A virtual rodent predicts the structure of neural activity across behaviors

Authors: Diego Aldarondo*, Josh Merel, Jesse D. Marshall, Leonard Hasenclever, Ugne Klibaite, Amanda Gellis, Yuval Tassa, Greg Wayne, Matthew Botvinick & Bence P. Ölveczky*

Speaker: Ziyuan Ye

Thursday, July 25, 2024

Aldarondo, D., Merel, J., Marshall, J. D., Hasenclever, L., Klibaite, U., Gellis, A., ... & Ölveczky, B. P. (2024). A virtual rodent predicts the structure of neural activity across behaviors. *Nature*, 1-3.

Overview

1. Introduction to Bence P. Ölveczky
2. Background
3. Framework
4. Method & Results
5. Discussion

Overview

1. Introduction to Bence P. Ölveczky
2. Background
3. Framework
4. Method & Results
5. Discussion

Bence P. Ölveczky

Bence Ölveczky

FOLLOW

Professor, Harvard
Verified email at fas.harvard.edu - Homepage
Neuroscience

TITLE	CITED BY	YEAR
Temporal structure of motor variability is dynamically regulated and predicts motor learning ability HG Wu, YR Miyamoto, LNG Castro, BP Ölveczky, MA Smith Nature neuroscience 17 (2), 312-321	704	2014
Motor cortex is required for learning but not for executing a motor skill R Kawai, T Markman, R Poddar, R Ko, AL Fantana, AK Dhawale, ... Neuron 86 (3), 800-812	593	2015
Vocal experimentation in the juvenile songbird requires a basal ganglia circuit BP Ölveczky, AS Andalman, MS Fee PLoS biology 3 (5), e153	589	2005
Segregation of object and background motion in the retina BP Ölveczky, SA Baccus, M Meister Nature 423 (6938), 401-408	509	2003
The role of variability in motor learning AK Dhawale, MA Smith, BP Ölveczky Annual review of neuroscience 40 (1), 479-498	462	2017

Cited by VIEW ALL

	All	Since 2019
Citations	6613	3875
h-index	32	27
i10-index	41	38

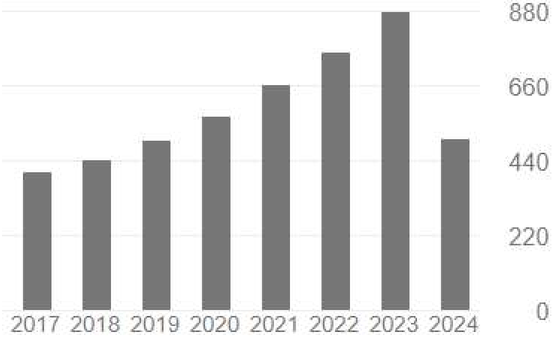

Public access VIEW ALL

1 article 25 articles

not available available

Based on funding mandates

Overview

What is AI Agent?

AI Agent Landscape

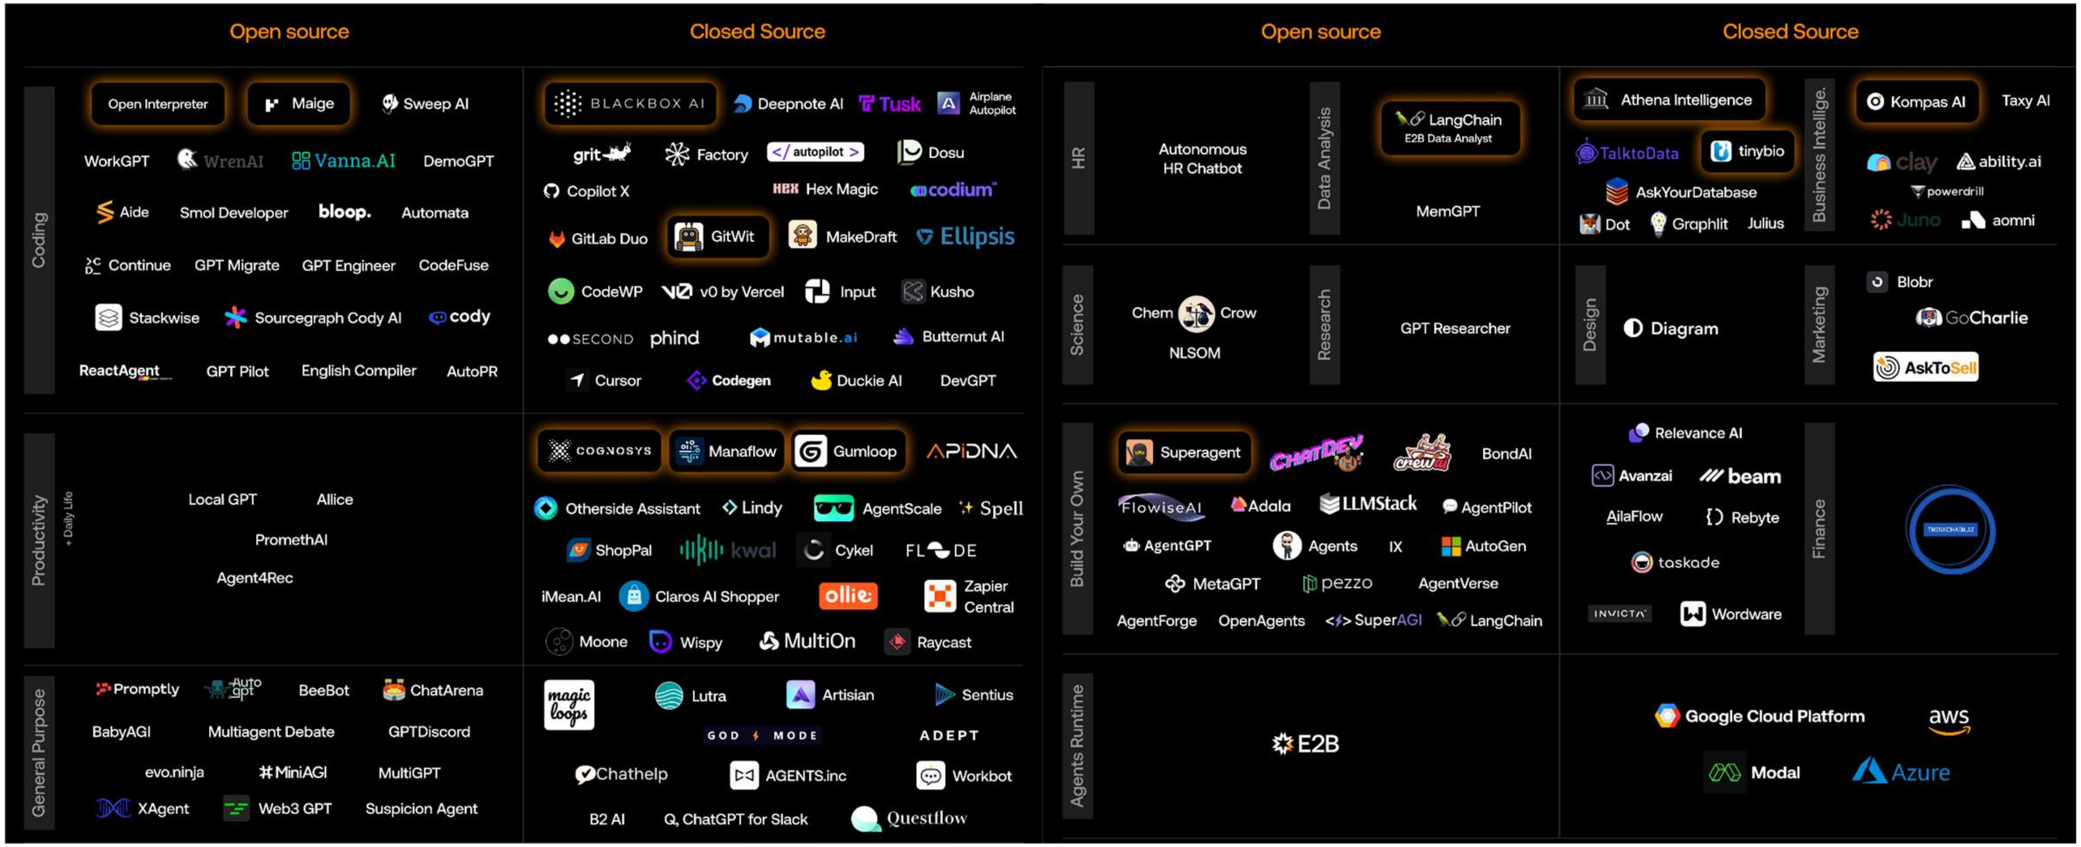

Most of the listed AI agents are mainly based on LLMs and do not have real-time perception of the physical world. They cannot directly perceive visual, auditory, or tactile information unless this information is explicitly input to them in text form.

Credits: https://github.com/e2b-dev/awesome-ai-agents

What is AI Agent?

