A quantitative model of conserved macroscopic dynamics predicts future motor commands

Brennan, Proekt ‘19

Geelon So, agso@eng.ucsd.edu
BGGN246 Week 7 — February 19, 2021

eLIFE
elifesciences.org
RESEARCH ARTICLE

A quantitative model of conserved macroscopic dynamics predicts future motor commands

Connor Brennan[1†], Alexander Proekt[2†*]

[1]Departmentof Neuroscience, University of Pennsylvania, Philadelphia, United States; [2]Department of Anesthesiology and Critical Care, University of Pennsylvania, Philadelphia, United States

Abstract In simple organisms such as *Caenorhabditis elegans*, whole brain imaging has been performed. Here, we use such recordings to model the nervous system. Our model uses neuronal activity to predict expected time of future motor commands up to 30 s prior to the event. These motor commands control locomotion. Predictions are valid for individuals not used in model construction. The model predicts dwell time statistics, sequences of motor commands and individual neuron activation. To develop this model, we extracted loops spanned by neuronal activity in phase space using novel methodology. The model uses only two variables: the identity of the loop and the phase along it. Current values of these macroscopic variables predict future neuronal activity. Remarkably, our model based on macroscopic variables succeeds despite consistent inter-individual differences in neuronal activation. Thus, our analytical framework reconciles consistent individual differences in neuronal activation with macroscopic dynamics that operate universally across individuals.

DOI: https://doi.org/10.7554/eLife.46814.001

Background

What we know

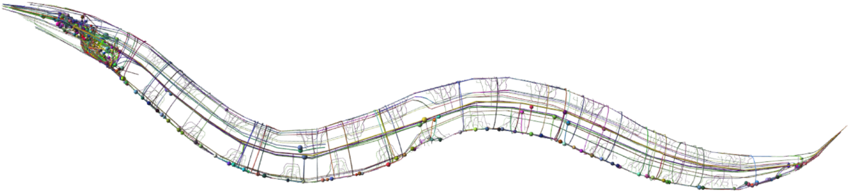

- all 302 neurons in *C. elegans* and their connections are known
- certain neurons can be individually identified across animals
- we can simultaneously record activity of a large fraction of neurons *in vivo*

Neuronal activity to locomotor behavior

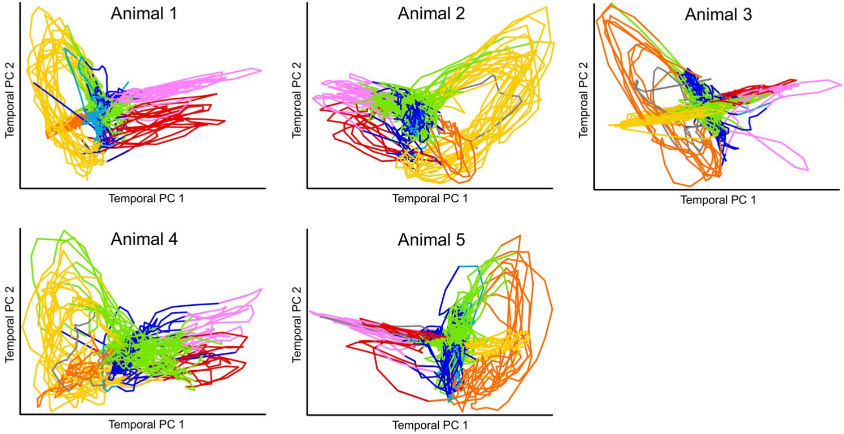

Figure 1: Neuronal activity form loops corresponding to a different behavior. However, microscopic dynamics differ across individuals.

Conservation of activities across individuals?

The microscopic activity of an individual can be reliably mapped to certain locomotion behaviors, allowing us to **predict** the behavior of the animal.

- These microscopic activities **are not conserved** across individuals, so microscopic activity cannot be directly used to predict behavior across animals.
- **Question:** the loops formed appear qualitatively similar across animals. Can we study the **macroscopic dynamics**?

