Available online at www.isas.org.in/jisas
JOURNAL OF THE INDIAN SOCIETY OF AGRICULTURAL STATISTICS 68(2) 2014 293-301

Big Data Comes to Functional Neuroscience: New Vistas for Statisticians

Mark Reimers
Departments of Psychiatry and Biostatistics, VCU, Richmond, VA 23298, USA

Received 17 April 2014; Revised 18 June 2014; Accepted 22 June 2014

SUMMARY

High-throughput data from new imaging technologies is coming to experimental neuroscience, as big data came to genomics a decade earlier. This represents a significant opportunity for new methodology in statistics. However many of the challenges facing statisticians will be quite different from those in genomics. This talk will introduce some issues in analysis of high-throughput functional neuroscience data, and illustrate them with recently published work, mostly drawn from animal studies. First many new technologies are burdened by significant noise so signal extraction techniques need development. Classical statistical dimension reduction strategies seem capture very limited fractions of variance in neuroscience data, and yet multivariate predictions and decoding have yielded some biological insight. Some alternative multivariate strategies have been proposed, but none are entirely satisfactory. How to characterize plasticity from neural activity data remains unclear. Finally we may anticipate a convergence of theoretical neuroscience with detailed experimental observations, as the heretofore unobservable dynamics of neural networks becomes visible. This emerging field presents an exciting opportunity to statisticians who are willing to learn neuroscience and engage with the field's questions.

Keywords: Neuroscience, Big data, Time series.

1. INTRODUCTION

Modern functional neuroscience increasingly draws on time series of neural activity to shed light on behavior and learning, and to give insight into clinical conditions such as psychiatric disorders or brain injury. New high-throughput technologies in neuroscience are now delivering new types of dynamic 'big data': long time series of measures on hundreds or thousands of individual neurons (brain cells) or many brain regions at high resolution. The BRAIN (Brain Research through Advancing Innovative Neurotechnologies) initiative in the United States and analogous initiatives in Korea, Japan and Europe, all point to a near future in which big data, and their associated analytic challenges, become the norm in neuroscience. The aim of this article is to introduce the reader to some of the new technologies and the characteristics of their data, and to point at some of the statistical challenges that will be faced in the next few decades of functional neuroscience.

Statistics has long played a role in interpreting neural activity data, both time series of events from individual neurons, and more recently, whole-brain activity data gathered on human subjects by functional magnetic resonance imaging (fMRI). In the former case a few neurons are sampled at high time resolution; in the latter case most of the brain is sampled, but at low resolution in time and averaging over millions of neurons. The promise of the new technologies is to gather information at a high resolution in both space and time over a large region of an animal's brain.

Neuroscience has many facets but a major agenda is to understand how brain activity works to guide behavior. This agenda calls for 'big data', for it is the

E-mail address: mreimers@vcu.edu

interactions among large numbers of connected brain cells that guide behavior. Until very recently neuroscientists have lacked the tools necessary to observe network activity at high resolution, and so have restricted their inquiries to questions that can be addressed by comparing the activity of individual neurons, or of one brain region at a time, between different situations.

In traditional electrophysiology researchers insert conducting electrodes into a brain region of interest in a live animal. The electrodes pick up fluctuations in voltage, which are recorded and separated into background and the sharp waves, called 'spikes' characteristic of action potentials. Statistics has contributed a great deal to the interpretation of long time-series of spikes, and shown how subtle changes in the timing of spikes may reflect perception or behavior. (Brown *et al.* 2004) review classical issues in spike train analysis in the context of measuring a small number of concurrent neurons.

Although fMRI has been used for almost twenty years to capture global images of brain activity, and generates large data sets, fMRI technology can only resolve events separated by several seconds, more than a hundred times slower than the time scale of neural activity; this has limited fMRI primarily to studies of localization of brain function; it cannot resolve the dynamics of brain activity; its average activity measures can tell us only that a particular piece of brain tissue is working hard, but do not tell us how that tissue is doing its work. Nevertheless fMRI has some characteristics of the new 'big data' technologies in that each experiment yields thousands of activity time series.

New technologies for measurement through imaging portend a dramatic increase in big data. Calcium imaging will soon give simultaneous activity on tens of thousands of individual neurons at 10ms resolution over many hours. Current whole-brain single-cell resolution data sets by 'light-sheet' calcium-indicator technology (Ahrens *et al.* 2013b) are already 1TB in size, and expected to grow to 50TB by the end of this year (Misha Ahrens, personal communication). Wide-field imaging using voltage-sensitive dyes (VSDs) can record activity over large swaths of cortex at 50 micron and 10ms resolution. One early VSD data set the author worked with contained 45 billion numbers, covering just over an hour of total recording