AI Agent Landscape

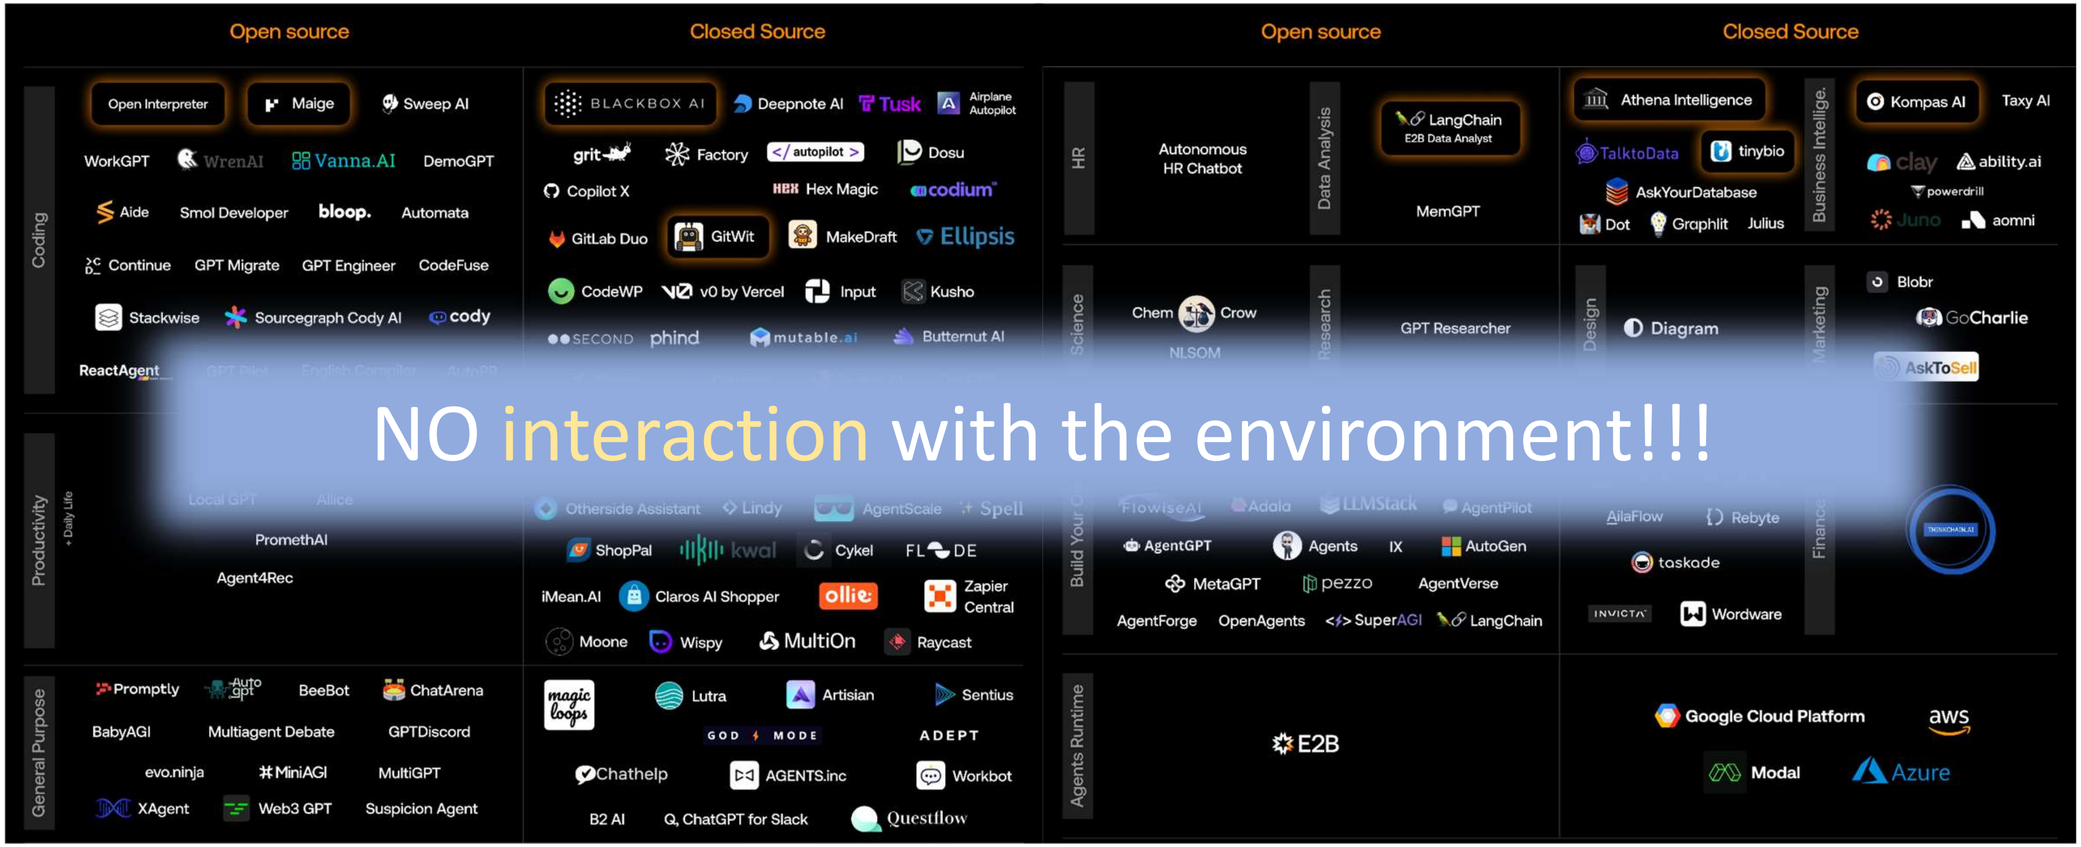

NO interaction with the environment!!!

Most of the listed AI agents are mainly based on LLMs and do not have real-time perception of the physical world. They cannot directly perceive visual, auditory, or tactile information unless this information is explicitly input to them in text form.

Credits: https://github.com/e2b-dev/awesome-ai-agents

Will AI Agent with Perception Module Allow Us to Well Model Real-world Animals?

DEEP NEUROETHOLOGY OF A VIRTUAL RODENT

Josh Merel[⋆1], **Diego Aldarondo**[⋆2,3], **Jesse Marshall**[⋆3,4], **Yuval Tassa**[1],
Greg Wayne[1], **Bence Ölveczky**[3,4]
[1]DeepMind, London, UK.
[2]Program in Neuroscience, [3]Center for Brain Science, [4]Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA 02138, USA.

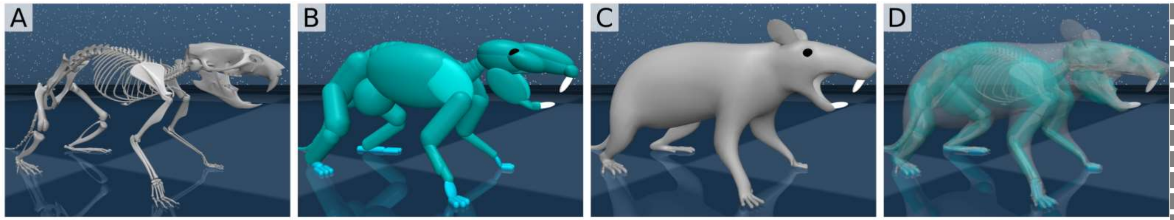

Figure 1: (A) Anatomical skeleton of a rodent (as reference; not part of physical simulation). (B) A body designed around the skeleton to match the anatomy and model collisions with the environment. (C) Purely cosmetic skin to cover the body. (D) Semi-transparent visualization of (A)-(C) overlain.

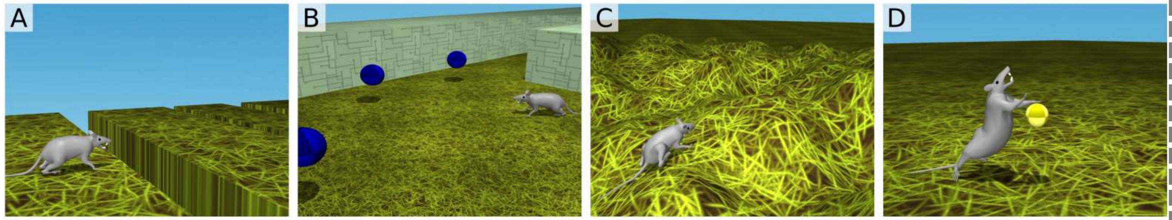

Figure 2: Visualizations of four tasks the virtual rodent was trained to solve: (A) jumping over gaps ("gaps run"), (B) foraging in a maze ("maze forage"), (C) escaping from a hilly region ("bowl escape"), and (D) touching a ball twice with a forepaw with a precise timing interval between touches ("two-tap").

Input: Perceptual information / Visual information

Goal:
Through deep reinforcement learning methods to train an ANN (Perception: Small-ResNet; Core layer: LSTM; Policy layer: LSTM) to complete four different tasks in virtual environment.

Results:

1. Some behavioral representations were shared across tasks, especially in the policy layers.
2. Policy layers encoded more low-level motor features, while core layers encoded higher-level task variables.
3. Networks with fewer layers showed more sharing of representations across tasks.

Limitations:

1. Cannot well mimic the real-world actions and behaviors.
2. Four specific tasks cannot cover the whole action space.

Merel, J., Aldarondo, D., Marshall, J., Tassa, Y., Wayne, G., & Olveczky, B. (2020). Deep neuroethology of a virtual rodent. In *International Conference on Learning Representations*.

Overview

1. Introduction to Bence P. Ölveczky
2. Background
3. Framework
4. Results
5. Discussion

How to replicate the movement and behavior of real animals?

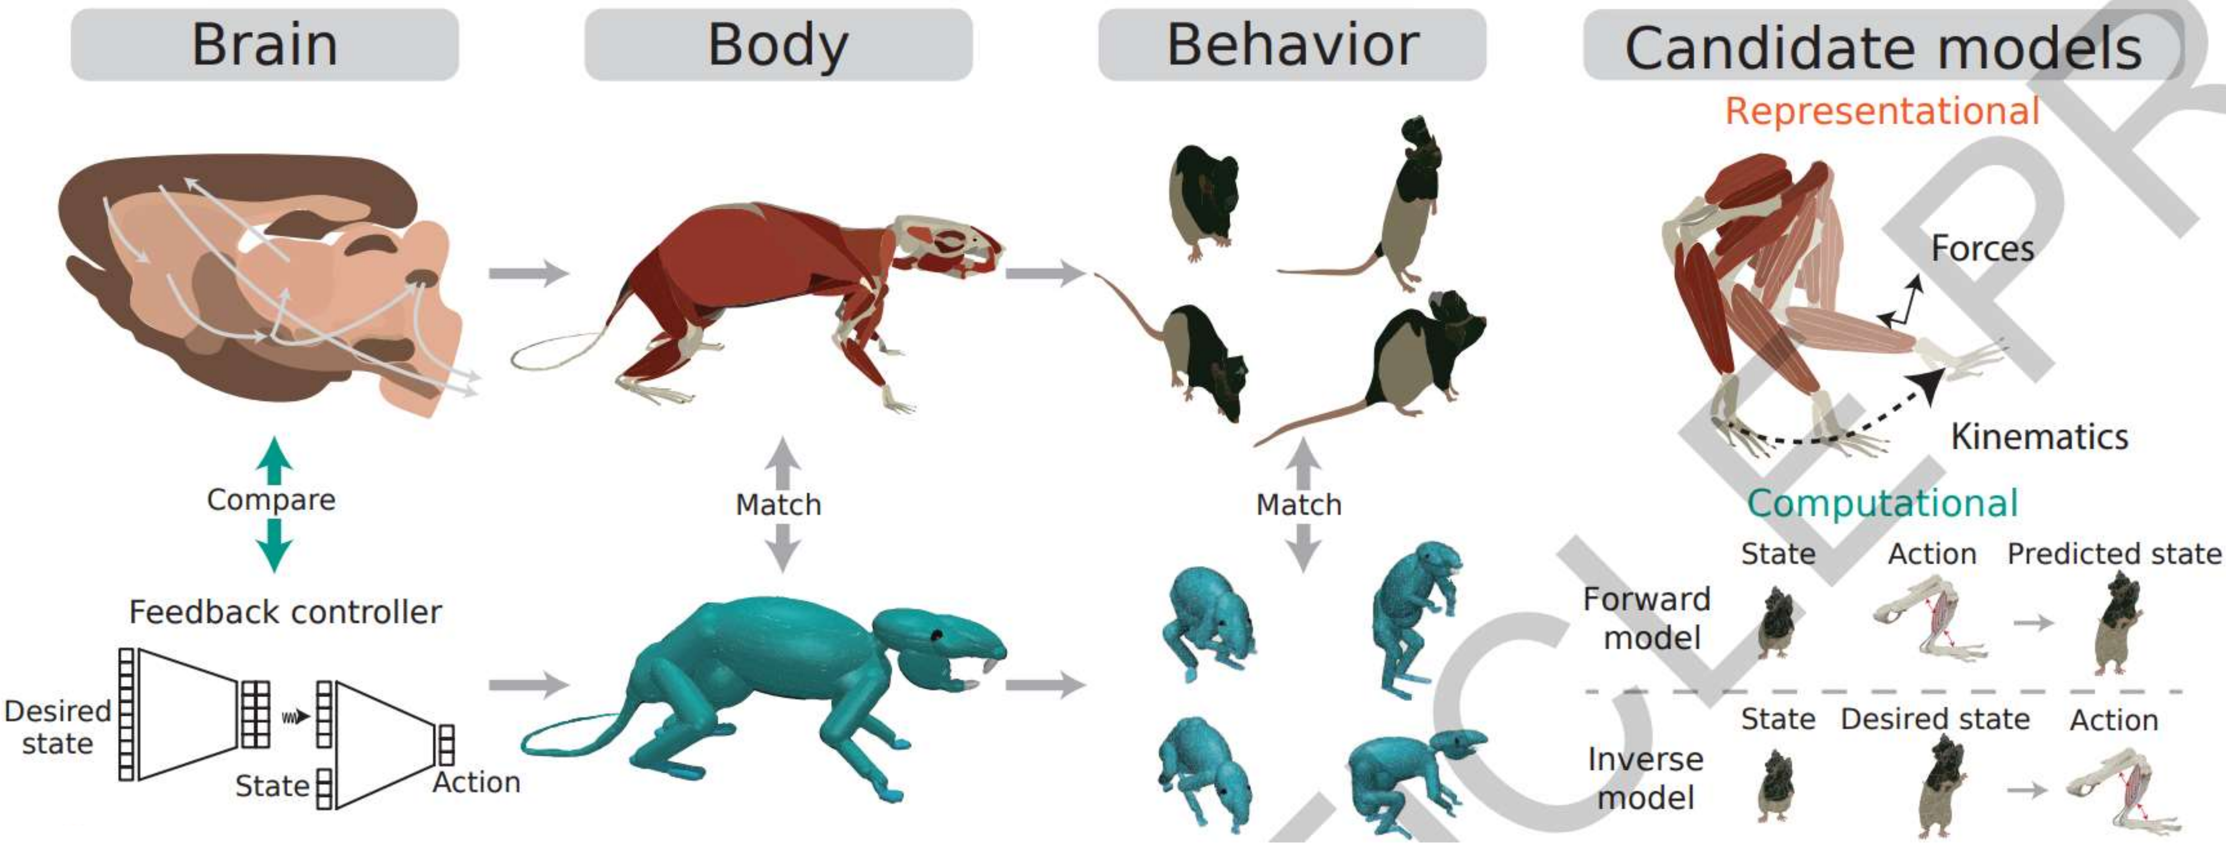

Four key aspects of replication.