Biological scales

- **Neuronal activity** (microscopic): activity of individual neurons, synapses, etc.
- **Neuronal dynamics** (macroscopic): laws of motion that govern neuronal activity
- **Locomotion behavior** (organism): behavior observed at the organism level

Microscopic v. macroscopic dynamics

> “Many disparate microscopic configurations lead to almost **indistinguishable macroscopic behavior**.
>
> *Brennan and Proekt (2019)*

Problem statement

Hypothesis: evolutionary pressure selects at the macroscopic (and not at the microscopic) level because any microscopic dynamics that generates the right macroscopic dynamics yields the same behavior.

- **Question:** can we use the observed neuronal activity to deduce neuronal dynamics? What are the relevant variables to describe neuronal dynamics?

Key results of paper

1. Neuronal activity across genetically-identical organisms are consistently different.
2. A new method is developed (asymmetric diffusion map modelling) to construct neuronal dynamics from neuronal activity.
3. The constructed neuronal dynamics is conserved across individuals.
4. Prediction of behavior is possible using neuronal dynamics across individuals.

Roadmap

1. Variable activation of identified neurons in *C. elegans*
2. Underlying neuronal dynamics give rise to neuronal activity
3. Simulations of neuronal dynamics predict behavioral switches
4. Macroscopic dynamics are conserved among animals
5. Theoretical explanations from stochastic differential equations
6. Asymmetric diffusion map modeling

Variable activation of identified neurons in *C. elegans*

Is neuronal activity different across individuals?

Experiment: 100 neurons simultaneously recorded in 5 animals (immobilized)[1]

- certain neurons closely associated to locomotion used to infer motor commands
    - fictive locomotor behavior inferred for each animal
- 15 neurons could be consistently and unequivocally identified in each individual
    - e.g. AIB, ALA, AVA, AVB, RIM, RME, etc.

Goal: quantify and predict motor commands across individuals

[1] Data collected by (Kato et al., 2015).

Neuronal activity

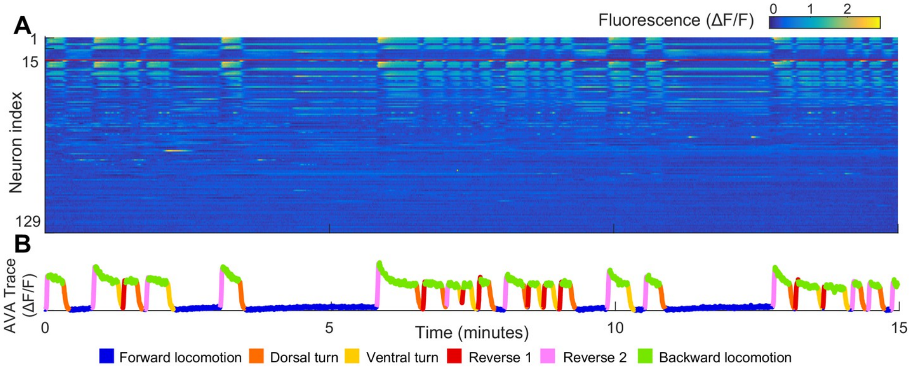

Figure 2: (A) Calcium signals for one animal for $\sim$15 min. (B) Activity for AVA neuron colored by behavioral state.

Variability of neuronal activity

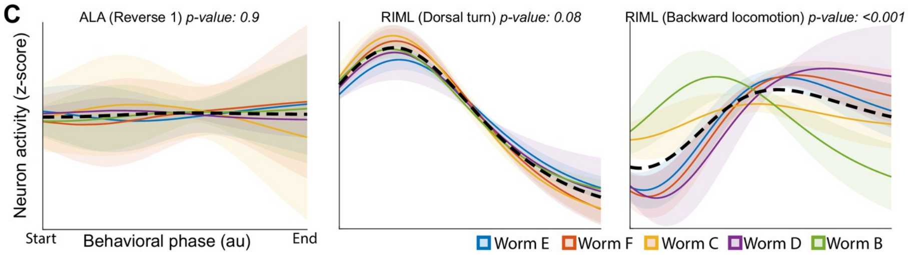

Figure 3: (C) Average neuron activity (ALA and RIML) during specific locomotion behavior (reverse, dorsal turn, backward locomotion). Each color corresponds to one of five animals, showing 95% confidence interval.