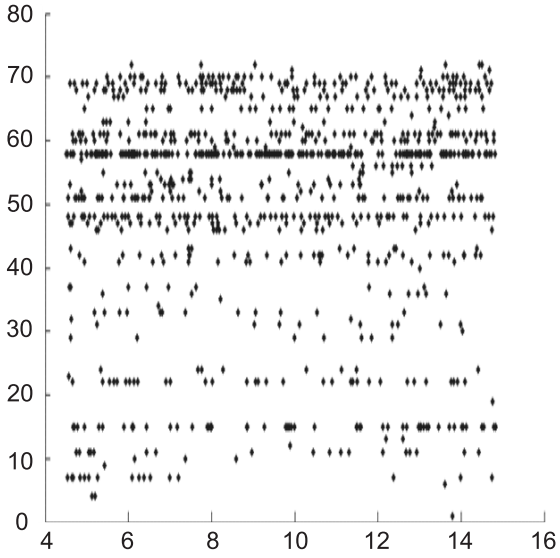

Fig. 1. Typical set of spike trains from electrode recordings circa 2005. Time in seconds. Unpublished data, courtesy of David Euston.

time (Mohajerani *et al*. 2013). Even electrophysiology is steadily expanding its capacities: currently it is possible to record from hundreds of neurons over many days. These technologies are generating long series of records (called 'spike trains') of the firing of many individual neurons; Fig. 1 shows a tiny slice of such a data set.

Such 'big data' poses many problems, not least of which is sufficient computational infrastructure. New data analysis frame works built on industrial data mining infrastructure, are being developed, for example the Thunder package (http://freeman-lab.github.io/thunder). Moreover such big data also affords unprecedented opportunities to ask the kinds of questions about how neural networks work, questions that formerly neuroscientists could only dream of, and statisticians have an important role to play in ensuring that these questions are answered well.

Big data in functional neuroscience differs from many other kinds of big data sets in engineering, network technologies, marketing, or social science. First functional neuroscientists deal with multivariate time-series, typically comprising hundreds or thousands of measures at thousands or millions of time points (fMRI data sets typically have tens of thousands of measures at hundreds of time points). Second, brain systems are

interconnected and neural activity is expected to be coherent and related to behavior, unlike many data mining situations, in which individual measures are mostly independent. Therefore integrative analysis approaches are important for neuroscience.

Big data in neuroscience also differs from the other recent expansion of statistics into physiology – into big genomic data – in several important ways. First, genes have distinct identities and functions and many have a rich annotation of function, whereas neurons in a recording of vertebrate brains are largely indistinguishable, and have no previously annotated known functions. This makes comparisons between individuals more difficult: for the most part corresponding genes function the same way in different individuals, but specific neurons in a specific region of one individual's brain will not correspond in function to specific cells in the corresponding region of another. fMRI is somewhat more like genomic data because functions for many brain regions have been broadly characterized, and many brain regions can be roughly aligned between individual brains; nevertheless there are consistent inter-individual differences in usage of some corresponding brain regions. Second, it is rare to gather data from the new technologies on many individuals, because experiments in living subjects are costly and require painstaking preparation, unlike genomics where the assay technology may be expensive but the sample preparations are fairly routine. Again fMRI data occupies an intermediate position, the technology is expensive, but the protocols are less demanding than the new imaging and recording technologies.

2. NEW TECHNOLOGIES AND PRE-PROCESSING ISSUES

Most of the revolutionary new technologies in functional neuroscience depend on imaging in some way. A number of new molecular tools enable researchers to generate visible light from invisible changes in membrane voltages, or from changes in concentrations of key signalling molecules, such as calcium orglutamate. The raw data are usually light intensity at many points in a series of images; the biologically interesting events, such as neurons firing, or membrane voltage, have to be inferred from the raw data.

The new high-throughput methods for functional neuroscience depend on pushing technology to its limits. Therefore issues about processing raw data play a much bigger role than in fields using established and well-characterized instruments. Many new technologies have strong, but poorly understood, artifacts, and to address these problems requires both sophisticated statistical technique and a deep understanding of the technology.

New imaging technologies are not the only methods which need statistical attention. Even methods for pre-processing raw data from long-standing technologies leave much to be desired. Recordings from extra-cellular electrode have been the standard way to identify and count 'spikes' (the firings of neurons) for fifty years: each electrode picks up voltage fluctuations from many nearby neurons. Usually some form of multivariate analysis is used to identify distinctive voltage fluctuation patterns resulting from repeated firing of the same neurons (this process is called 'spike sorting'). However current methods can reliably distinguish only one or two neurons by their distinctive wave-forms out of the many neurons surrounding the electrode. This problem introduces a strong bias toward highly active neurons to the downstream analysis because most low-firing rate neurons don't have enough exemplars to be reliably identified and separated from each other or from the more active neurons. Furthermore changing conditions may induce changes in the wave-form from a single neuron, and sometimes spikes from a single neuron get assigned as two separate neurons (Hill *et al.* 2011, Neymotin *et al.* 2011). Thus spike sorting remains more art than science and there is considerable scope for advances in pre-processing methods.

