Master's Thesis Defense

Towards the neural basis of belief state computation in the brain

Célia Benquet

Current unreliable observation

Current incomplete observation

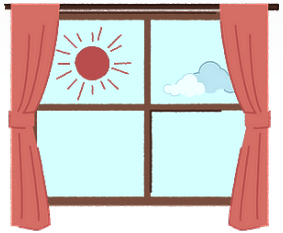

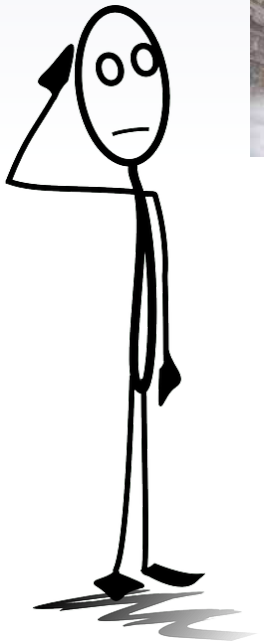

Past experiences and actions

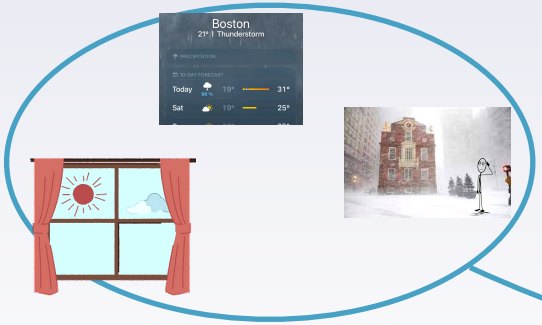

$P(storm \mid observations)$

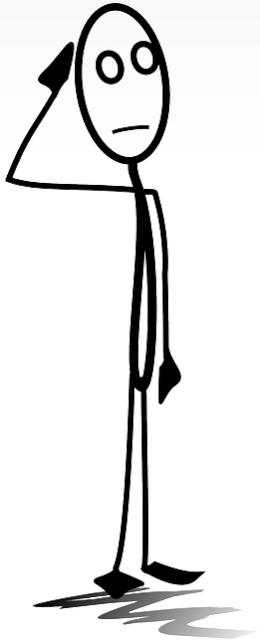

Markov Decision Process (MDP)

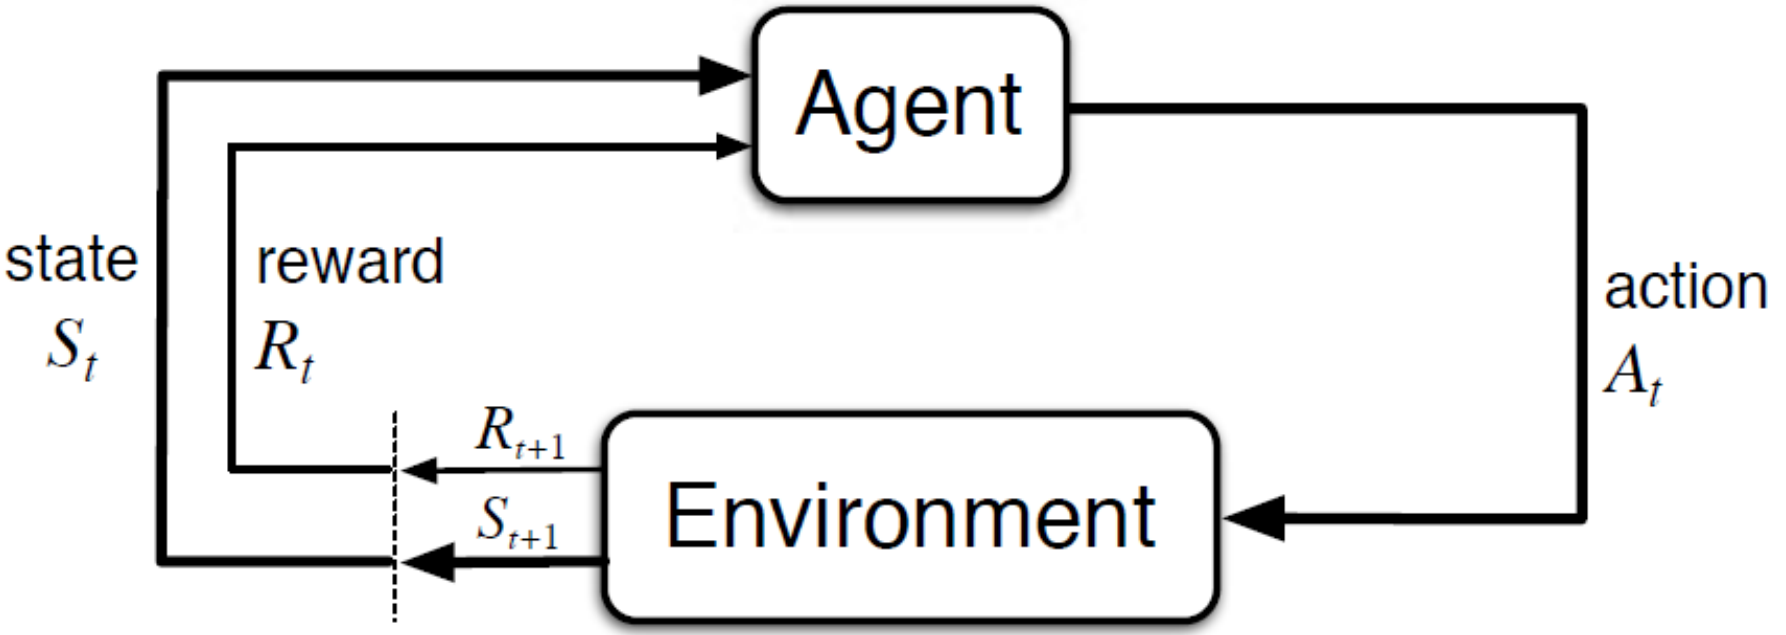

Sutton & Barto, 1998

Starkweather & Uchida, 2020
Schultz et al., 1997, Science

Temporal Difference learning (TD)

Value estimate

$$V(s_t) = \sum_{\tau=t}^{\infty} \gamma^{\tau-t} r(\tau)$$

$$\hat{V}(t) = \sum w_i x_i(t) = w(t)^T x(t)$$

(CSC features representation)

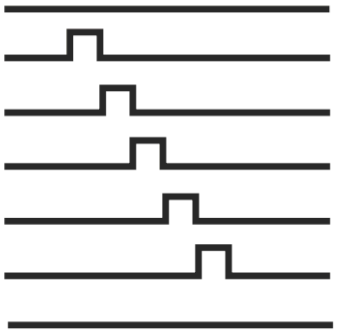

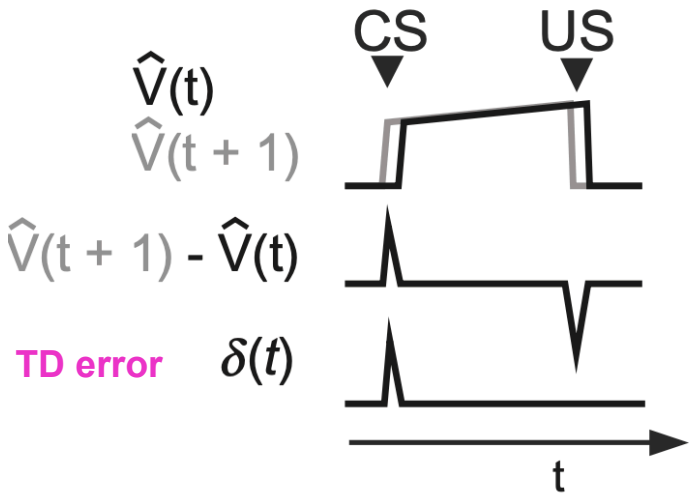

CSC: Complete Serial Compounds

Partially-observable Markov decision process (POMDP)

$$\hat{V}(t) = \sum w_i x_i(t) = w(t)^T x(t)$$

$$\hat{V}(t) = \sum w_i b_i(t) = w(t)^T b(t)$$

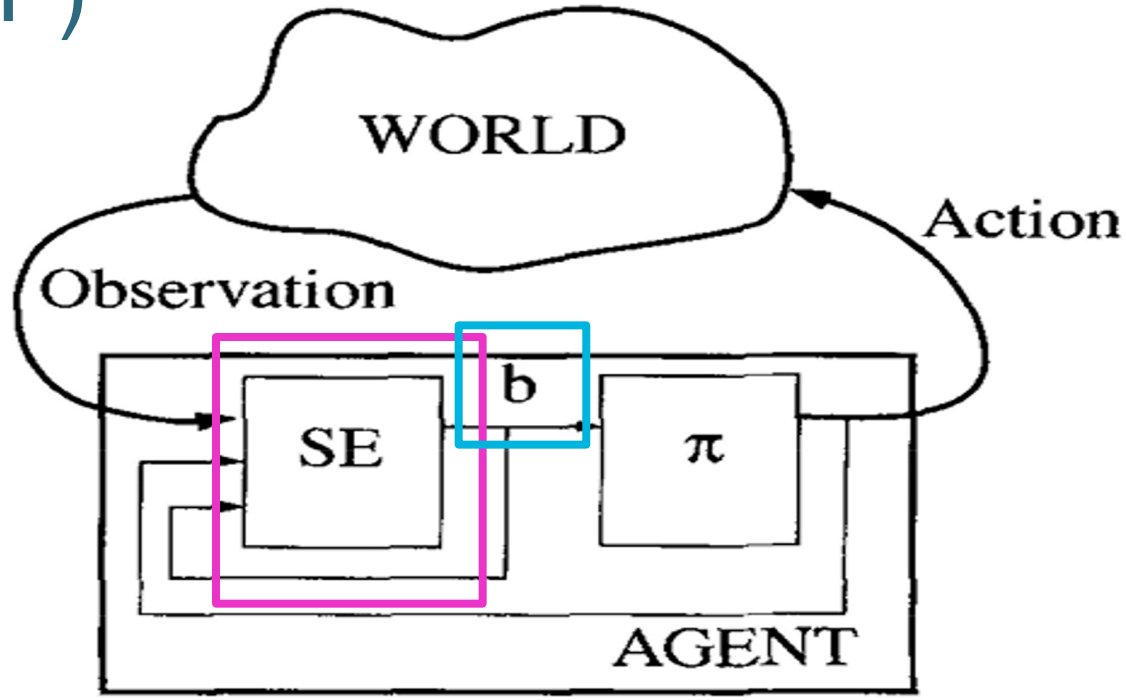

SE: state estimator
π: policy

Kaelbling et al., 1998, Artif. Intell.
Daw et al., 2006, Neural Comput.

Variable reward delay paradigm

Deterministic task

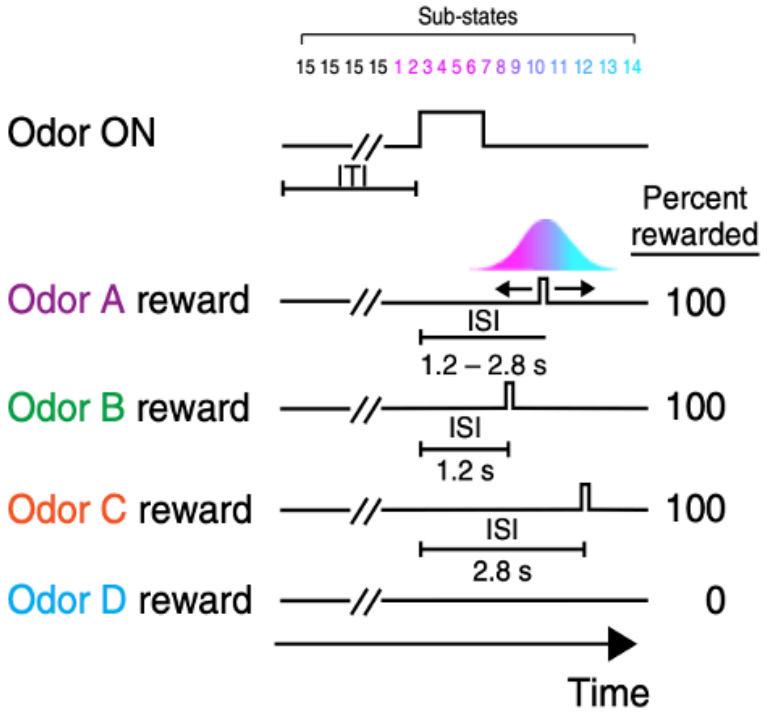

Stochastic task

(90% rewarded)

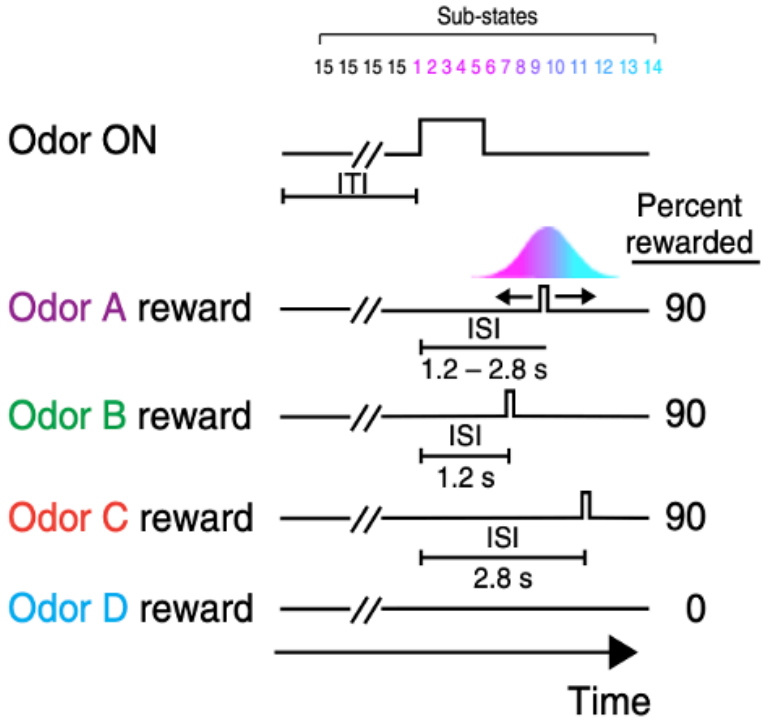

Starkweather et al., 2017, Nat. Neurosci.

Variable reward delay paradigm

Deterministic task

Fully observable states

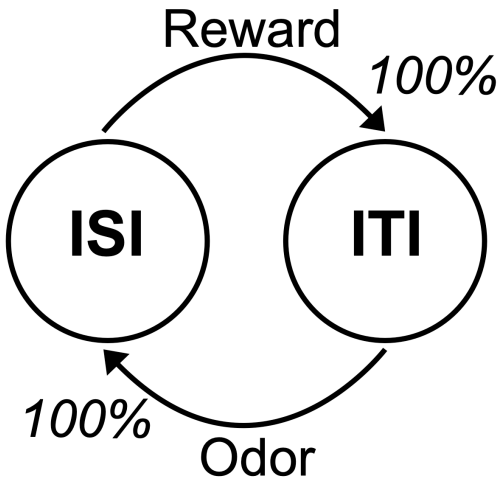

Stochastic task

(90% rewarded)

Hidden states

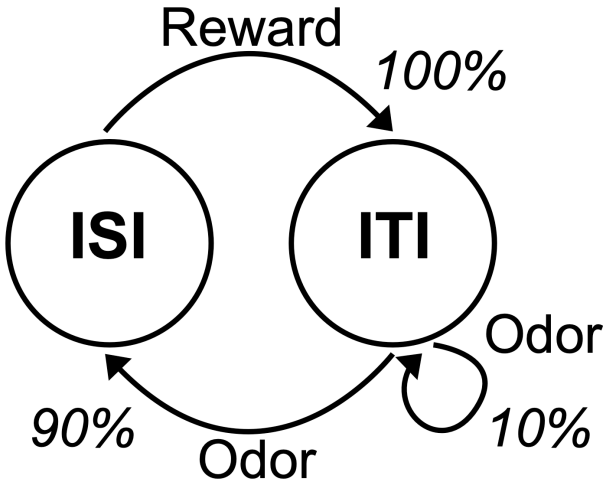

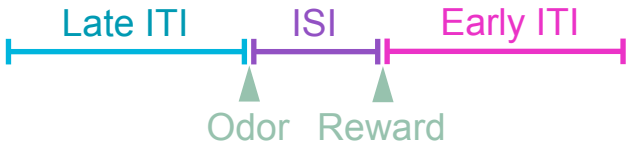

ISI: Inter-stimuli interval
ITI: Inter-trial interval

Starkweather et al., 2017, Nat. Neurosci.

TD model with belief state

Deterministic task **Stochastic task**

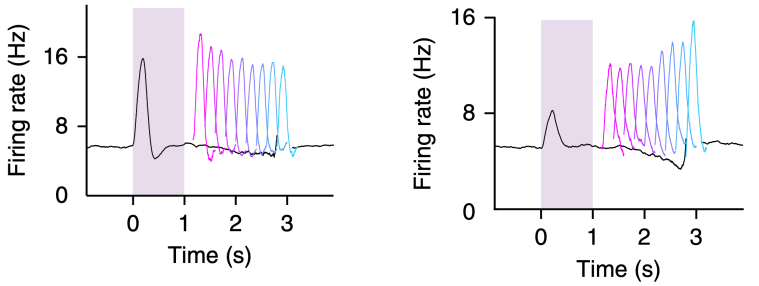

Experimental data in the VTA

TD model with CSC

Deterministic task **Stochastic task**

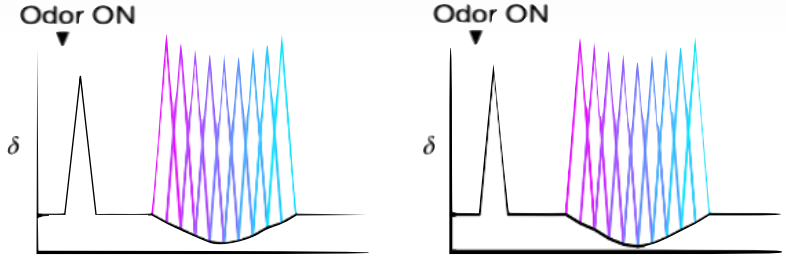

TD model with belief state

Deterministic task **Stochastic task**

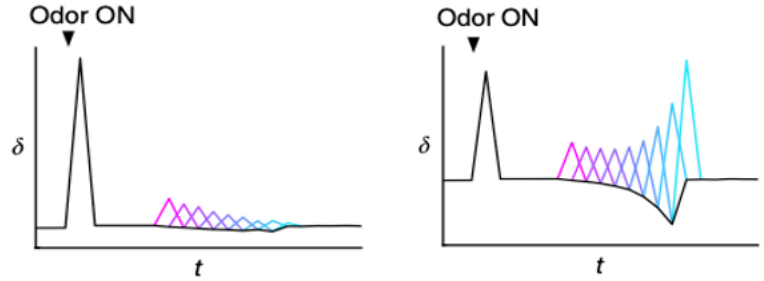

VTA: ventral tegmental area

Starkweather et al., 2017, Nat. Neurosci.

Belief state representation in the brain

Stochastic task, mPFC inactivation

mPFC: medial pre-frontal cortex
SalB: salvinorin B

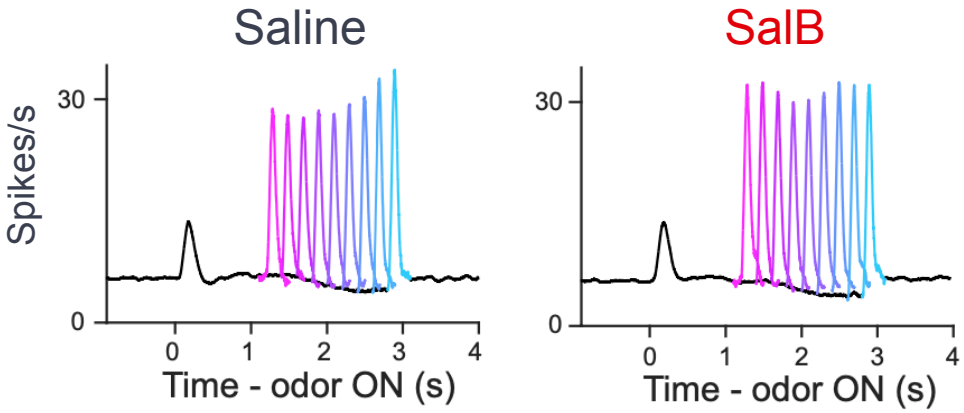

TD model with CSC

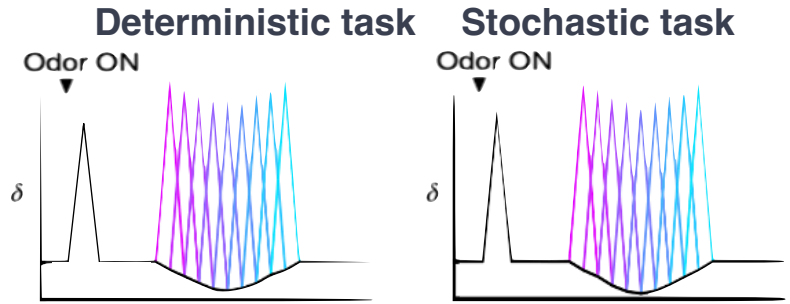

A different region to compute belief state representation?

Starkweather et al., 2018, Neuron

Starkweather et al., 2017, Neuron

Belief state representation in the brain

Stochastic task, mPFC inactivation

mPFC: medial pre-frontal cortex
SalB: salvinorin B

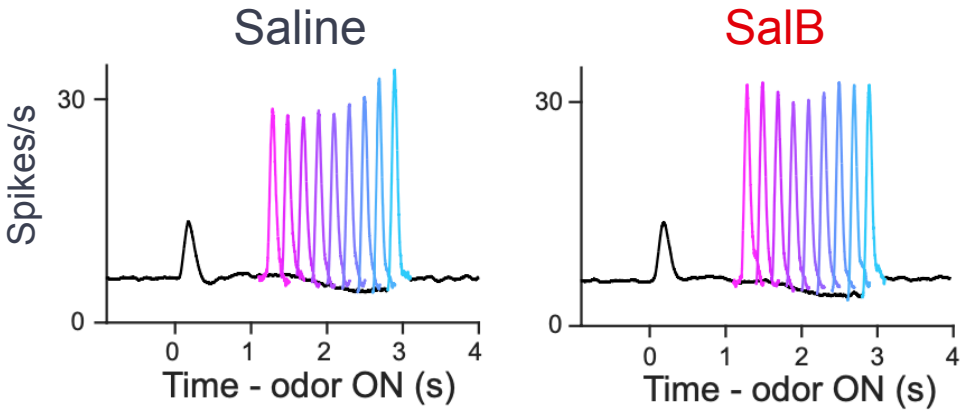

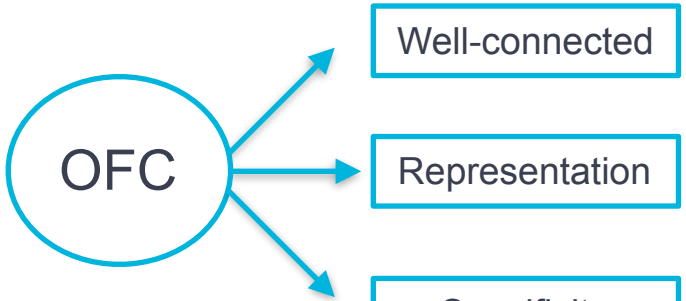

OFC: Orbito-Frontal Cortex

A different region to compute belief state representation?

Starkweather et al., 2018, Neuron

Wilson et al., 2014, Neuron

Is belief state representation computed in the OFC, and if so, how?

Is belief state representation computed in the OFC, and if so, how?

Individual neurons investigation

Neural activity encoder

State decoder

Evolution of the neural activity at the population level

Deterministic Task

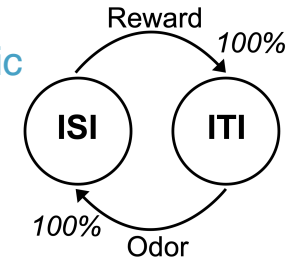

Sandra Romero Pinto

Is belief state representation computed in the OFC, and if so, how?

