

SIMONS FOUNDATION

Global Brain



Credit: Ricardo Bessa/Folio Art

Searching for the Hidden Factors Underlying the Neural Code

Three new methods to identify hidden structure in neural activity.

By Richard Gao

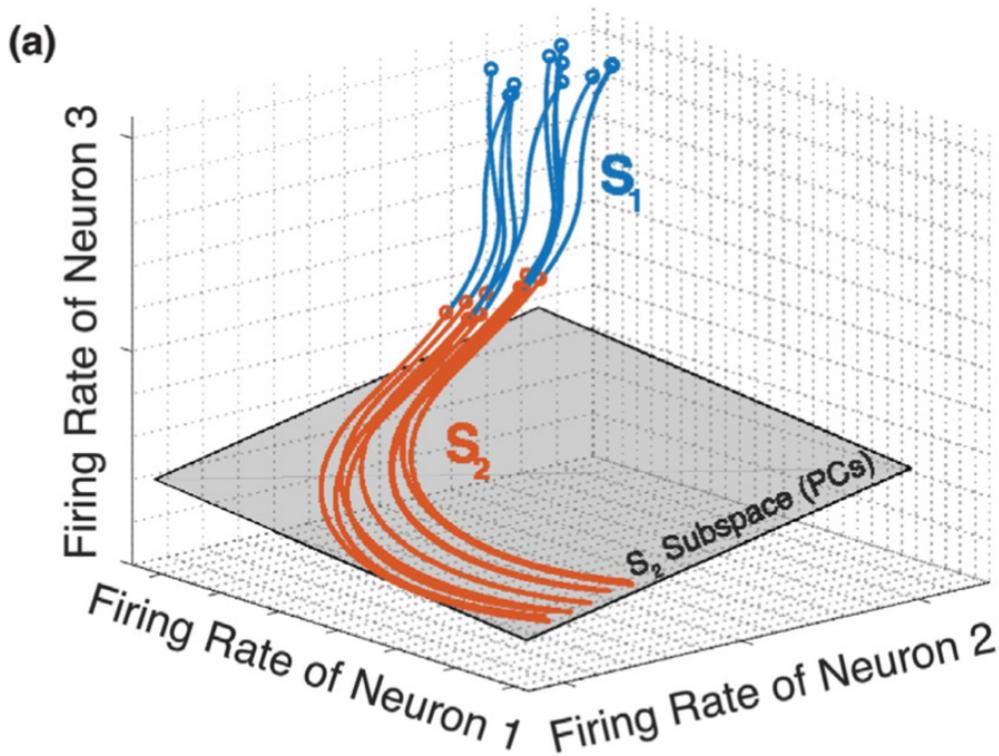
July 31, 2019

Series

Cosyne 2019

Aside from unenthusiastically playing the flute during band class in middle school, I've never really had any relationship with orchestral music. I do find symphonic orchestras inspiring, though, because these highly skilled musicians chose to work with up to 100 others rather than on their own, sacrificing their individual expression to be a part of a collective capable of epic performances of works like the *Avengers* theme.

Our neurons take a similar tack: Rather than singing their own tunes independently, they form groups of neurons that fire in complex, coordinated patterns. New technologies capable of recording from hundreds of neurons at once, such as high-density electrodes and wide-field imaging, have revealed an elegant structure underlying this cooperative activity. It's as if these neurons were an ensemble of orchestral musicians, coming together to perform an unknown composition. Our goal as neuroscientists is to uncover these compositions and to better understand the organizing principles behind cooperative neural activity — and how they drive complex behaviors.