Neuronal activity and behavior across animals

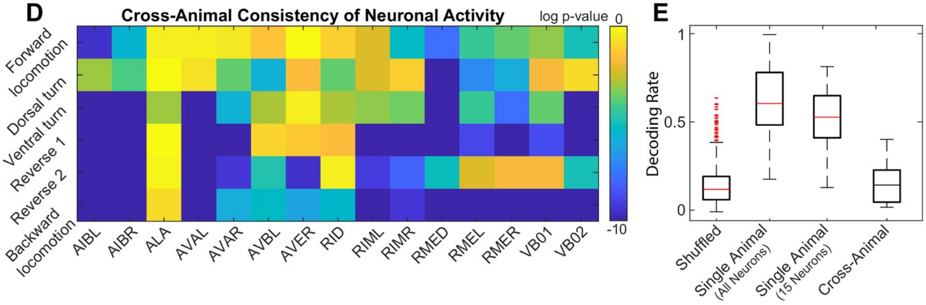

Figure 4: (D) Probabilities that neuronal activity from different individuals drawn from the same distribution (darker ⇒ distributions likely different). (E) Decoding rate for backward locomotion from neuronal activity.

Summary

The **same behavior** across individuals arise from **different firing patterns**.

- Perhaps the brains of different individuals are essentially incommensurable, in which case we really just need to study each brain in isolation.
- Or, there is some sort of equivalence between individuals, in which case we would like to recover the equivalence from the neuronal activity.

Underlying neuronal dynamics give rise to neuronal activity

Configurations of neural networks

A spiking neural network with 100 neurons has 2^{100} configurations. However, even in a much more complex biological system, we might expect that:

- most configurations never actually occur
- the transition from one configuration to another is not arbitrary

Neuronal dynamics

The laws of motions of a neural network can be mathematically specified by:

- a **phase space**: the set of all biologically possible configurations
- a **time evolution function**: a description of how the configuration of the neural network changes over time

Asymmetric diffusion map modeling

This paper introduces asymmetric diffusion map modeling, which is a method to **construct a neuronal dynamics** from observed neuronal activity data.

- Are the dynamics recovered by this method a salient representation?
- Do the dynamics generalize across individuals?

Reconstructed neuronal dynamics

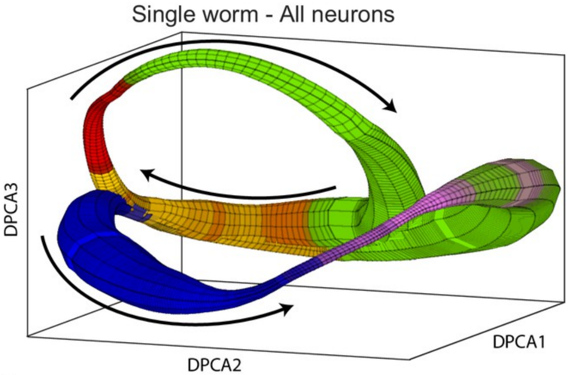

Figure 5: Visualization of a phase space constructed by asymmetric diffusion map modeling, colored by locomotion behavior.