Individual neurons investigation

Neural activity encoder

State decoder

Evolution of the neural activity at the population level

OFC neurons encode stimuli presence in the task … which are also state-relevant

Odor-tuned

Reward-tuned

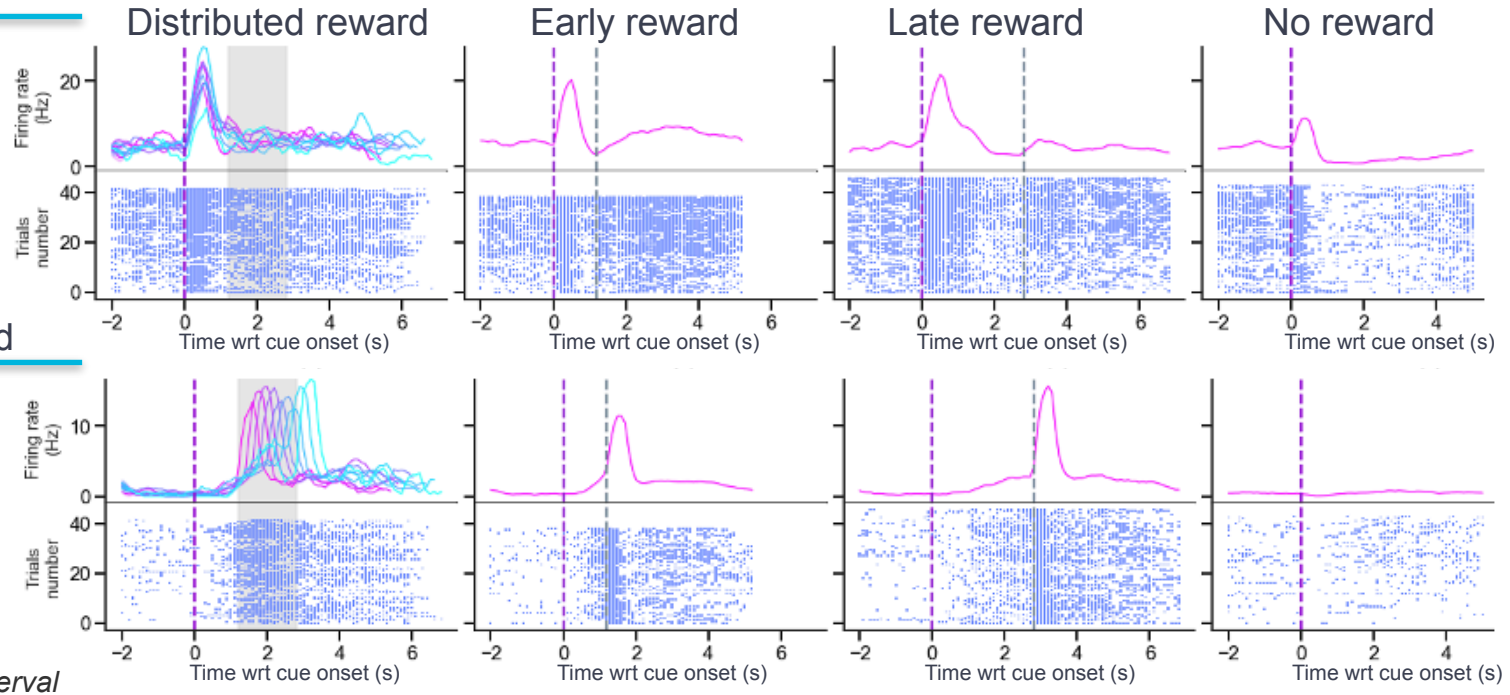

ISI: Inter-stimuli interval
ITI: Inter-trial interval

OFC neurons encode the task structure

Background-tuned

Bursting delay-tuned

Sustained delay-tuned

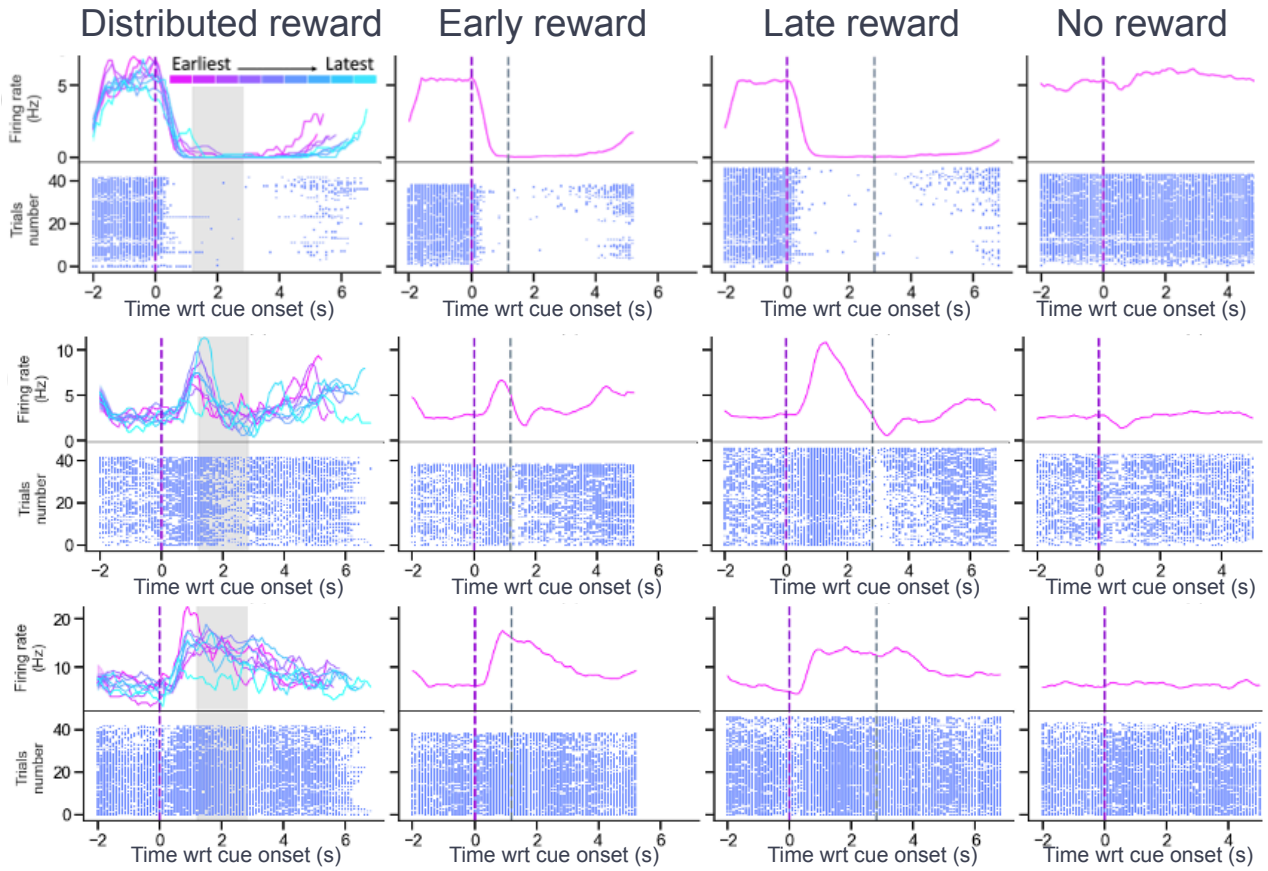

Activity map in the OFC

- Trial -averaged activity
- Neurons sorted by timing of their peak activity
- Same neurons order for all trial type

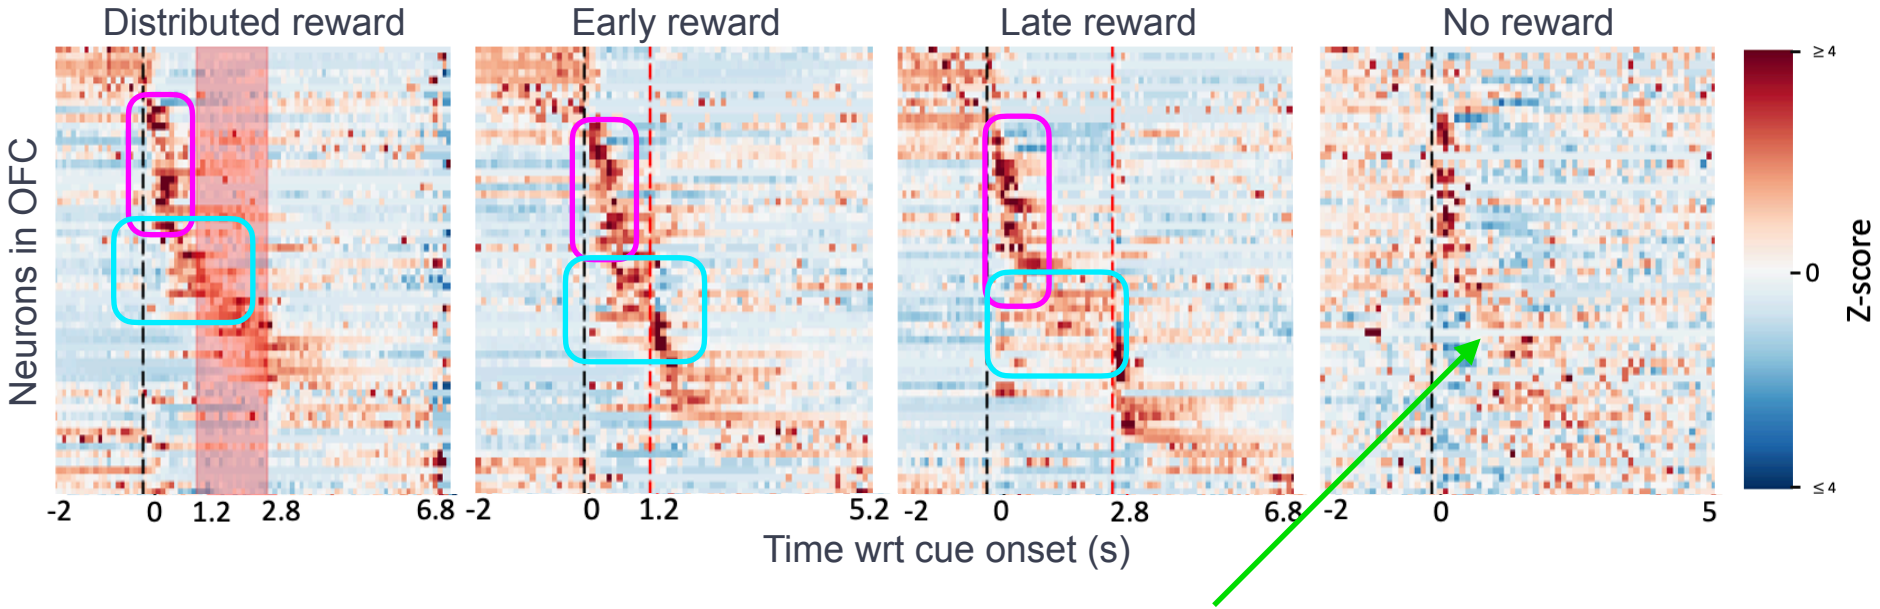

Is belief state representation computed in the OFC, and if so, how?

Individual neurons investigation

Single-neuron activity shows specific task-variables tuning in the OFC

Neural activity encoder

State decoder

Evolution of the neural activity at the population level

Is belief state representation computed in the OFC, and if so, how?

Individual neurons investigation

Single-neuron activity shows specific task-variables tuning in the OFC

Neural activity encoder

State decoder

Evolution of the neural activity at the population level

Poisson GLM model

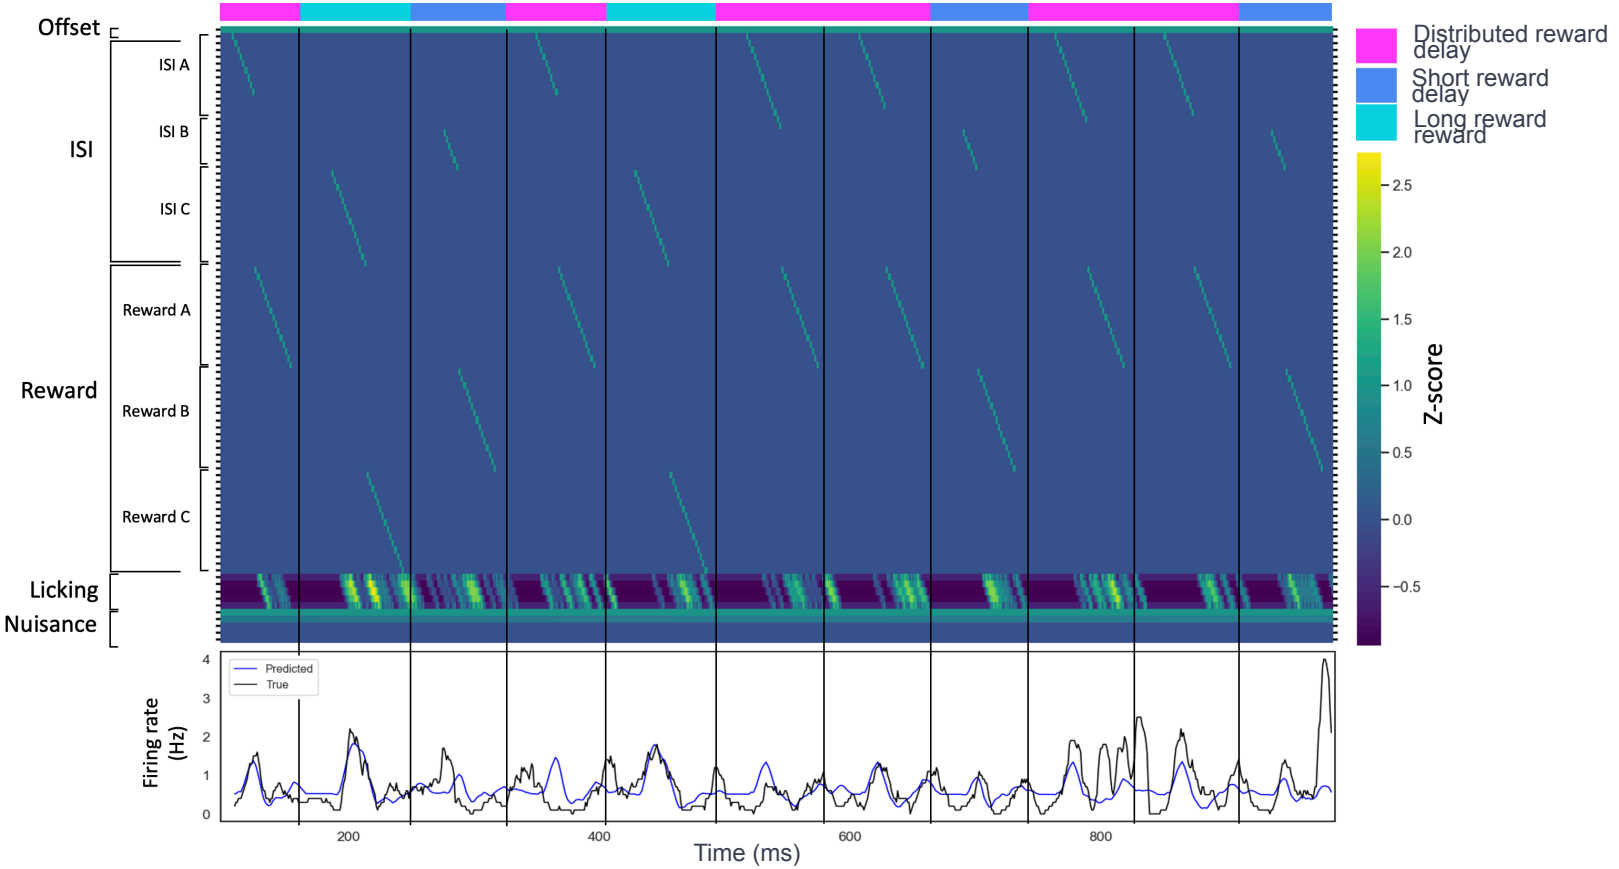

A large fraction of neurons in the OFC encode state-relevant variables

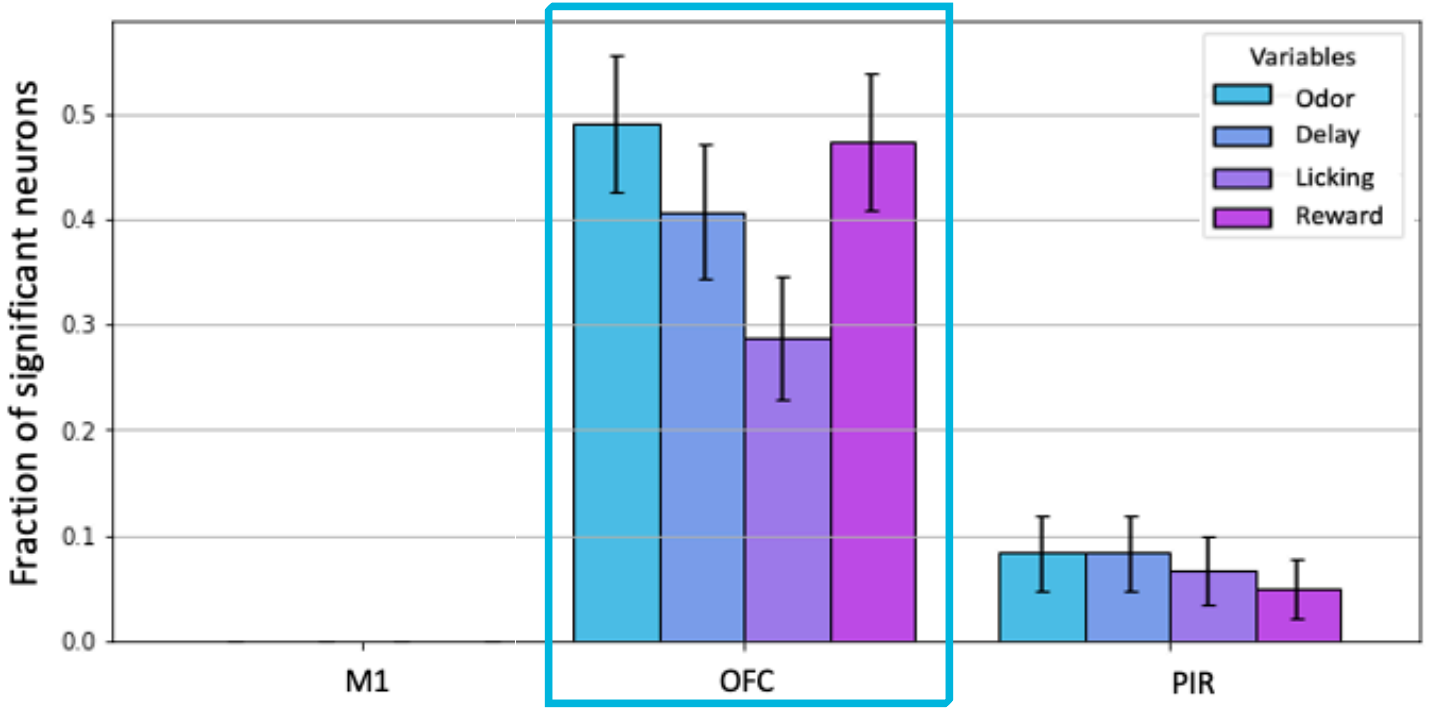

M1: primary motor cortex
PIR: piriform cortex

Is belief state representation computed in the OFC, and if so, how?

Individual neurons investigation

Single-neuron activity shows specific task-variables tuning in the OFC

Neural activity encoder

OFC neurons encode variables that are essential to state representation

State decoder

Evolution of the neural activity at the population level

Is belief state representation computed in the OFC, and if so, how?

Individual neurons investigation

Single-neuron activity shows specific task-variables tuning in the OFC

Neural activity encoder

OFC neurons encode variables that are essential to state representation

State decoder

Evolution of the neural activity at the population level

States classifier

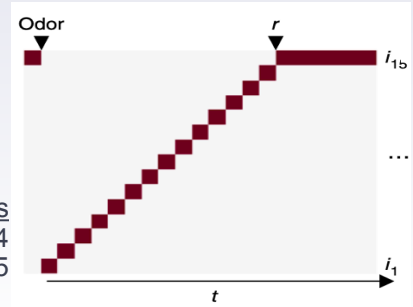

Micro-states
ISI: 1-14
ITI: 15

Macro-states classifier

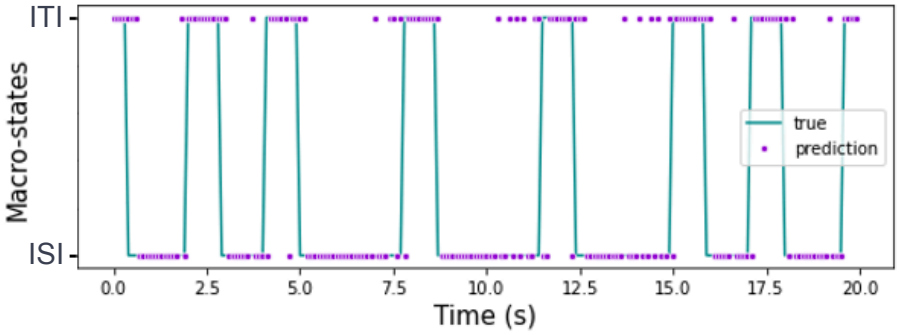

Micro-states classifier

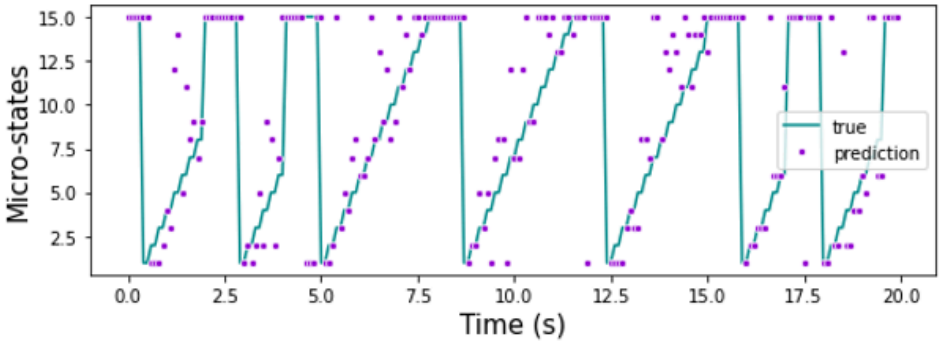

States classifier

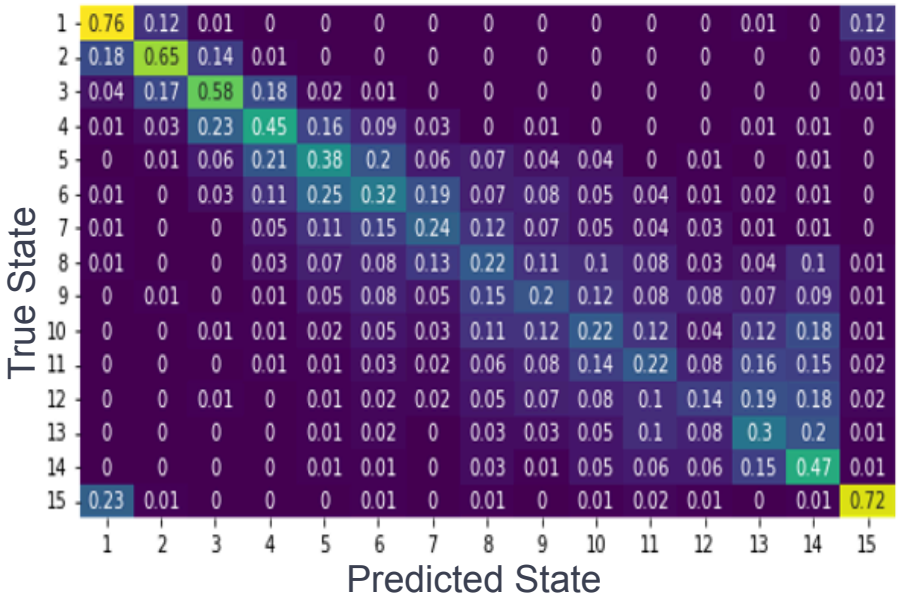

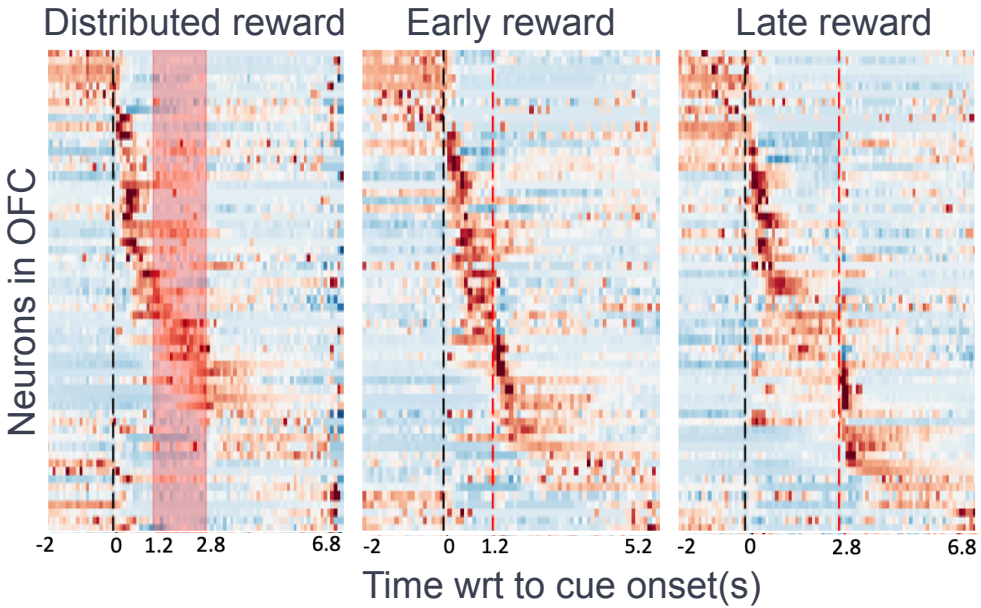

Is belief state representation computed in the OFC, and if so, how?