MIMIC Framework: Motor IMItation and Control

Step 1: Data acquisition

- 6 cameras for behavior recording
- 2 128-channel tetrode for brain neural activities recording

Step 2: 3D pose estimation

- Using DANNCE to track 23 anatomical landmarks (key points) on the rat

Step 3: Skeletal registration

- Register a 74 degree-of-freedom / 38 controllable degrees-of-freedom skeletal rat model to the key points

Step 4: Imitation

- Trained ANNs to achieve inverse dynamics models using DRL
- Input: reference trajectories and current body state
- Outputs: joint torques

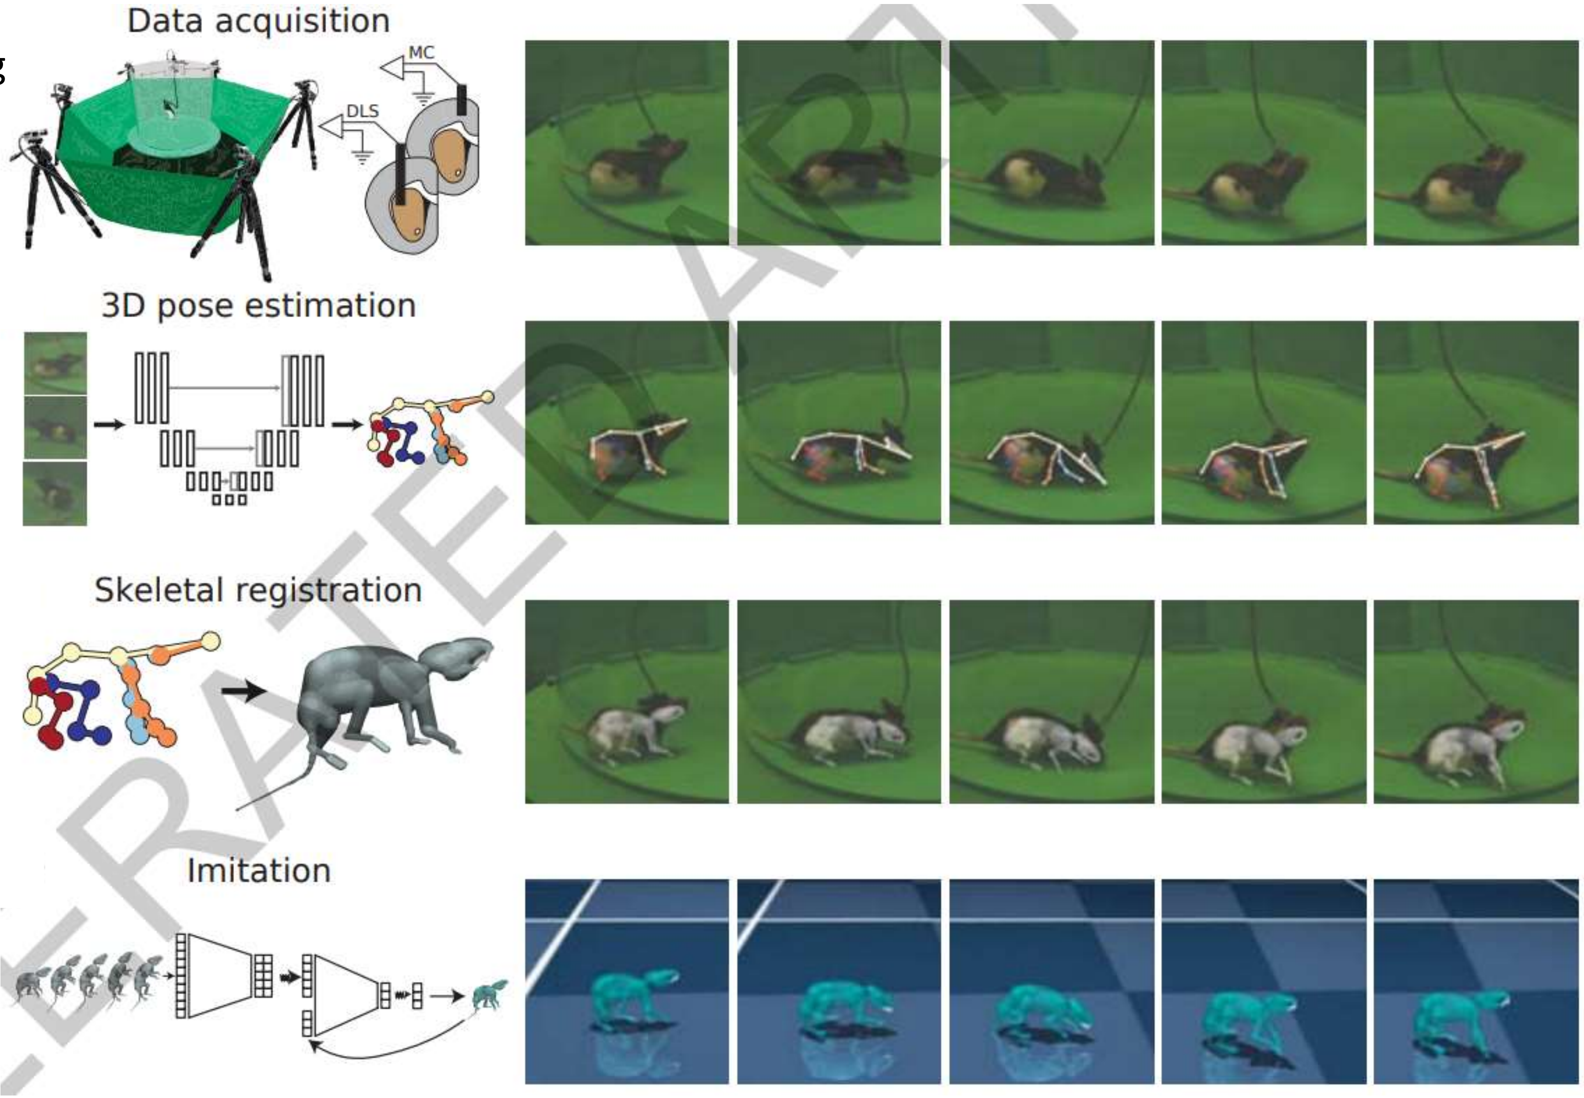

Overview

Training Artificial Agents to Imitate Rat Behavior with MIMIC

A. Training framework
The neural network implements an inverse dynamics model to generate the desire action based on the current state and the reference trajectory.

B. Behavior similarity comparison
The anterior-posterior (AP) position of the key points; Key points' height from floor.

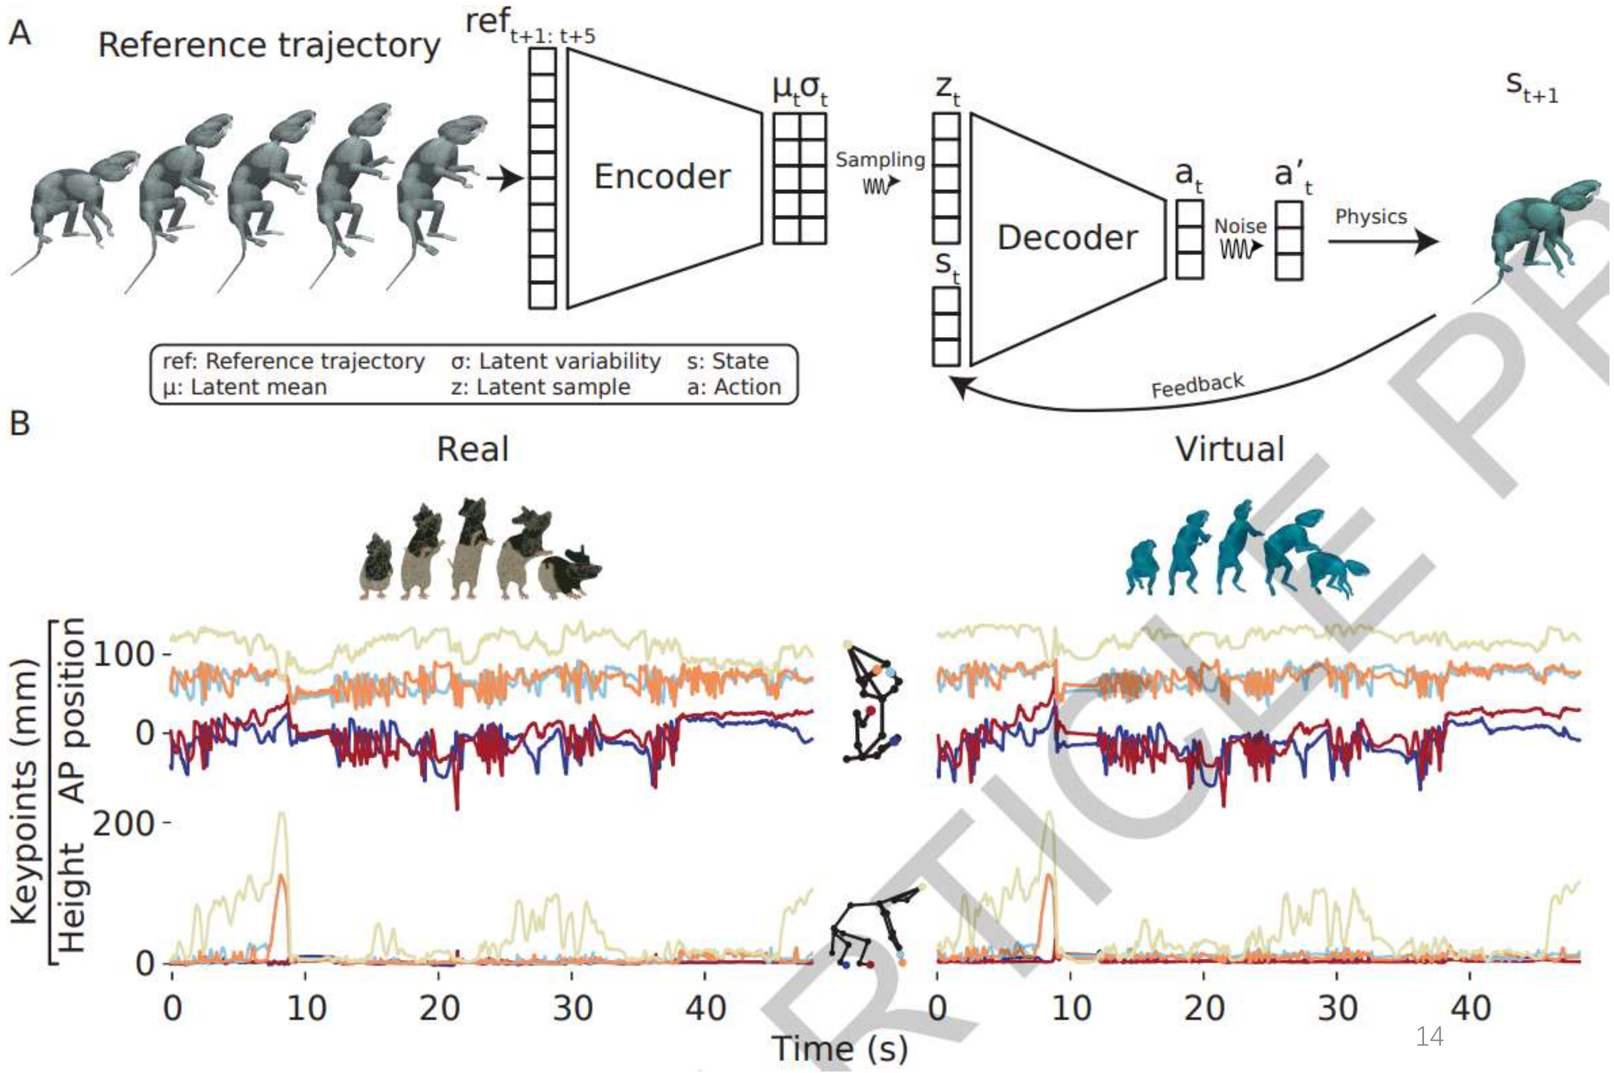

Training Artificial Agents to Imitate Rat Behavior with MIMIC

C. Behavior demonstration
Rear (后肢站立)
Walk (自然行走)

D. Simulation error of different body parts

E. Simulation error of different behavior

F. Illustration of error accumulation over time

G. Robustness evaluation: Demonstrate how long the virtual mouse model can continue to accurately mimic the behavior of a real mouse without triggering termination and reset conditions

