Neurons & Behavior: a dynamical perspective

Antonio Costa, Greg J Stephens

Introduction Recently, advanced imaging methods have made it possible to simultaneously record whole-brain activity in a behaving nematode, *C. elegans* [1, 2]. This means that we are able to get measurement of the brain $g(x)$ and behavior $h(y)$ simultaneously (g and h are measurement functions). In abstract terms, we can write a set of first order differential equations that describe the dynamics of the brain and behavior and their coupling,

$$\dot{x} = F(x, y, t) \tag{1}$$
$$\dot{y} = G(x, y, t), \tag{2}$$

where F and G are possibly non linear functions of x, y and time t. A natural question arises: how do we understand the mapping between neural and behavioral dynamics? Can we find an expression for F and G? Or can we infer the relationship between x and y without having a functional representation of the dynamics? Which principles can we use (from dynamical systems and information theory) to guide our analysis? In this project, we will build upon dynamical systems theory, ergodic theory, inference and information theory, in an attempt to find answers to these questions.

Warm-up. A useful starting point is to review low-order, deterministic dynamical systems and the Lorenz equation (originally derived as a simpler description of atmospheric turbulence) provides a canonical example:

$$\dot{x} = \sigma(y - x),$$
$$\dot{y} = \rho x - y - xz,$$
$$\dot{z} = xy - bz,$$

where $\{\sigma, b, \rho\}$ are parameters. Though perhaps most famous for its "butterfly" chaotic attractor, the Lorenz equations are a wonderful laboratory in which to explore many phenomena in nonlinear dynamics and you might start by demonstrating different attracting trajectories (fixed points, limit cycles, chaos) depending on parameter choices and initial conditions. For analytical understanding, fix two of the three parameters $\{\sigma = 10, b = \frac{8}{3}\}$ and see happens as a function of ρ. First, verify that the origin is a fixed point. Use linear stability analysis to deduce the asymptotic dynamics for $\rho < 1$? What happens to the stability of the origin fixed point at $\rho > 1$? For $\rho > 1$ (but not too large) are there new fixed points? For more complicated attractors such as limit cycles and chaos you will need to develop a numerical integration scheme. We also have some very-finely sampled high resolution simulation data for the Lorenz dynamics on their standard chaotic attractor if you're interested in comparing.

Randomness in deterministic dynamics. In the Lorenz system, once we increase ρ above a certain value, the dynamics become chaotic. What are the signatures of a chaotic dynamical system? Try to follow the evolution of a pair of neighbors in phase space. How do their trajectories diverge? One of the properties of a chaotic system is the existence of positive Lyapunov exponents [3, 4, 5]. Try to estimate the largest Lyapunov exponent (or the entire spectrum of exponents if you'd like!) of the Lorenz system in different dynamical regimes.

Bonus: Symmetries and Lyapunov Exponents. Lyapunov exponents reflect the symmetries underlying the dynamics. For example, continuous symmetries lead to the existence of zero exponents. Can you prove this? Use the calculations in [6] as a guide. In addition, discrete symmetries such as time-reversal symmetry, or the existence of a symplectic structure can lead to additional structure in the Lyapunov spectrum. Use the calculations in [7, 6, 8] to prove the existence of a conjugate pairing of Lyapunov exponents in Hamiltonian systems. What are the consequences of this property? Can you justify Lioville's theorem from this point of view? How would the Liouville's theorem be extended in the presence of dissipation?

Entropy & Determinism. Chaotic dynamics has many similarities with stochastic processes and one example is the equivalence of the Kolmogorov-Sinai entropy with the Shannon entropy of an appropriately constructed symbolic sequence. In simple systems, it possible to find a discretization that preserves a one-to-one mapping with the underlying phase space trajectories. Explore the notion of a generating partition in simple 1D maps (e.g. tent map, logistic map) and try to estimate the entropy of the generated symbolic sequence in the chaotic regime. How is this related with the Ising model? Explore this equivalence in continuous chaotic dynamical systems and how the partition process affects entropy estimation [9, 10]. Can you derive the asymptotic behavior of the entropy estimates as a function of the partition/resolution? Is there a relationship between that and the notion of predictive information in the context of learning [11, 12, 13]?

Inverse Physics: Learning Equations of Motion from Data. What if we didn't know the equations (the usual case)? Use current ideas from the literature such as [14, 15] to explore how you can learn equations of motion directly from data. Do it for a system(s) of your choosing! What can inferring these equations tell you about the system? What are the difficulties and important principles? Why not simply use a deep network for prediction (e.g. [16, 17])

State Space Reconstruction. In practice, we rarely have access to the original equations of motion, or even the state of the system. We simply have measurement data from an underlying dynamical system. In this case, the first step towards understanding the dynamics is to reconstruct the state space (or phase space as is known in physics) from the available measurements. The concept of state is one of the most influential ideas in modern physics, almost all the current theories are framed in the phase space, leading to a geometrization of physics, which can be very useful (for a history of the idea see here [18]). As an example, consider a simple harmonic oscillator. Working with the phase space variables (x,p) allows for a description of the system in terms of a simple first order differential equation, which would be second order if we thought only about x. By constructing the phase space, you get access to a set of variables that are maximally predictive.