Simulated neuronal activity

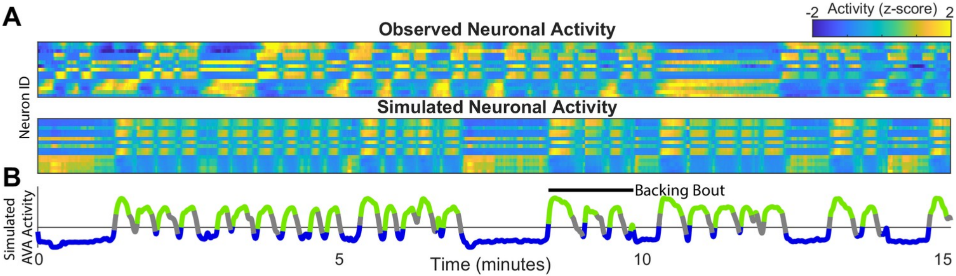

Figure 6: (A) Experimentally observed (top) and simulated (bottom) activity of 15 shared neurons. (B) Simulated trace of AVA neuron colored by behavior (blue/green are forward/backward locomotion).

Neuronal activity trace spectral residues

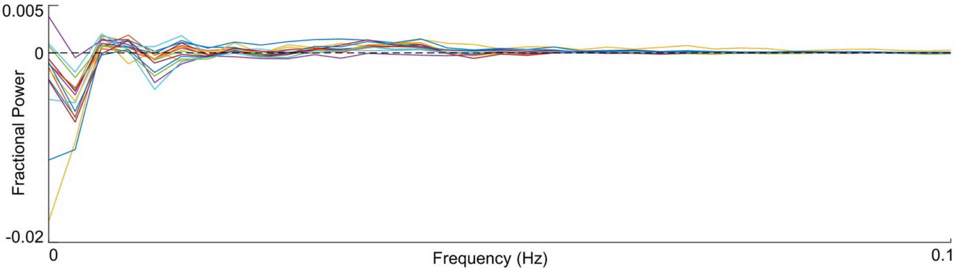

Figure 7: Difference in the fractional power of the spectrum of the observed neuron and its simulation. The spectra over .05 Hz of simulated neurons are statistically indistinguishable from experimentally observed activity.

Dwell time distributions

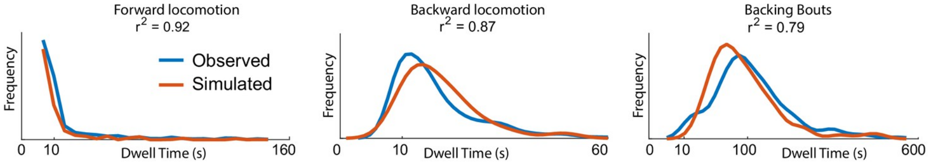

Figure 8: Amount of time spent in different behavioral states.

Summary

The neuronal dynamics constructed by asymmetric diffusion map modeling can yield **simulations of neuronal activity**.

- Simulated and observed neuronal activity have **various matching statistics**.

Simulations of neuronal dynamics predict behavioral switches

Behavioral switches

Switching between different modes of locomotion is stochastic, likely depending entirely on the dwell time within the current behavior.

- **Question:** can the neuronal dynamics model predict future changes in behavior?

Prediction setup

Experiment: constructed model used to predict time to behavior switch from forward to backward locomotion.

- Prediction task data for 11 new animals under similar conditions to model animals.[2]
- Compare prediction with dwell time statistics model:

$$\text{prediction} = \mathbb{E}[\text{time in forward motion}] - \mathbb{E}[\text{time to reach current state}].$$

[2] Prediction task data collected by (Nichols et al., 2017). Model data collected by (Kato et al., 2015).

Prediction on new animals

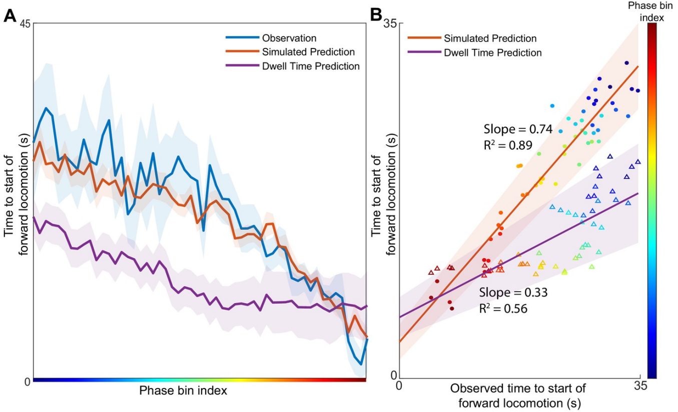

Figure 9: Comparing predictions based on simulated dynamics and behavioral dwell time statistics in the region where most backward locomotion terminates and forward locomotion begins.

