Fudan University
Institute of Science and Technology
for Brain-inspired Intelligence

NEURAL INFORMATION PROCESSING SYSTEMS

CAML

NeurIPS 2026

A Theory of Continuous Attractor Manifold Learning for Spatial Cognition

Chong Li Xiangyang Xue Jianfeng Feng Taiping Zeng

BIRLab
Brain-inspired Intelligent Robotics Lab

https://birlabs.github.io/CAML/

Stable spatial codes must also track movement

a **Classical CANN**

stable bumps ✓

hard to design ✗

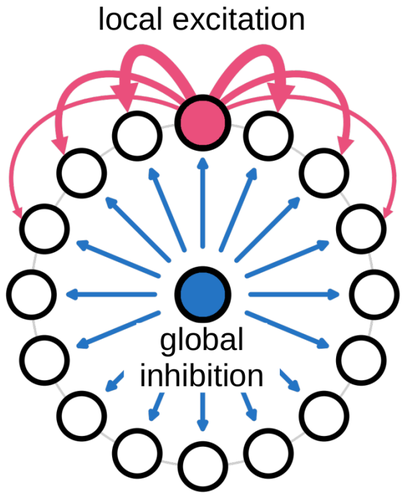

b **Task-trained RNN**

emergent spatial code ✓

no stability guarantee ✗

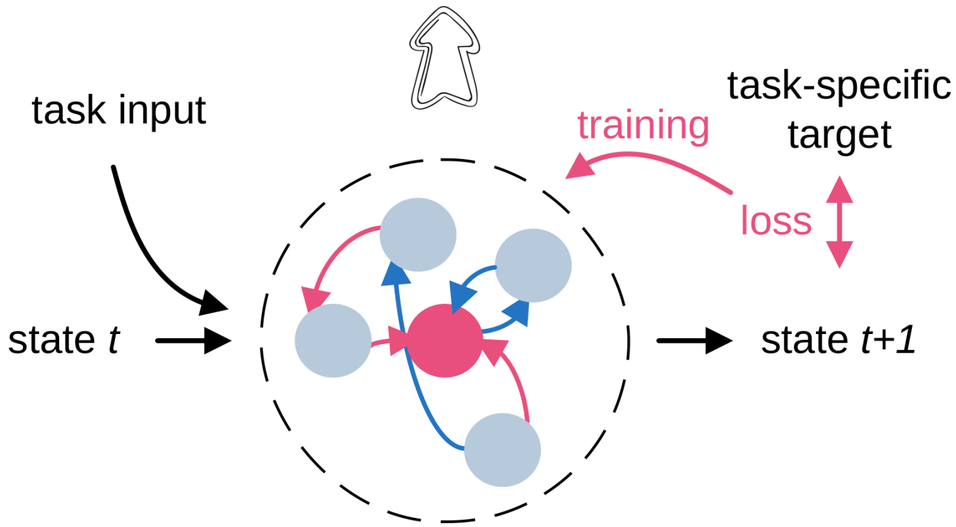

c **Classical Hopfield**

energy stability ✓

isolated attractors ✗

$$\mathbf{E}_{\mathbf{Hopfield}}(\mathbf{v}) = -\frac{\mathbf{1}}{\mathbf{2}}\mathbf{v}^{\top}\mathbf{W}\mathbf{v} - \mathbf{b}^{\top}\mathbf{v} + \sum_{\mathbf{i}}\int_{\mathbf{0}}^{\mathbf{v_i}} \mathbf{g}^{-\mathbf{1}}(\mathbf{s})\,\mathbf{ds}$$

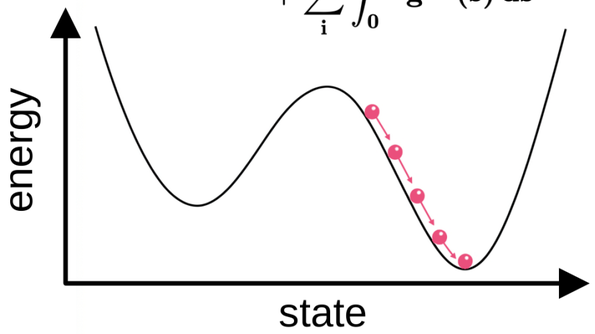

energy decreases toward point attractors

CAML learns the dynamics of a supplied population manifold

A continuous set of stable states, with motion driven by velocity inputs.

