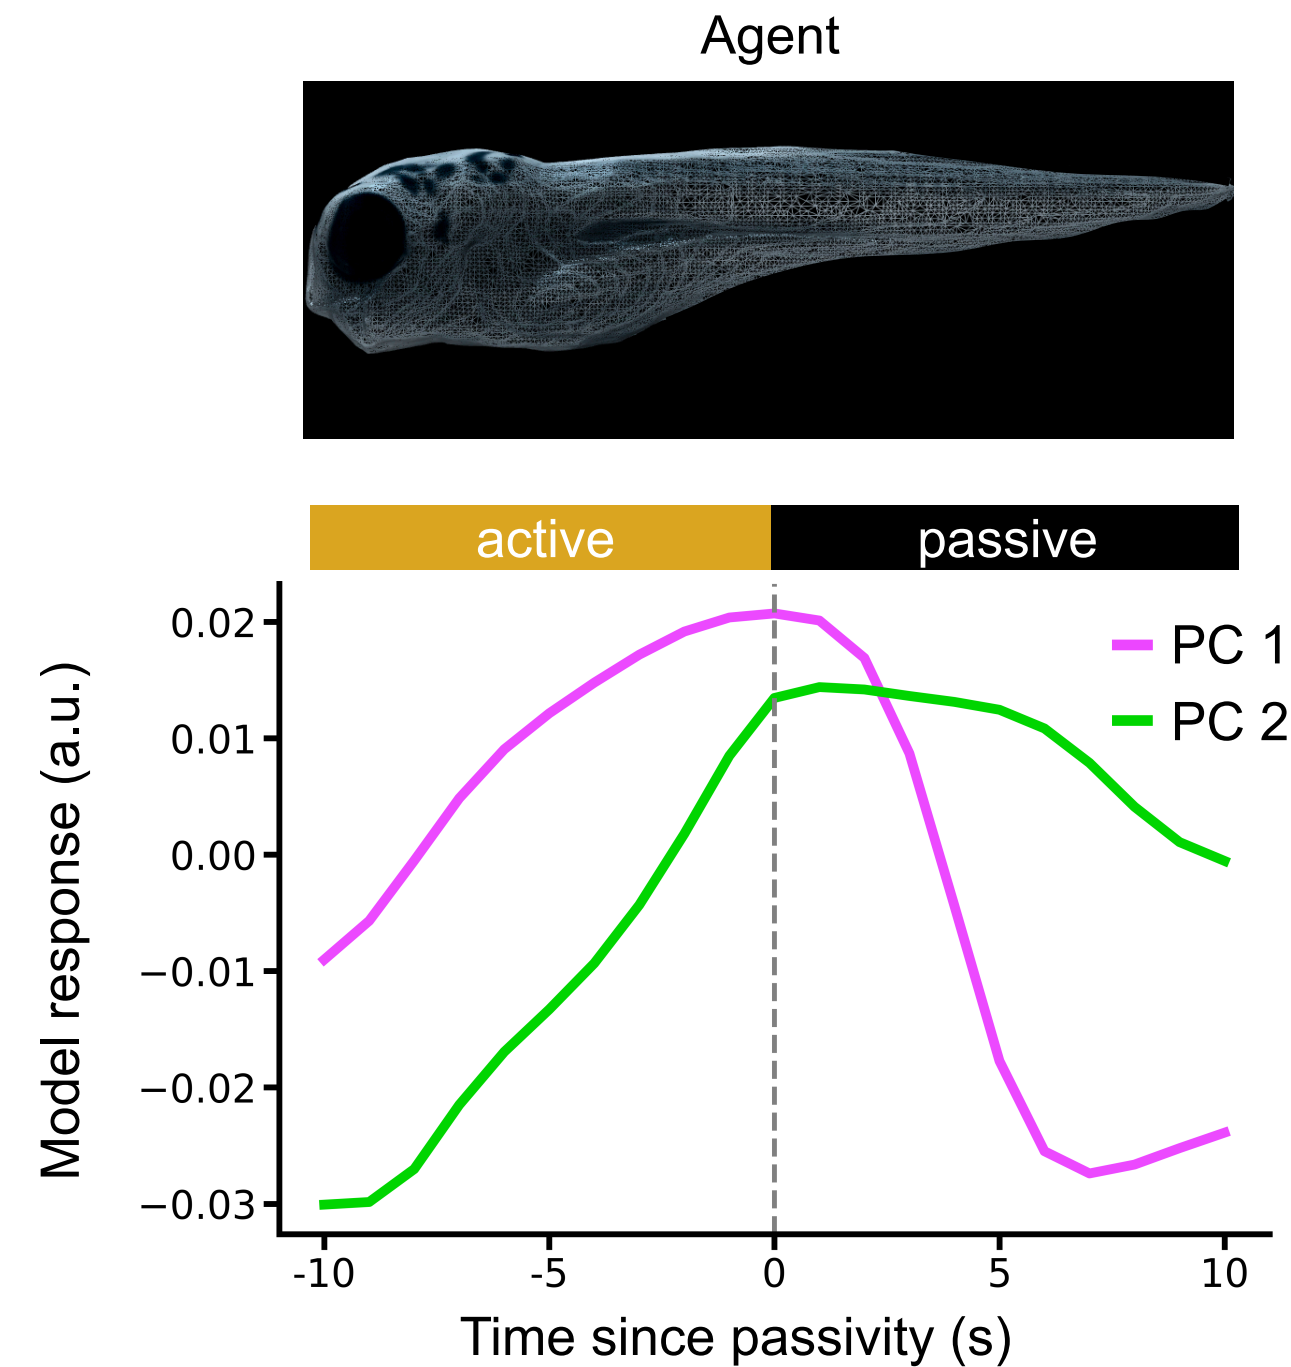
Neural-Glial Circuits Emerge in 3M-Progress Agents

The promise of Embodied NeuroAI

- Propose and validate behaviorally relevant circuit hypotheses at macroscopic scale in closed loop.



- Algorithms with Animal-like sensorimotor intelligence



NeurIPS 2025 Paper



<https://arxiv.org/abs/2506.00138>



Reece Keller



Alyn Kirsch



Felix Pei



Xaq Pitkow



Leo Kozachkov