Prediction using single neuron

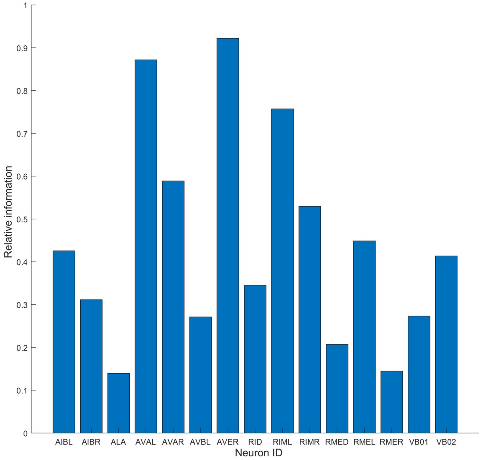

Figure 10: Dwell time prediction using single neuron. The quality of prediction varies depending on how coupled the neuron is to behavior.

Summary

Asymmetric diffusion map modeling constructs a neuronal dynamics that is able to predict behavior switches on new animals (data collected years apart).

- Prediction was based only on activity from 8–13 neurons.
- Prediction is possible based even only on one neuron.

Macroscopic dynamics are conserved among animals

Dynamics across animals

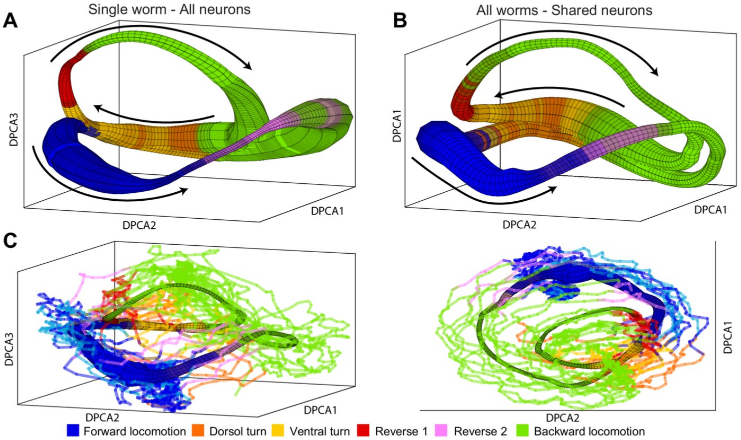

Figure 11: Visualization of phase space of neuronal dynamics. Fatter width ⇒ lower phase velocity (greater stochasticity); smaller width to greater determinism.

Summary

The neuronal dynamics are highly correlated to locomotion behaviors across animals.

- Where the loops overlap, dynamics tend to be dominated by stochasticity.
- Within the individual loops, the dynamics are more deterministic.

Theoretical explanations from stochastic differential equations

Topology of phase space

Question: why does the phase space to look like a superposition of closed loops?

Oscillatory behavior in neural networks

> A feature that has often been observed in neural networks is oscillation. The corresponding potential landscape shows a **closed-ring shape topology**.
>
> *Yan et al. (2013)*

Closed-ring topology

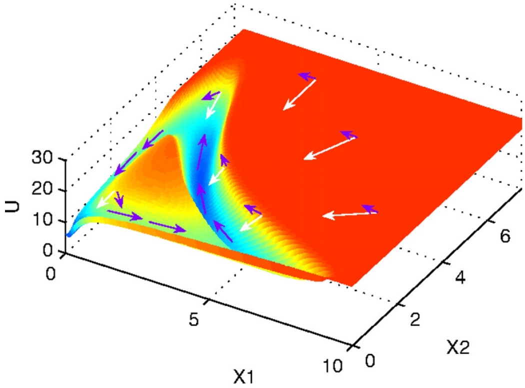

Figure 12: Potential landscape (white arrows show negative gradient) with a divergence-free flux (blue arrows), Wang et al. (2008).

