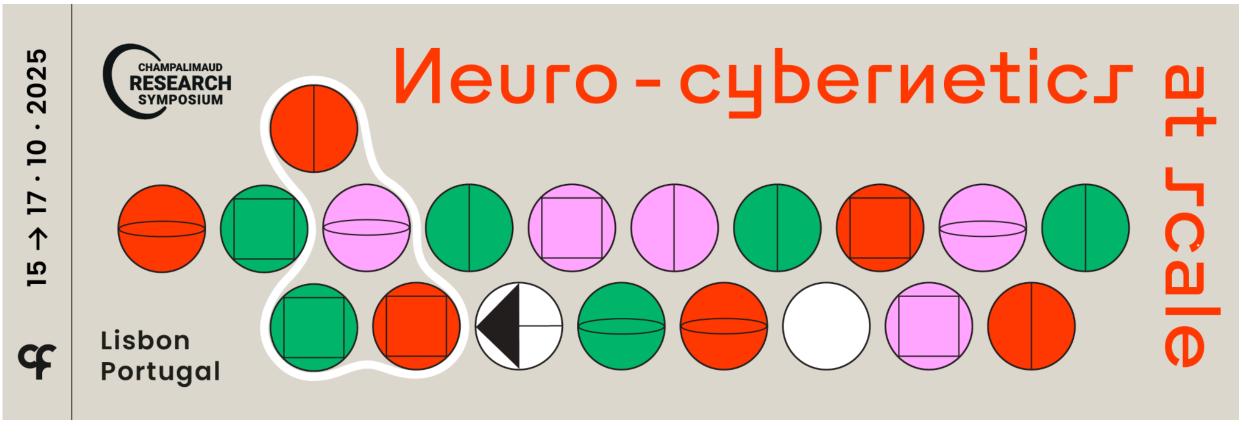

Planning to Learn: Active Inference of Latent Dynamics with Fisher Information

Hyungju Jeon 1, Isabel Rosário 1,2, Il Memming Park 1
1. Champalimaud Foundation, Lisbon, PT
2. Instituto Superior Técnico, Lisbon, PT

Understanding the latent dynamical structure that governs behavior and neural activity is a central challenge in computational neuroscience and many complex system domains. However, performing joint state and system identification for high-dimensional nonlinear dynamics from partial, noisy observations remains data-intensive and often sample-inefficient under naïve data acquisition strategies.
We propose a computationally efficient active learning framework for inferring latent dynamics models that plans optimized input sequences to maximize future information gain. Our method combines local Fisher Information metrics, which quantify parameter sensitivity in probabilistic latent dynamics, with predictive uncertainty reduction strategies derived from ensemble disagreement or posterior entropy. To address the compounding effect of early model inaccuracies, we introduce discount factors that regularize information gathered under high epistemic uncertainty in the latent posterior, ensuring robustness in belief updates over time. We validate our approach through computational experiments in various nonlinear dynamics environments from classical control, and simulated neural population dynamics, showing improved sample efficiency over myopic and uncertainty-driven baselines. Our method offers a flexible and general framework for structured system identification and active experimental design in partially observed dynamical systems.

neural dynamics, active learning, fisher information, state-space models