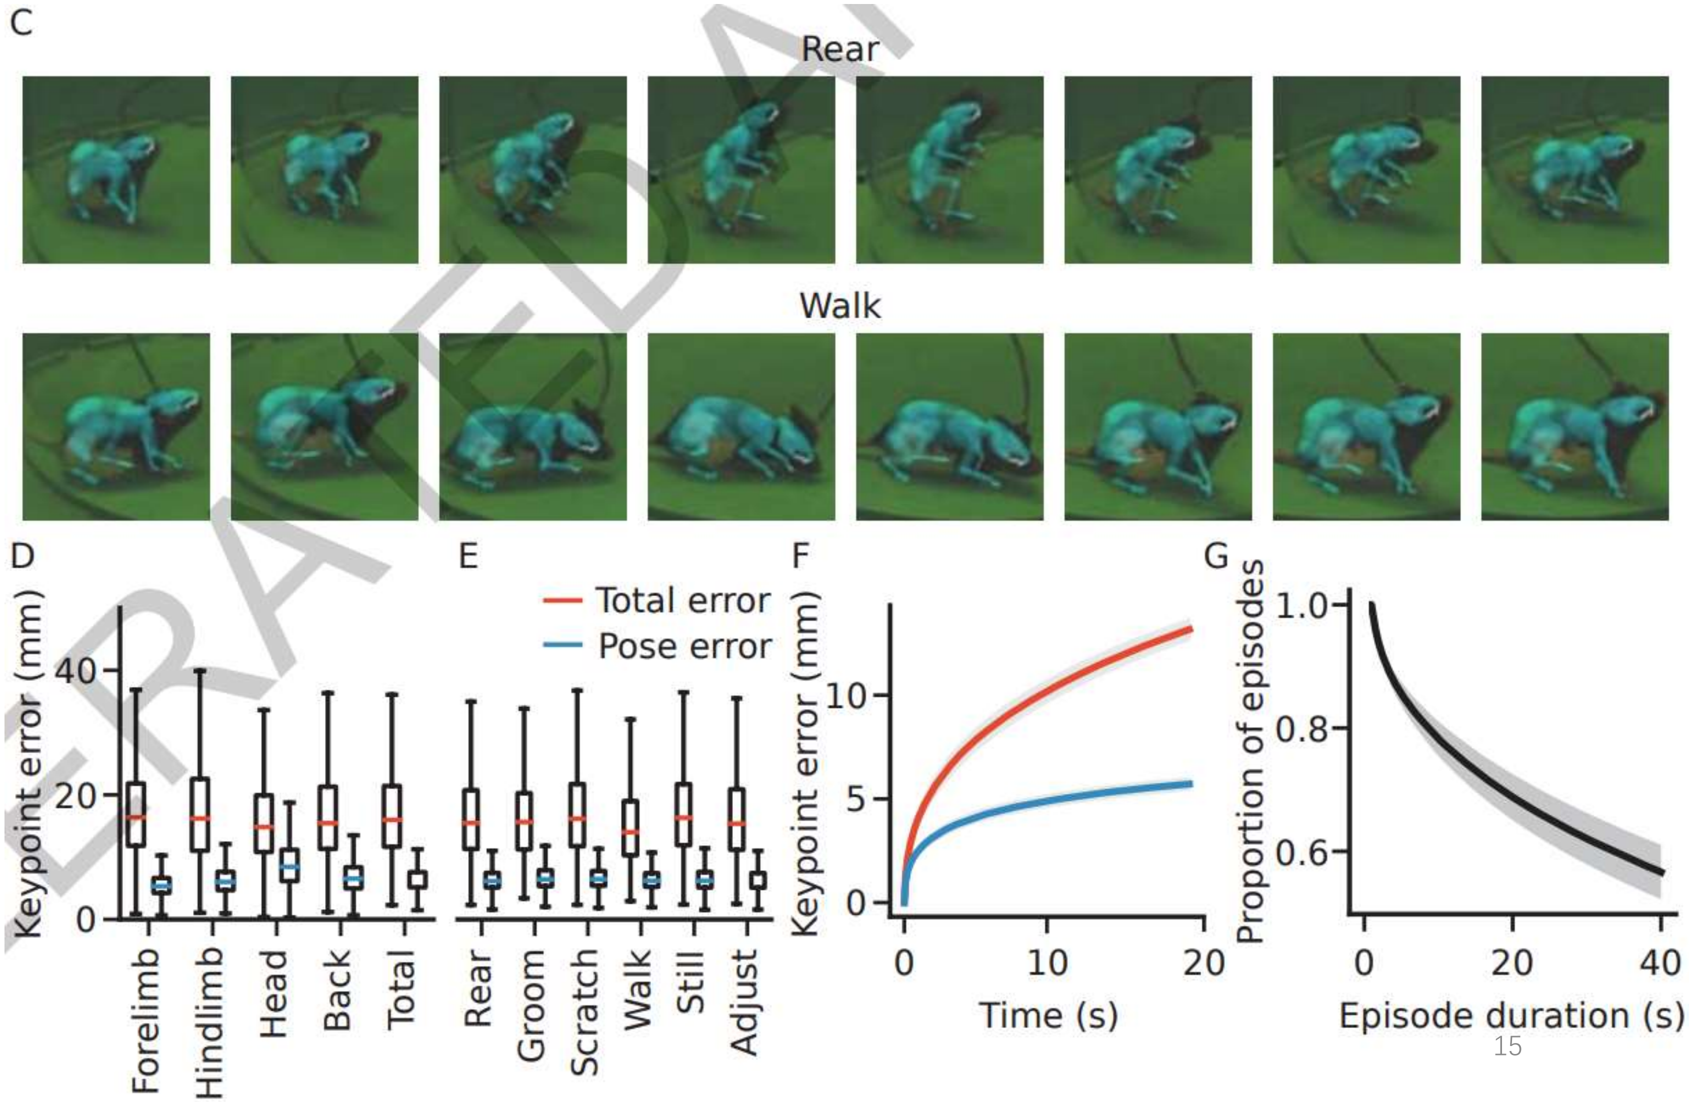

Neural Activity of Virtual and Real Rat

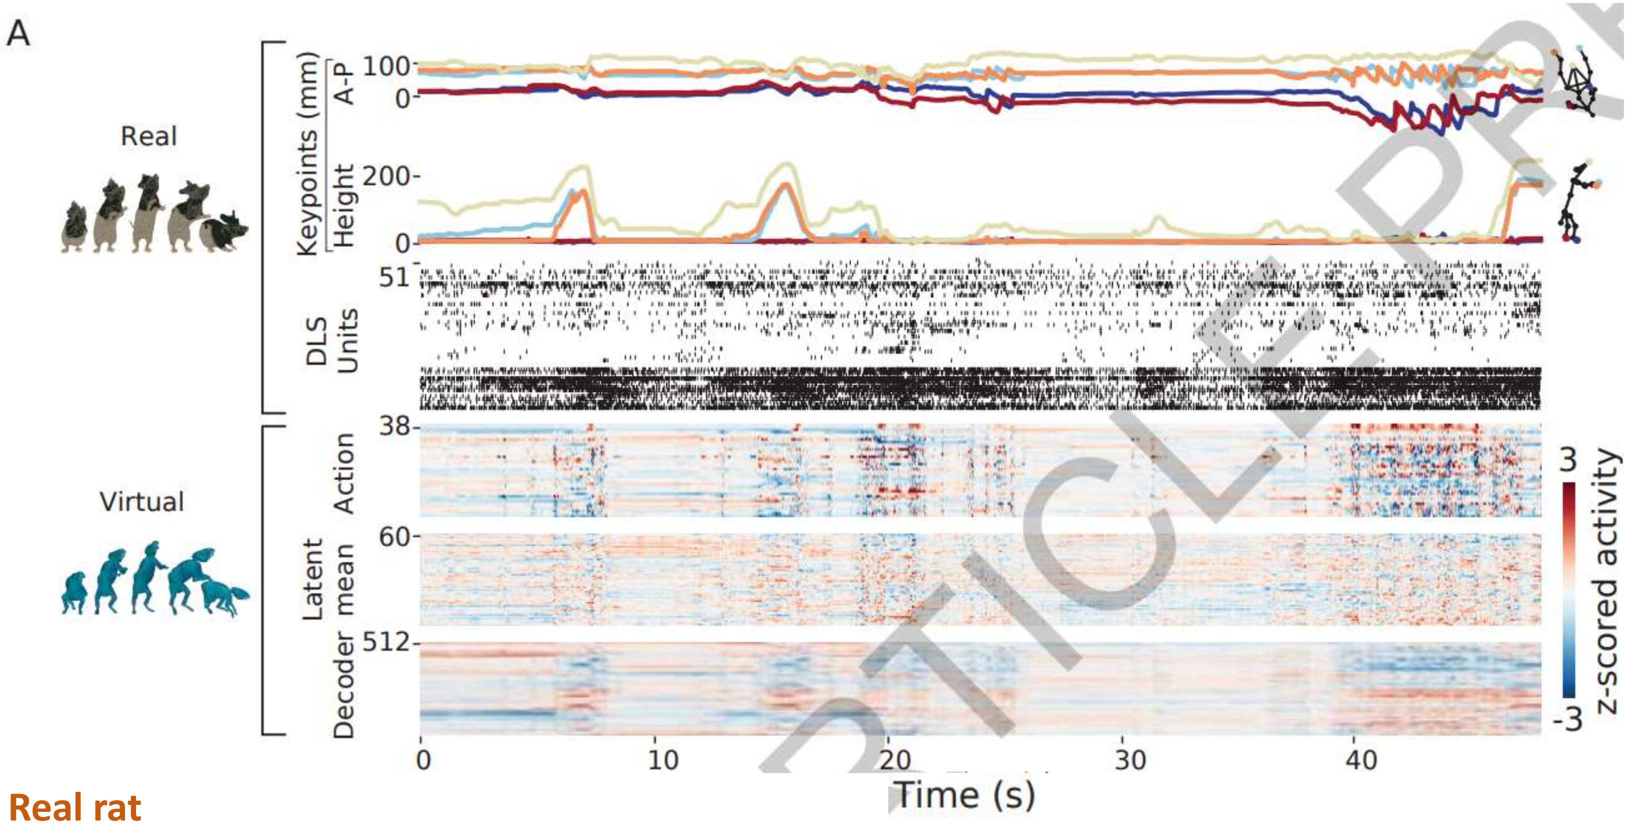

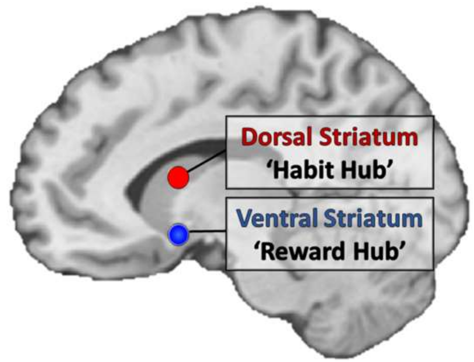

Real rat

Top: A-P: Distance between each anterior-posterior (A-P) position of the key points and centroid position of rat; Height: Key points' height from floor.