Individual neurons investigation

Single-neuron activity shows specific task-variables tuning in the OFC

Neural activity encoder

OFC encode variables that are essential to belief state representation

State decoder

OFC population activity is sufficient to discern between states

Evolution of the neural activity at the population level

Is belief state representation computed in the OFC, and if so, how?

Individual neurons investigation

Single-neuron activity shows specific task-variables tuning in the OFC

Neural activity encoder

OFC encode variables that are essential to belief state representation

State decoder

OFC population activity is sufficient to discern between states

Evolution of the neural activity at the population level

Representing the belief states in the brain

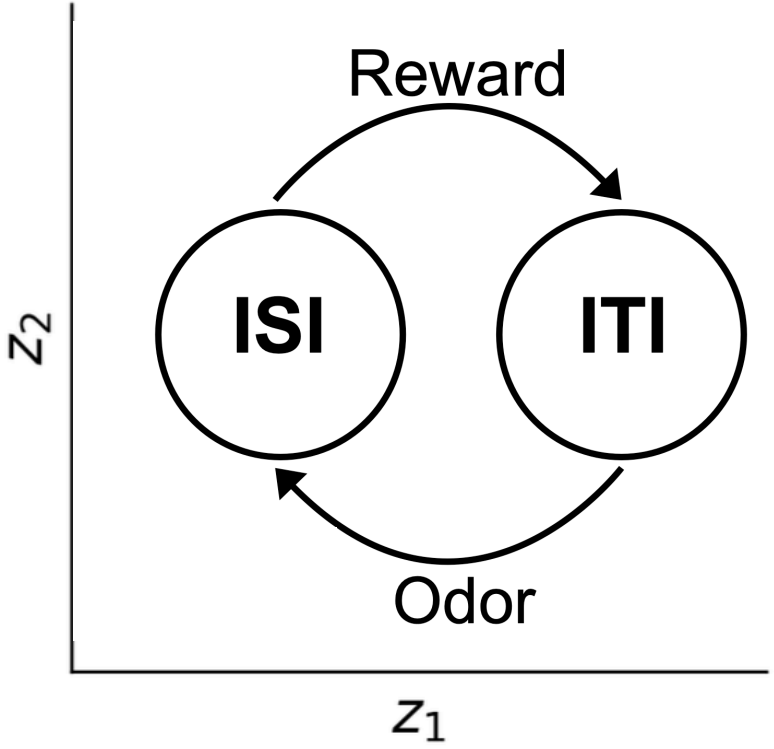

GRU network modeling

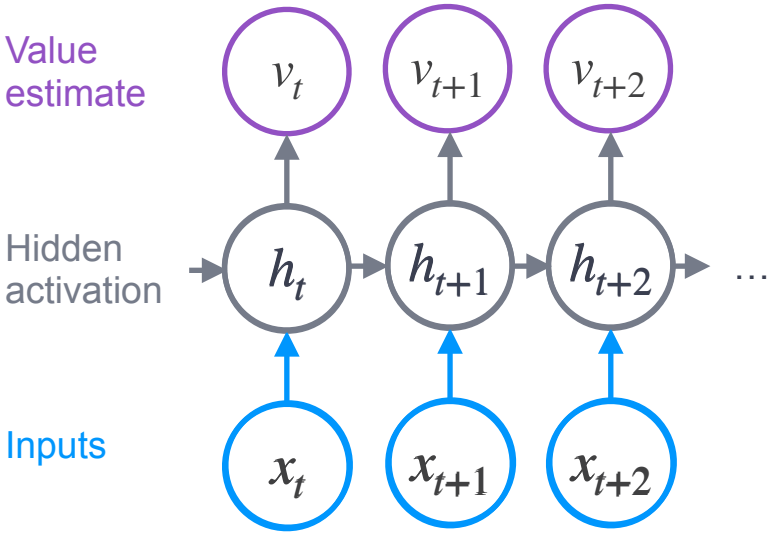

Activation layers state decoder

Stochastic task

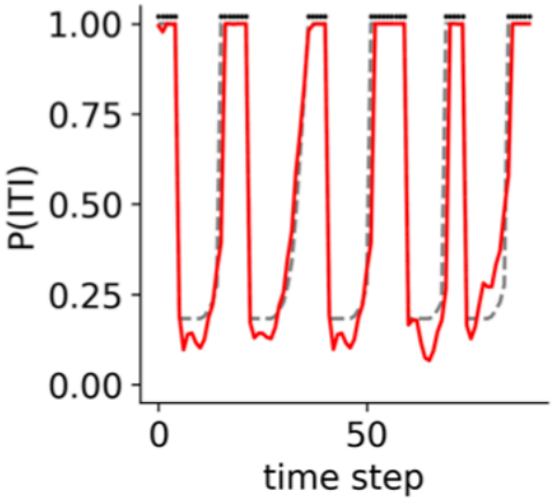

GRU: gated recurrent unit

GRU network modeling

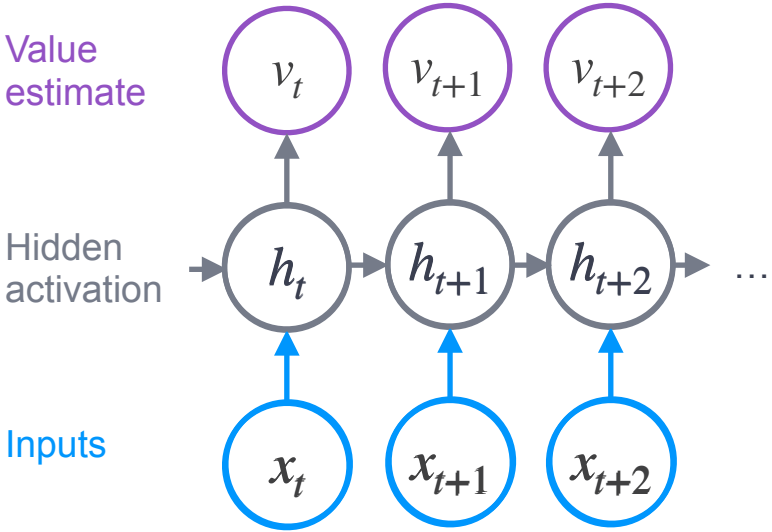

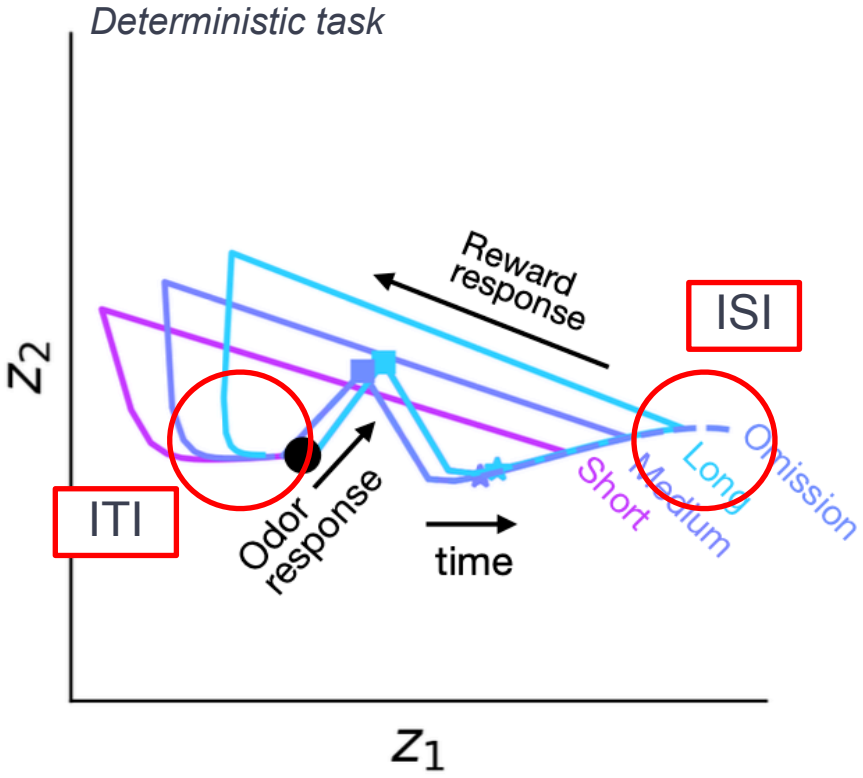

GRU network modeling

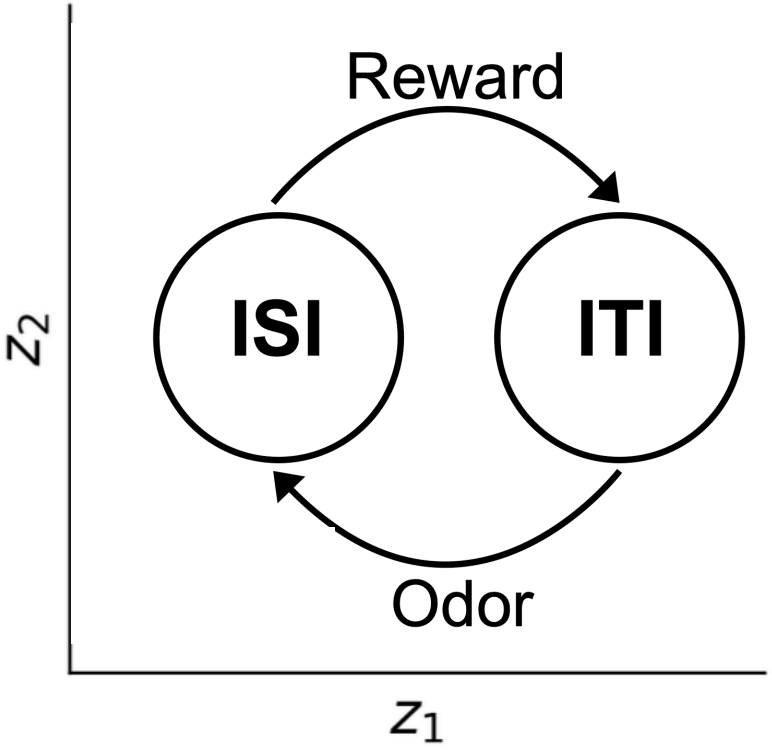

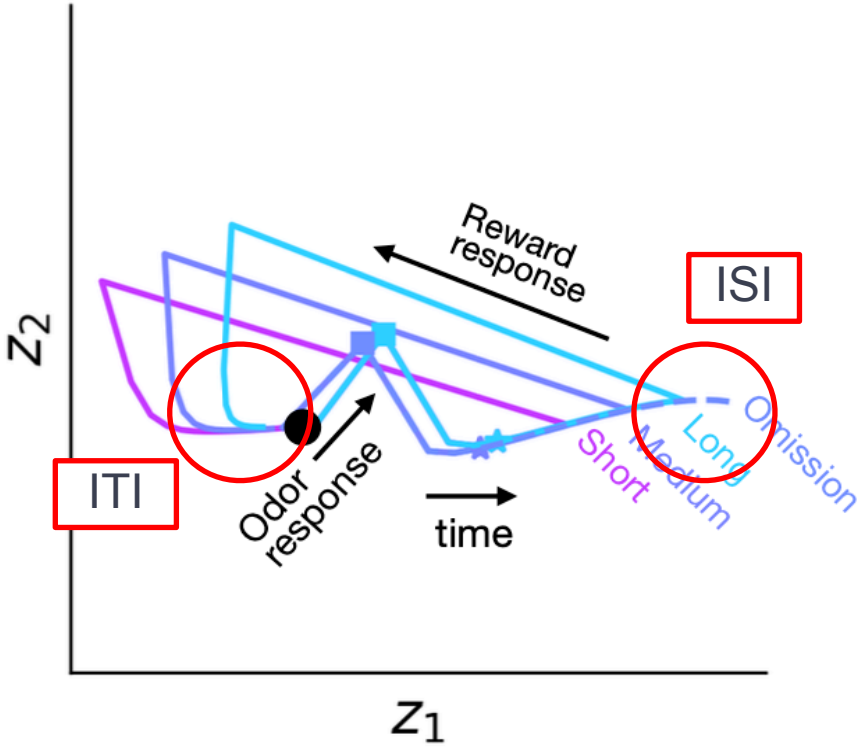

Geometrical representation (PCA)

- Low-dimensional trial-averaged activity trajectory

Distributed reward

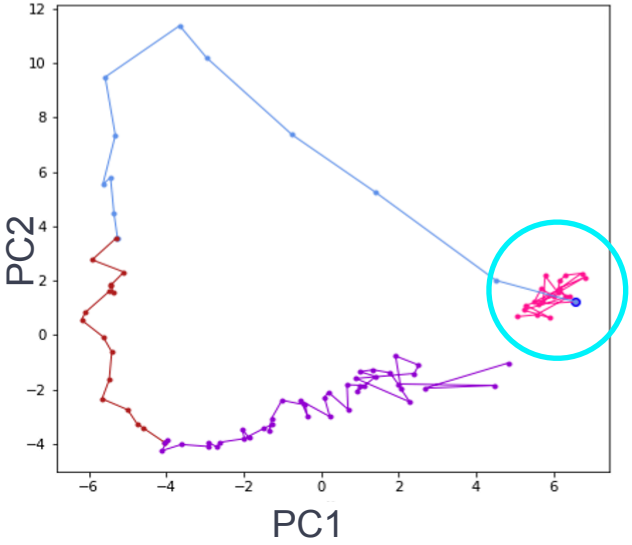

Early reward

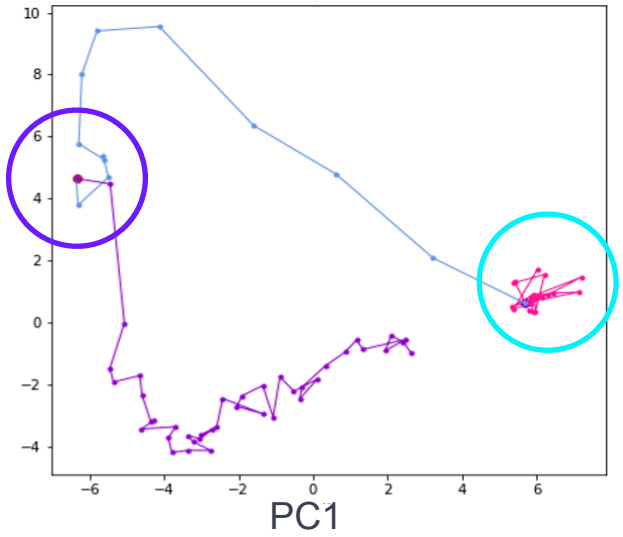

Late reward

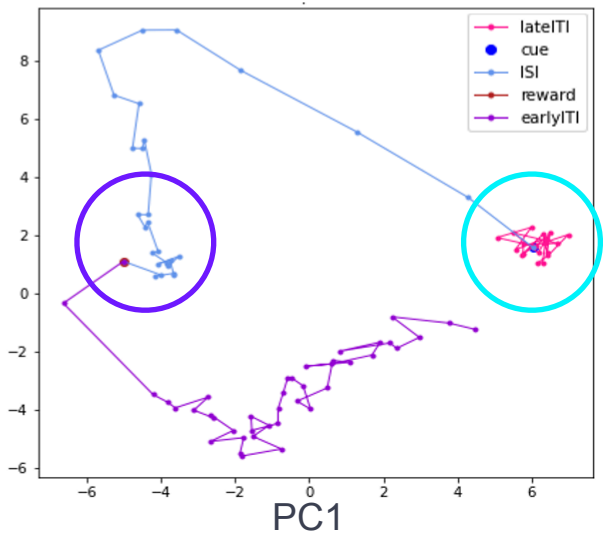

Geometrical representation (PCA)

Distributed reward

Early reward

Late reward

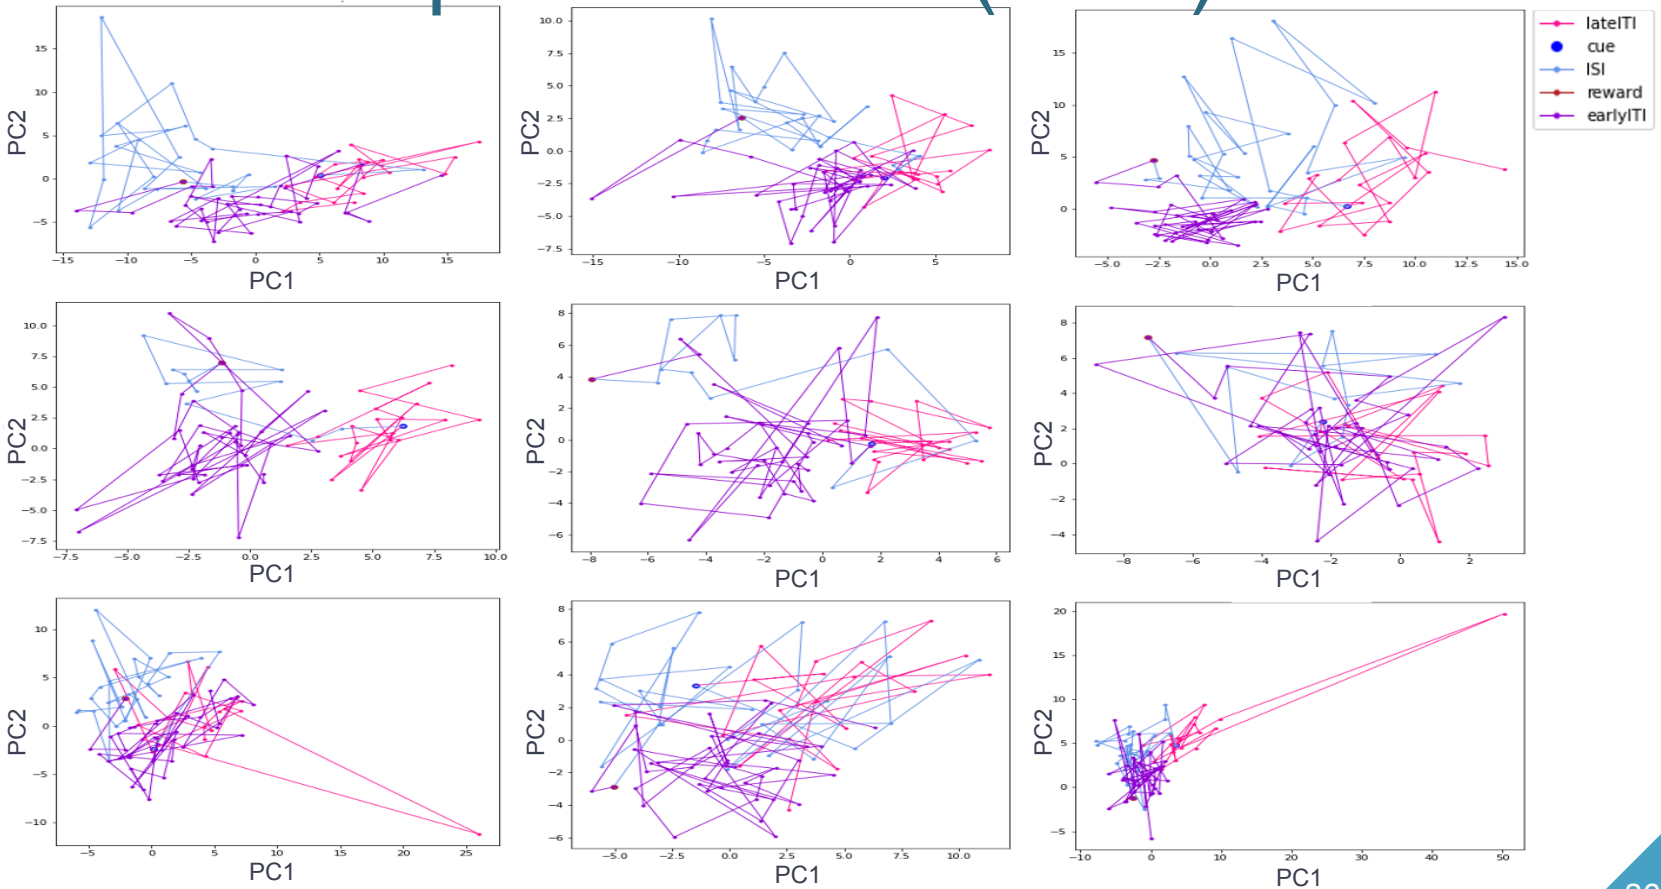

PC: principal component

Maheswaranathan et al., 2019, Adv Neural Inf Process Syst.

Recurrent switching linear dynamical systems (rSLDS)