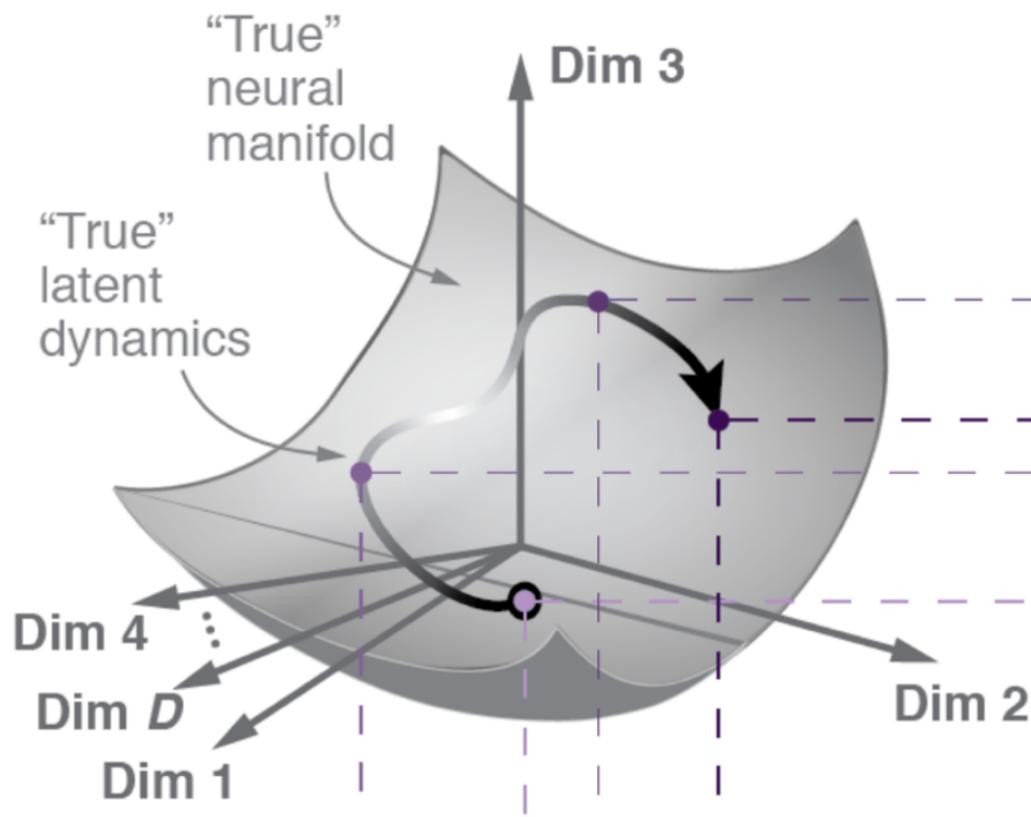


Collective neural activity over time can be visualized as trajectories through a 'neural state space.' Similar behaviors often generate orderly traces. Saxena and Cunningham, COIN 2019.

Traversing musical landscapes and neural manifolds

In an orchestra, the full range of music the group can perform is defined by what each individual musician can play, with one dimension for each musician. In physics, this is known as the state space. A musical performance traces out a trajectory in this space. If each of the 100 musicians played completely independently, this trajectory would look like a random meandering path jumping all over the place. The result? Total cacophony. Fortunately, orchestra members cooperate from beginning to end, generating a beautiful and orderly trajectory threading through the realm of all that is musically possible.

Similarly, recording from 100 neurons, we could theoretically observe completely independent activities, where each neuron takes full advantage of its individual degree of freedom. But we rarely do. When our arm reaches in different directions, for example, neural activity in the motor cortex often follows structured trajectories through an analogous neural state space, where each dimension (or axis) is defined by the firing rate of one neuron (see figure, right). Strikingly, these paths reside on a lower-dimensional surface, or manifold, of the full state space. In other words, collective neural activity underlying the arm movement seems to be constrained to a smaller region of all that's possible, which is specific to and defined only for that particular population of neurons.



The trajectory representing coordinated population activity through time is shown above to be constrained to a lower-dimensional

