Available online at www.sciencedirect.com

ScienceDirect

Current Opinion in Neurobiology

On simplicity and complexity in the brave new world of large-scale neuroscience

Peiran Gao[1] and Surya Ganguli[2]

Technological advances have dramatically expanded our ability to probe multi-neuronal dynamics and connectivity in the brain. However, our ability to extract a simple conceptual understanding from complex data is increasingly hampered by the lack of theoretically principled data analytic procedures, as well as theoretical frameworks for how circuit connectivity and dynamics can conspire to generate emergent behavioral and cognitive functions. We review and outline potential avenues for progress, including new theories of high dimensional data analysis, the need to analyze complex artificial networks, and methods for analyzing entire spaces of circuit models, rather than one model at a time. Such interplay between experiments, data analysis and theory will be indispensable in catalyzing conceptual advances in the age of large-scale neuroscience.

Addresses
[1] Department of Bioengineering, Stanford University, Stanford, CA 94305, United States
[2] Department of Applied Physics, Stanford University, Stanford, CA 94305, United States

Corresponding author: Gao, Peiran (prgao@stanford.edu)

Current Opinion in Neurobiology 2015, **32**:148–155

This review comes from a themed issue on **Large-scale recording technology**

Edited by **Francesco P Battaglia** and **Mark J Schnitzer**

For a complete overview see the Issue and the Editorial

Available online 29th April 2015

http://dx.doi.org/10.1016/j.conb.2015.04.003

'*Things should be as simple as possible, but not simpler.*'
– Albert Einstein.

Introduction

Experimental neuroscience is entering a golden age marked by the advent of remarkable new methods enabling us to record ever increasing numbers of neurons [1–5,6•], and measure brain connectivity at various levels of resolution [7–10,11•,12–14], sometimes measuring both connectivity and dynamics in the same set of neurons [15•,16]. This recent thrust of technology development is spurred by the hope that an understanding of how the brain gives rise to sensations, actions and thoughts will lurk within the resulting brave new world of complex large-scale data sets. However, the question of how one can extract a conceptual understanding from data remains a significant challenge for our field. Major issues involve: (1) What does it even mean to conceptually understand 'how the brain works?' (2) Are we collecting the right kinds and amounts of data to derive such understanding? (3) Even if we could collect *any* kind of detailed measurements about neural structure and function, what theoretical and data analytic procedures would we use to extract conceptual understanding from such measurements? These are profound questions to which we do not have crisp, detailed answers. Here we merely present potential routes towards the beginnings of progress on these fronts.

Understanding as a journey from complexity to simplicity

First, the vague question of 'how the brain works' can be meaningfully reduced to the more precise, and proximally answerable question of how do the connectivity and dynamics of distributed neural circuits give rise to specific behaviors and computations? But what would a satisfactory answer to this question look like? A detailed, predictive circuit model down to the level of ion-channels and synaptic vesicles within individual neurons, while remarkable, may not yield conceptual understanding in any meaningful human sense. For example, if simulating this detailed circuit were the *only* way we could predict behavior, then we would be loath to say that we *understand* how behavior emerges from the brain.

Instead, a good benchmark for understanding can be drawn from the physical sciences. Feynman articulated the idea that we understand a physical theory if we can say something about the solutions to the underlying equations of the theory without actually solving those equations. For example, we understand aspects of fluid mechanics because we can say many things about specific fluid flows, without having to numerically solve the Navier–Stokes equations in every single case. Similarly, in neuroscience, understanding will be found when we have the ability to develop simple coarse-grained models, or better yet a hierarchy of models, at varying levels of biophysical detail, all capable of predicting salient aspects of behavior at varying levels of resolution. In traversing this hierarchy, we will obtain an invaluable understanding of which biophysical details matter, and more importantly, which do not, for any given behavior. Thus our goal should be to find simplicity amidst complexity, while of course keeping in mind Einstein's famous dictum quoted above.

How many neurons are enough: simplicity and complexity in multineuronal dynamics

What kinds and amounts of data are required to arrive at simple but accurate coarse grained models? In the world of large scale recordings, where we do not have access to simultaneous connectivity information, the focus has been on obtaining a state-space description of the dynamics of neural circuits through various dimensionality reduction methods (see [17] for a review). This body of work raises a key conceptual issue permeating much of systems neuroscience, namely, what precisely can we infer about neural circuit dynamics and its relation to cognition and behavior while measuring only an infinitesimal fraction of behaviorally relevant neurons? For example, given a doubling time of about 7.4 years [18] in the number of neurons we can simultaneously measure at single cell, single spike-time resolution, we would have to wait more than 100 years before we can observe $O(10^6$–$10^9)$ neurons typically present in full mammalian circuits controlling complex behaviors [19]. Thus, systems neuroscience will remain for the foreseeable future within the vastly undersampled measurement regime, so we need a *theory* of neuronal data analysis in this regime. Such theory is essential for firstly guiding the biological interpretation of complex multivariate data analytic techniques, secondly efficiently designing future large scale recording experiments, and finally developing theoretically principled data analysis algorithms appropriate for the degree of subsampling.

