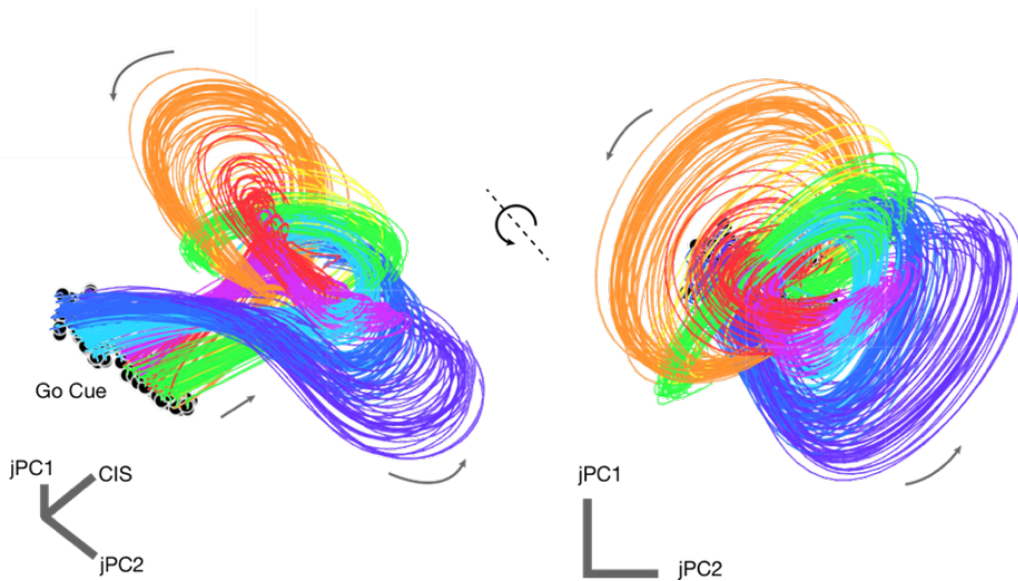

Global Brain

Noise-Cancelling Method Enhances Ability to Predict Behavior from

A new technique that uses machine learning to mute noise inherent in single-trial neural recordings will make it easier to decipher the neural code

Grace Lindsey

October 26, 2018



LFADS finds the underlying patterns in how neurons interact over time. Here, each line represents the progression of neural activity in motor cortex as an animal reaches in different directions. Credit: Pandarinath et al Nature Methods 2018

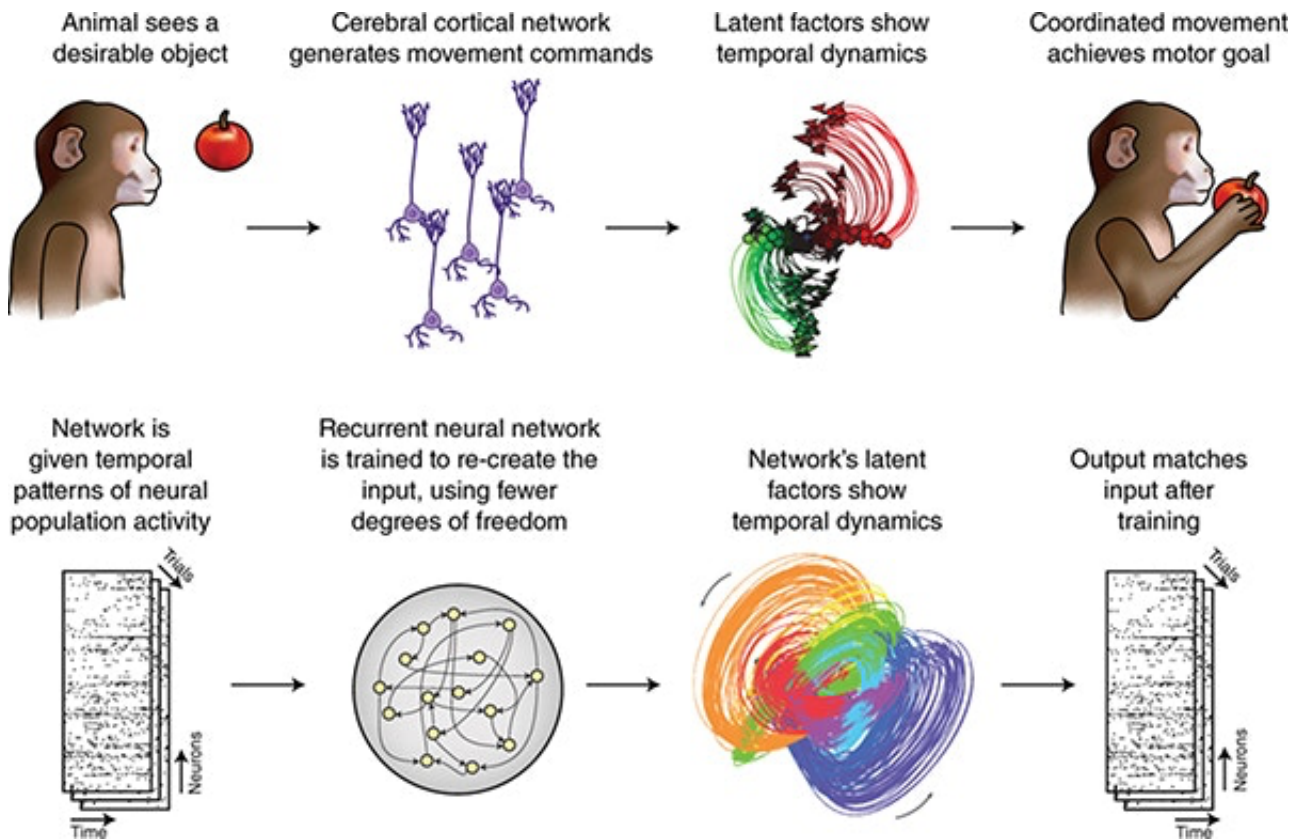
Like the height, width and depth of a shipping container, data from neural recordings are frequently