parameters

discrete states

continuous states

data

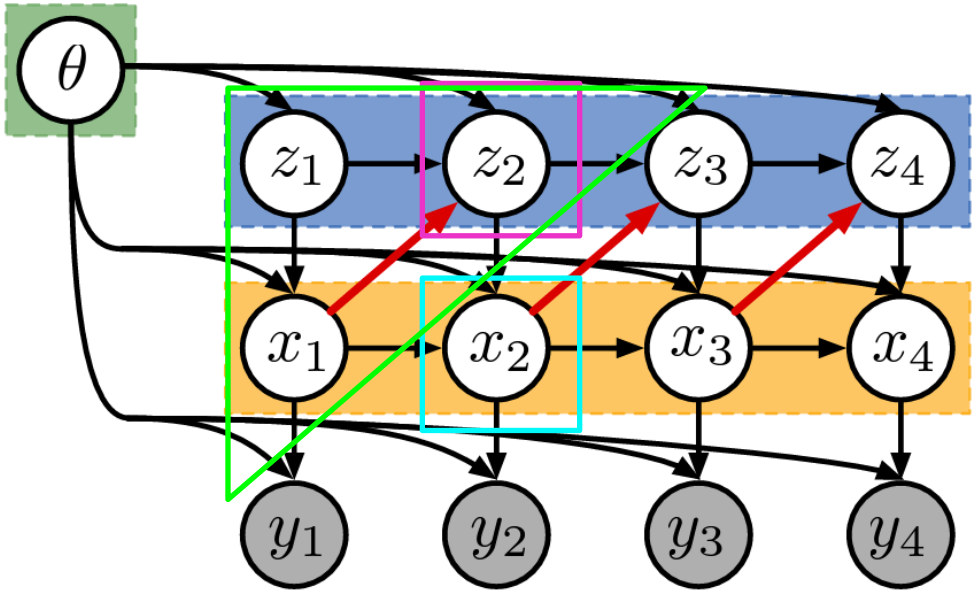

Linderman et al., 2016, preprint

Dynamics of the GRU network

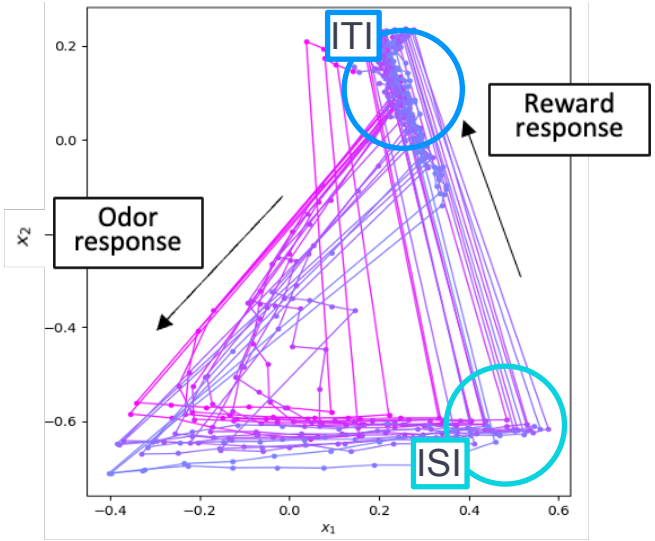

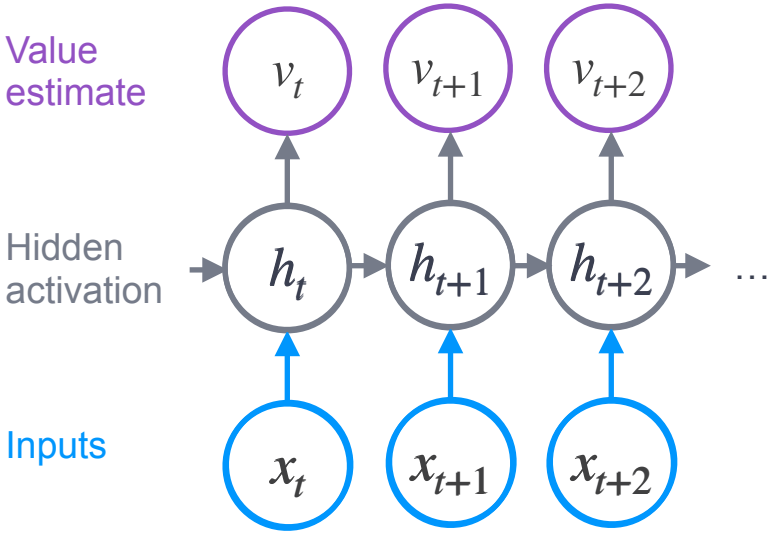

Dynamics of the GRU network

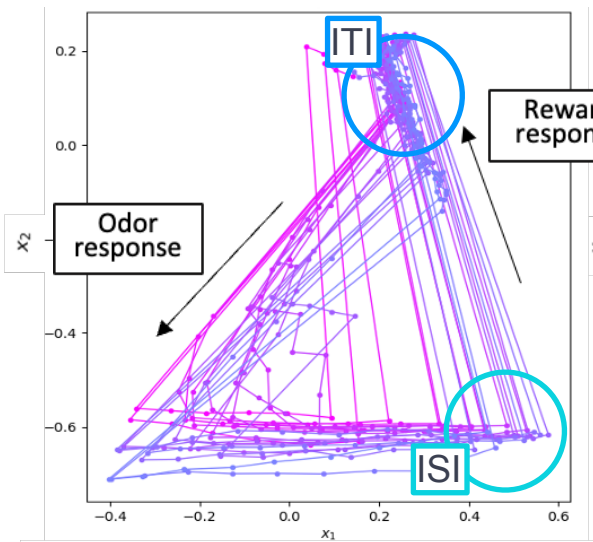

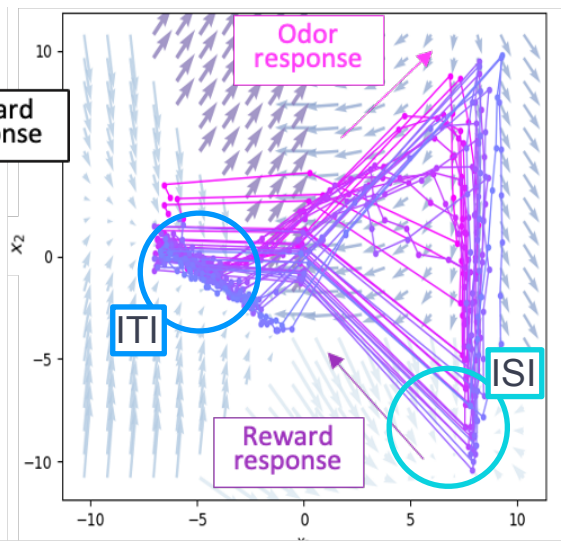

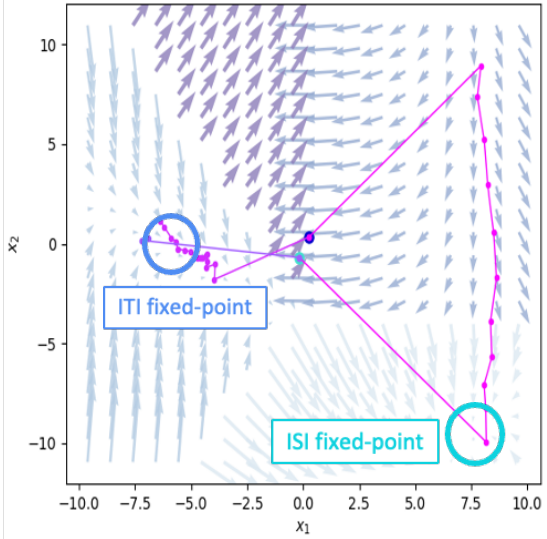

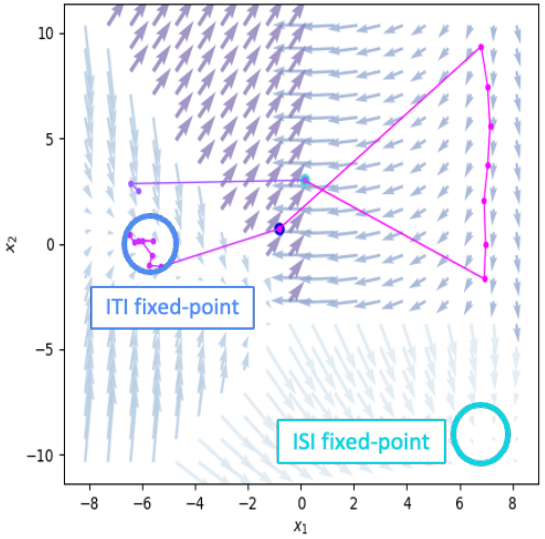

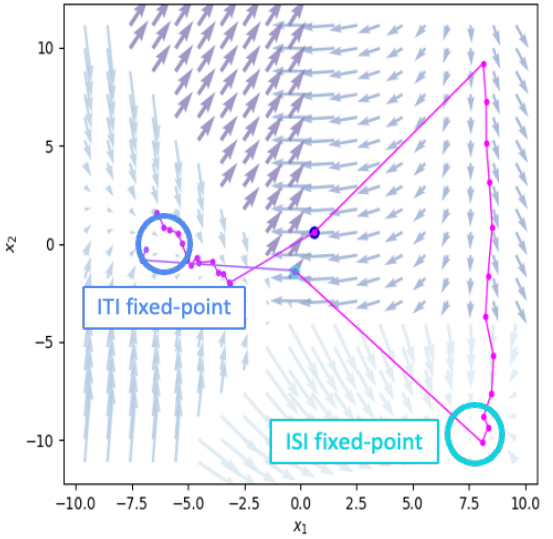

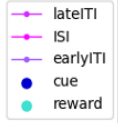

Dynamics of the neural population activity

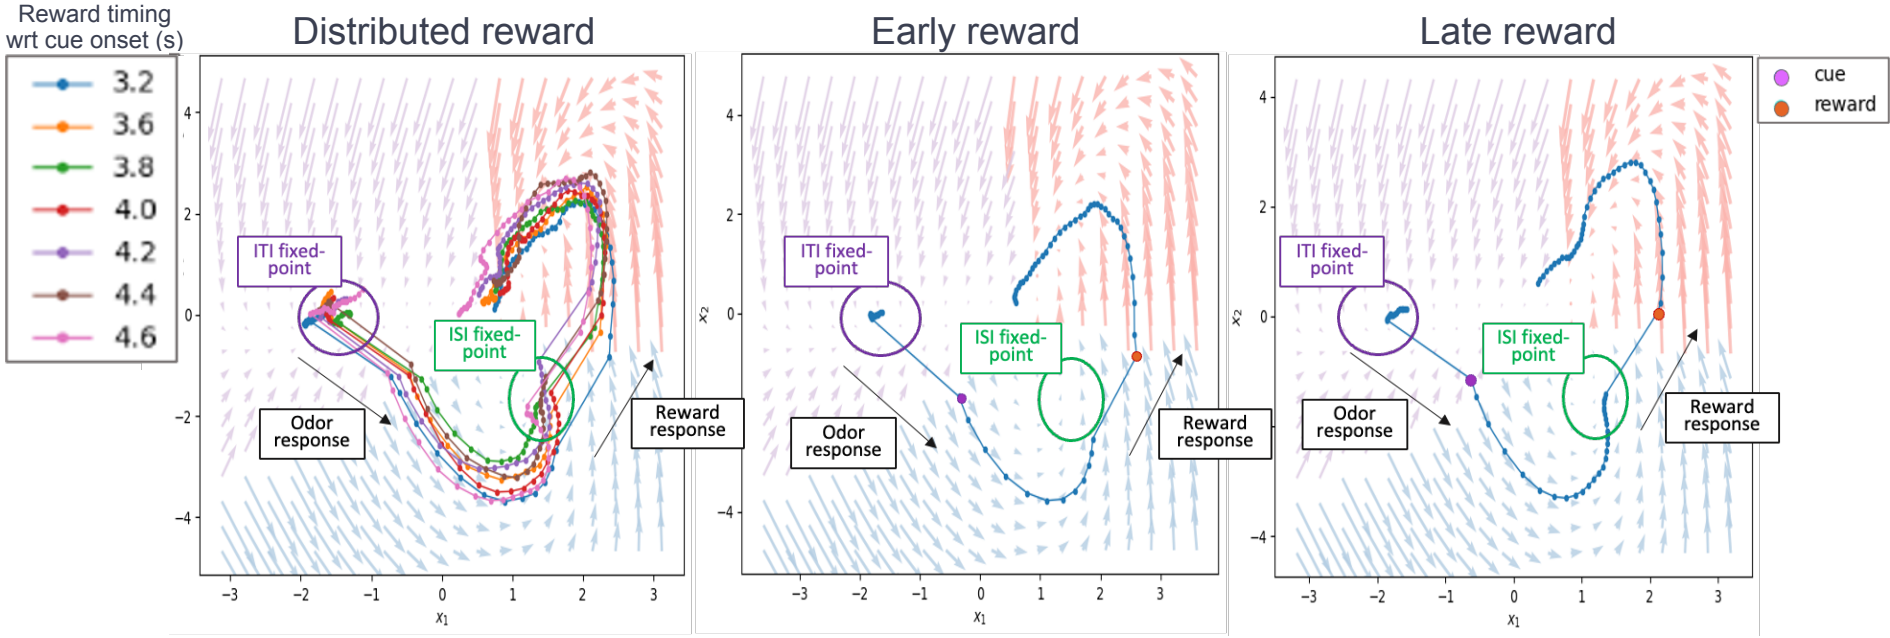

Dynamics of the neural population activity

Distributed reward

Early reward

Late reward

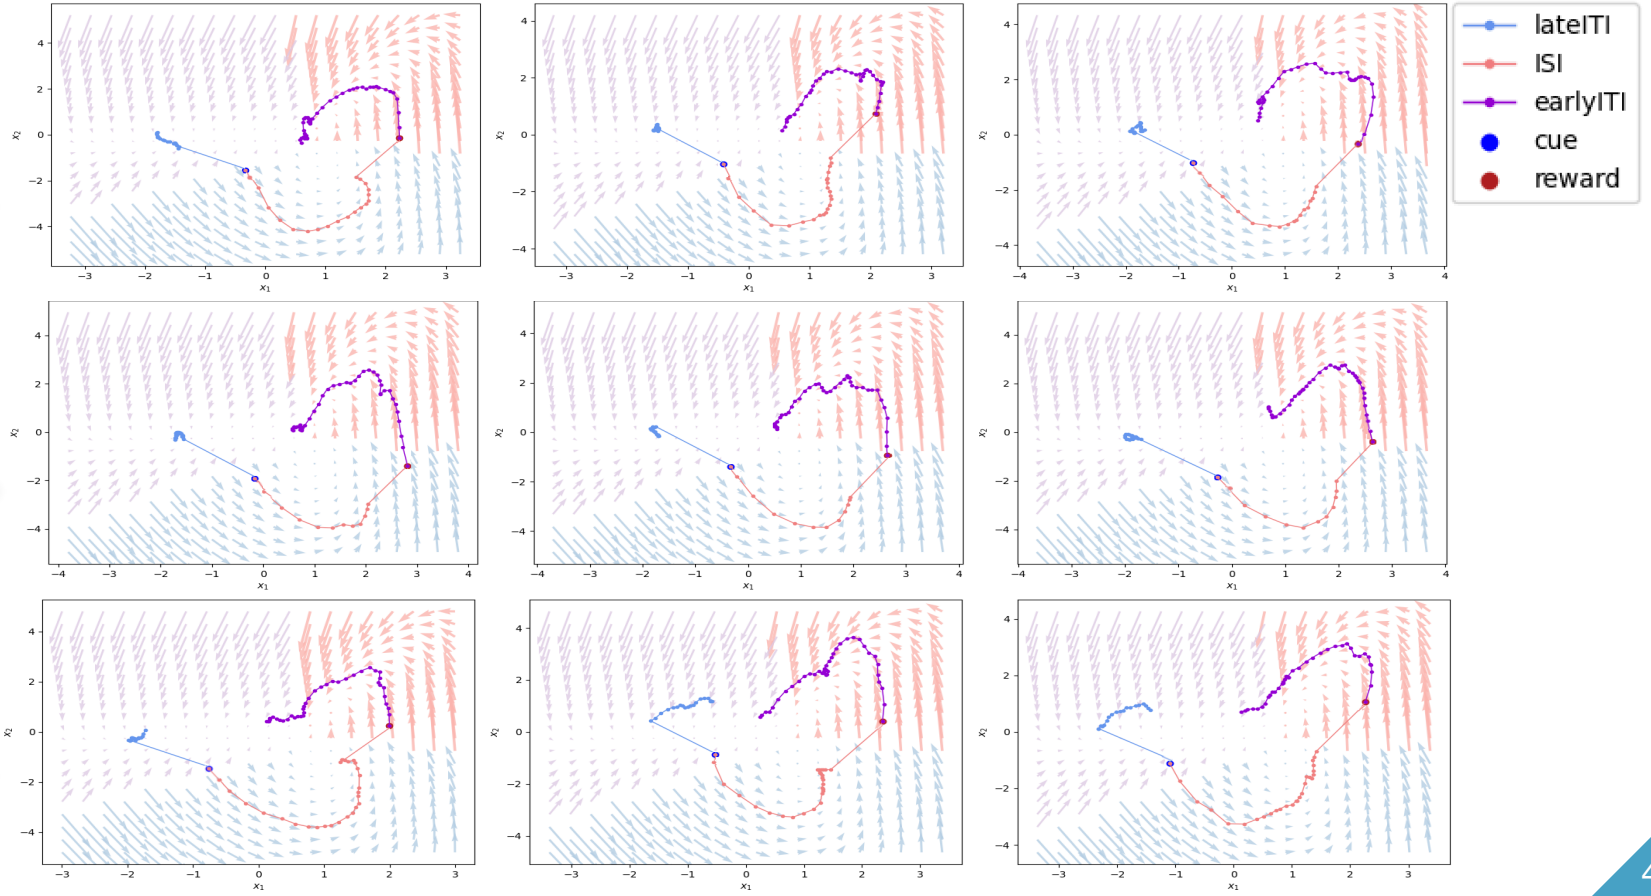

Dynamics of the neural population activity

Distributed reward

Early reward

Late reward

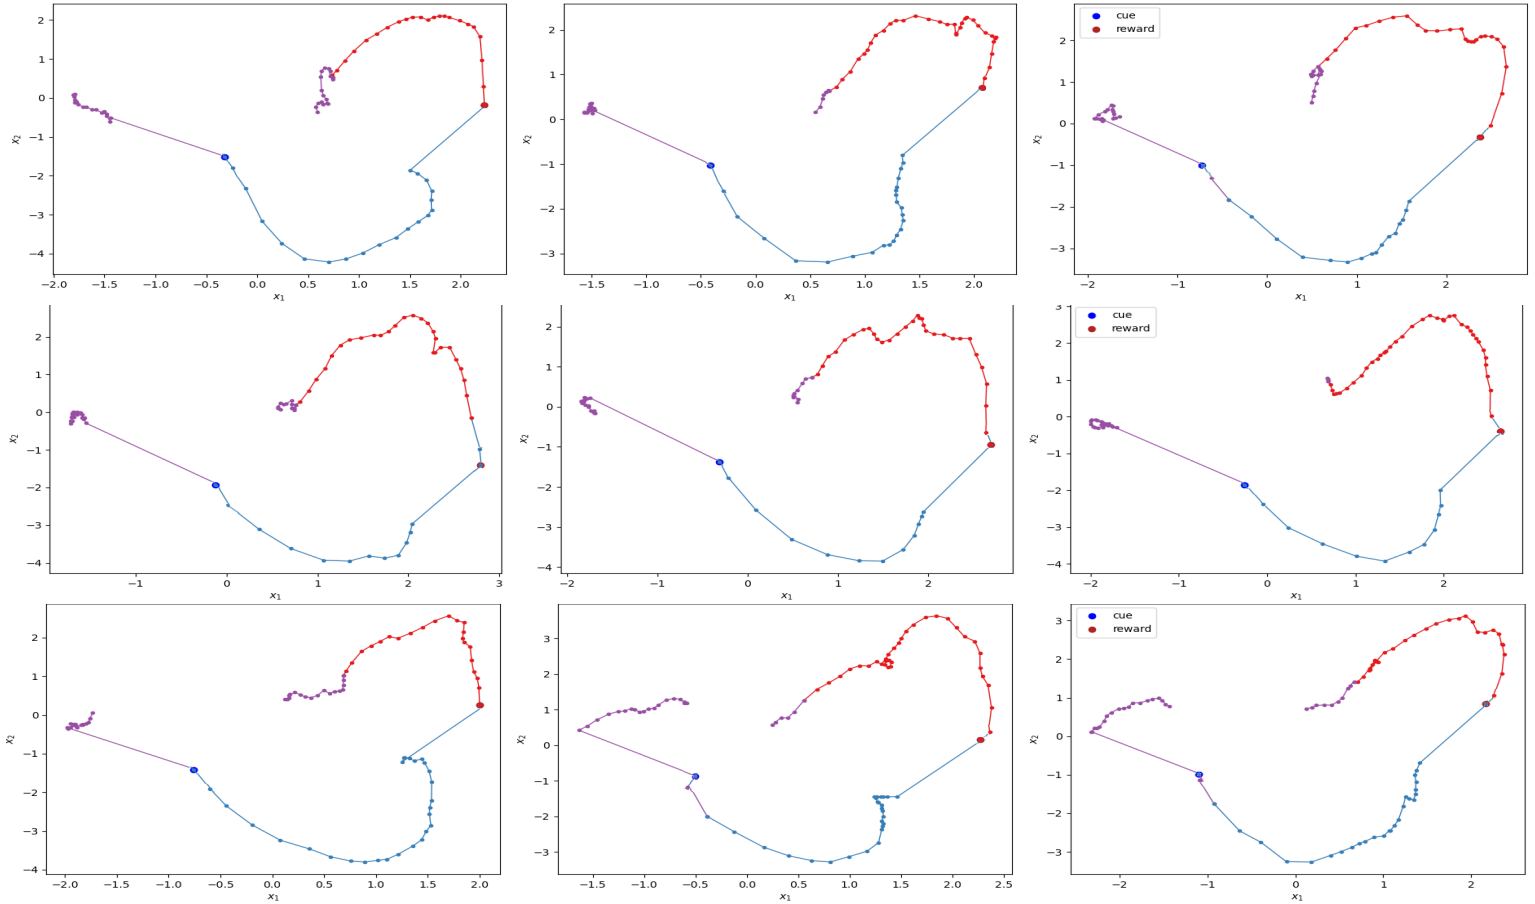

Dynamics of the neural population activity

Distributed reward

Early reward

Late reward

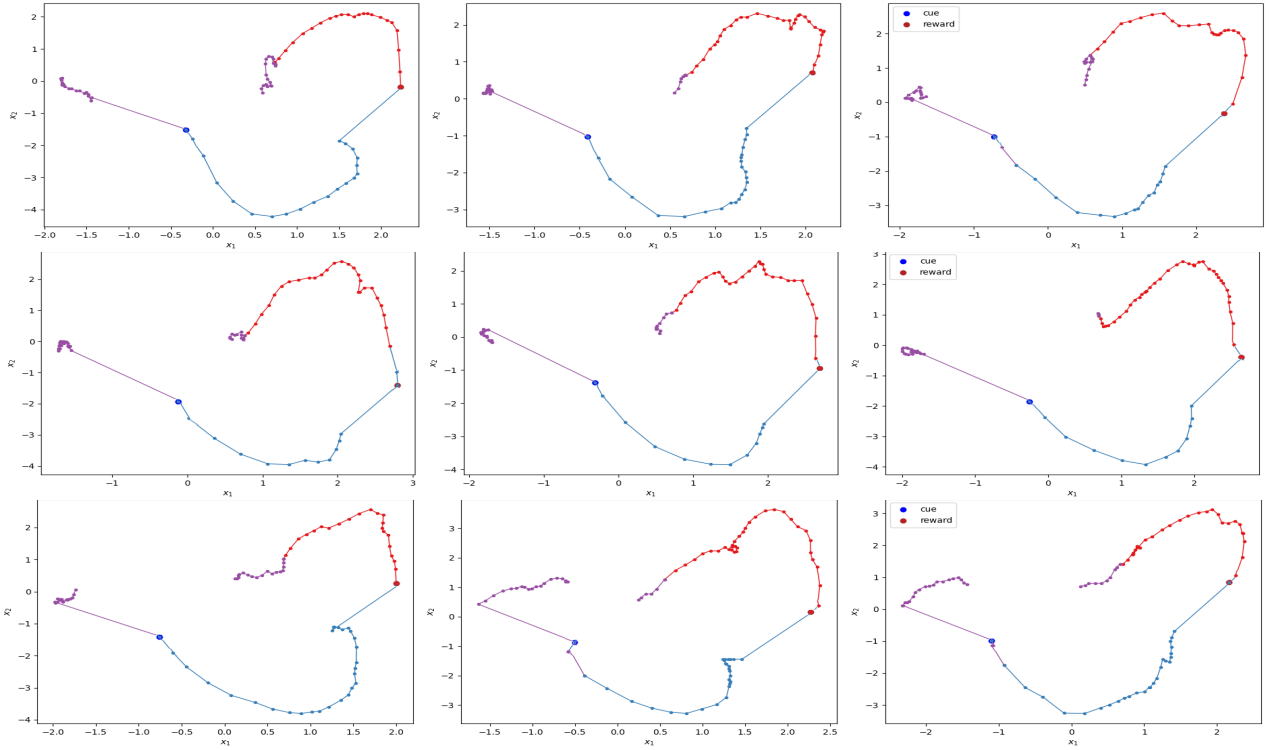

$$x_t = A_{z_t} \cdot x_{t-1} + V_{z_t} \cdot u_t + b_{z_t}$$

Is belief state representation computed in the OFC, and if so, how?