Bottom: DLS (Dorsolateral striatum, 背外侧纹状体, responsible for motor control and coordination) spike train

Virtual rat

Different representation value trains

ref: Reference trajectory σ: Latent variability s: State
μ: Latent mean z: Latent sample a: Action

Neural Activity in DLS and MC is Well Predicted by An Inverse Dynamics Model

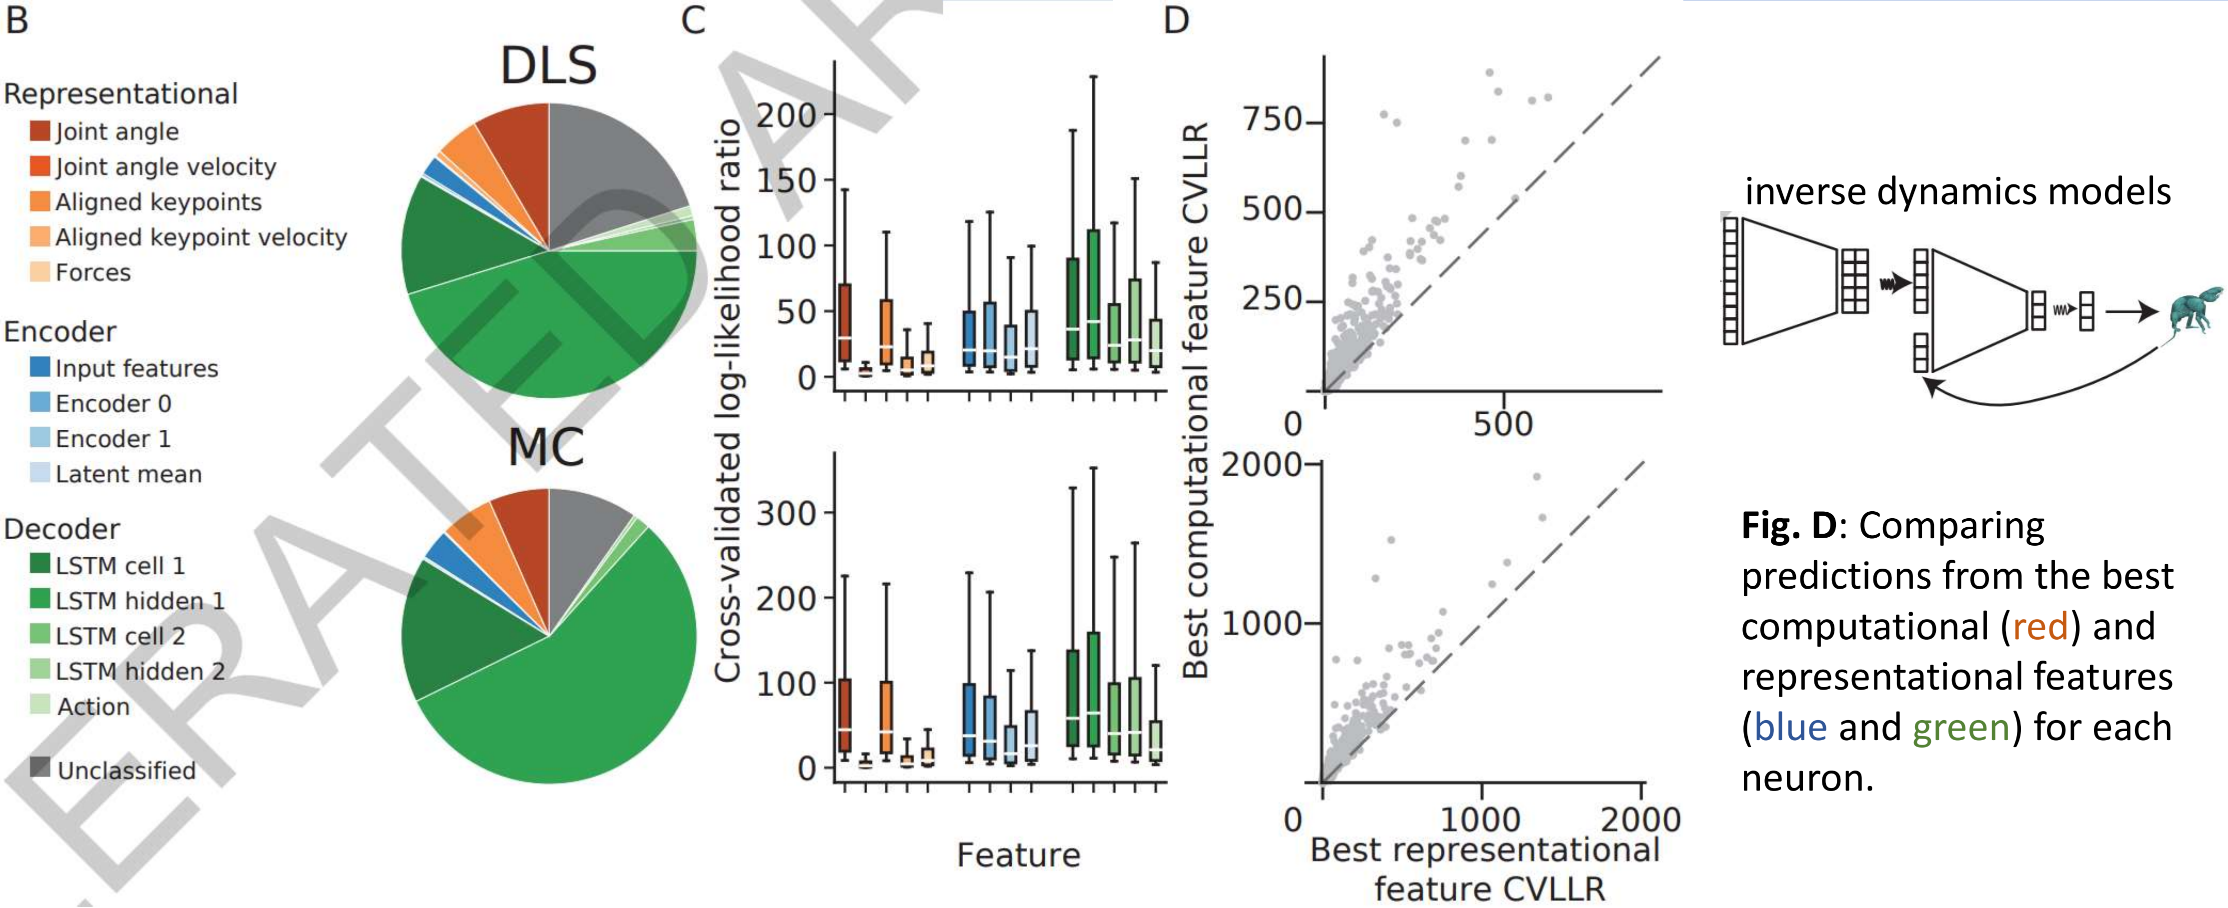

Fig. D: Comparing predictions from the best computational (red) and representational features (blue and green) for each neuron.

Fig. B-C: Using Poisson generalized linear models (GLM) to construct the relationship between each real-world neuronal activity and different virtual rodent representations.

Key finding: GLMs based on the inverse dynamics models outperform those based on representational features for the majority of classified neurons in both DLS and MC.