Modern ideas of state space reconstruction build upon seminal results by the mathematician Floris Takens, who proved that in fairly general conditions it is possible to reconstruct a one-to-one copy of the original state space using the delays of our measurements alone. Readable references on state space reconstruction can be found in [19, 4]. When possible, the reconstructed state space preserves all fixed points, limit cycles, and other limiting sets. As a toy example, consider that we only have noisy measurements of the x variable of the Lorenz system. Given this measurement reconstruct the state space of the Lorenz system. Try the exercise again with other variables and report the findings. In particular how does the reconstruction vary with different measurement functions and reconstruction parameters? How would you measure the quality of a state space reconstruction? **Hint**: Think about prediction, and what does a "good reconstruction" mean in terms of predictability. Can you use, for example, predictive information as a guide for the reconstruction?

Back to the project questions. Data from [1] can be downloaded from `https://osf.io/79ghf/`. Using ideas from dynamical systems theory and information theory, explore the relationship between the dynamics of the brain and behavior. Essentially, there are two main approaches you can use:
Parametric approach: Find some functional form for F and G and use it to estimate a relationship between neural and behavioral dynamics. Big questions: Scholz et al. [1] showed that the posture dynamics can be predicted using the neural activity using a simple linear regression method. Can you recover these results using a more complete model? Can behavior predict the neurons as well? If so, what is the functional relationship?
Non-parametric approach: Reconstruct the state space of a single neuron (e.g. AVA) and a body point (e.g. head), and try to infer the relationship between their dynamics. One way in which you can do that is through mutual information,

$$I(X,Y) = H(X) + H(Y) - H(X,Y). \tag{3}$$

What are the challenges of estimating mutual information in the context of a continuous phase space? How does that depend on the way you choose to discretize? Why should we expect there to be an instantaneous relationship between neural and behavioral dynamics? Shouldn't there be a lag time between them? Explore the possibility of estimating the mutual information between $X_{t+\tau'}$ and $Y_{t+\tau}$. Big questions: How does the estimate of mutual information using a state space representation differ from a simple estimate of the mutual information between the measurements themselves, g(x) and h(y)? In other words, how important is the fine scale dynamics and a proper definition of state to neural computation? How does the estimate of the mutual information depend on the lag times? Given the maximal mutual information between the past behavior and the current neural activity $I_{\max}(X_t, Y_{t-\tau})$ and the maximal mutual information between the past neural activity and the current behavior $I_{\max}(X_{t-\tau}, Y_t)$, what is the relative difference between the two? What are the challenges of using the entire network dynamics and the entire posture? How do we expect the network dynamics to impact our estimates of mutual information? In other words, at what level of description is neural computation being done?

References

[1] Scholz M, et al. (2018) Predicting natural behavior from whole-brain neural dynamics. *bioRxiv*.

[2] Nguyen JP, et al. (2016) Whole-brain calcium imaging with cellular resolution in freely behaving caenorhabditis elegans. *Proc. Natl. Acad. Sci. USA* 113(8).

[3] Eckmann JP, Ruelle D (1985) Ergodic theory of chaos and strange attractors. *Reviews of Modern Physics* 57(3):617–656.

[4] Abarbanel HD, Brown R, Sidorowich JJ, Sh Tsimring L (1993) The analysis of observed chaotic data in physical systems Chaotic time series data are observed routinely in experiments on physical systems and in observations in the field. The authors review developments in the extraction of information of physical imp. *Rev. Mod. Phys.* 65.

[5] Strogatz S (1994) ME 406 The van der Pol Oscillator. pp. 1–15.

[6] Haken H (1983) At least one lyapunov exponent vanishes if the trajectory of an attractor does not contain a fixed point. *Physics Letters A* 94(2):71–72.

[7] Gupalo D, Kaganovich AS, Cohen EGD (1994) Symmetry of Lyapunov spectrum. *Journal of Statistical Physics* 74(5-6):1145–1159.

[8] Hoover W, Hoover C (2017) Instantaneous Pairing of Lyapunov Exponents in Chaotic Hamiltonian Dynamics and the 2017 Ian Snook Prizes. *Computational Methods in Science and Technology* 23(1):73–79.

[9] Grassberger P, Schreiber T, Schaffrath C (1991) Nonlinear time sequence analysis. *International Journal of Bifurcation and Chaos* 1:43.

[10] Grassberger P, Procaccia I (1983) Characterization of Strange Attractors. *Rev. Lett* 50(5):1448.

[11] Bialek W, Tishby N (1999) Predictive Information. (3).

[12] Bialek W, Nemenman I, Tishby N (2001) Complexity through nonextensivity. *Physica A: Statistical Mechanics and its Applications* 302(1-4):89–99.

[13] Bialek W (2001) Predictability , Complexity , and Learning. 2463(1949):2409–2463.

[14] Brunton SL, Proctor JL, Kutz JN (2016) Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl. Acad. Sci. USA* 113(15):3932–3937.

[15] Daniels BC, Nemenman I (2015) Automated adaptive inference of phenomenological dynamical models. *Nature Communications* 6:1–8.

[16] Pathak J, Lu Z, Hunt BR, Girvan M, Ott E (2017) Using machine learning to replicate chaotic attractors and calculate lyapunov exponents from data. *Chaos* 27(12):121102.

[17] Champion K, Lusch B, Kutz JN, Brunton SL (2019) Data-driven discovery of coordinates and governing equations. pp. 1–26.

[18] Nolte DD (2010) The tangled tale of phase space.

[19] Casdagli MEJDFGJ (1991) The theory of state space reconstruction in the presence of noise. *Information Dynamics* 0:61–95.