Individual neurons investigation

Single-neuron activity shows specific task-variables tuning in the OFC

Neural activity encoder

OFC encode variables that are essential to belief state representation

State decoder

OFC population activity is sufficient to predict states

Evolution of the neural activity at the population level

Population dynamics in the OFC show two fixed-points corresponding to states

Discussion

Belief state representation in the OFC

Deterministic task

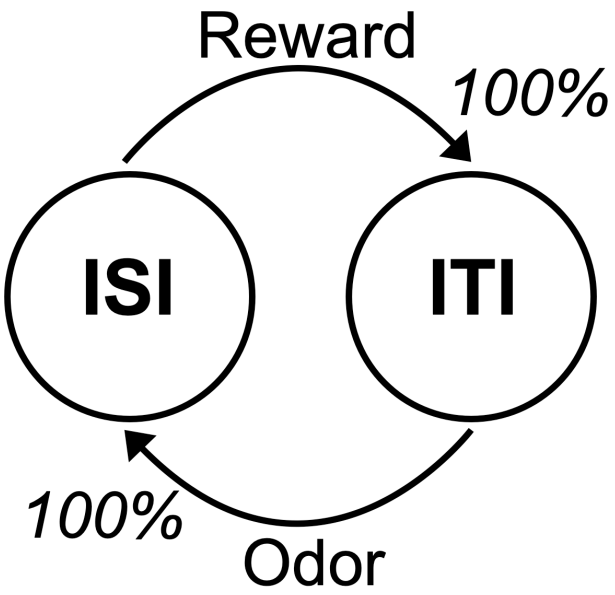

Stochastic task *(90% rewarded)*

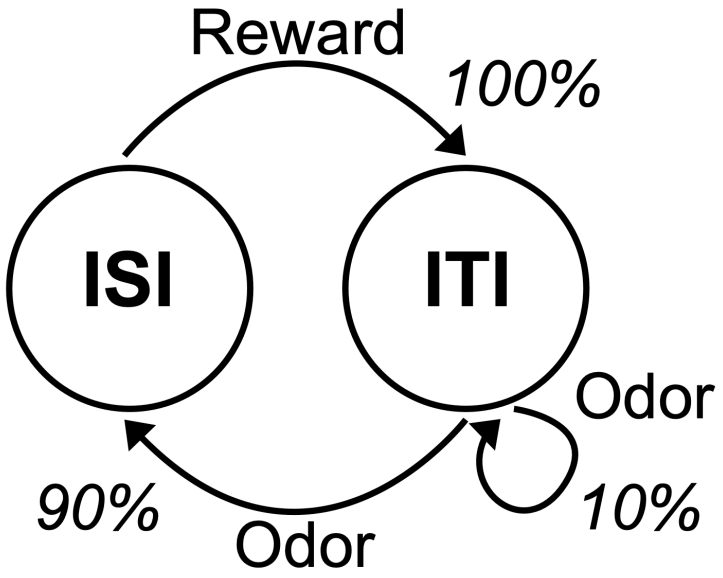

Discussion

Belief state representation in the OFC

Background-tuned

Sustained delay-tuned

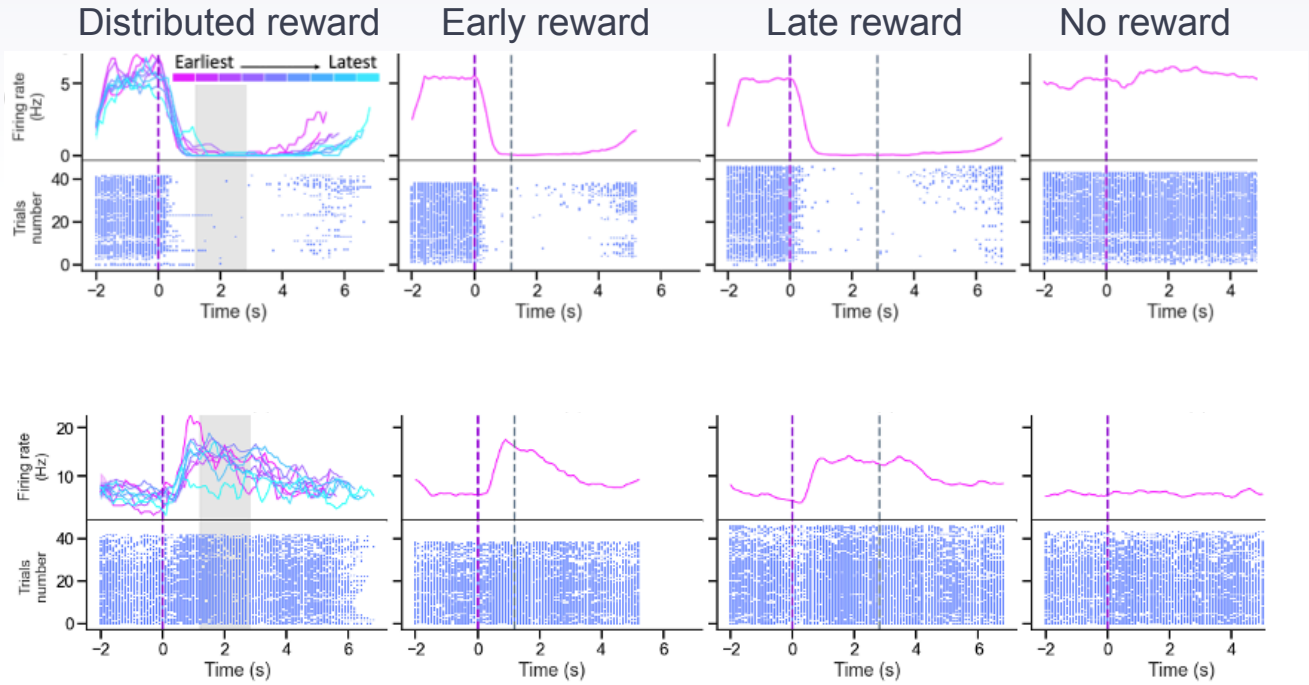

Discussion

Further investigation on the stochastic task

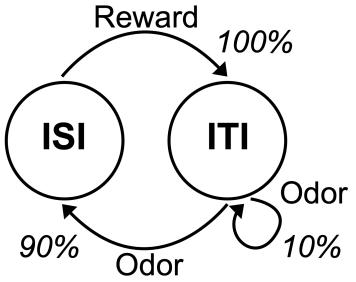

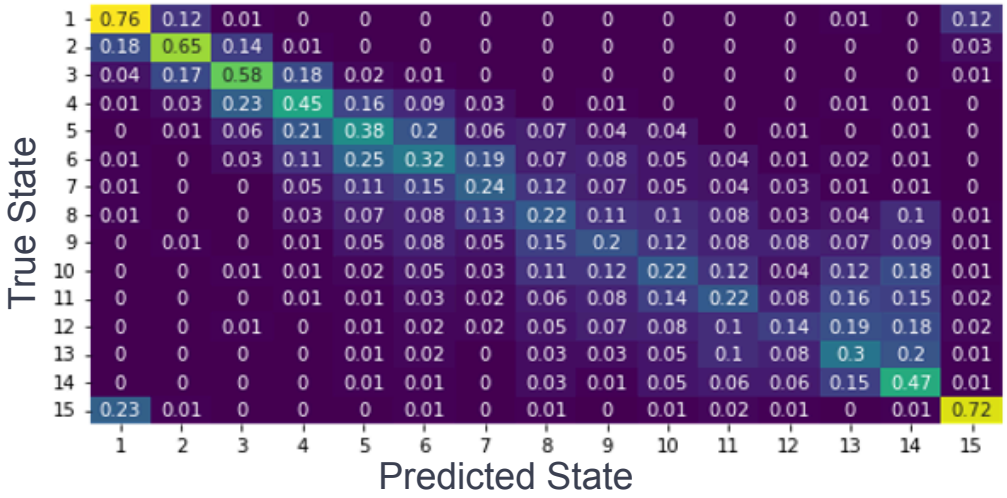

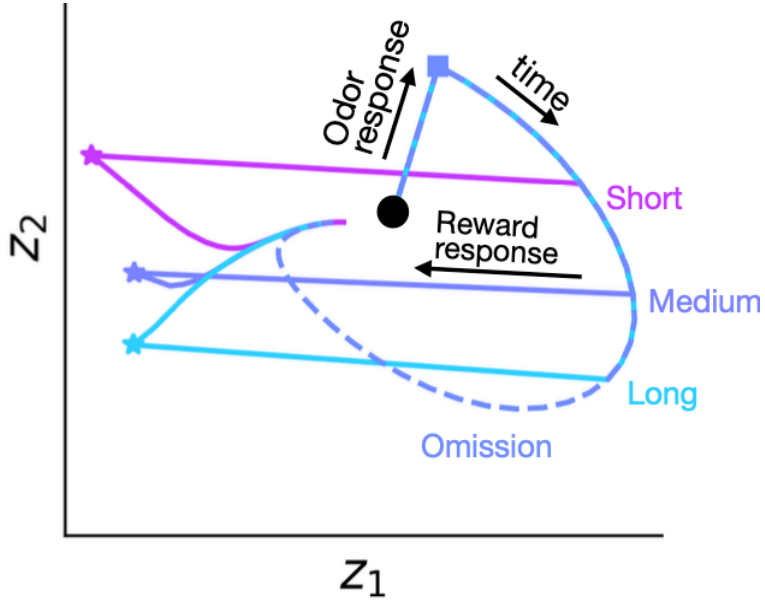

Thank you Uchida lab!

Supervision
Sandra Romero Pinto
Jay Hennig
Naoshige Uchida
Mackenzie Mathis

But also
Mitsuko Uchida
Isobel Green
Adam Lowet
Lyle Kingsbury
Malcom Campbell
Nacho Sanguinetti
Mengzi Yun
Paul Masset
Sara Pinto Dos Santos
Shudi Xu
Iku Kimura
Ryu Amo
For welcoming me to the lab!

Fundings
EPFL • WISH FOUNDATION
WOMEN IN SCIENCE AND HUMANITIES

Appendix

Supplementary Figures

Discussion

Limits of the rSLDS & further experiments

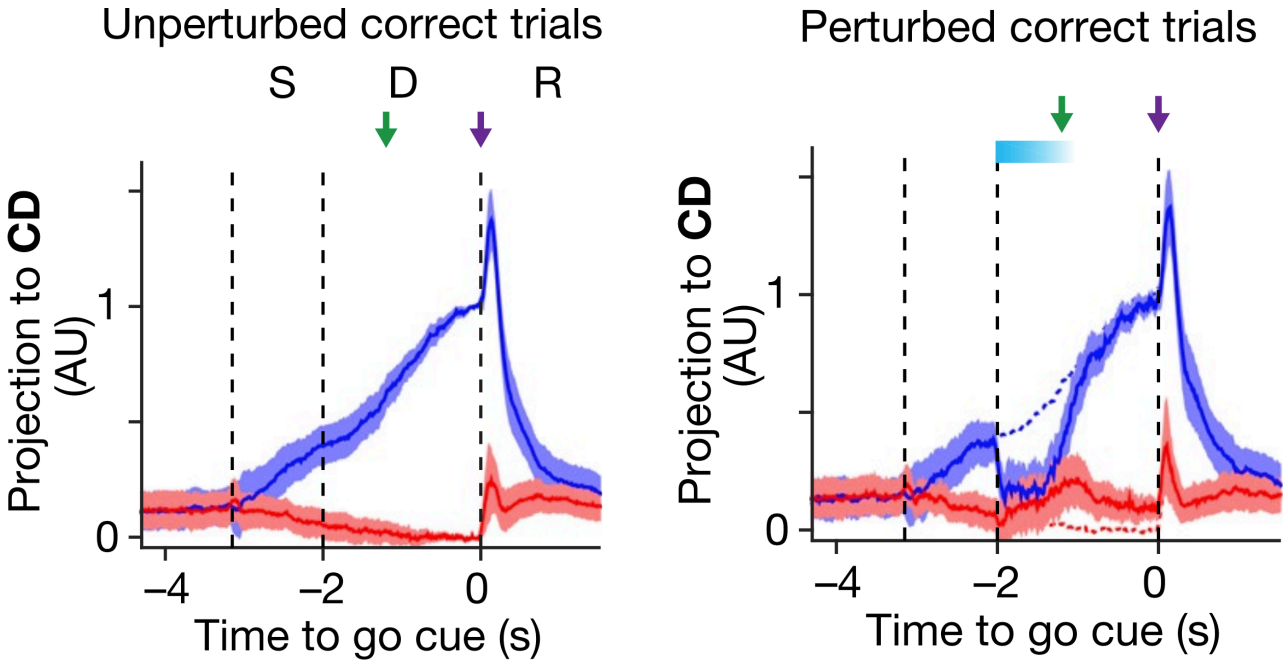

For a "discrete" attractor

- **Perturbation small enough**
  → goes back to the fixed-point
- **Perturbation bigger**
  → goes to the other fixed-point