A clue to the beginnings of this theory lies in an almost universal result occurring across many experiments in which neuroscientists tightly control behavior, record many trials, and obtain trial averaged neuronal firing rate data from hundreds of neurons: in such experiments, the dimensionality (i.e. number of principal components required to explain a fixed percentage of variance) of neural data turns out to be much less than the number of recorded neurons (Figure 1). Moreover, when dimensionality reduction procedures are used to extract neuronal state dynamics, the resulting low dimensional neural trajectories yield a remarkably insightful dynamical portrait of circuit computation (e.g. [20,21,22•]).

These results raise several profound and timely questions: what is the origin of the underlying simplicity implied by the low dimensionality of neuronal recordings? How can we trust the dynamical portraits that we extract from so few neurons? Would the dimensionality increase if we recorded more neurons? Would the portraits change? Without an adequate theory, it is impossible to quantitatively answer, or even precisely formulate, these important questions. We have recently started to develop such a theory [41,42]. Central to this theory is the mathematically well-defined notion of neuronal task complexity (NTC). Intuitively, the NTC measures the volume of the manifold of task parameters (see Figure 2a for the special cases of simple reaches) measured in units of the neuronal population autocorrelation scale across each task parameter. Thus the NTC in essence measures how many neuronal activity patterns could possibly appear during the course of an experiment given that task parameters have a limited extent and neuronal activity patterns vary smoothly across task parameters (Figure 2b). With the mathematical definition of the NTC in hand, we derive that the dimensionality of neuronal data is upper bounded by the NTC, and if the neural data manifold is sufficiently randomly oriented, we can accurately recover dynamical portraits when the number of observed neurons is proportional to the log of the NTC (Figure 2c).

These theorems have significant implications for the interpretation and design of large-scale experiments. First, it is likely that in a wide variety of experiments, the origin of low dimensionality is due to a small NTC, a hypothesis that we have verified in recordings from the motor and premotor cortices of monkeys performing a simple 8 direction reach task [43]. In any such scenario, simply increasing the number of recorded neurons, without a concomitant increase in task complexity will not lead to richer, higher dimensional datasets — indeed data dimensionality will be independent of the number of recorded neurons. Moreover, we confirmed in motor cortical data our theoretically predicted result that the number of recorded neurons should be proportional to the

Figure 1

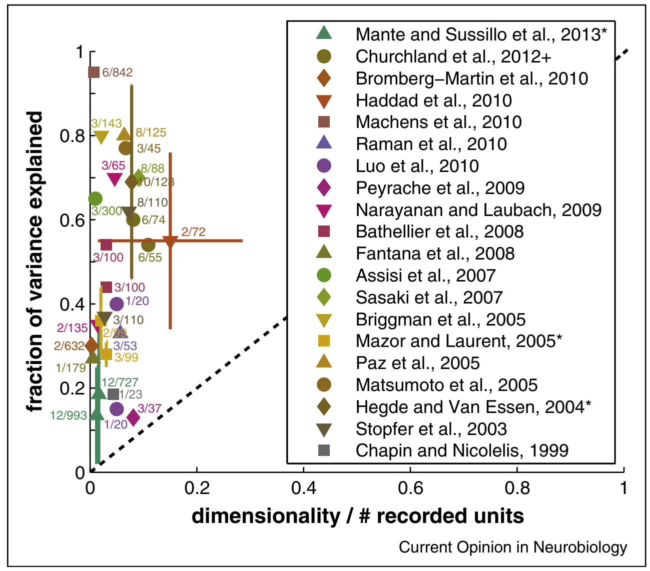

In many experiments (e.g. in insect [20,23–26] olfactory systems, mammalian olfactory [26,27], prefrontal [21,22•,28–30], motor and premotor,[31,32], somatosensory [33], visual [34,35], hippocampal [36], and brain stem [37] systems) a *much* smaller number of dimensions than the number of recorded neurons captures a large amount of variance in neural firing rates.