Learning restoration and bump shifting

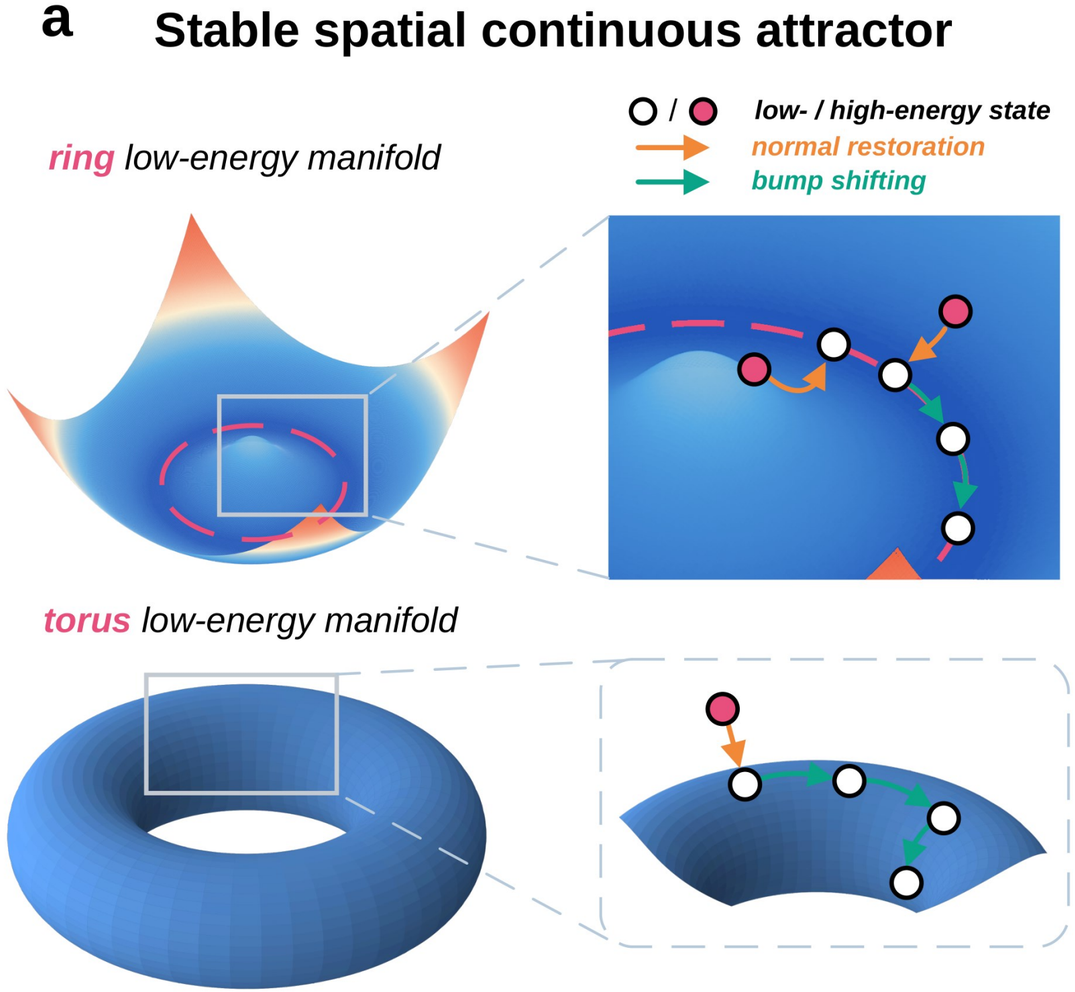

$$\dot{\mathbf{h}} = \underbrace{-\nabla_{\mathbf{h}} E_{\theta}(\mathbf{h})}_{\text{normal restoration}} + \underbrace{\mathbf{B}(\mathbf{h}, \mathbf{u})}_{\text{bump shifting}}$$

1 Learn normal restoration

Learn an energy valley that draws noisy activity back toward the manifold.