Inagaki et al., 2019, Nature

Discussion

New data in OFC

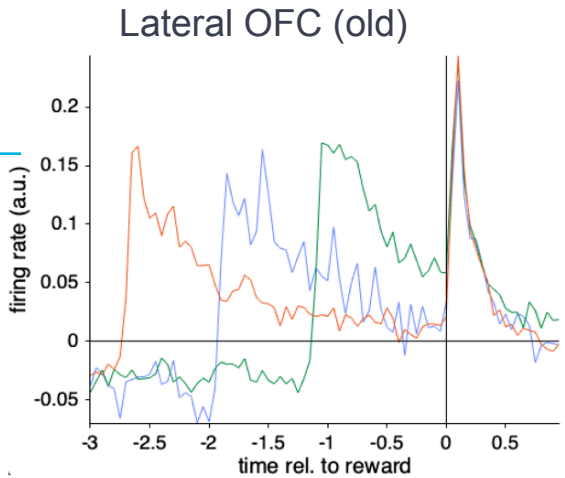

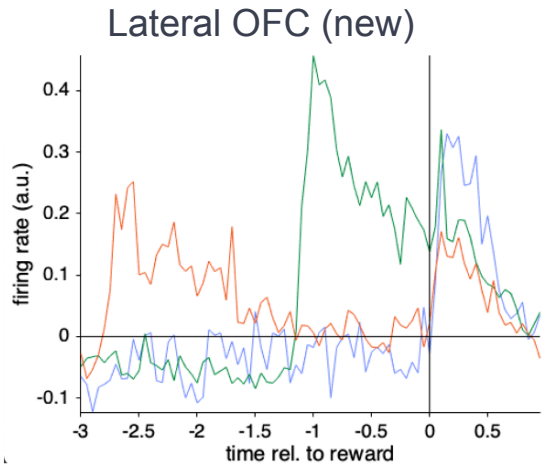

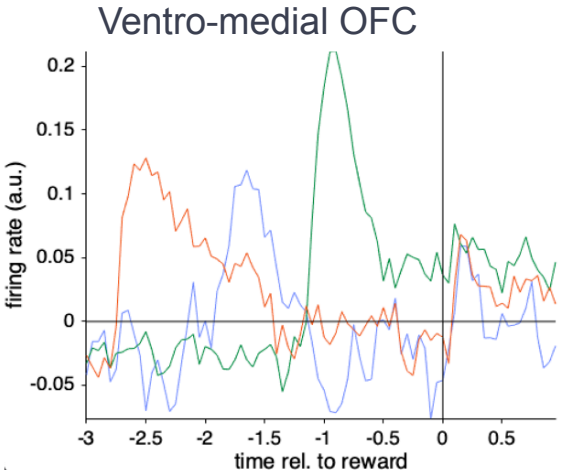

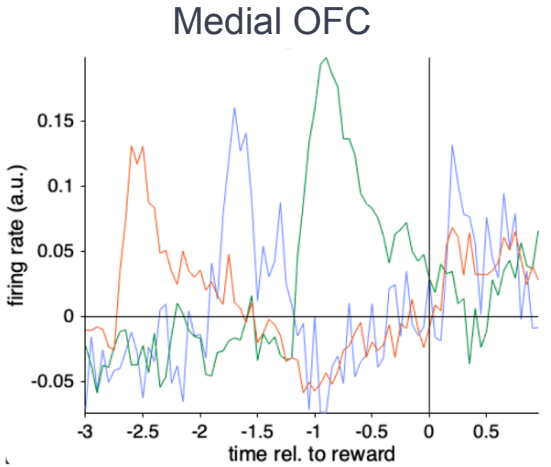

Trial-averaged licking rate

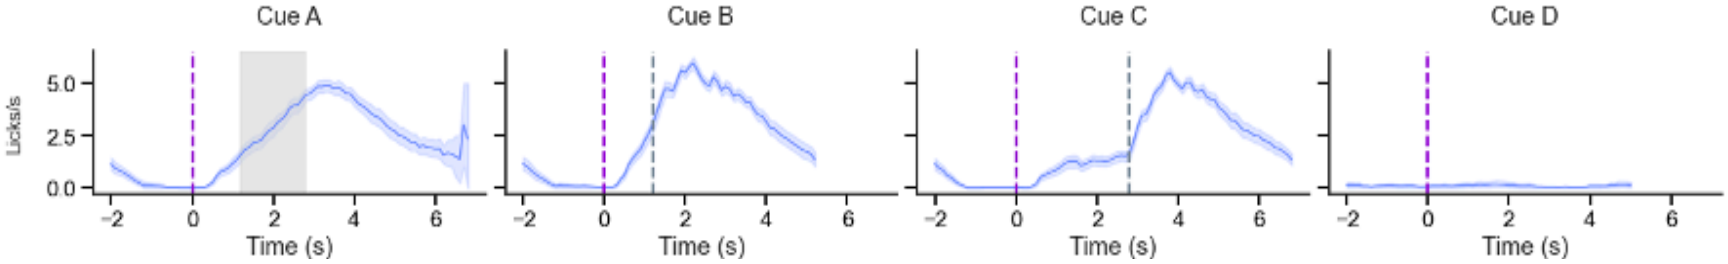

Example neuron in piriform cortex

Odor-tuned

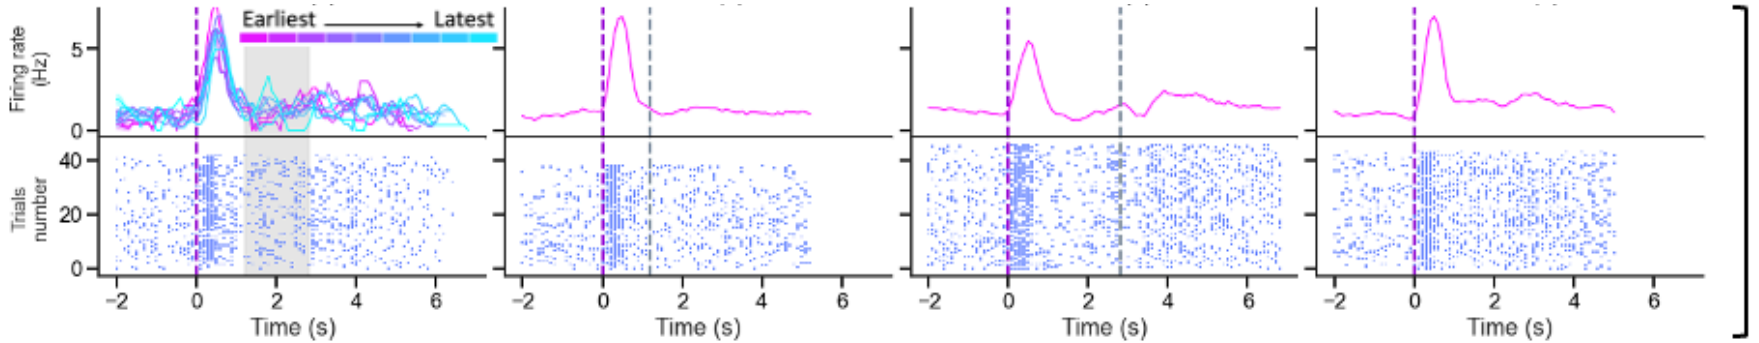

Example neuron in M2

Reward-tuned

Sustained delay-tuned

Background-tuned

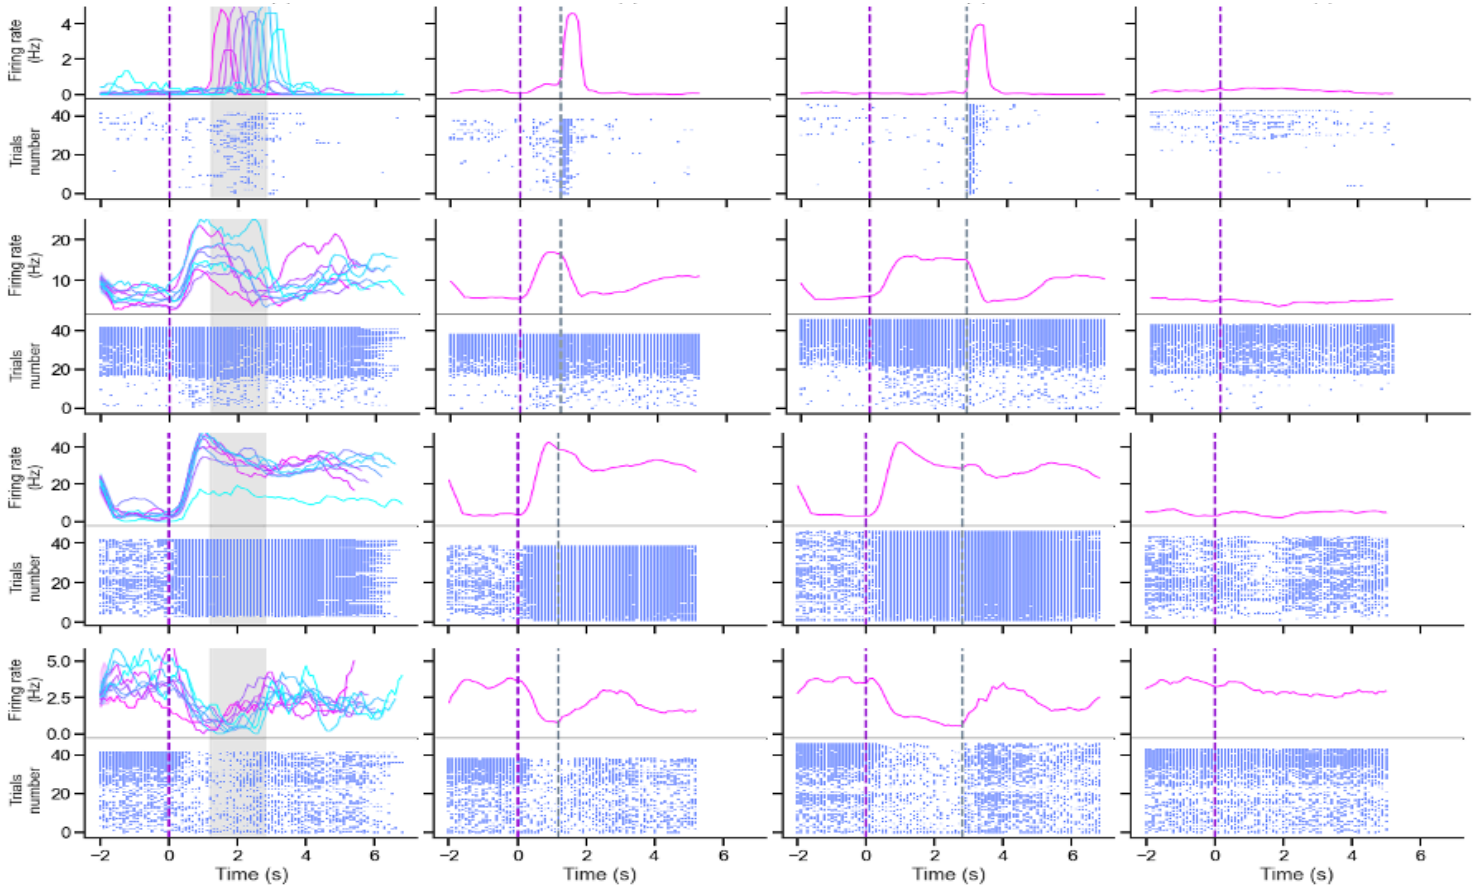

Fractions of neurons encoding each variable in M2

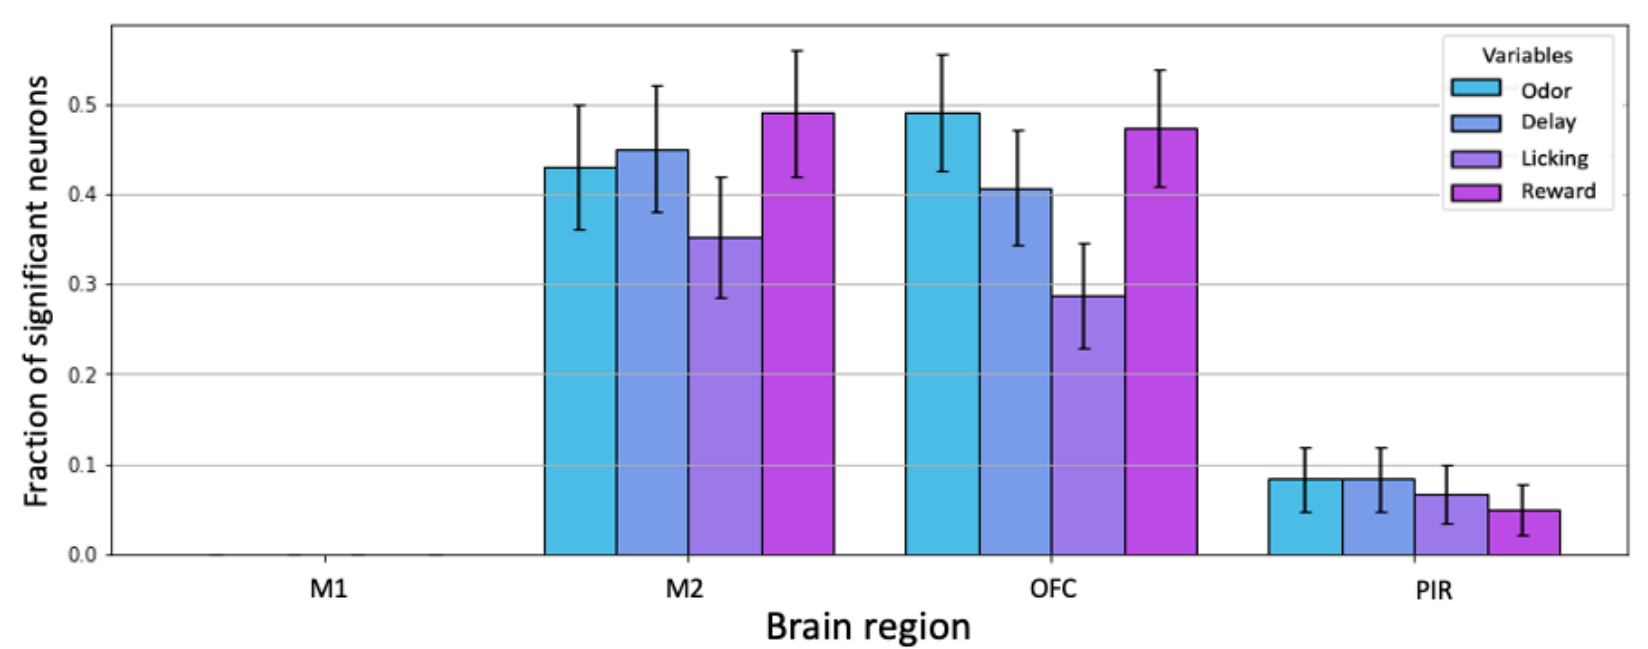

Activity map in other regions

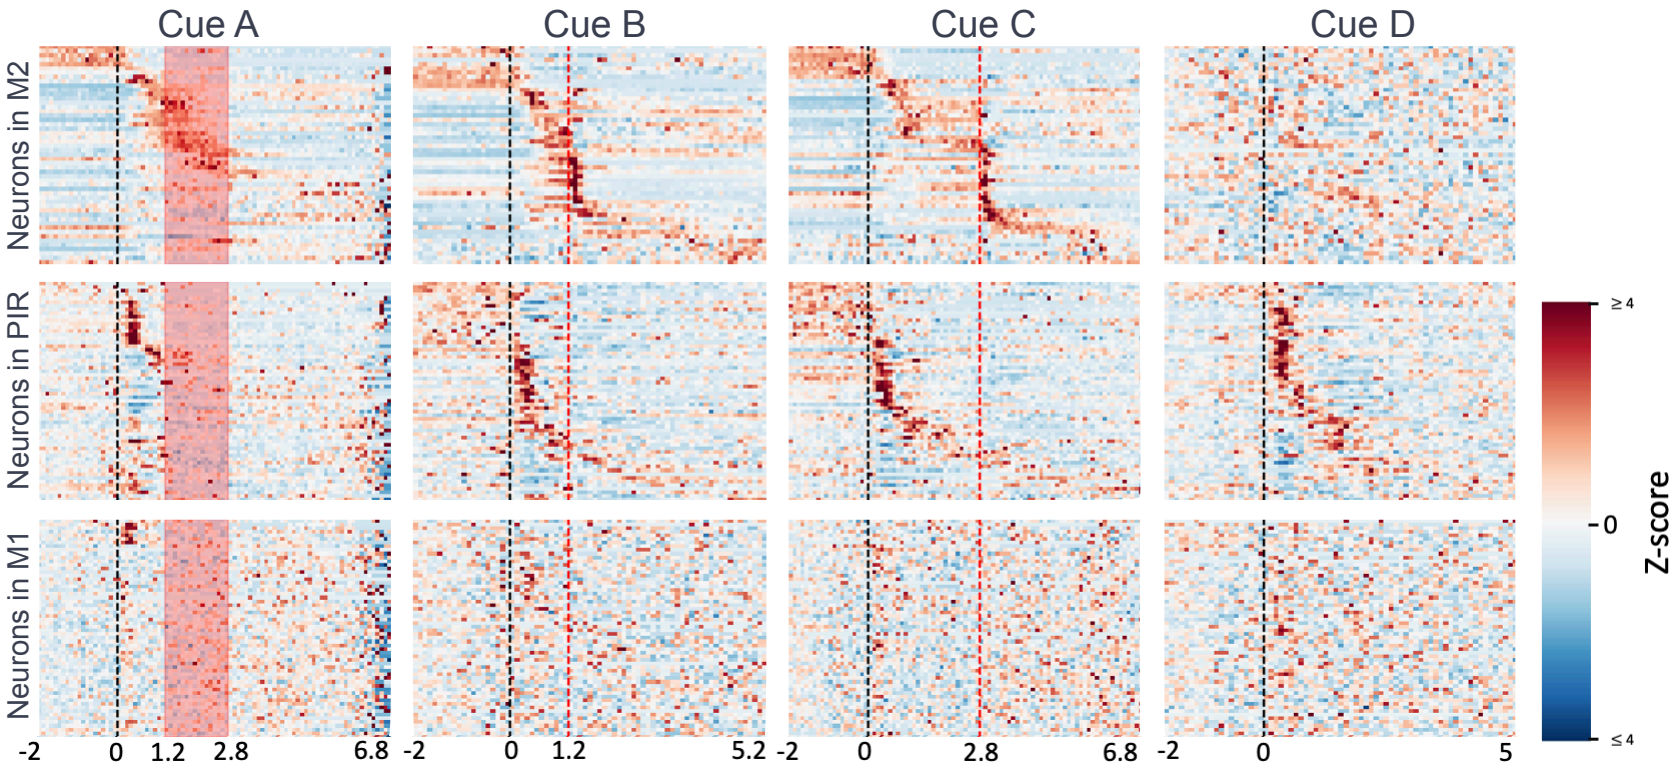

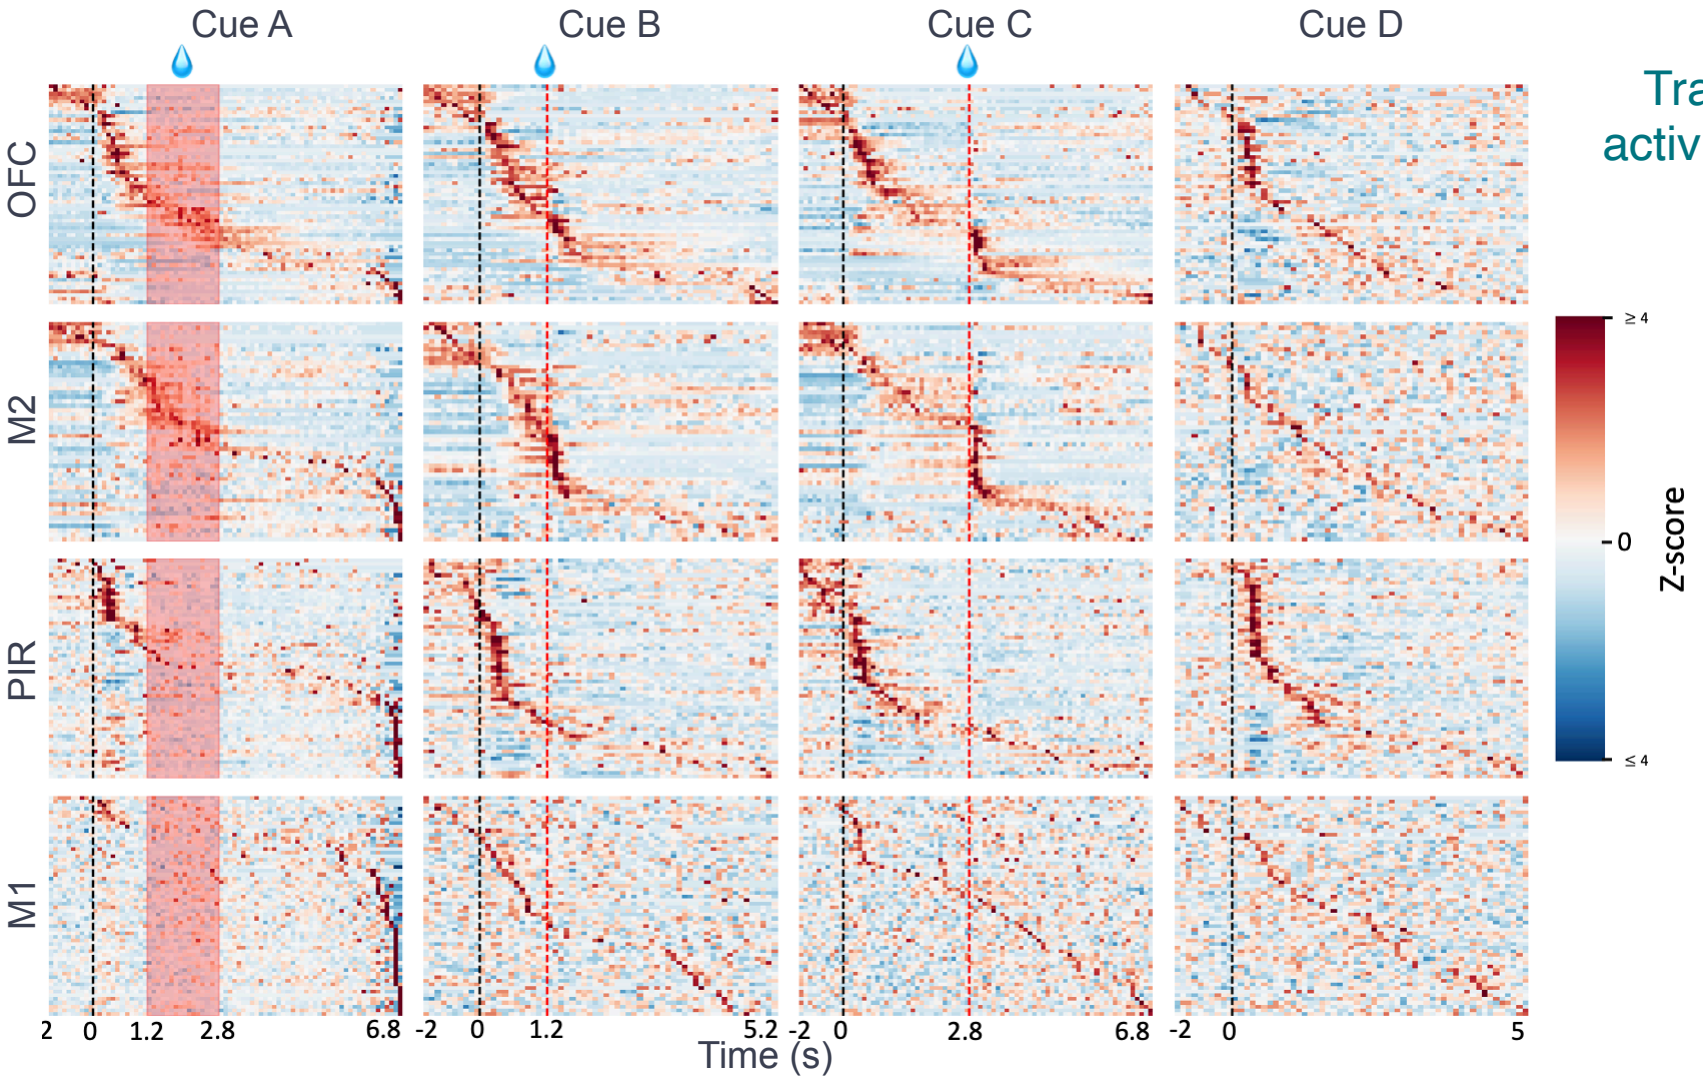

Training activity map

Activity map on a different animal

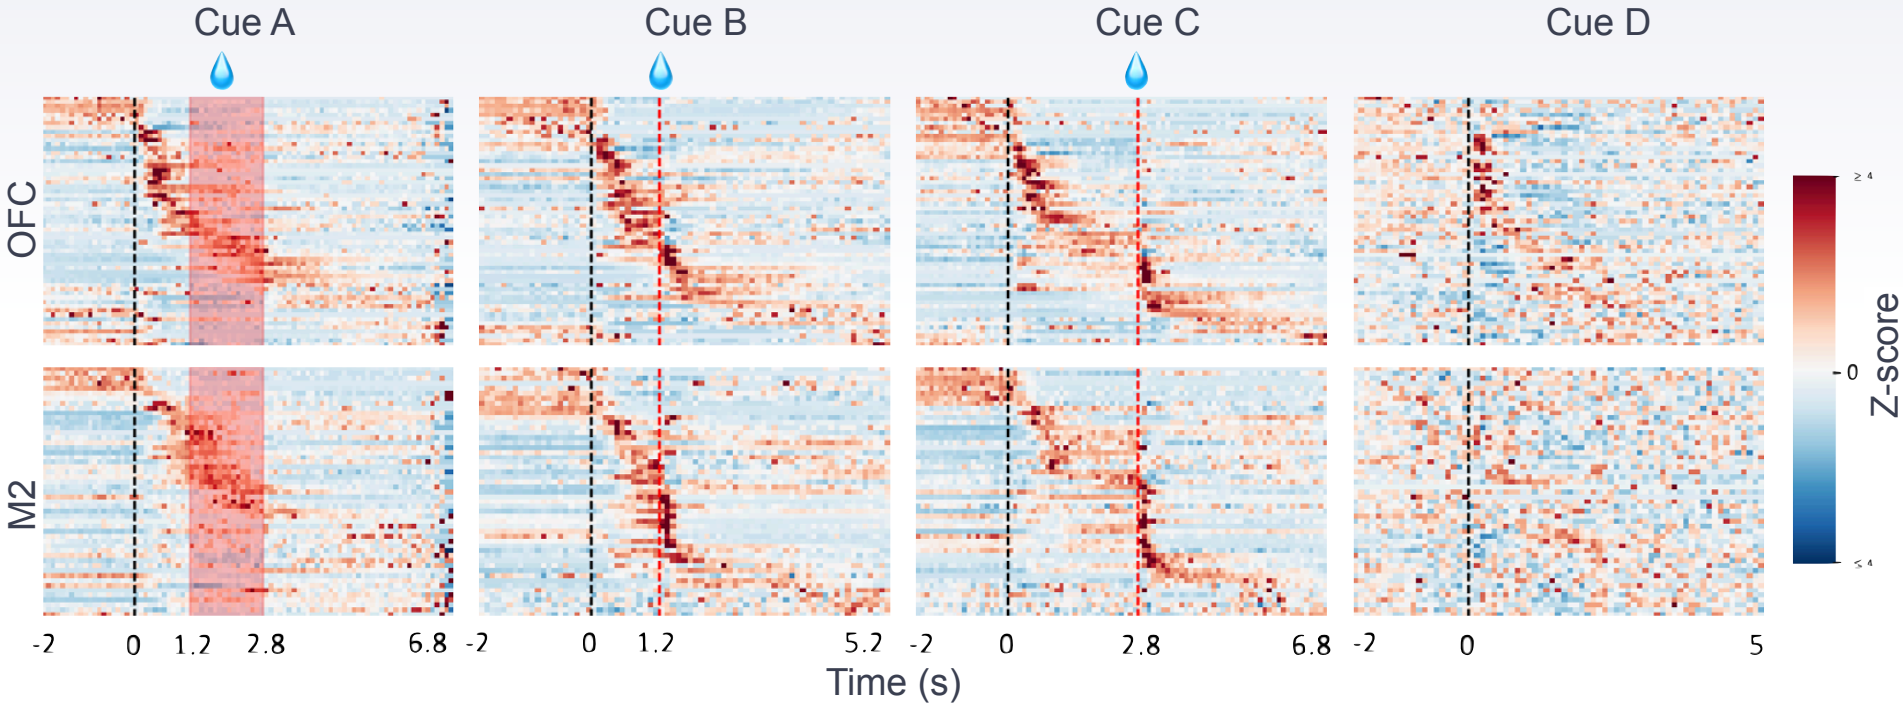

Performances on full/reduced matrices with Poisson GLM model