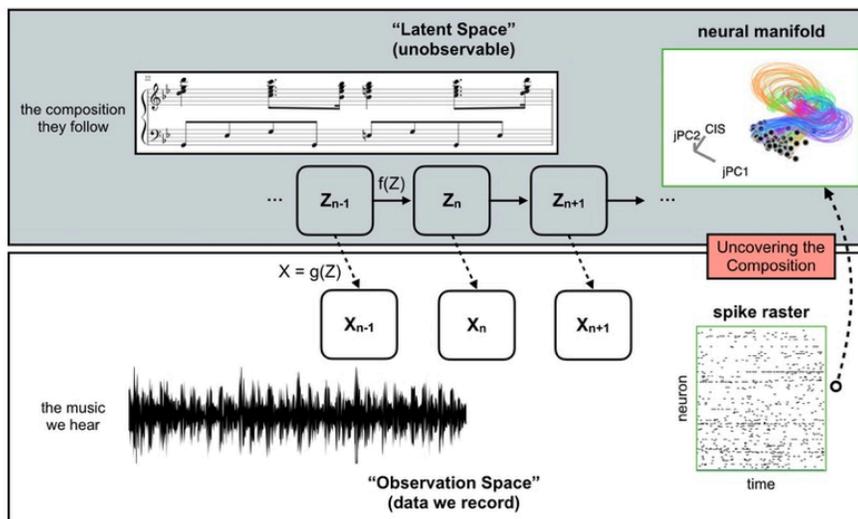
(nonlinear) surface in the neural state space, otherwise known as a 'neural manifold.' More accurately, the manifold is the hypothetical surface that contains all possible trajectories. Gallego et al. *bioRxiv* 2019

Just how low-dimensional are these manifolds, and why? For an orchestra, each composition (like Beethoven's Symphony No. 9) defines a single trajectory, and a collection of compositions makes up the musical landscape. These musical manifolds are low dimensional because compositions and the orchestra itself are divided into sections — the woodwinds play one part, the strings another, and the percussion a third — forming smaller units that harmonize within themselves to create layers of organization. As we tune into the orchestra in the brain by recording from large populations, we also observe regularities in how neurons cooperate with each other, but we don't quite know why. A major goal for computational neuroscience is to understand the organizational principles behind coordinated neural activity, relying on statistical inferences to uncover and characterize these hidden neural manifolds.

There, now you can finally follow #neuralmanifold on Twitter.

New Tools for Uncovering Neural Manifolds

As technology evolves to record from more and more neurons, computational neuroscientists are developing more complex models to take advantage of the data we have, building richer and more detailed portraits of neural dynamics. A number of such methods were presented at this year's Cosyne meeting in Lisbon. Here, I will highlight three complementary approaches presented by three junior researchers — Lea Duncker, Daniel Hernandez and Dan O'Shea — that tackle similar problems in different ways. I can write a whole blog post for each of these projects on why they are uniquely cool, but for the rest of this post, I will focus on the commonalities and briefly touch on their unique features.



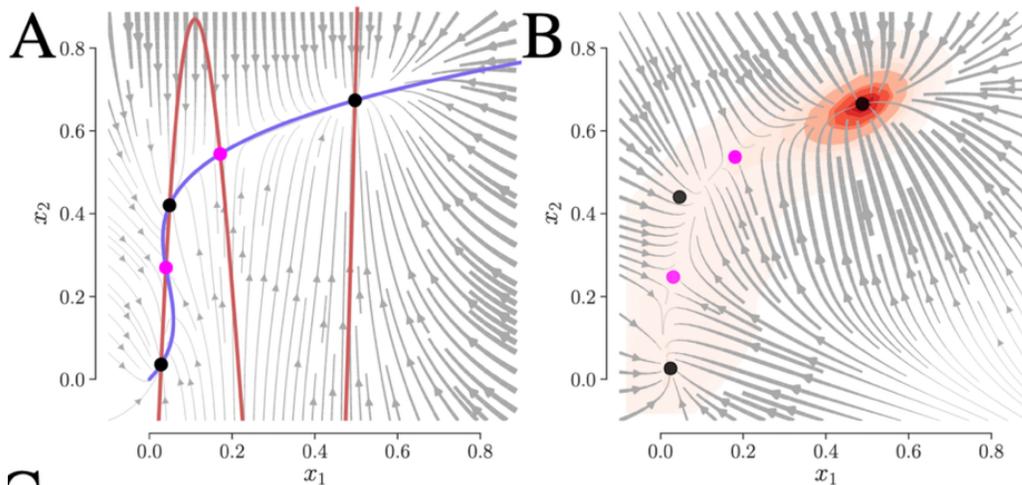
The musical composition is the set of rules, $f()$, that defines the evolution of dynamics of the latent variables, Z , the notes. As listeners, we are not privy to those rules, only the sounds, X , from a performed version of the music. The projection or observation function, $g()$, maps from the latent variables Z to the observed variables X , acting like the musician who translates notes to sounds. The tools presented all perform the reverse inference: Given the data we observe, X , we want to retrieve the (theoretical) latent variables Z and their dynamics $f()$, which often explicitly approximates the observation function $g()$.