Similar Patterns of Neuron Population between Virtual and Real Brain

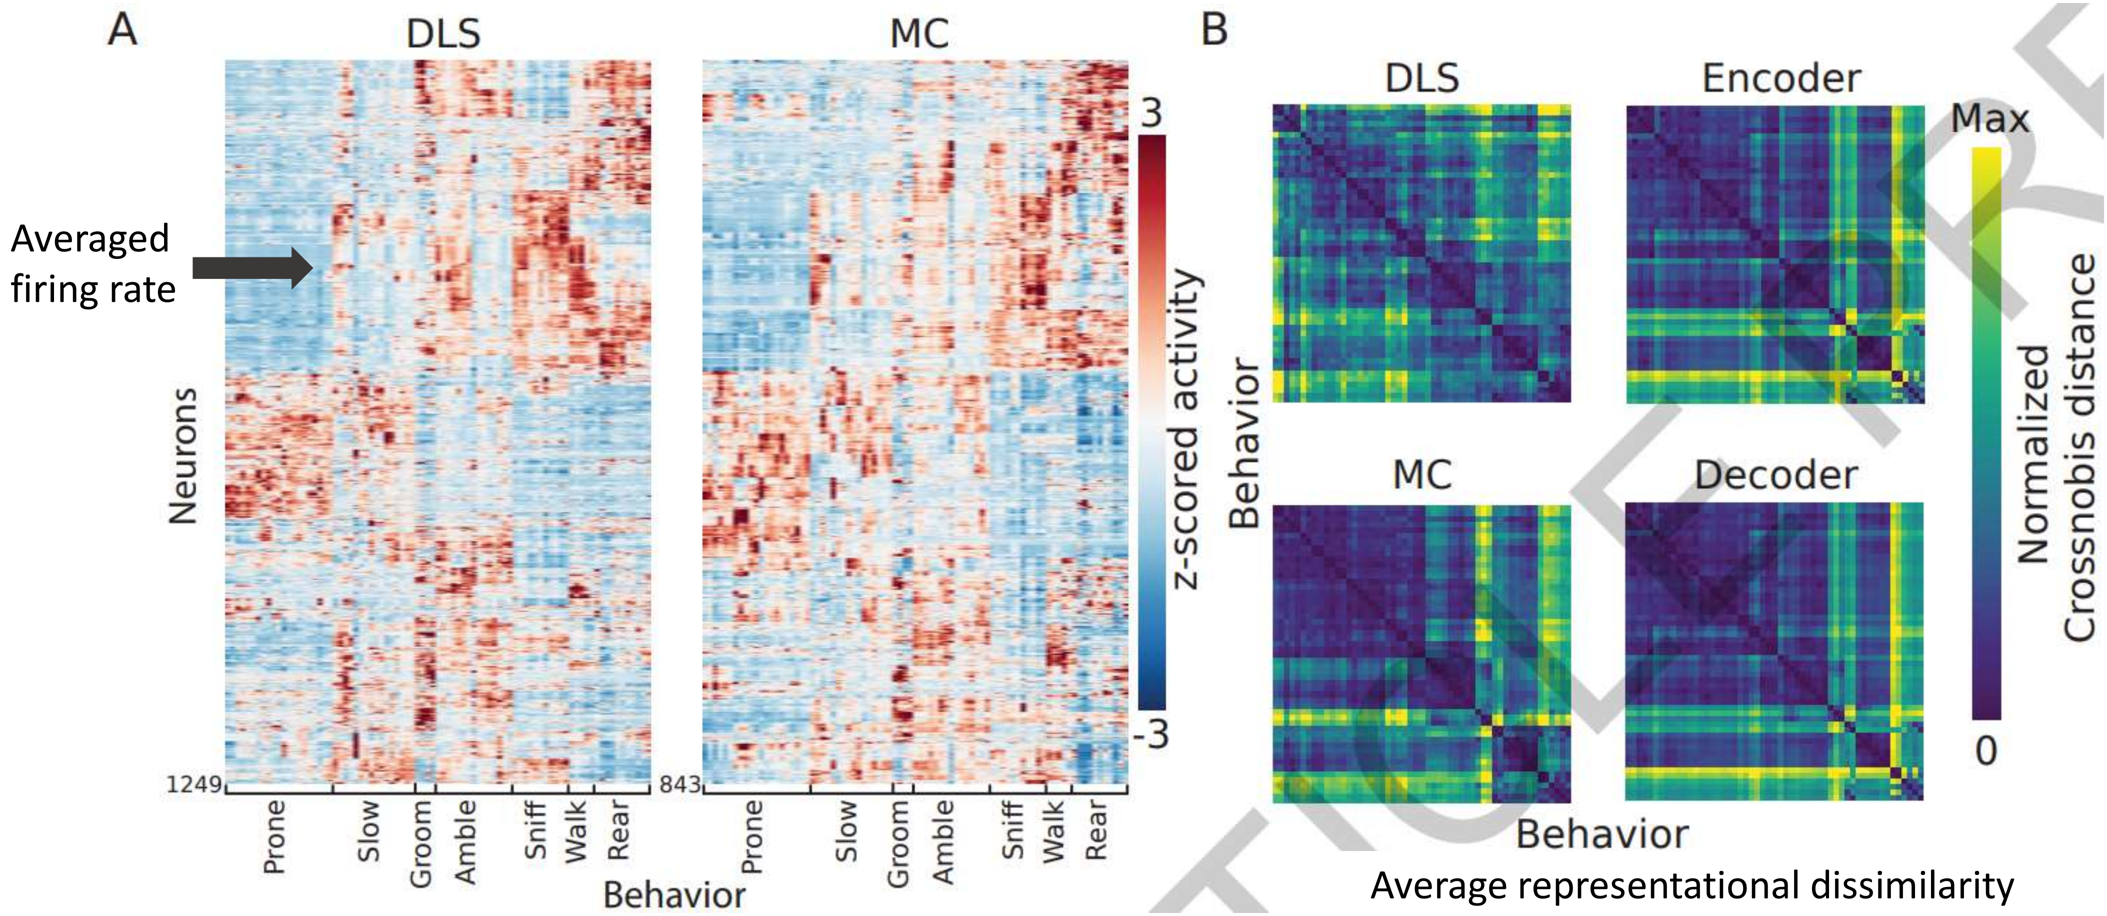

Average representational dissimilarity matrices (RDMs) for neural activity

Background knowledge:

- The DLS plays a key role in learning movement sequences, influencing the output of the motor cortex (MC).
- The MC is responsible for executing the final selected action (similar to decoder).

Similar Patterns of Neuron Population between Virtual and Real Brain

Fig. C-E: Across-subject average of whitened-unbiased cosine (WUC) similarity between RDMs of different computational and representational models and neural activity.
Results: Inverse dynamics models (Green & Blue) perform better than representational models (Red)

Fig. F-I: Similarity between neural activity and each of network layers.
Results: Neural dynamics in MC is easier to predict by artificial neural network.

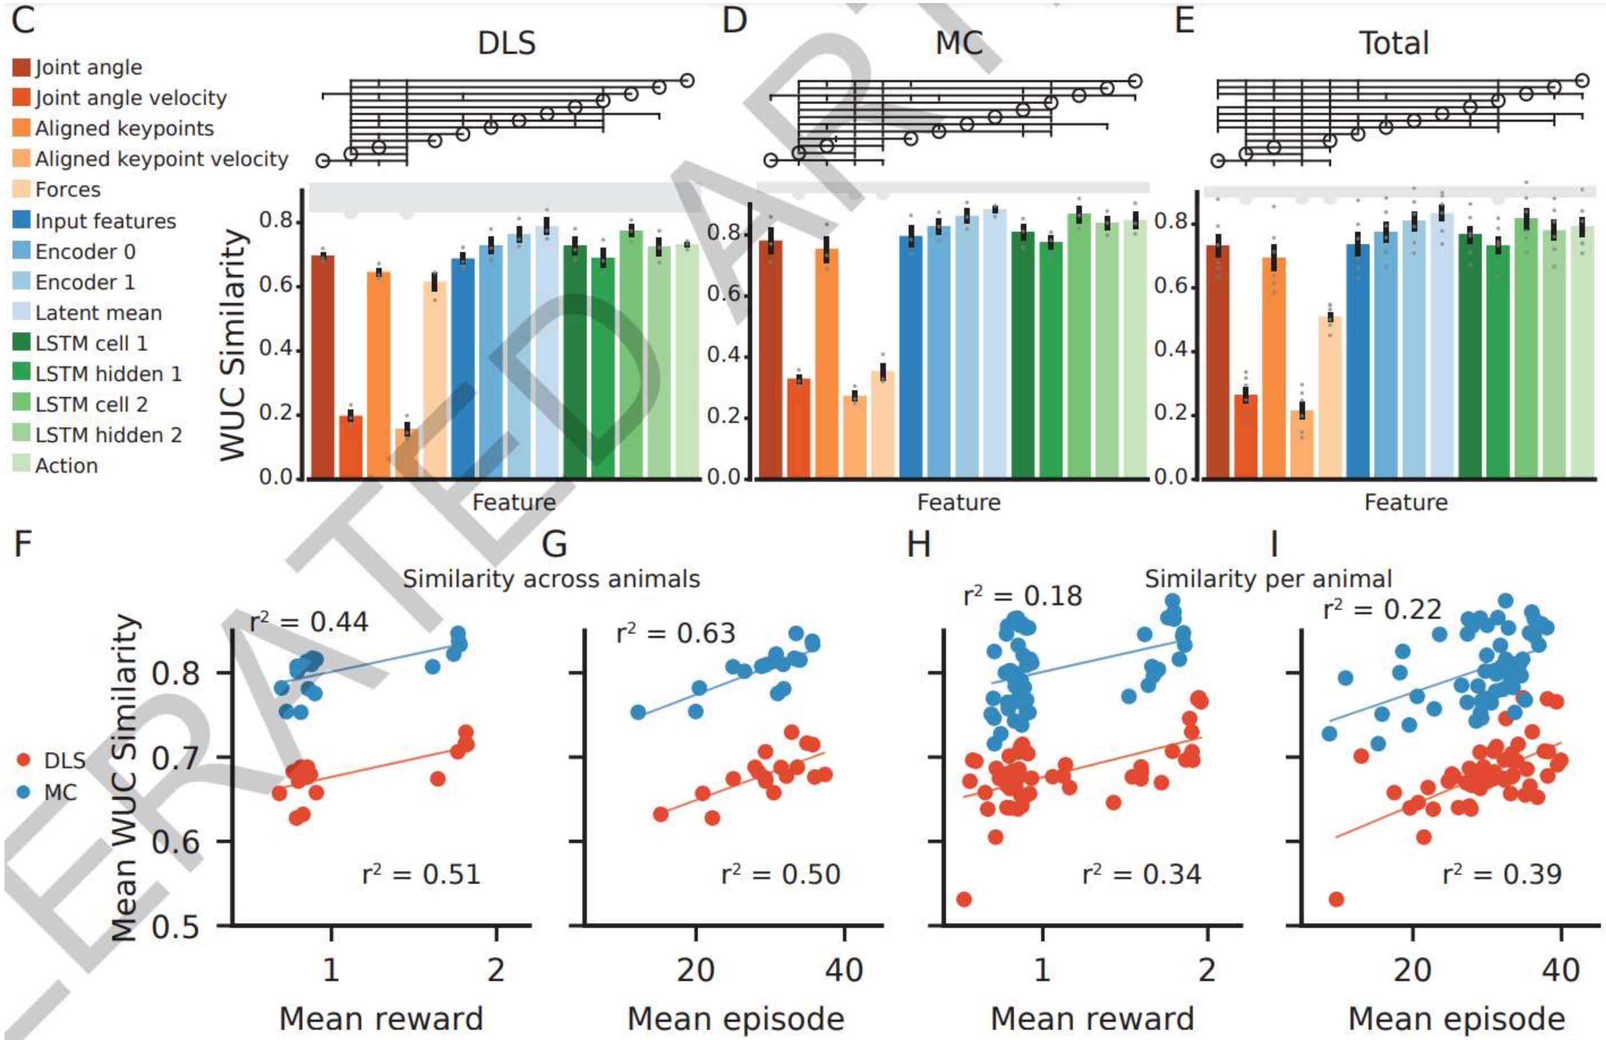

Action Variability & Latent Variability

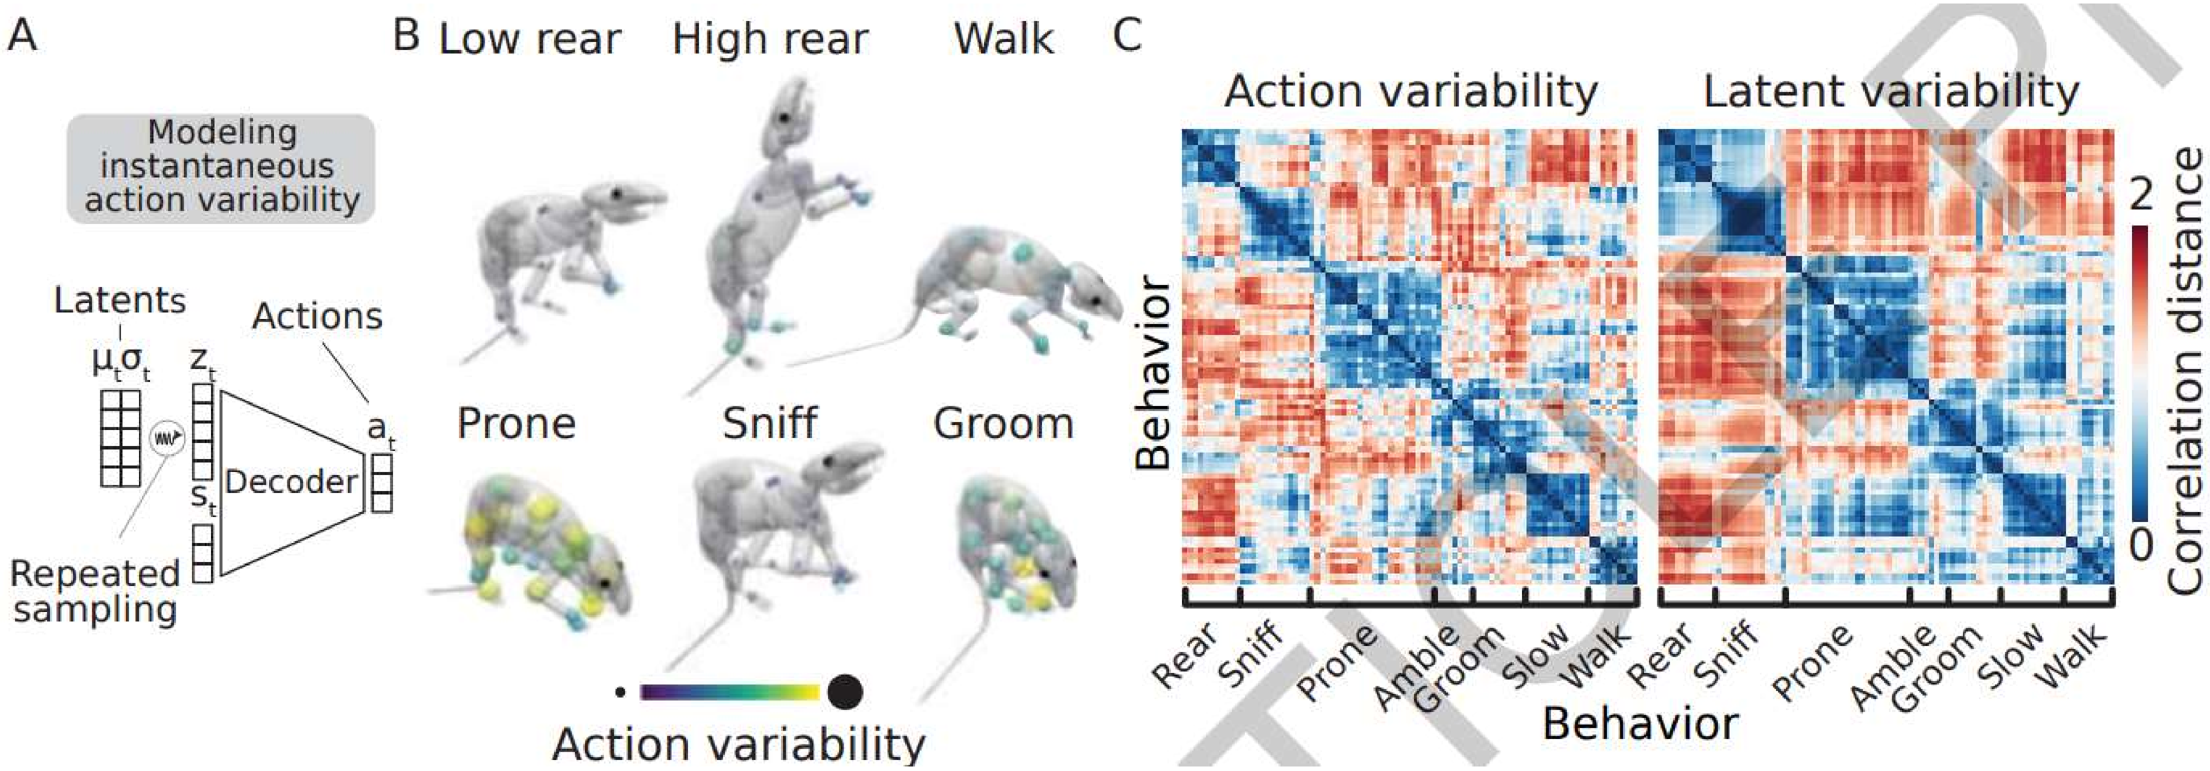

Variability calculation method: Resampling the latent space 50 times at every step to obtain instantaneous action variability and latent variability

Action Variability & Latent Variability

D. Trajectories of six latent dimensions along which variability was differentially regulated across behavior.

E. The population latent variability discriminates behaviors.

F-H. The deviations from normal variability structure <u>reduce</u> the model's <u>robustness</u> to noise.