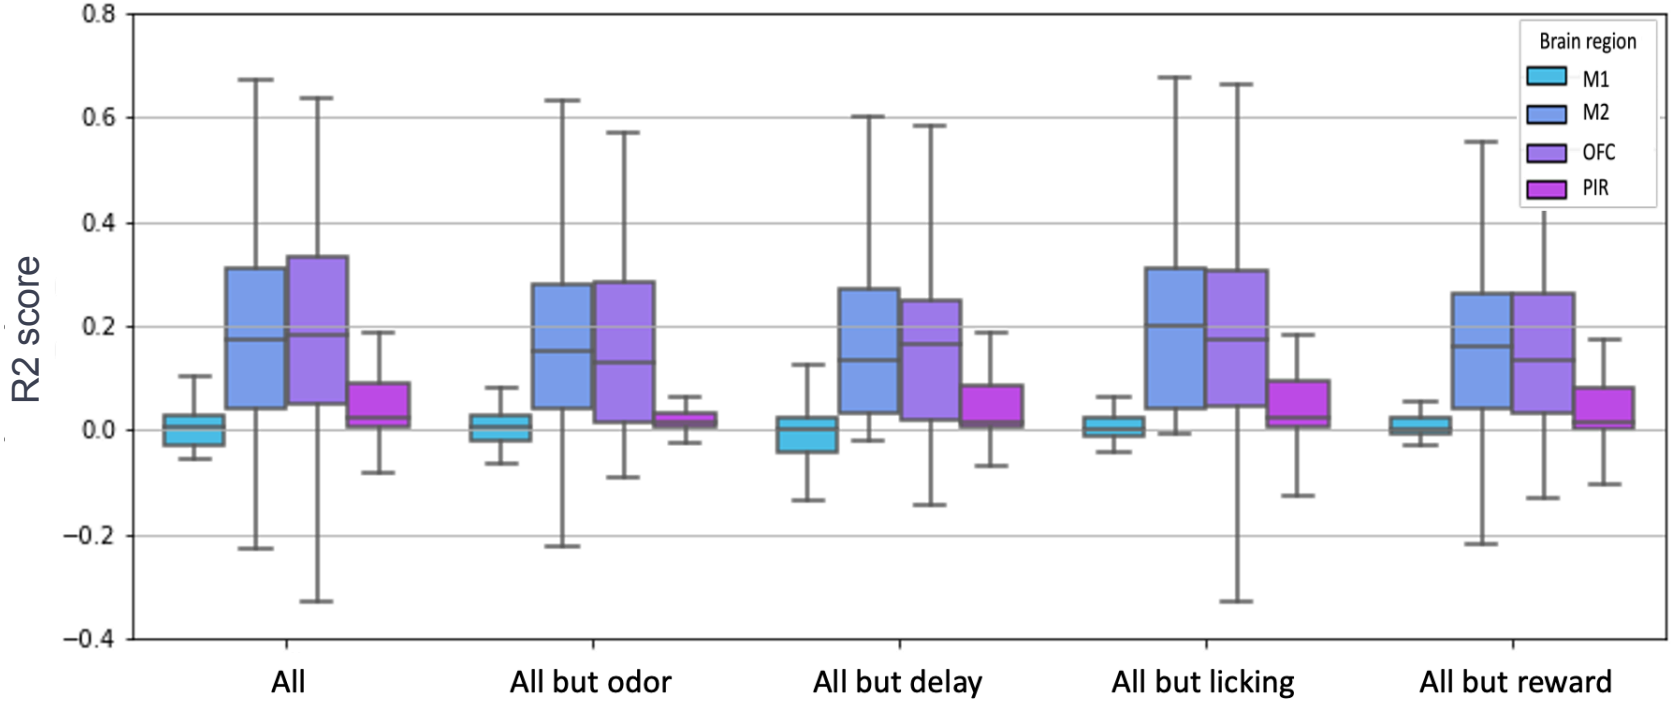

Comparison linear regressor vs Poisson GLM

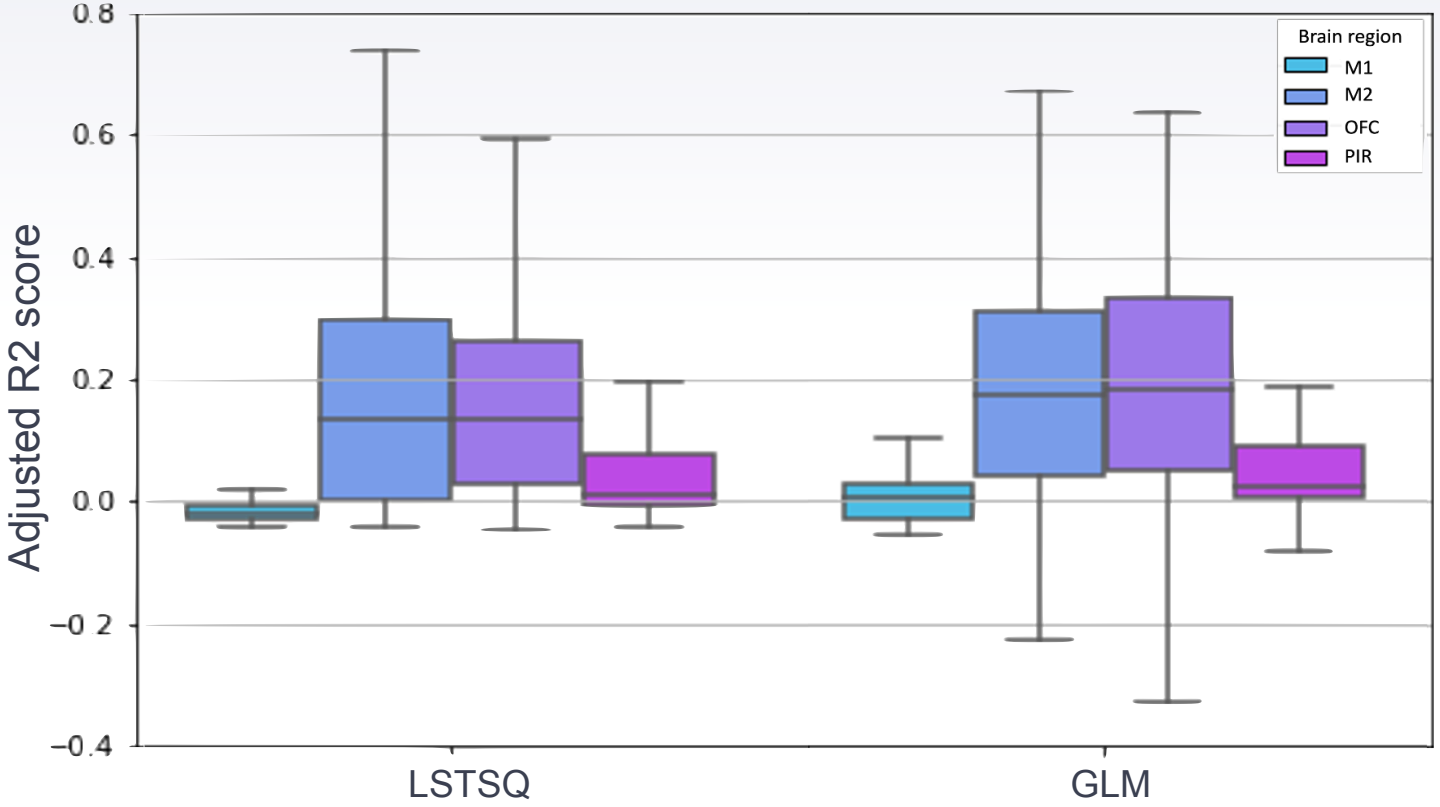

Performances on full/reduced matrices with linear regression model

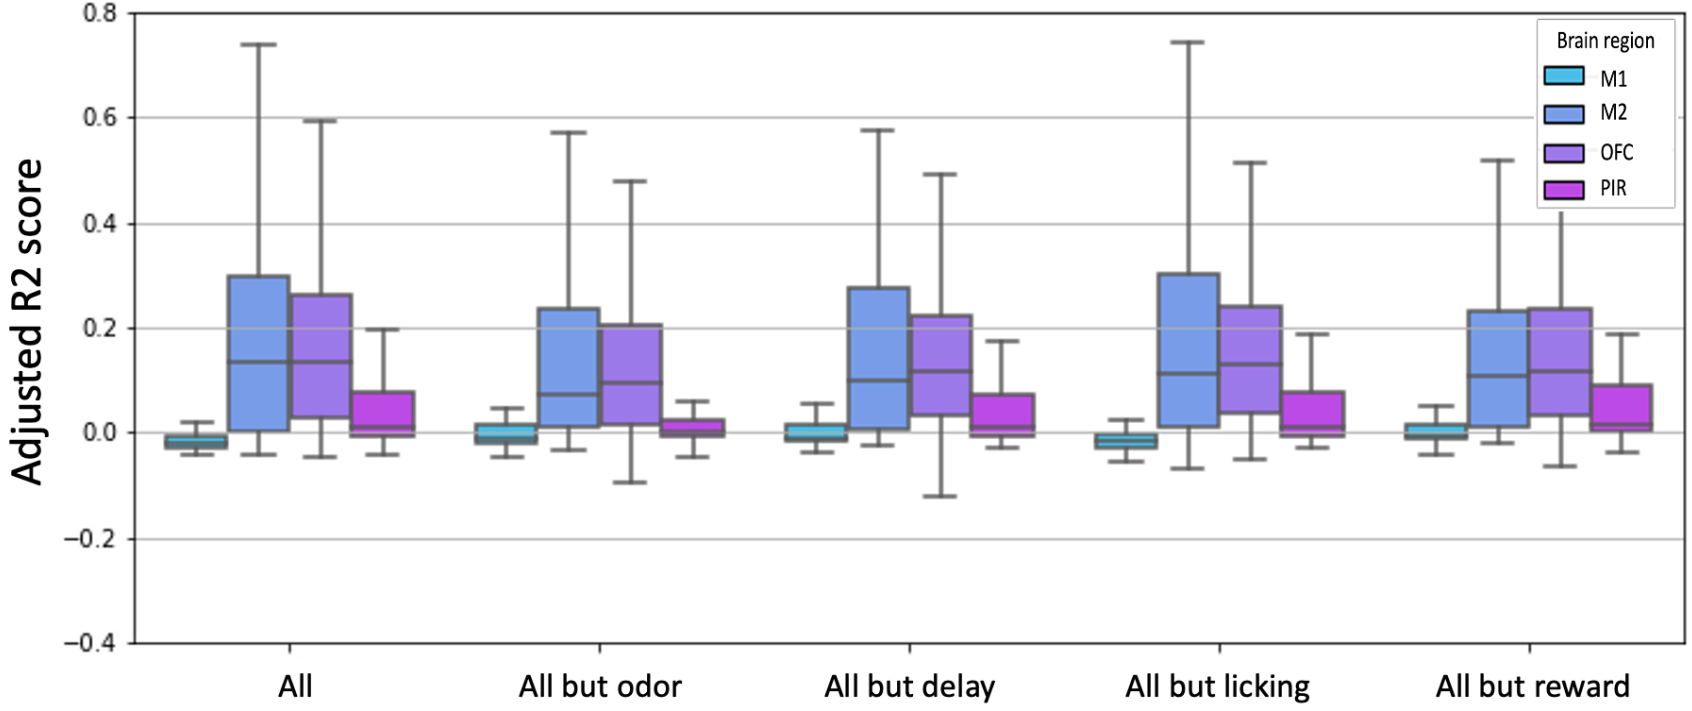

State classifier on M1 neural activity

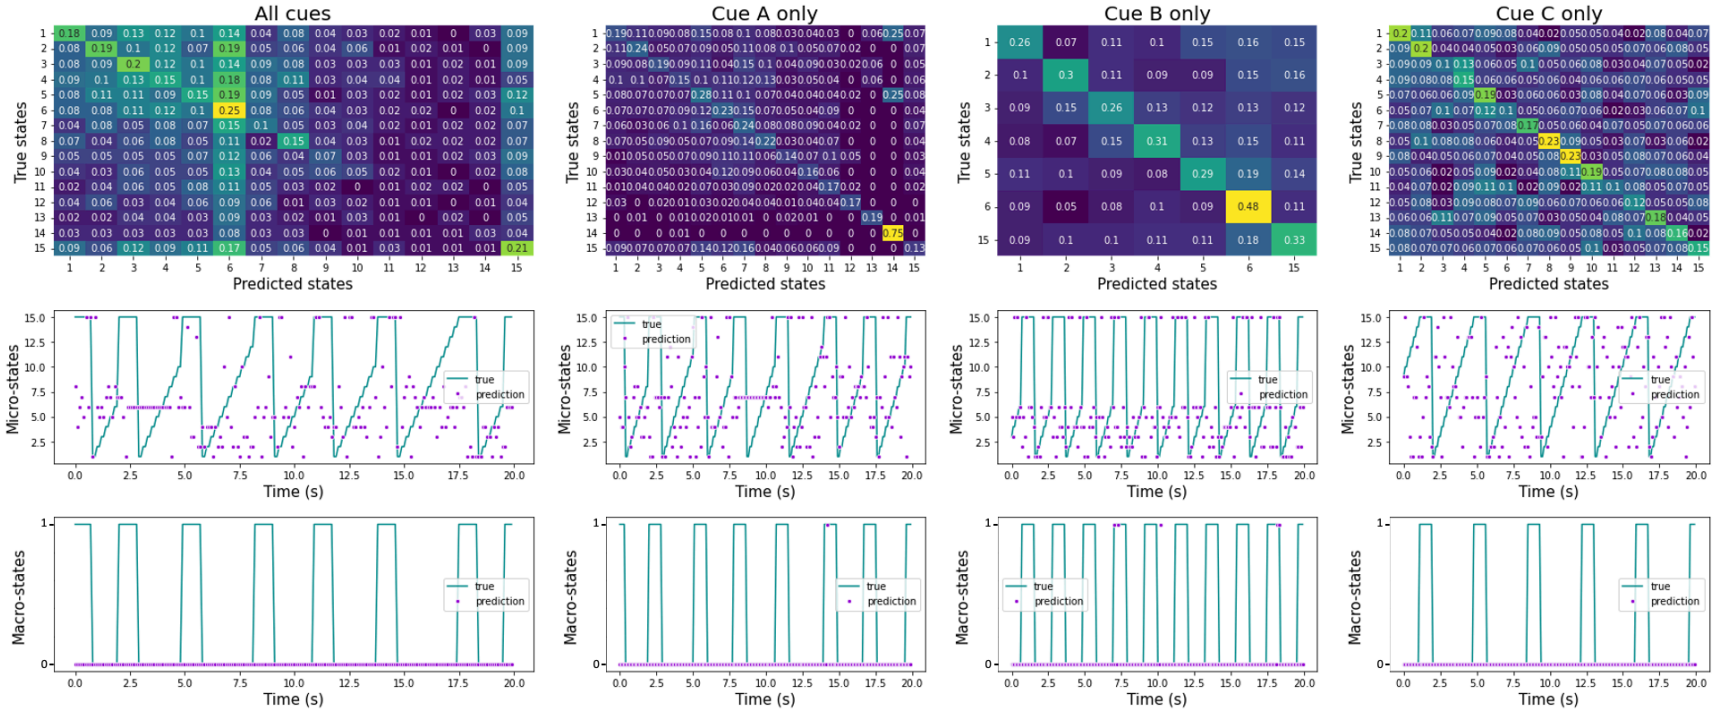

State classifier on PIR neural activity

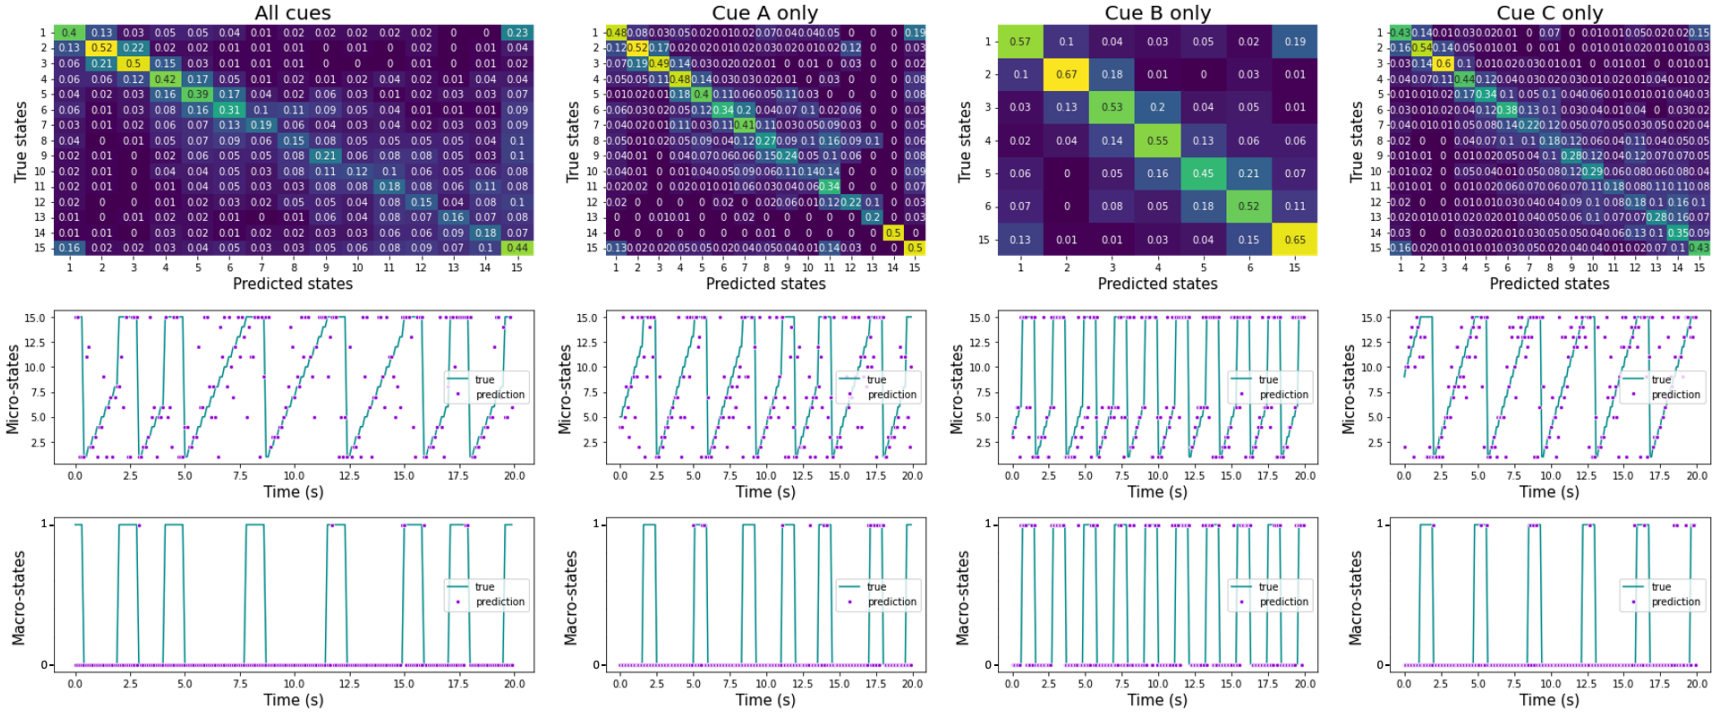

State classifier on M2 neural activity

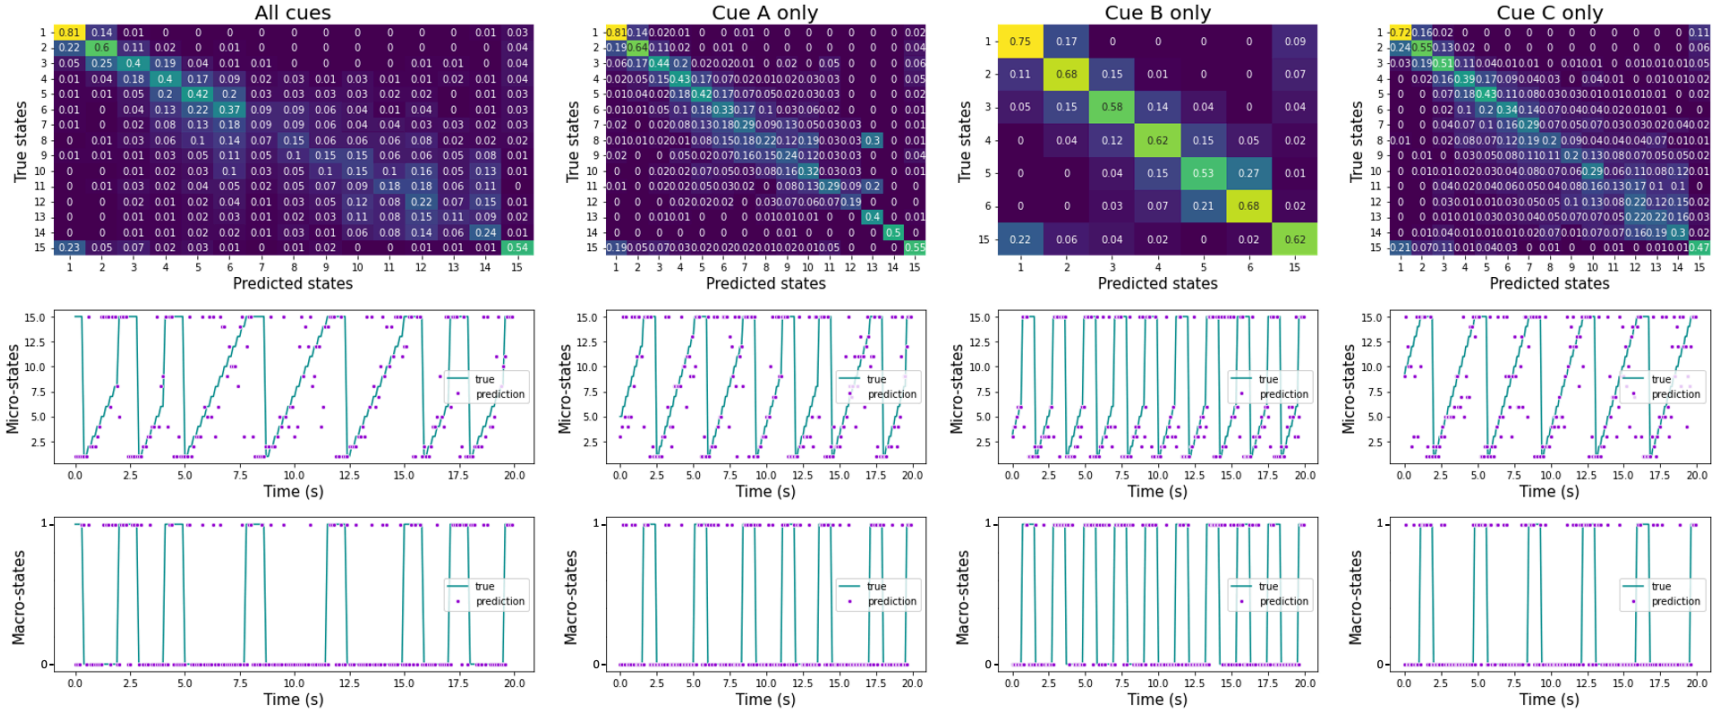

Global drift in the population activity

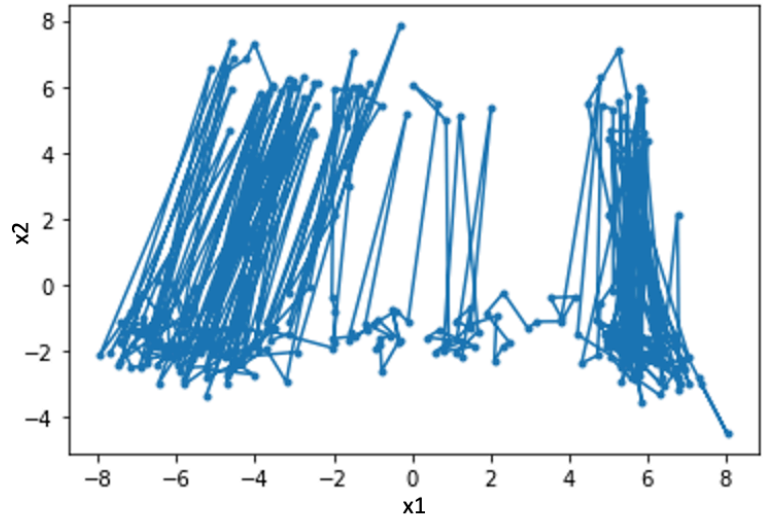

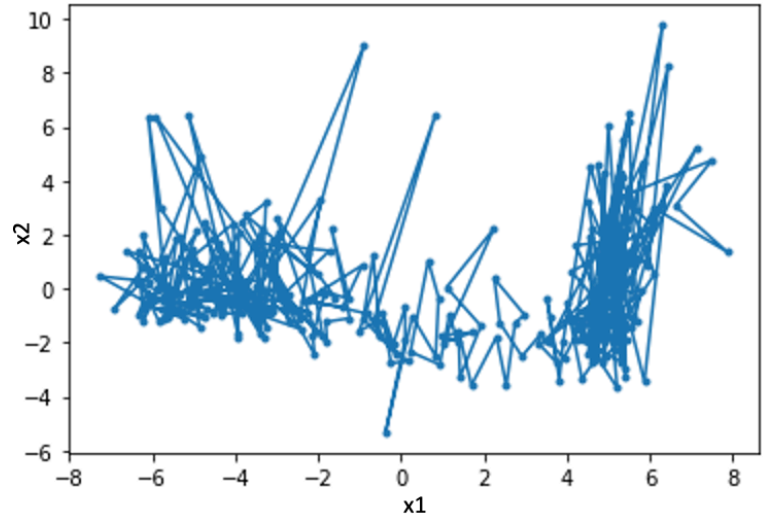

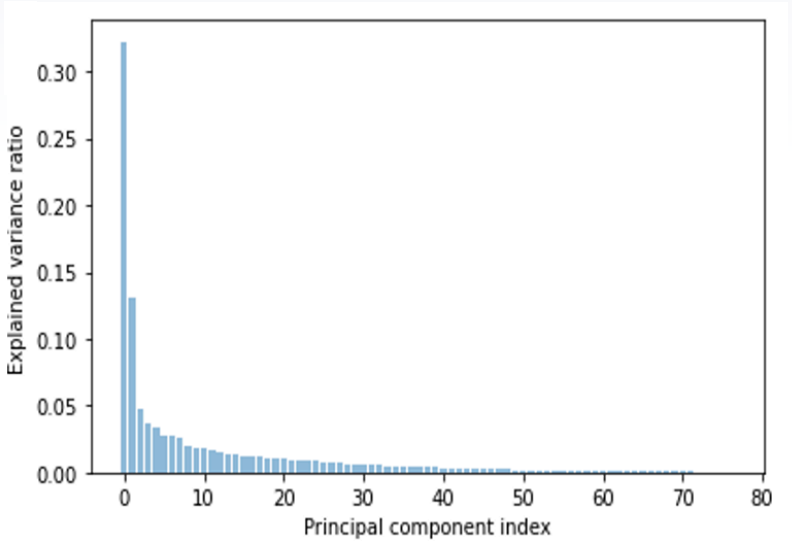

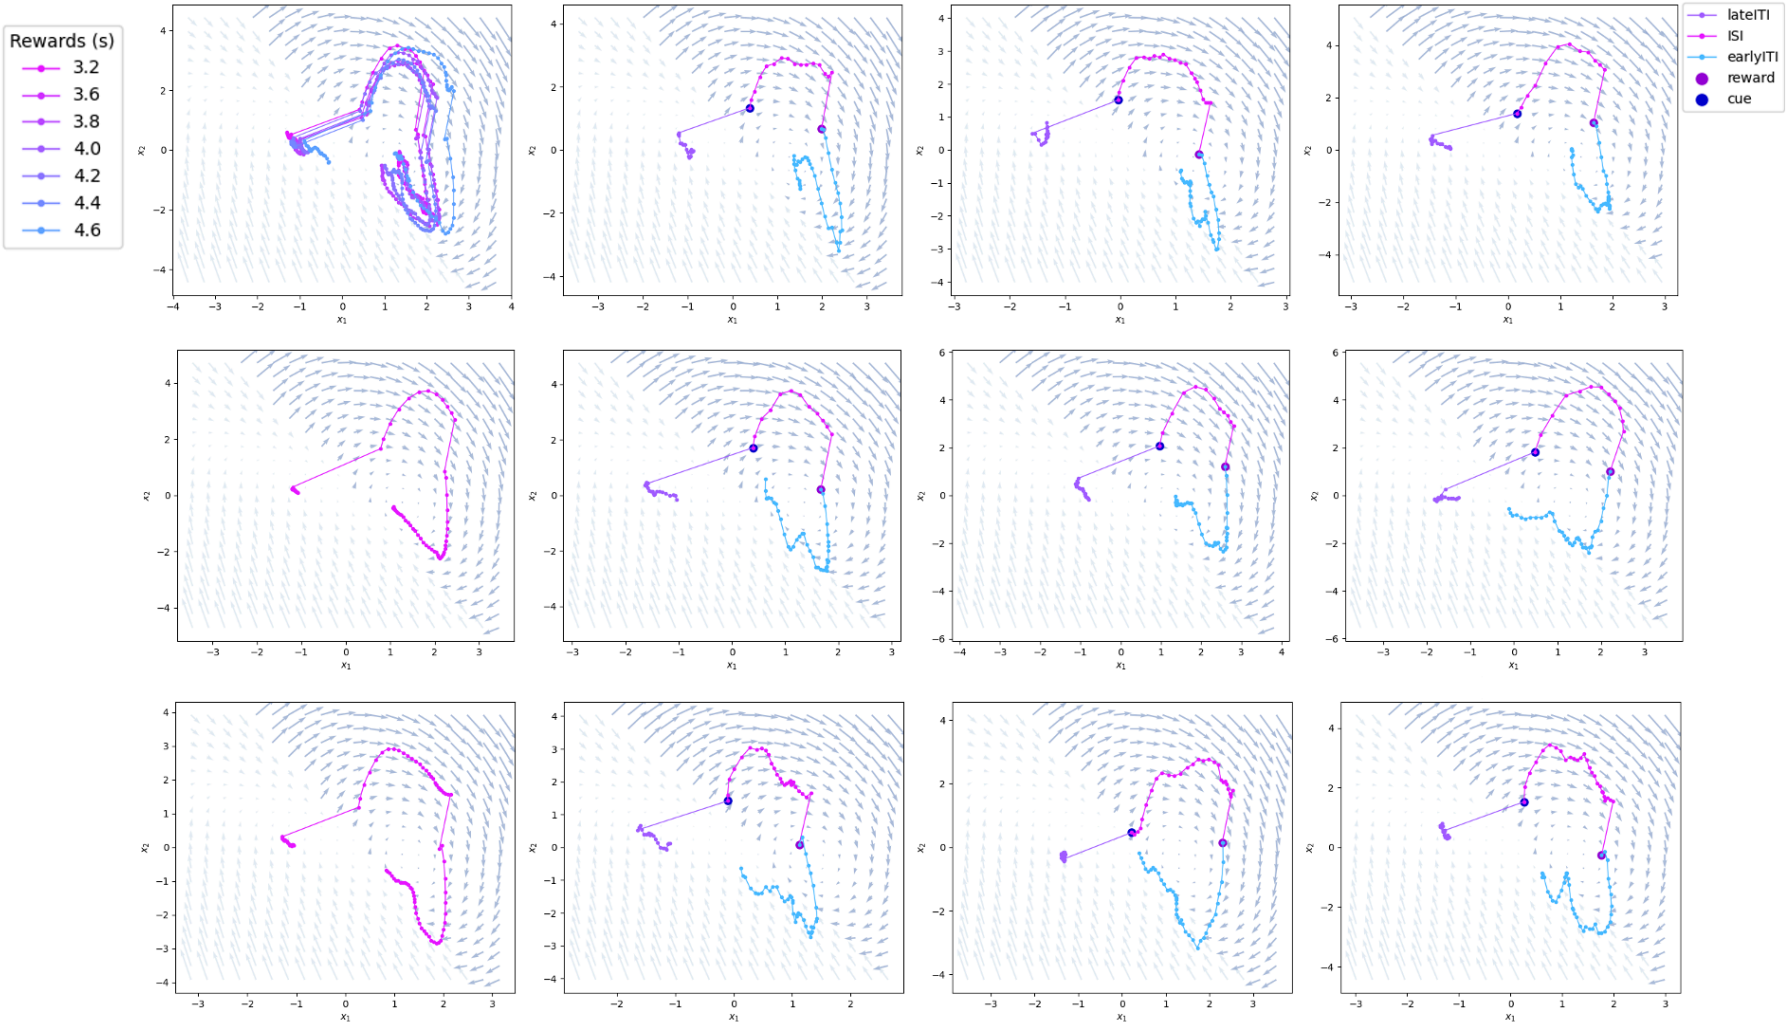
Rewards (s)
3.2
3.6
3.8
4.0
4.2
4.4
4.6
lateITI
ISI
earlyITI
reward
cue
x_1
x_2

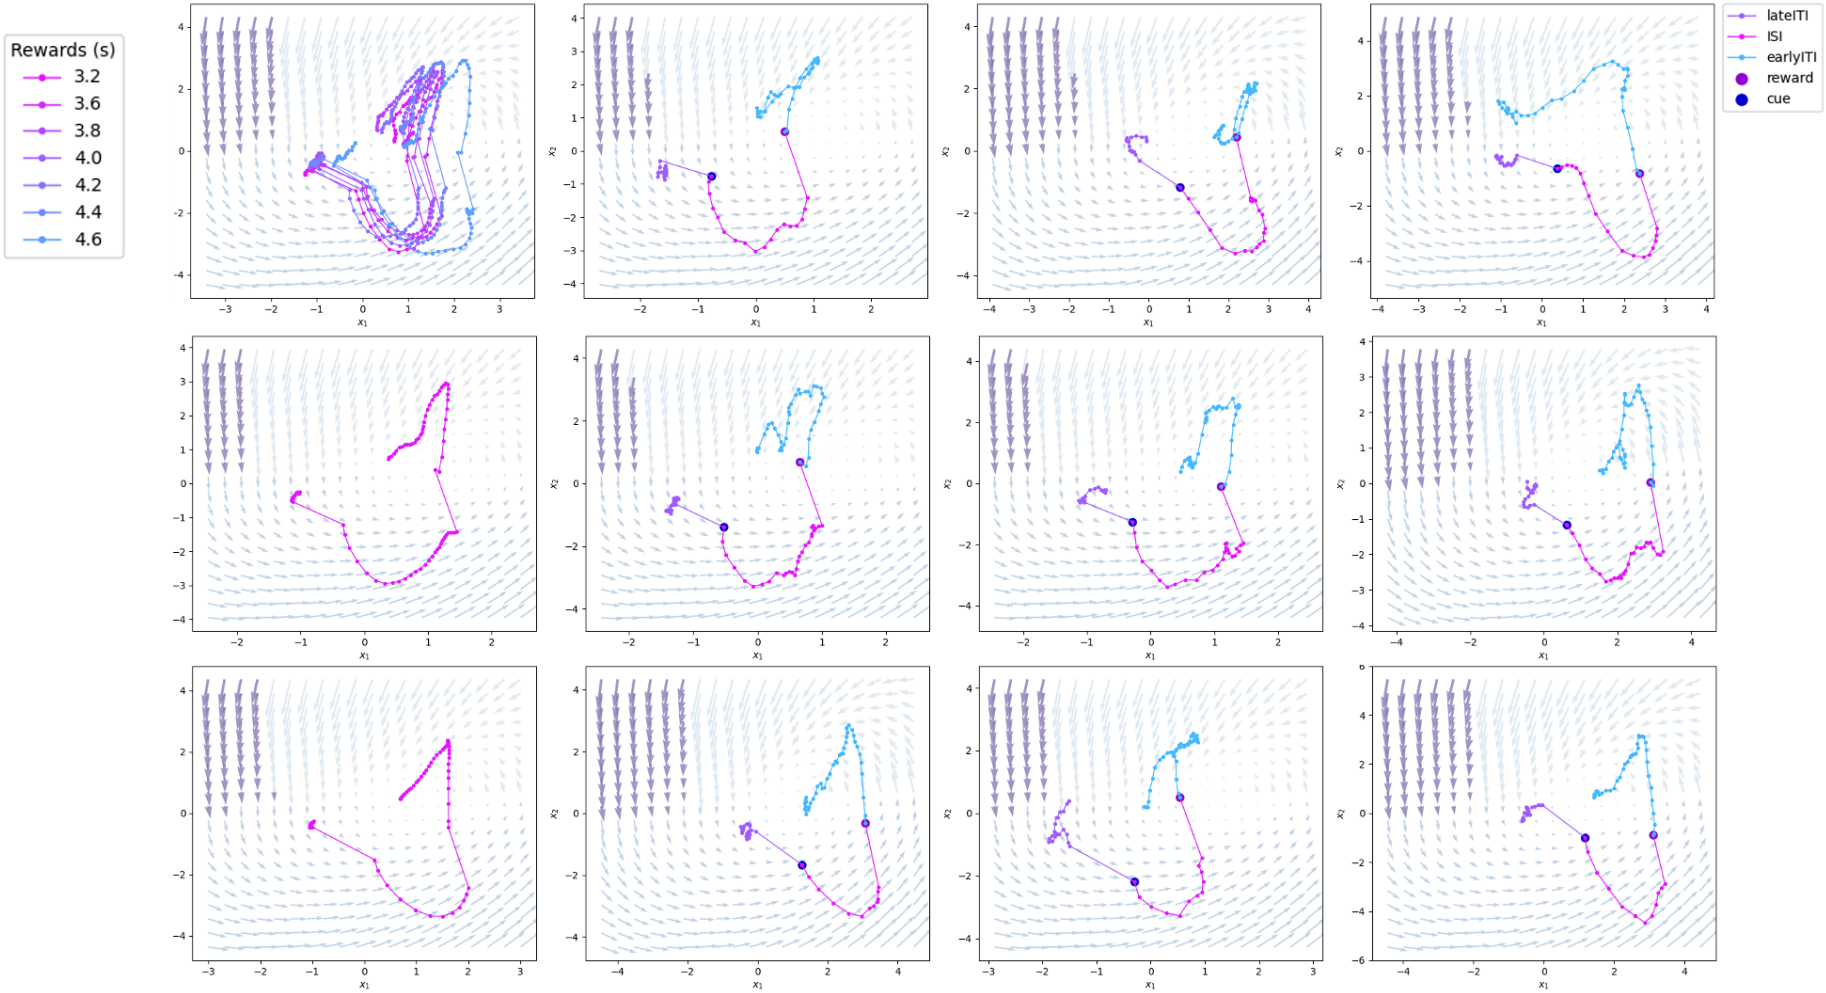
Rewards (s)
3.2
3.6
3.8
4.0
4.2
4.4
4.6
lateITI
ISI
earlyITI
reward
cue
x_1
x_2