Stochastic differential equation model

Let $\mathbf{X}(t) \in \mathbb{R}^d$ describe the neuronal activity at time t. Suppose there is some driving force $\mathbf{F}(\mathbf{X})$ corresponding to the deterministic aspect of the neuronal dynamics.

- The time-evolution of the system is described by:

$$\frac{d\mathbf{X}}{dt} = \mathbf{F}(\mathbf{X}) + \varepsilon,$$

where ε is independent noise introducing stochasticity into the system.

Evolution of probability mass

Define $P(\mathbf{X}, t)$ the probability that the system is in state $\mathbf{X}$ at time t. By the **law of conservation of probability mass**,

$$\frac{dP(\mathbf{X}, t)}{dt} = -\nabla \cdot \mathbf{J}(\mathbf{X}, t),$$

- $\mathbf{J}(\mathbf{X}, t) = \mathbf{F}(\mathbf{X})P - \mathbf{D}\nabla P$ is the probability flux, describing how much probability mass is passing through $\mathbf{X}$ at time t
- $\mathbf{D}$ is the diffusion coefficient describing how noisy the neural network is

Steady-state solution

If the neural network is at a steady-state solution (i.e. it is not learning), then $\frac{dP}{dt} = 0$, which implies that $\mathbf{J}$ is divergence-free:

$$\nabla \cdot \mathbf{J}(\mathbf{X}, t) = 0.$$

- Equilibrium systems have $\mathbf{J} = 0$.
- But generally, $\mathbf{J} = \nabla \times \mathbf{A}$ can be a **curl flux field**.[3]

[3] Generalized to higher dimensions by the Hodge decomposition.

Dynamics of nonequilibrium systems

> The dynamics of a nonequilibrium network spirals (from flux) down the gradient (from potential) instead of only following the gradient as in the equilibrium case.
>
> *Wang et al. (2008)*

Summary

We might expect **the configuration of a neural network** at a steady state (i.e. not learning) to **evolve along loops** in some phase space.

- Loops are 1-dimensional objects, with a single phase parameter θ.
- A reasonable representation of neuronal dynamics is the pair (α, θ)
  - α is the identity of the loop and θ its phase

Asymmetric diffusion map modeling

Constructing the phase space

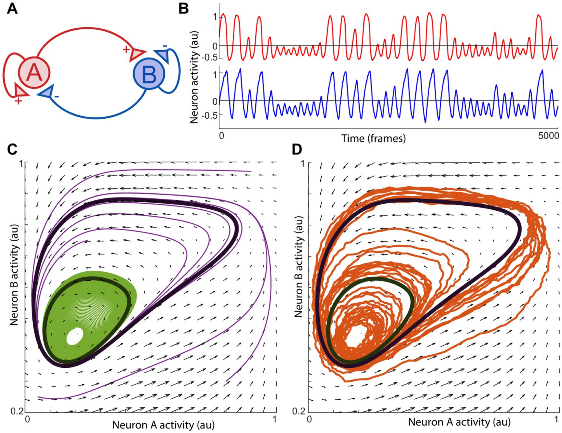

Figure 13: Takens' theorem states that the **delay embedding** can recover the intrinsic phase space of a dynamical system.

Manifold reconstruction

The delay embedding allows us to **recover a point cloud** that approximates the underlying phase space:

$$\text{neuronal activity} \mapsto \text{delay embedding}.$$

- **Question:** can we apply a *manifold learning* technique to recover the manifold? What about the dynamics on the phase space?