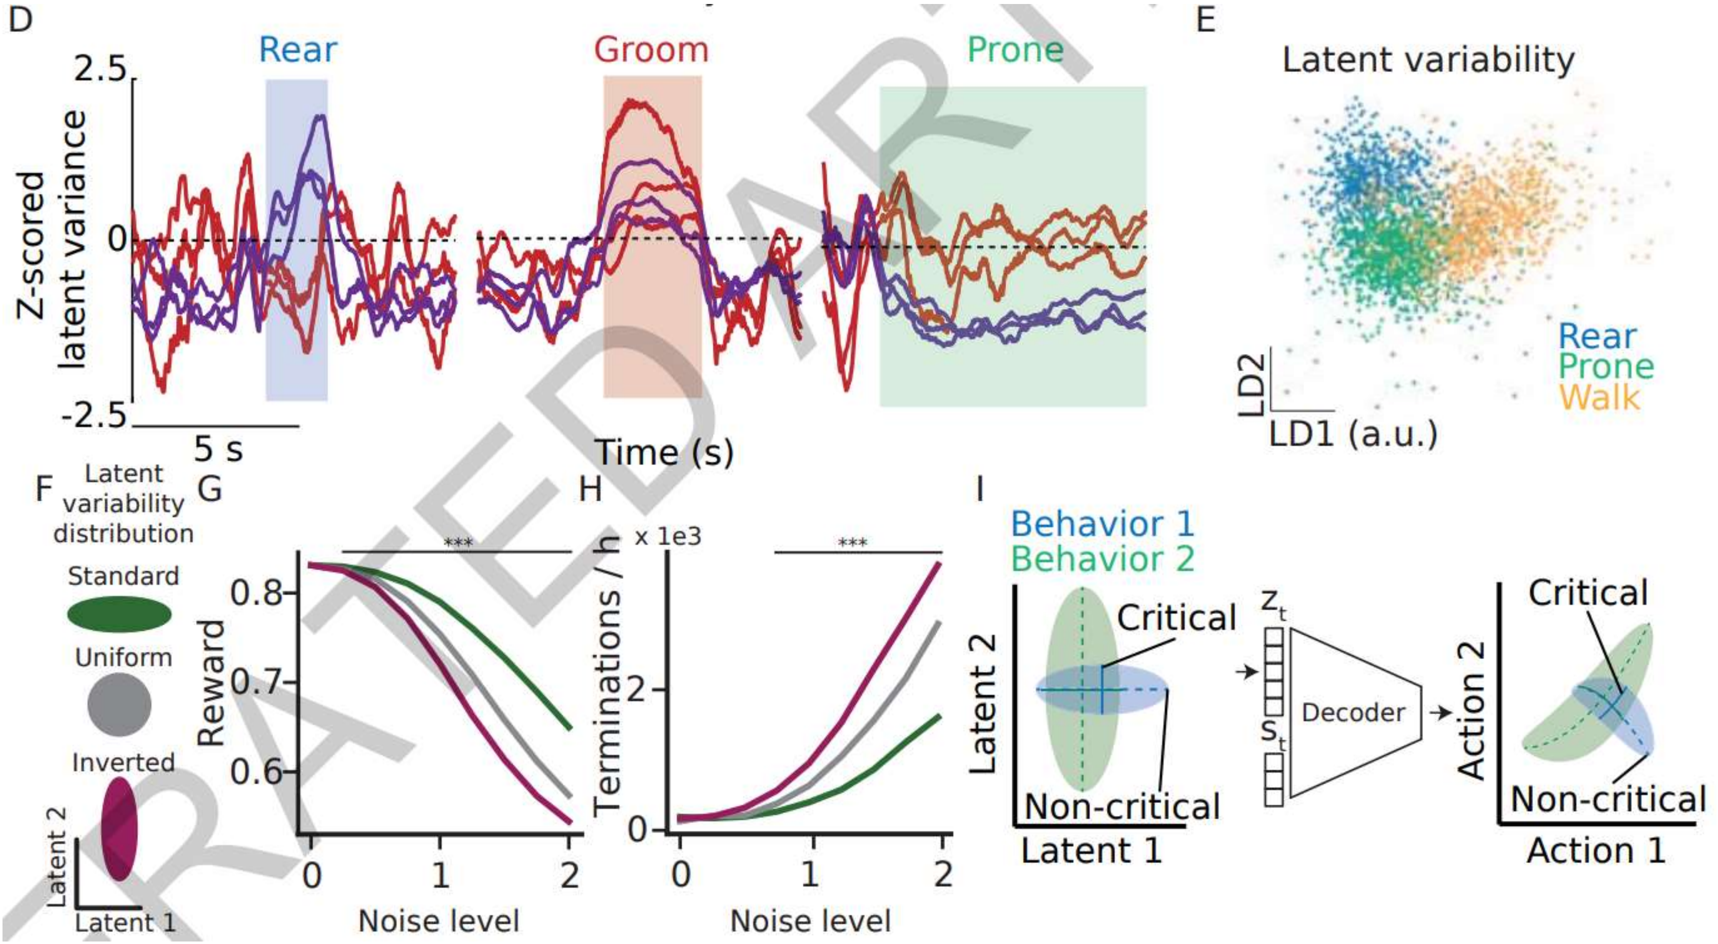

I. The virtual rodent model is able to adaptively adjust the variability in the latent space according to different behavioral demands (Fig.I-right). Different behaviors may require precise control in different latent dimensions (Fig.I-left). This is reflected in the reduction of variability in critical dimensions in the latent space.

Overview

Bridging The Relationship between Neural Activity and Behavior via AI Agent

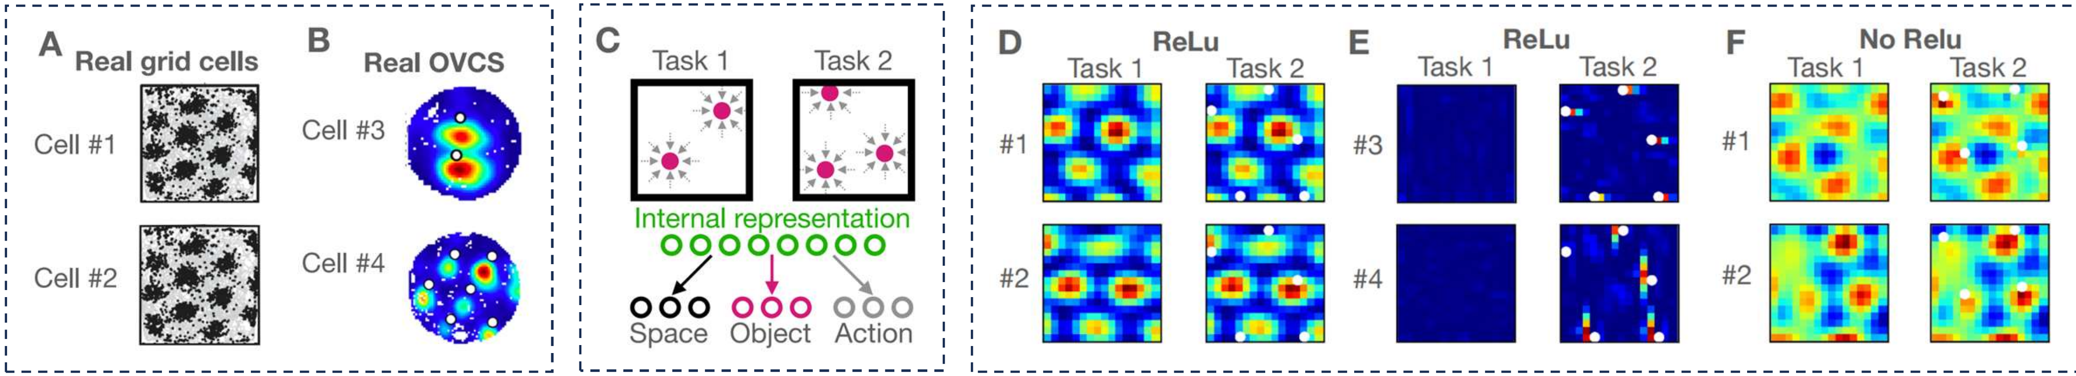

Whittington et al. (2023). ICLR.

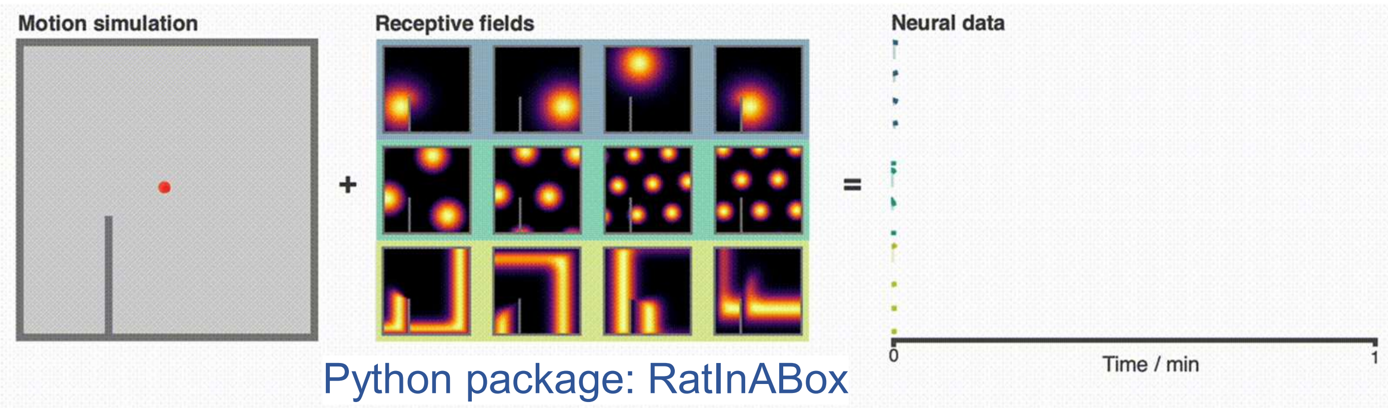

User-generated Environment

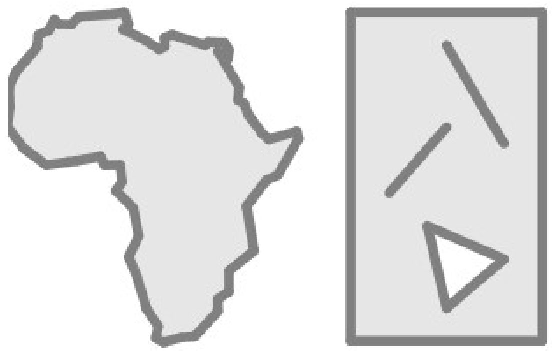

George et al. (2024). Elife.