Richard Gao

In all three studies, it's assumed that the neural data represent noisy observations from what is termed a 'nonlinear dynamical system.' Conceptually, a dynamical system is defined by a set of equations that dictate, at every point in a defined state space, where you should go next — a mapping from one place to the next. With some starting points and this set of rules, the state of the system will trace out trajectories over time in the state space — precisely like the orchestra following a composition from beginning to end. These equations also constrain the possible trajectories to a low-dimensional manifold: The rules define the landscape.

Even though these equations are mathematical abstractions of the physical system, uncovering these rules allows researchers to model the system's behavior and make predictions. However, the data we use to perform that inference is rarely a direct measurement of latent variables in the system; rather, it's a transformation or 'projection' of them. The spiking data we gather is analogous to the music we hear. But we're most interested in inferring the exact sequence of notes in the composition itself — the latent variables governing neural activity. Uncovering these variables and the rules they follow requires reverse inference using computational techniques to model our noisy population spiking data. The following works all present new tools for achieving that goal.

Lea Duncker's work with Maneesh Sahani, a theorist at University College London and an investigator with the Simons Collaboration on the Global Brain, uses Gaussian processes to learn the equations of a continuous-time stochastic (noise-driven) dynamical system. Gaussian processes are very flexible in their ability to model unknown and continuous functions, providing not only an estimate of any point on the function, but also confidence bounds on those estimates. Gaussian processes further narrow the learned function based on important properties of the system, like fixed points and the local behavior around them (is the fixed point an attractor, a repeller or a saddle?), which the authors exploit to obtain human-interpretable descriptions of the complex nonlinear dynamics.



The ground truth landscape of the simulated system (A, left), where arrows denote direction of flow at that location and dots represent fixed points of the system. The learned landscape (B, right) has very similar features Dunker et al arXiv 2019

The authors use a number of well-defined dynamical systems, such as the van der Pol oscillator, as the hidden rules to simulate spike train data, and they demonstrate the approach's ability to retrieve the dynamics defined by those equations. The key feature of this approach is that the Gaussian process description allows them to obtain an estimate of the entire landscape (or phase portrait), even where no observations were directly made. It's like being able to distinguish the entire body of Mozart's works after hearing just a handful of his compositions. Because of this, we can directly compare the simulated ground truth and the learned dynamics visually, noting a close similarity between the two, especially in important regions of the state space such as fixed points (see figure, right). It will be really exciting to see what the full manifolds learned from real spiking data look like, especially in regions where data is not directly

Daniel Hernandez, working with John Cunningham and Liam Paninski, SCGB investigators at Columbia University, developed Variational Inference for Nonlinear Dynamics (VIND) as another tool to infer hidden structure in neural population activity. VIND takes the same approach in assuming that there is a set of latent or hidden variables whose evolution traces out the manifold, and that the recorded neural activity is

a projection of those hidden variables. The algorithm accordingly learns two components: how the hidden variables evolve, and how spikes are generated from them. For the first part, VIND approximates a nonlinear dynamical system as a 'locally linear' one, which means the equations are different depending on the location in state space and are learned using an artificial neural network. For the second part, VIND assumes that the spiking data we observe are randomly drawn from complicated probability distributions parameterized by the hidden variables, which are often intractable to learn directly. In that sense, the spikes we observe have an almost improvisational feel, like jazz — the musicians are following a concrete composition, but how they play it varies from one show to another. That variation makes the underlying composition more difficult to decipher, but VIND approximates the complicated probability distributions with Gaussian or Poisson distributions (i.e., variational inference).

Validating their approach on simulated data from the classic Lorenz system (the butterfly attractor), as well as various electrophysiological datasets, they are able to retrieve and denoise the underlying low-dimensional latent dynamics, showing that the method has better performance and speed than several others. Moreover, VIND is able to learn the rules of a system (e.g., single-neuron voltage dynamics) and predict many steps into the future given a different starting point, which is very useful. Its ability to predict future behavior also implies that the system does, to some degree, follow a set of dynamical equations, cashing in on the original assumptions. To the detriment of my analogy, symphonies (at least human-composed ones) do not actually follow a set of dynamical equations as they evolve. Nevertheless, after listening to your favorite artist over and over again, you might have the experience of being able to guess the chord sequence of their new song after hearing just the first few seconds.