2 Learn bump shifting

Learn how velocity moves activity along the manifold to update direction or position.

Energy descent holds for autonomous restoration.

Restoration toward ring and torus manifolds

Perturbed population states evolve under the learned energy gradient.

Head-direction ring

2,000 population states

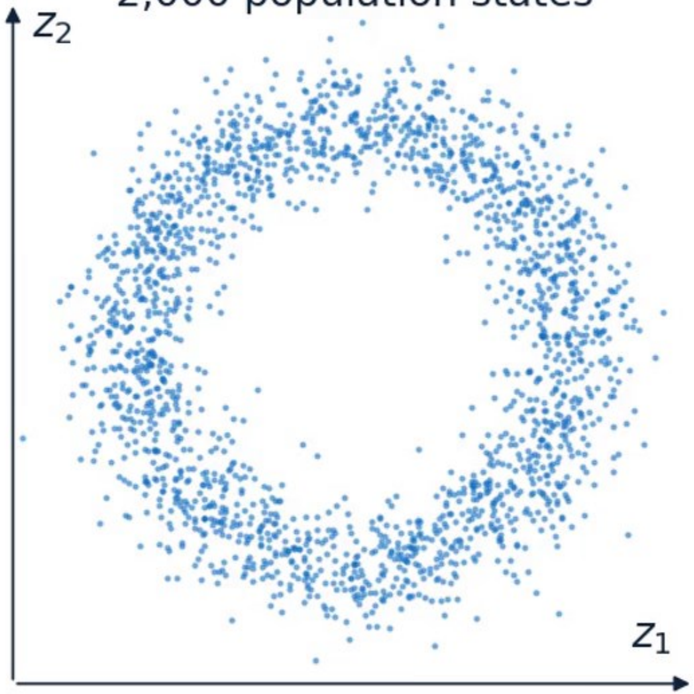

t = 0.00 s

Grid-cell torus

2,000 population states

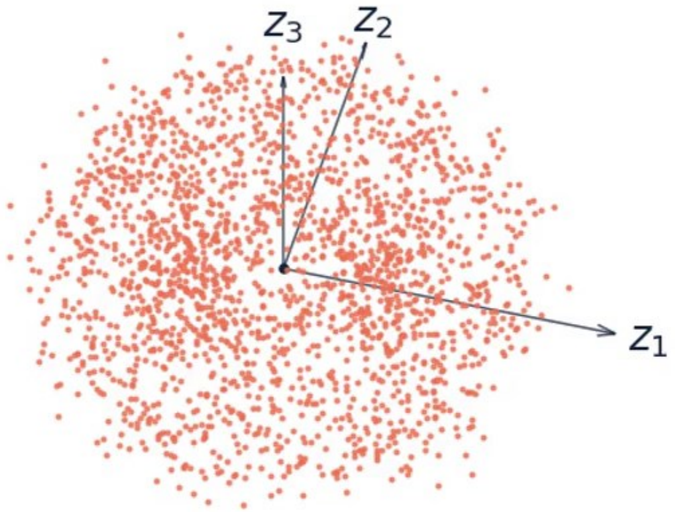

t = 0.00 s

Each point represents one neural population state.

Rotating view

Head-direction bump recovery

The energy gradient restores perturbed population activity toward the manifold.