Diffusion map

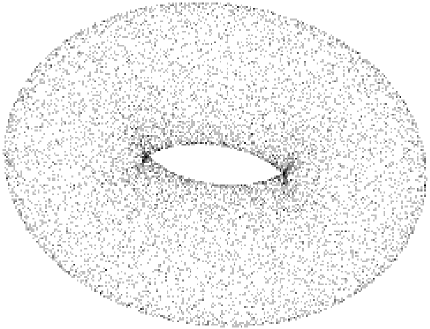

Figure 14: Diffusion map is a manifold learning technique to recover a low-dimensional representation from a high-dimensional point cloud (Nadler et al., 2005).

Diffusion map: main idea

Let $\mathbf{X} \in \mathbb{R}^{N \times D}$ be a point cloud of size N in D dimensions.

- Simulate a random walk on point cloud (think of as a diffusion process).
- Goal is to construct a low-dimensional representation $\widetilde{\mathbf{X}} \in \mathbb{R}^{N \times d}$ of point cloud
- Diffusion map is a technique that constructs $\widetilde{\mathbf{X}}$ so that the distance between $\widetilde{\mathbf{X}}_i$ and $\widetilde{\mathbf{X}}_j$ are related to *how long it would take a diffusion process starting at i to reach j.*

Asymmetric diffusion map

We don't merely have a point cloud—we have **time series data**. We can reconstruct the manifold *while learning the underlying dynamics*.

- Simulate a random walk with drift on the point cloud
  - Use dynamics data to obtain drift, use local distances to obtain diffusion
- Obtain both a low-dimensional phase space and a time-evolution law.

Overall pipeline

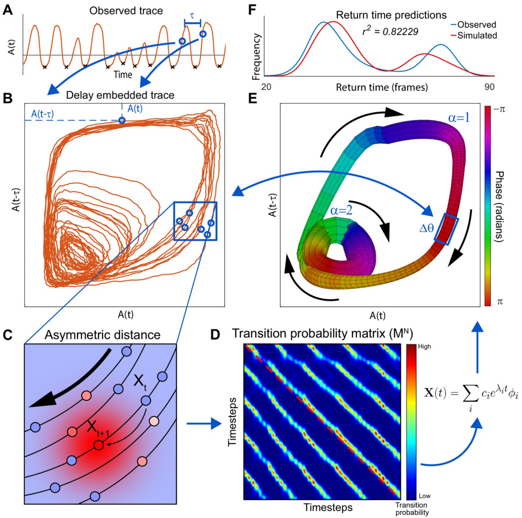

References

Connor Brennan and Alexander Proekt. A quantitative model of conserved macroscopic dynamics predicts future motor commands. *Elife*, 2019.

Saul Kato, Harris S Kaplan, Tina Schrödel, Susanne Skora, Theodore H Lindsay, Eviatar Yemini, Shawn Lockery, and Manuel Zimmer. Global brain dynamics embed the motor command sequence of caenorhabditis elegans. *Cell*, 163(3):656–669, 2015.

Boaz Nadler, Stephane Lafon, Ronald R Coifman, and Ioannis G Kevrekidis. Diffusion maps, spectral clustering and eigenfunctions of fokker-planck operators. *arXiv preprint math/0506090*, 2005.

Annika LA Nichols, Tomáš Eichler, Richard Latham, and Manuel Zimmer. A global brain state underlies c. elegans sleep behavior. *Science*, 356(6344), 2017.

Jin Wang, Li Xu, and Erkang Wang. Potential landscape and flux framework of nonequilibrium networks: robustness, dissipation, and coherence of biochemical oscillations. *Proceedings of the National Academy of Sciences*, 105(34):12271–12276, 2008.

Han Yan, Lei Zhao, Liang Hu, Xidi Wang, Erkang Wang, and Jin Wang. Nonequilibrium landscape theory of neural networks. *Proceedings of the National Academy of Sciences*, 110(45):E4185–E4194, 2013.