defined by three values: the number of neurons, the number of trials, and the temporal length of a trial. These recordings are inherently noisy measures of neural activity, so scientists frequently average over one of these dimensions: the activity of a neuron over multiple trials, or of multiple neurons on one trial, or even of one neuron over time within a trial.

But this averaging comes at a cost: it muddies the precise timing information contained in a single trial. By sacrificing these detailed dynamics, researchers may be missing out on clues to how brain regions interact with each other or generate behavior. A new approach — published in *Nature Methods* in September 2018 by Chethan Pandarinath at Georgia Tech and collaborators from Stanford, Brown, Emory, and Columbia Universities and Google — is designed to rescue these dynamics and provide clean and precise estimates of a neuron's firing on a trial-by-trial basis.

The method — latent factor analysis via dynamical systems (LFADS) — uses a machine learning technique known as a variational autoencoder. Autoencoders are models trained to reproduce their input. In this case, the model is fed a sequence of spike counts from the data and produces a similar sequence of spike counts as output. The utility of this seemingly trivial transformation is hidden in the machinery that carries it out.

In LFADS, the data from the neural population on a given trial is used to generate the initial state for a recurrent artificial neural network. This recurrent network takes input and interacts with itself over time, generating latent dynamics that are used to produce a moment-by-moment estimate of the firing rate for each neuron in the population. The structure and training of the model predispose it to reproduce the real neural dynamics but not the noise. Better than naive averaging, this method also utilizes the state of the population as a whole to make estimates for individual neurons.



An artificial neural network (bottom) can capture the dynamical structure present in neural population activity (top). Credit: Kim Caesar/Springer Nature

“LFADS is a very creative application of the variational auto-encoder idea,” says Jonathan Pillow, an associate professor in the department of psychology at Princeton University and an investigator with the Simons Collaboration on the Global Brain, who was not involved in the work.

Having accurate firing-rate information for individual trials makes mapping the relationship between neural activity and behavior more precise than ever. Using their new method on data from motor and premotor cortices, the team showed that their de-noised firing-rate estimates could predict the velocity of hand movements on individual trials better than either the averaged neural data or a previous method for estimating rates on individual trials.

The same model can be used even if the specific neurons being recorded change (a common situation for many experimentalists). By adjusting the precise way in which the spike counts are read into and out of the model, data from different neurons recorded on different days can be accurately stitched together and help make the model better. “The stitching really does work,” says David Sussillo, a neuroscientist at Stanford University and Google and an investigator with the SCGB, who led the work.

The recurrent neural network at the heart of LFADS also brings another advance. “Most previous work in

this area has tended to focus on linear dynamical systems,” says Pillow. “But we know the brain’s dynamics are nonlinear.” According to Sussillo, allowing the latent population to interact in more complex, nonlinear ways provided a model of the dynamics “that was natural to the data.”

Sussillo and collaborators have developed an LFADS software package that is freely available for the neuroscience community. The researchers put considerable effort into making it user-friendly, but it may still require technical expertise to use. When it comes to training, Sussillo says, these models have “all the strengths and weaknesses that deep learning models have.” The technique requires a fair amount of data to train the model, and the training process can be difficult to troubleshoot.

Sussillo and collaborators will use LFADS in their ongoing SCGB project, ‘Computation-through-dynamics as a framework to link brain and behavior,’ which aims to build models that can dynamically execute cognitive tasks in a neurally plausible way. To do this, researchers need to compare the dynamics found in neural data with those of the models. LFADS will be instrumental in that part of the project, ensuring that the dynamics in the data are as accurate as possible.

© Simons Foundation