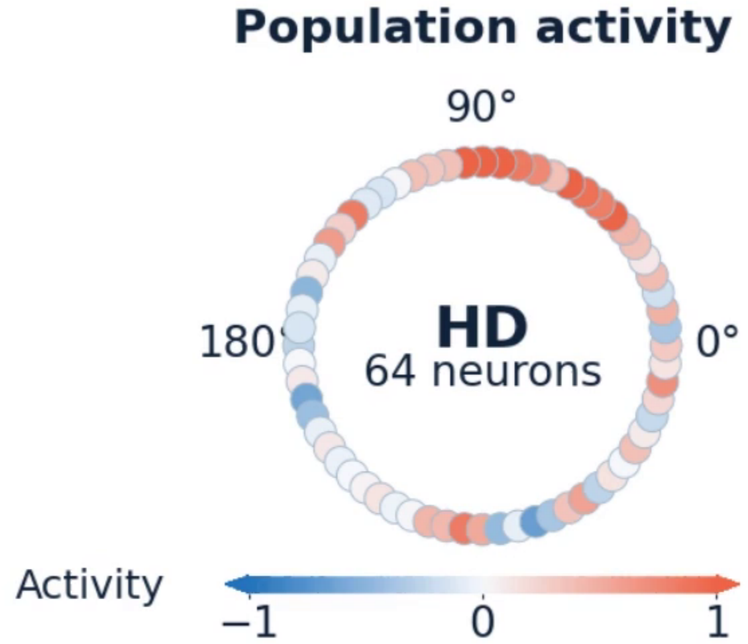

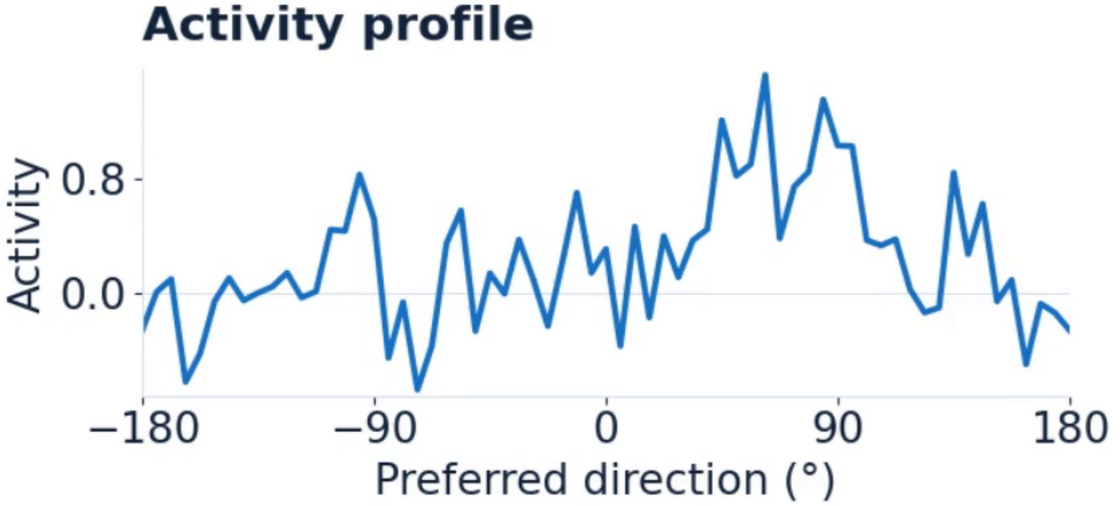

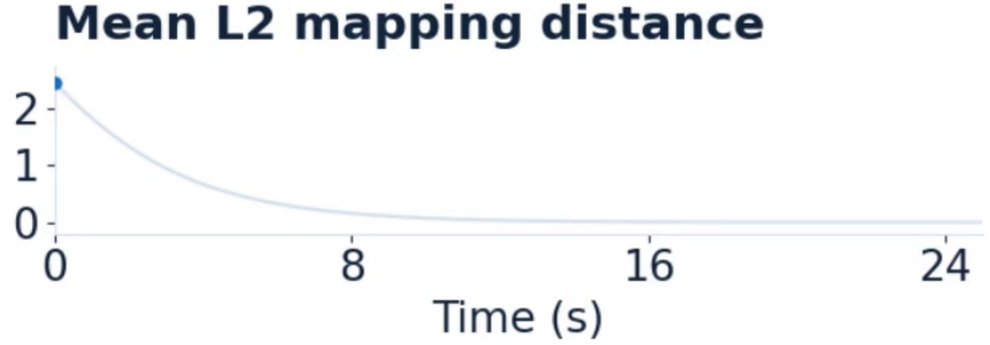

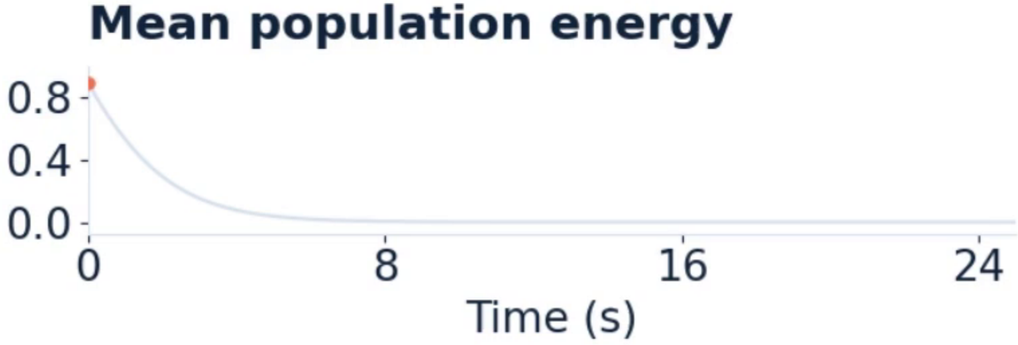

Recovery across noise levels

Normalized population RMSE, mean ± SD across test states. Lower is better.

Population	Noise σ	CAML	DMHN-derived	MHN	CANN†
HD	0.1	**0.014 ± 0.009**	0.474 ± 0.046	1.093 ± 0.001	0.120 ± 0.010
HD	0.5	**0.051 ± 0.038**	0.482 ± 0.060	1.093 ± 0.001	0.155 ± 0.054
Grid	0.1	**0.078 ± 0.032**	1.000 ± 0.000	0.995 ± 0.000	0.153 ± 0.114
Grid	0.5	**0.079 ± 0.032**	1.000 ± 0.000	0.995 ± 0.000	0.156 ± 0.116

Lower recovery error across both populations and noise levels

† CANN uses its own reference states. Other methods use prescribed codes.

Angular velocity drives head-direction activity

Learned bump shifting moves activity along the manifold to track self-motion.