Finally, Dan O'Shea, working with David Sussillo and Krishna Shenoy, SCGB investigators at Stanford University, presented his work extending the functionality of LFADS (latent factor analysis via dynamical systems) and its application on a massive mouse whole-cortex Neuropixels dataset. LFADS, previously discussed in a SCGB research summary by Grace Lindsay, similarly assumes that latent neural variables evolve according to a nonlinear dynamical system. LFADS uses techniques from machine learning to train a recurrent neural network (RNN) that can reproduce single-trial neural dynamics. At its core, an RNN is a complex but flexible nonlinear dynamical system, whose connection weights can be tuned through optimization to approximate other dynamical systems without learning the explicit rules. In addition, LFADS is capable of 'dynamical neural stitching,' a feature that allows researchers to learn shared neural dynamics for a behavior in datasets where separate groups of neurons were recorded in multiple, sequential experimental sessions. This is like recording the wind, percussion and brass sections playing the same piece on different days and then stitching the tracks together to infer the hidden composition. Check out Dan's explanatory Twitter thread, with beautiful animations of stitched neural trajectories.

More from simonsfoundation.org



Decoded single trajectories of latent variables (projected onto 3 dimensions) across 42 recording sessions were consistent when the underlying shared dynamics is learned via neural stitching. Each trace is a single trial, and color represents the trial type. Trajectories are similar within the same trial type and differences across trial types. Pandarinath et al. Nature Methods 2018

While we can now record many neurons in a session using high-density probes like Neuropixels, performing neural stitching across several brain regions simultaneously still poses technical challenges because the local dynamics within each brain region can vary. In his Cosyne talk, Dan presented a more powerful version of LFADS dynamical stitching that addresses these technical issues. He described results from applying this technique to a dataset (recorded by Will Allen) combining over 80 recording sessions across 34 brain regions in 21 mice performing the same olfactory decision task. He was able to retrieve robust latent trajectories and decode odor identity from the model's inferred input. The ability to combine sequential datasets opens up many new opportunities for experiments and for re-analysis of existing neural data (see LFADS instructions and tutorial).

What's next?

These works presented at Cosyne demonstrate that many different approaches can tackle the same problem using the mathematical tools available to us today. All of these methods are generalizable to many dynamical systems and can be easily used in other fields, such as climate science or economics. In

neuroscience, however, employing these tools to uncover the latent neural manifold from large datasets allows us to ask even more questions. At the risk of overstretching the orchestra analogy, I will end on three questions with its help:

First, we know that an orchestra relies on the composition to bring about cooperation and harmony. But given that the brain operates in harmony as well, what exactly is its composition? In other words, what physical quantities do the latent variables and manifolds represent? No external composer dictates the symphony in the brain, but factors like circuit connectivity constraints, external and upstream inputs, and neuromodulators may be good candidates for physical instantiations of these latent variables.

The second question concerns the inferences we make about a system where the data we observe may not reflect its most crucial components. Standing in front of the orchestra, the conductor goes through the entire performance without making a single sound. Yet, remove the conductor and the orchestra falls apart. Do we need to search for the silent orchestrator in neural circuits that makes everything cooperate, perhaps with complementary imaging techniques that listen in on glial cells and other non-spiking contributors?

Lastly, macroscopic brain signals, such as neural oscillations detectable with electroencephalogram, reflect the total activity of many, many neurons, particularly when they are synchronous, like during the climax of a symphony. How can we use our knowledge about neural manifolds to better inform our analysis of these signals, to wrangle more information for clinical and translational applications? Hopefully we'll see some of these questions addressed at next year's Cosyne! For now, enjoy this [epic performance](#).

Richard Gao is a graduate student in the Department of Cognitive Science at the University of California, San Diego, working with Bradley Voytek on decoding population activity from field potential signals. You can read his blog [here](#).

© Simons Foundation