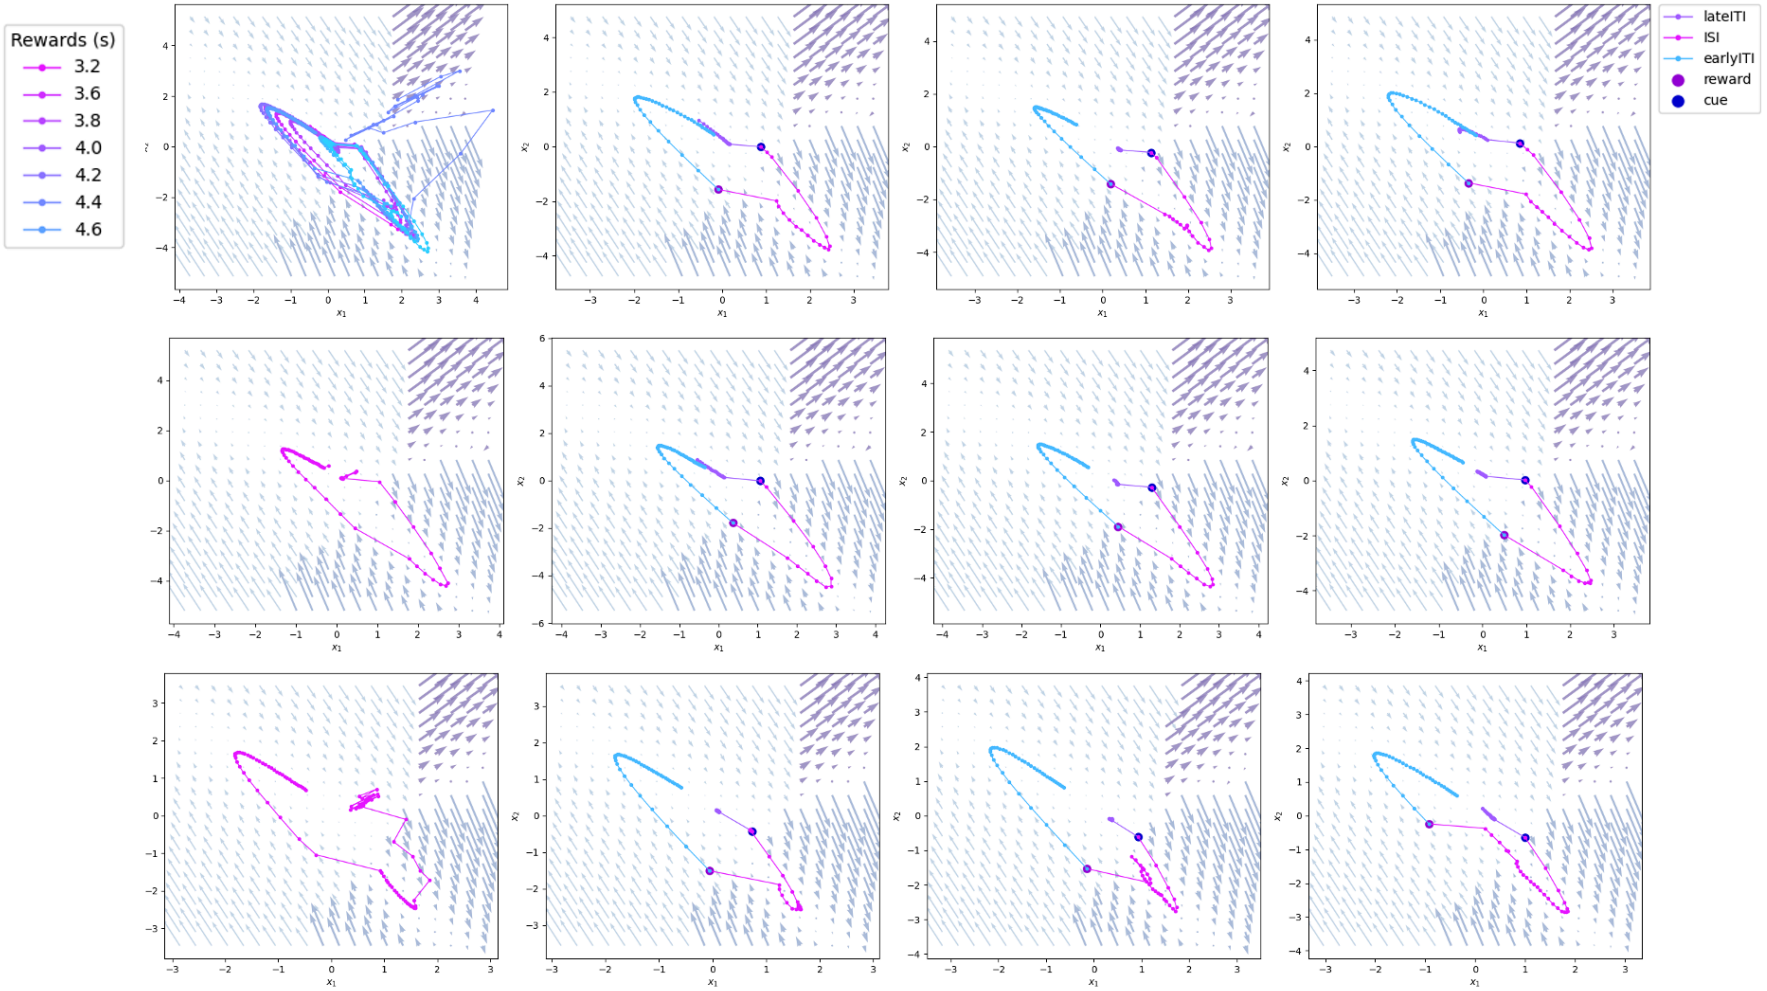
Rewards (s)
3.2
3.6
3.8
4.0
4.2
4.4
4.6
lateITI
ISI
earlyITI
reward
cue
x_1
x_2

Recurrent switching linear dynamical systems

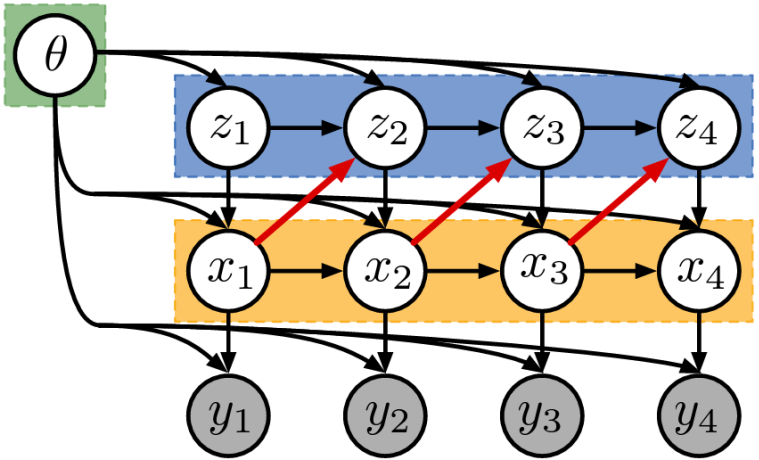

$$\theta = \{A_k, V_k, b_k, C, d, R, W, r\}$$ (Parameters)

$$p(z_t = i \mid z_t - 1 = j, x_{t-1}) \propto \exp(R x_{t-1} + W u_{t-1} + r)$$

$$z_t = \{1,2, \ldots K\}, z_t \mid u_t, z_{t-1}, x_{t-1}$$ (Discrete states)

$$x_t = A_{z_t} \cdot x_{t-1} + V_{z_t} \cdot u_t + b_{z_t}$$ (Latent dimensions)

$$y_t = C \cdot x_t + d$$ (High-dim data)

Gated Recurrent Unit (GRU)

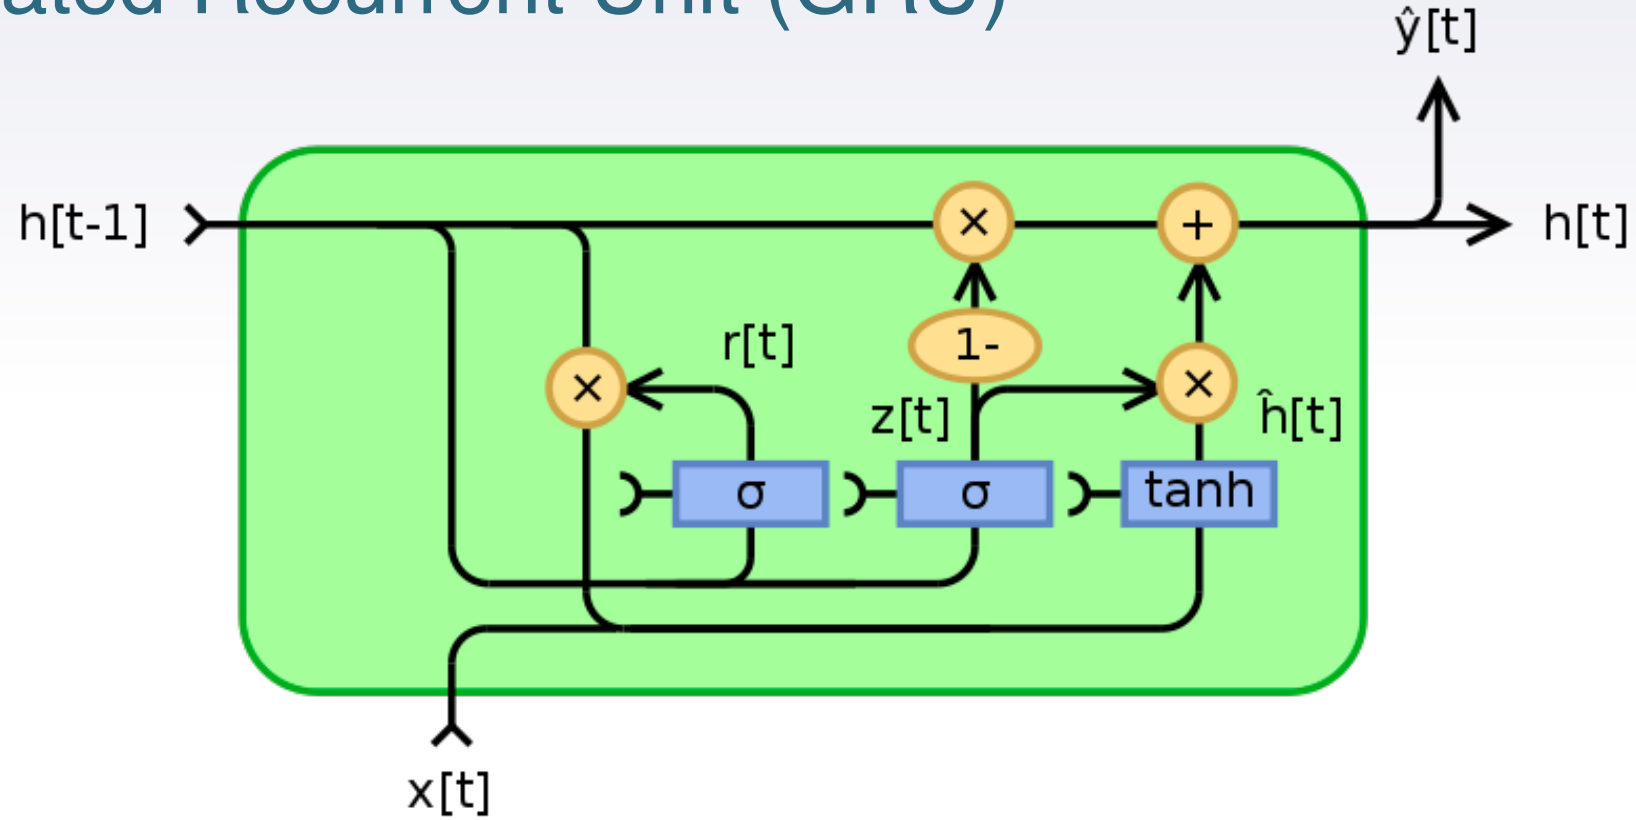

GLM variables structure

Variable	Event	Basis type	Basis Duration	Bases Number
ISI	Cue	unit	[1.2s, 2.8s]	14, cut at reward
Reward	Reward	unit	3s	15
Licking	Licking	Cosine	1s	5
Nuisance	Recording onset	Cosine	Full recording	5

CSC vs microstimuli features representation

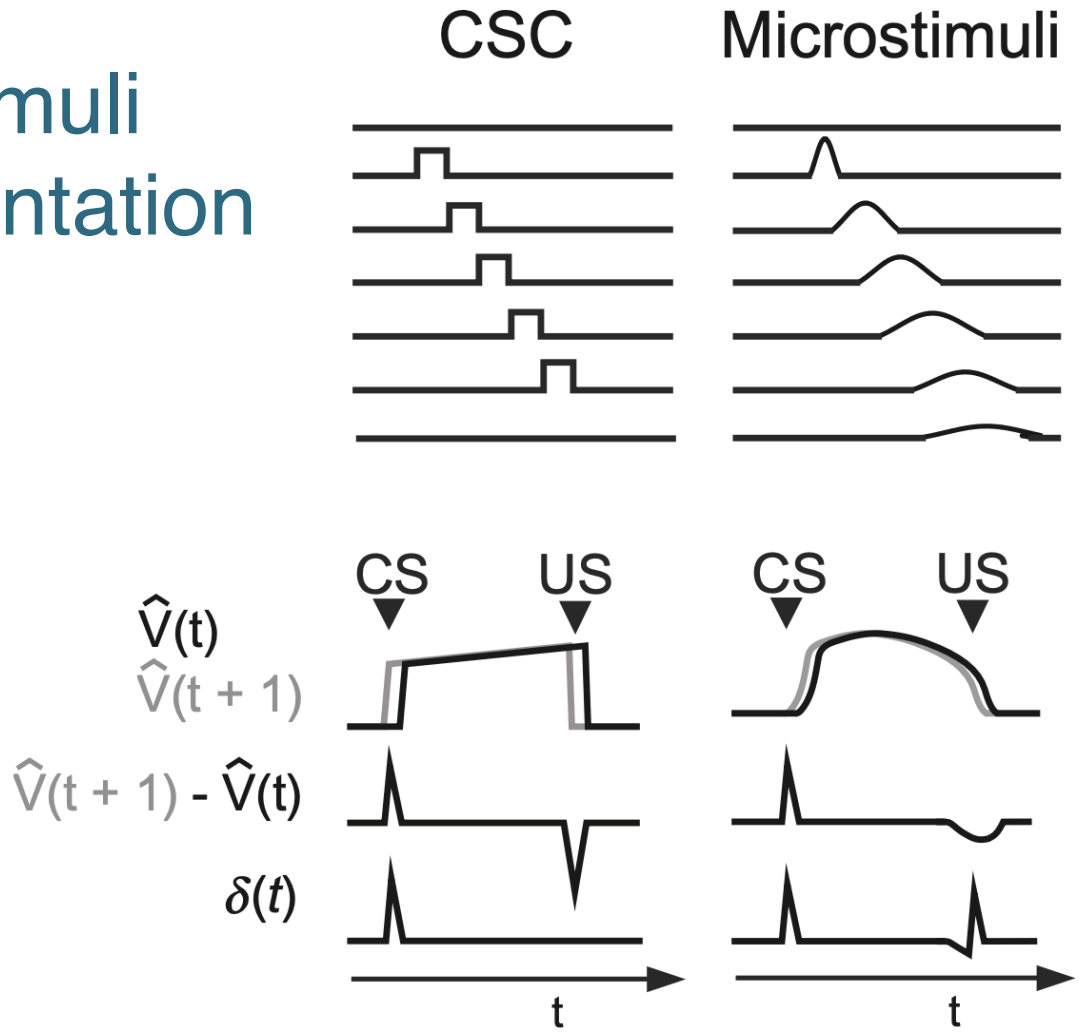

Starkweather & Uchida, 2020

Consistent EmBeddings of high-dimensional Recordings using Auxiliary variables (CEBRA)

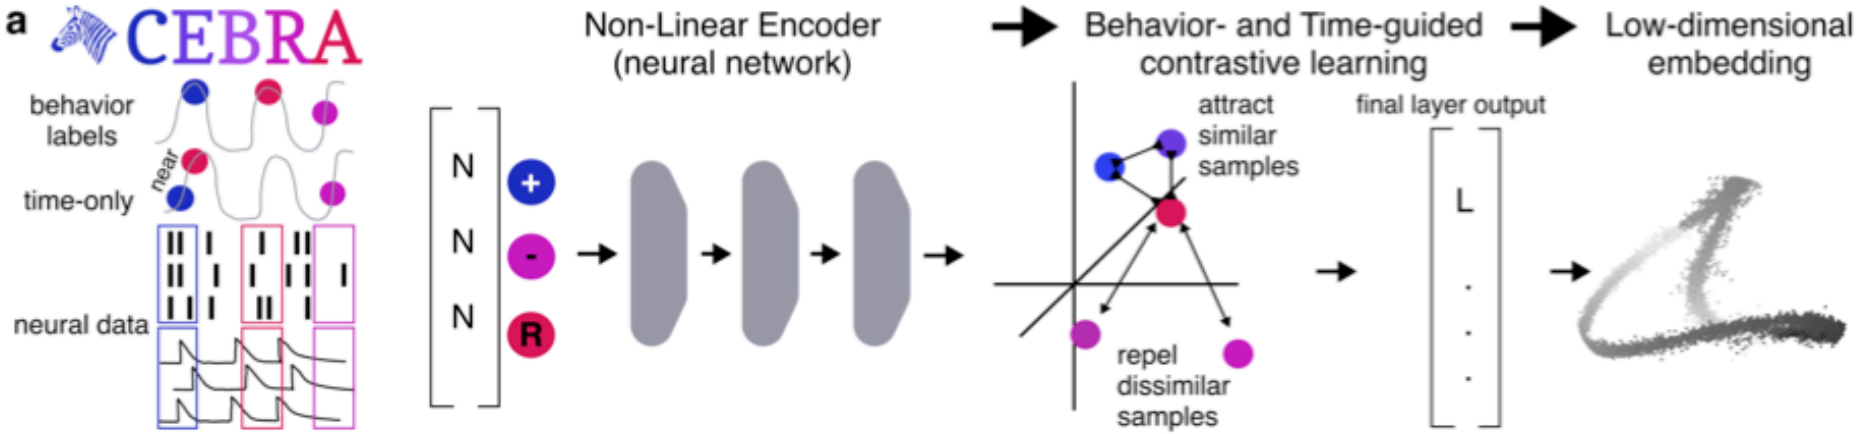

Schneider et al., 2022, preprint

Demixed PCA

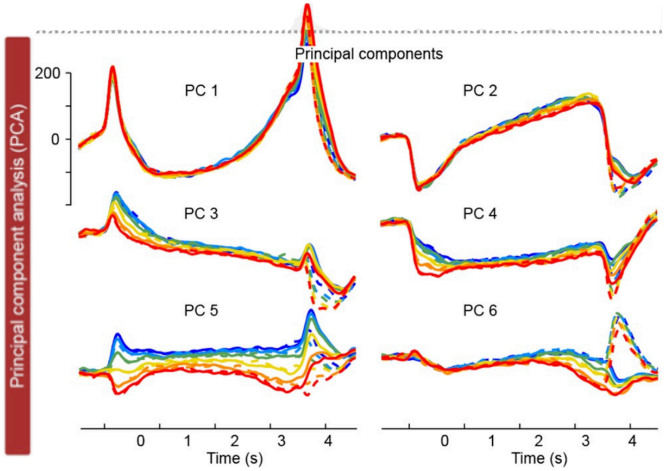

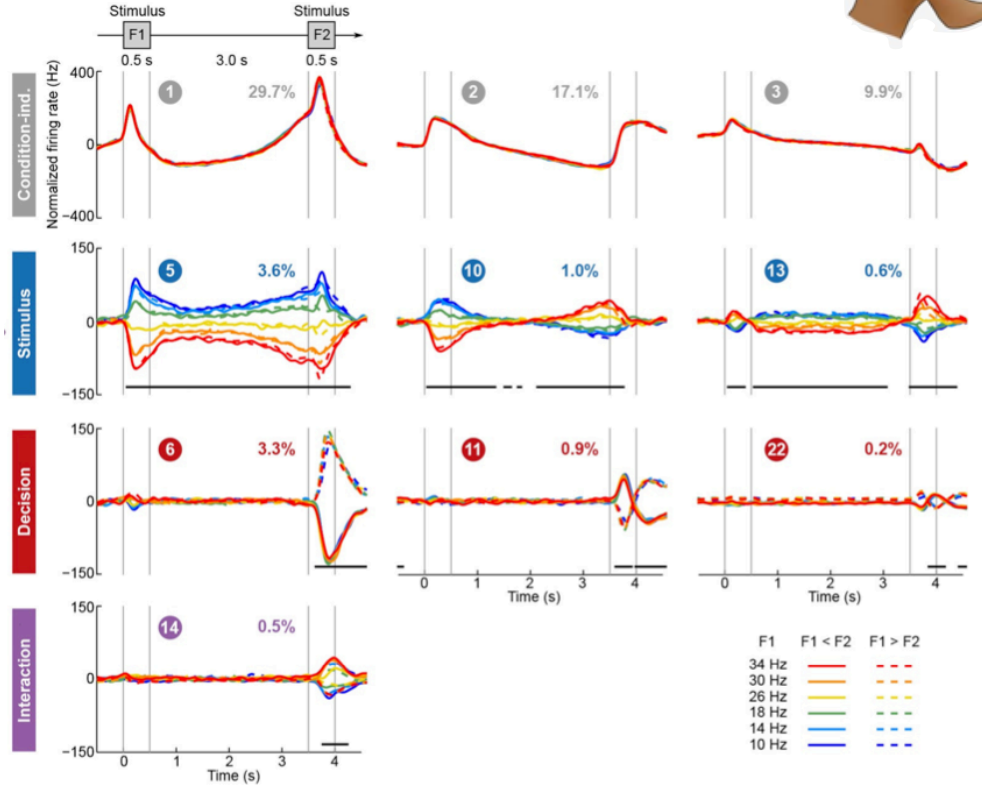

- Concise way of visualizing the data that summarizes the task-dependent features of the population response in a single feature.

Kobak et al., 2015, eLife