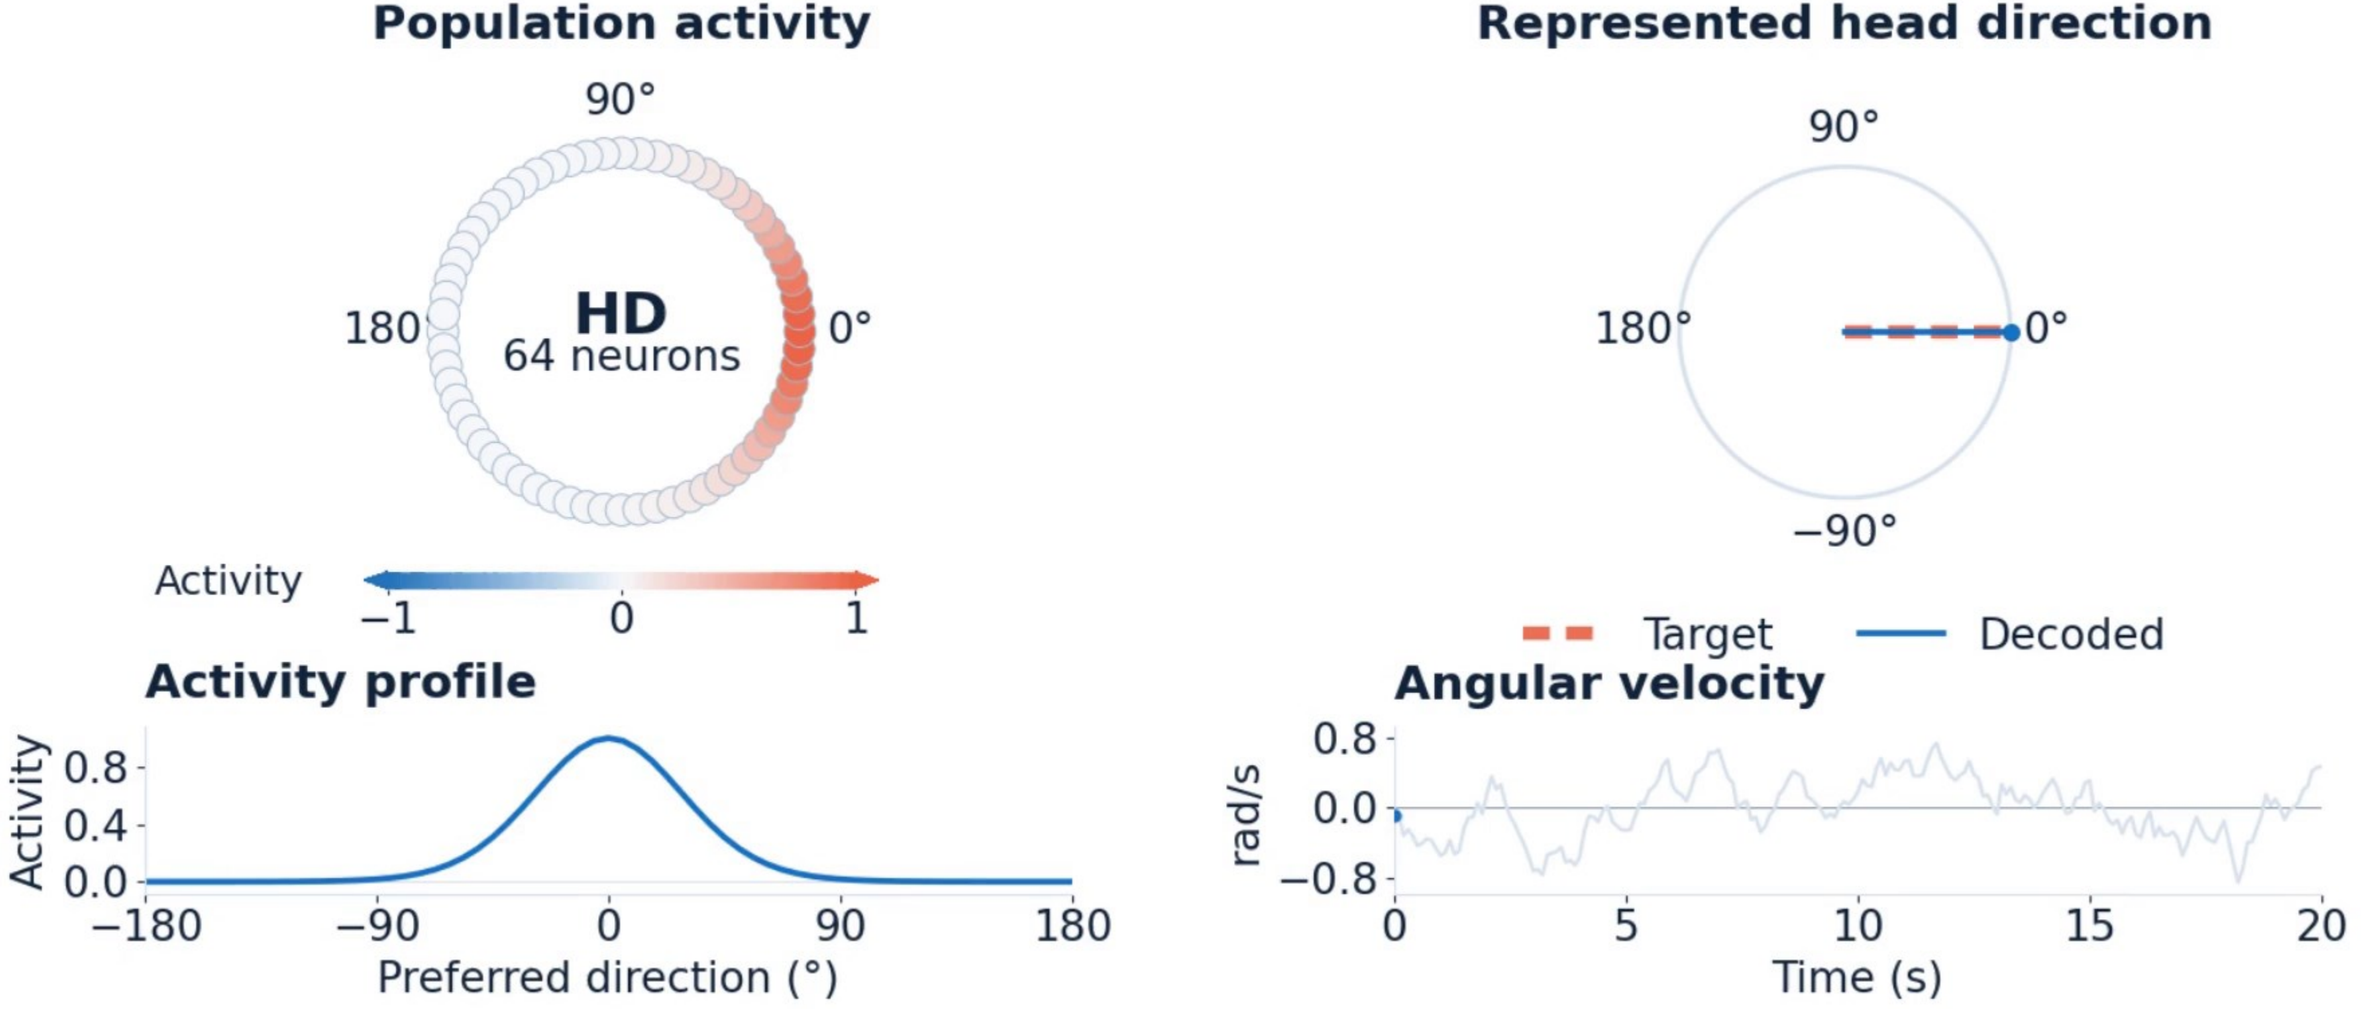

Learned shifting improves Grid and Place tracking

Tracking error, mean ± SD over 100 paired trajectories

Population	Velocity setting	Learned shifting	Fixed shifting
HD	0.5 rad/s	0.752 ± 0.375°	**0.654 ± 0.342°**
HD	1.0 rad/s	**1.425 ± 0.653°**	1.557 ± 0.825°
Grid	0.5 rad/s	**0.093 ± 0.037°**	0.715 ± 0.287°
Grid	1.0 rad/s	**0.157 ± 0.050°**	1.968 ± 0.731°
Place	Random walk	**3.9 ± 2.3 mm**	67.1 ± 70.0 mm

87–92% lower mean Grid tracking error

Grid velocity limits apply per axis. Grid errors average both phase axes.

