PHYSICAL MATH SEMINAR

Interpretable representations of non-linear (neural) dynamics using geometric deep learning

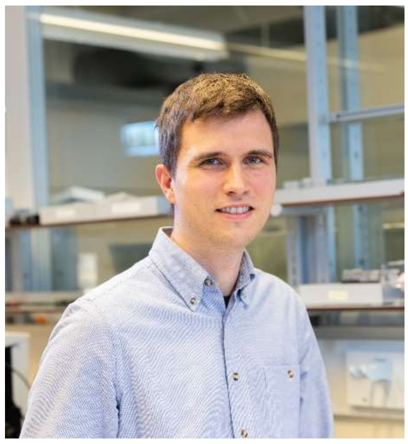

Adam Gosztolai

Medical University Vienna

ABSTRACT:

Large-scale neural population recordings permit unprecedented access to the dynamical processes that comprise neural computations. Yet, typical methods for interpreting these recordings through low-dimensional representations, such as PCA or UMAP do not explicitly model the underlying dynamics. Meanwhile, other methods that learn the dynamics in individual trials, such as recurrent neural networks, do not model the spatial context of these trials within the full dynamical landscape sampled by all trials across experimental conditions. Recent representation learning methods, such as CEBRA, have introduced user supervision with behavioural features — a 'dynamical template' to map neural states onto — to learn consistent latent representations across animals and systems. Yet, without a well-chosen supervision feature, low-dimensional representations are generally less interpretable and cannot act as a similarity metric across animals and experiments. Recent evidence suggests that task-relevant neural activity takes place on low-dimensional subspaces of the state space called neural manifolds, which provides an additional inductive bias for learning algorithms. In this talk, I will present a theoretical framework called MARBLE that leverages the manifold structure to represent non-linear dynamical systems from sparse time series data. Our approach relies on a geometry-informed decomposition of neural trajectories into local flow fields and then mapping them jointly into a shared latent space using geometric deep learning. I will show that MARBLE provides latent representations that give rise to a well-defined similarity metric between different neural systems to compare computations and detect fine-grained changes in dynamics due to task variables, e.g., decision thresholds and gain modulation. Being unsupervised, MARBLE is uniquely suited to biological discovery. Yet, I will show that it can discover more interpretable neural representations in several motor, navigation and cognitive tasks than generative models such as LFADS or (semi-)supervised models such as CEBRA. Intriguingly, this interpretability implies significantly higher decoding performance than state-of-the-art. Our results suggest that using the manifold structure yields a new class of algorithms with higher performance and the ability to assimilate data across experiments.

TUESDAY, NOVEMBER 12, 2024

2:30 PM – 3:30 PM

Building 2, Room 449

MIT MATHEMATICS

https://math.mit.edu/pms/