George, T. M., Rastogi, M., de Cothi, W., Clopath, C., Stachenfeld, K., & Barry, C. (2024). RatInABox, a toolkit for modelling locomotion and neuronal activity in continuous environments. *Elife*, *13*, e85274.

Whittington, J. C., Dorrell, W., Ganguli, S., & Behrens, T. (2023). Disentanglement with biological constraints: A theory of functional cell types. In *The Eleventh International Conference on Learning Representations*.

Bridging Neural Dynamics and Behaviour via Virtual Insect

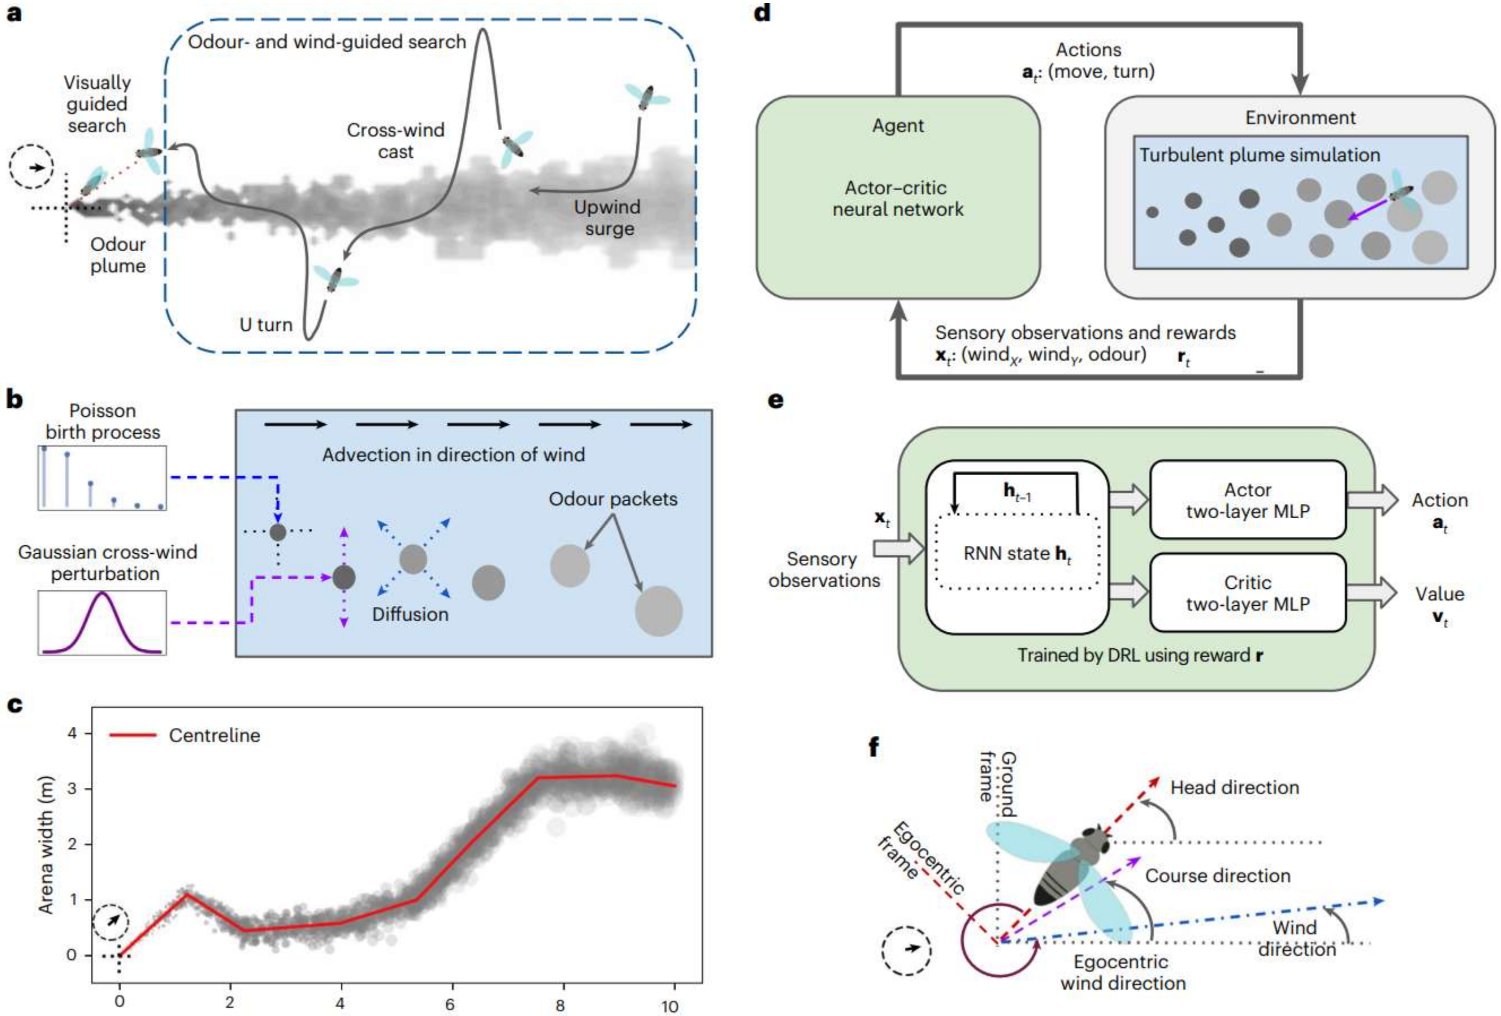

Goal: use deep reinforcement learning to train **RNN** agents to <u>locate the source of simulated odor plumes</u> in changing wind conditions.

Input: wind velocity [wind-X, wind-Y] and local odor concentration.
Output: move / turn

Results:
1. Well-trained RNN exhibited similar behaviors compared to those of real-world flying insects tracking odor plumes:
- Upwind surges when detecting odor
- Cross-wind casts and U-turns when losing the odor trail
- Different behavioral modules for tracking, recovering, and being lost

2. Long timescale memory is crucial for RNN to track odor plumes

Singh, S. H., van Breugel, F., Rao, R. P., & Brunton, B. W. (2023). Emergent behaviour and neural dynamics in artificial agents tracking odour plumes. *Nature machine intelligence*, *5*(1), 58-70.

Other virtual animals: Virtual Drosophila Simulator

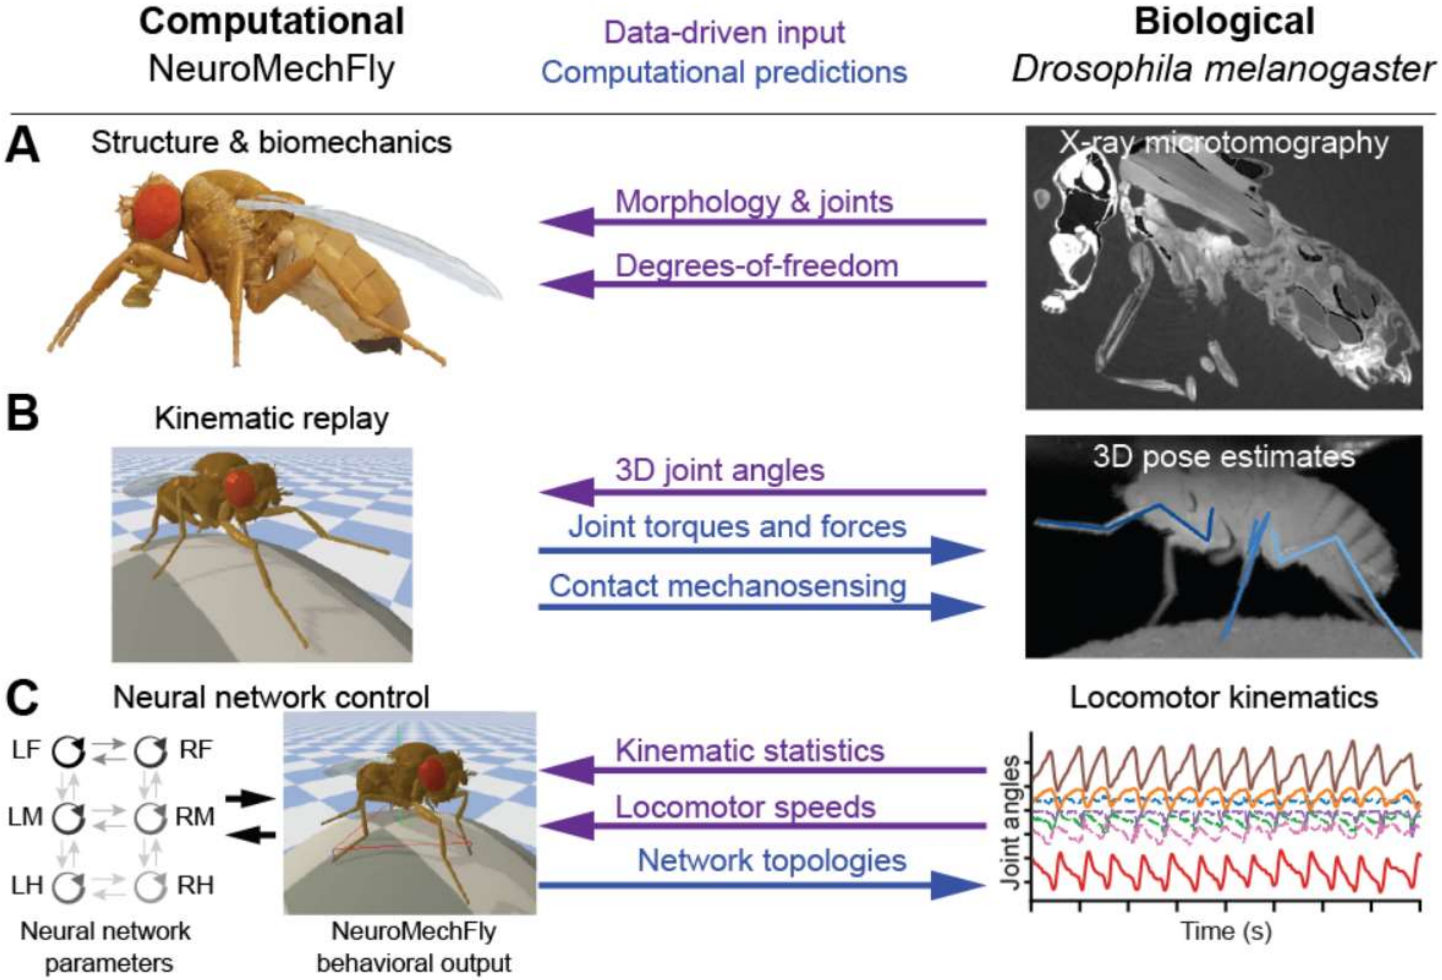

Overview:
NeuroMechFly is the first <u>comprehensive</u>, <u>morphologically</u> accurate neuromechanical simulation framework for the adult drosophila.

Simulation framework:
Obtaining <u>tactile information</u> when fruit flies come into contact with the environment through dynamic replay and inverse dynamics calculation.

Possible future work:
Construct a neural network as controller which take <u>tactile information as input</u>, and <u>output the kinematic statistics for motion</u>. This might help us to explore the <u>network dynamics behind environmental stimulus and behaviors</u>.

Lobato-Rios, V., Ramalingasetty, S. T., Özdil, P. G., Arreguit, J., Ijspeert, A. J., & Ramdya, P. (2022). NeuroMechFly, a neuromechanical model of adult Drosophila melanogaster. *Nature Methods*, 19(5), 620-627.

Thanks for your attention!