Local excitation and surround inhibition emerge

Effective interactions from the learned restoring field

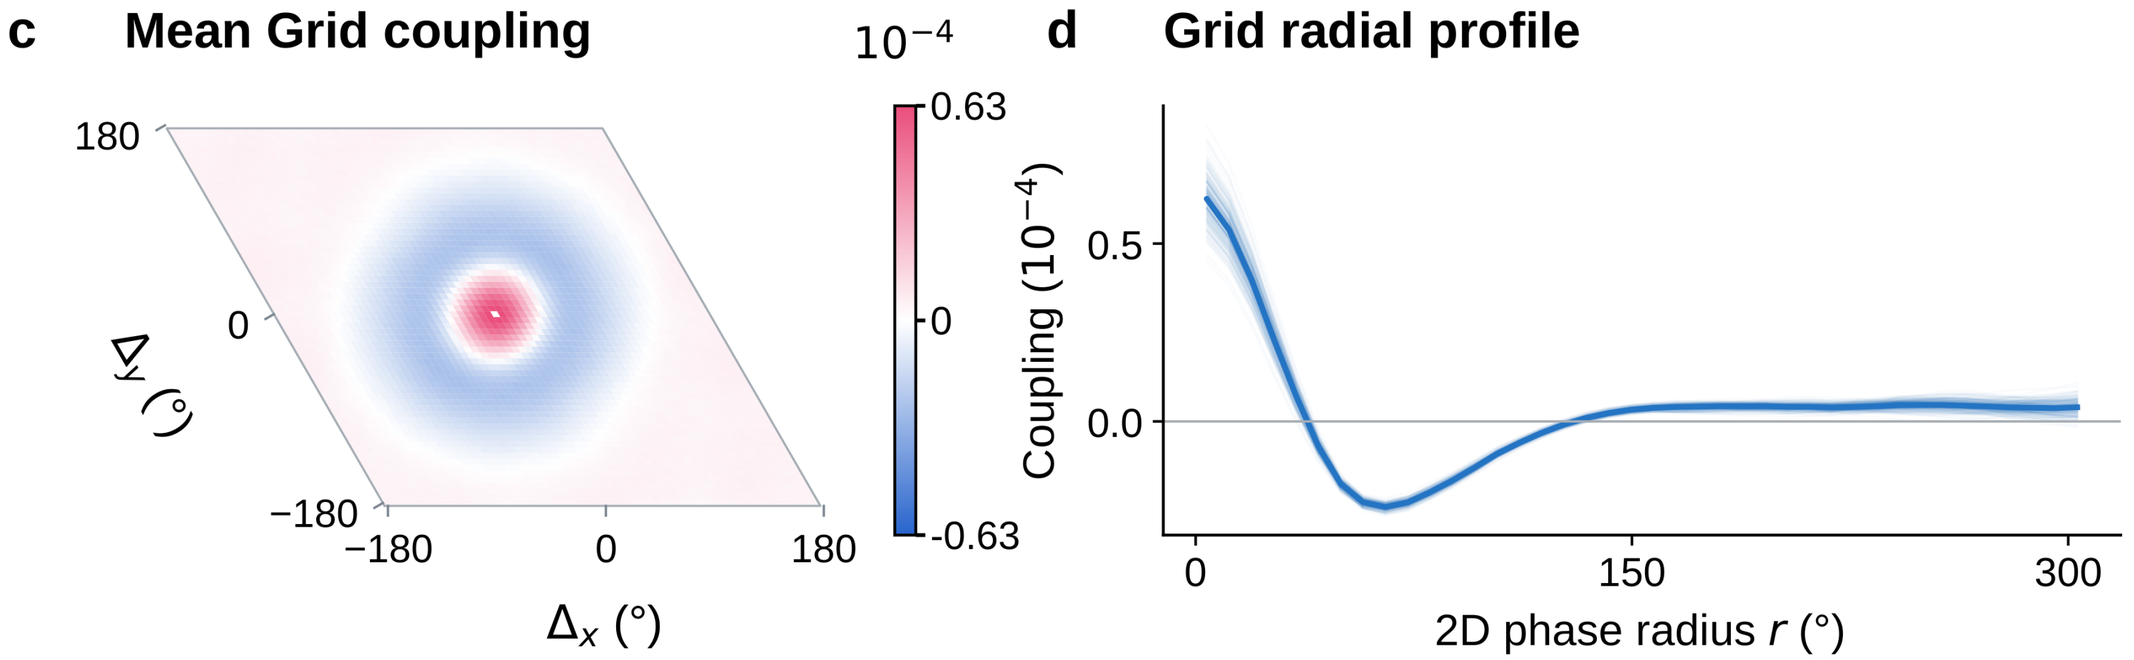

The learned energy induces structured coupling without a hand-designed kernel.

From spatial codes to learned attractor dynamics

CAML learns restoration and motion together

- Energy descent provides a restoring mechanism
- Learned bump shifting supports spatial tracking
- Interpretable neural interactions emerge from learning

Paper, demonstrations, and code

birlabs.github.io/CAML

Thank you