However the new imaging technologies are most obviously in need of better pre-processing. One such method is calcium imaging. Calcium is released into the cell cytoplasm when a neuron fires. In calcium imaging a fluorescent dye binds the released calcium and emits light that is detected by a camera (Grienberger and Konnerth 2012). In newer methods the natural calcium binding proteins are linked to a fluorophore by genetic engineering (Grienberger *et al*. 2012). However spikes occur at a time-scale of milliseconds, while the calcium response occurs over tens of milliseconds and the dye response over hundreds of milliseconds, thus blurring together individual spikes. Thus it becomes a problem

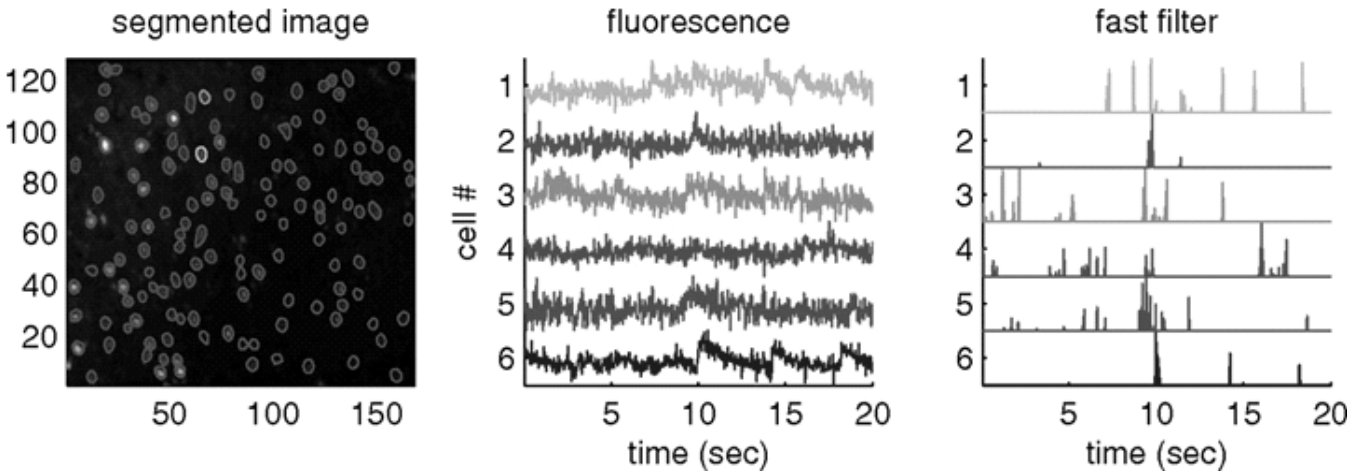

Fig. 2. Middle panel shows typical calcium traces from several individual cells shown at left, identified with colors, with inferred spikes after method of (Vogelstein *et al.* 2010) shown at right. [Copyright J Neurosci; permission applied for]

to identify individual spikes of a single cell, on a time scale of neural activity (see Fig. 2). It is a natural to think of the mapping from spikes to calcium signal as a convolution, and several deconvolution methods have been developed (Vogelstein *et al*. 2010). However the natural variability in neuron calcium signalling has frustrated the application of these methods. Furthermore in very dense neural tissue it is often hard to identify individual cells, which may appear in an image at several separated locations, which not be counted as separate cells. This kind of problem invites approaches using multivariate techniques such as Independent Components Analysis (ICA) (Mukamel *et al.* 2009).

Another emerging technology uses voltage-sensitive dyes (VSDs) or more recently genetically encoded voltage indicators, to translate voltage changes into changes in visible light emitted (Shoham *et al*. 1999). These indicators potentially can give very high resolution data over a very wide field, limited only by the resolution of the camera. Voltage-sensitive indicators have very rapid rise times, comparable to spike times, but somewhat slower decay times (tens of milliseconds) (Petersen *et al.* 2003); the dynamic response seems to vary somewhat over the surface of cortex. How best to pre-process VSD data remains an unsolved problem, and several artifacts are known. Other kinds of high-throughput measures have their own characteristic artifacts and biases, but little modeling of these has been done so far.

3. EXPLORATORY DATA ANALYSIS

In popular press interviews with neuroscientists, one refrain often heard is that "we don't yet know the neural code"; this means that no-one really knows how the firing of a particular neuron, or a group of neurons, might relate to behavior. Therefore for the foreseeable future exploratory data analysis will be an important aspect of what neural data analysts do.

When dealing with so many variables, a statistician's first instinct is to do some sort of dimension reduction. However one of the frustrations of most neuroscience analysis is that the first biological principal component (PC) typically explains less (often much less) than 20% of the variance; conversely if a first PC explains much of the variance, it is usually an artifact. Often even five biological principal components do not explain as much as 50% of the variance in a data set (Fig. 3). This may reasonably be

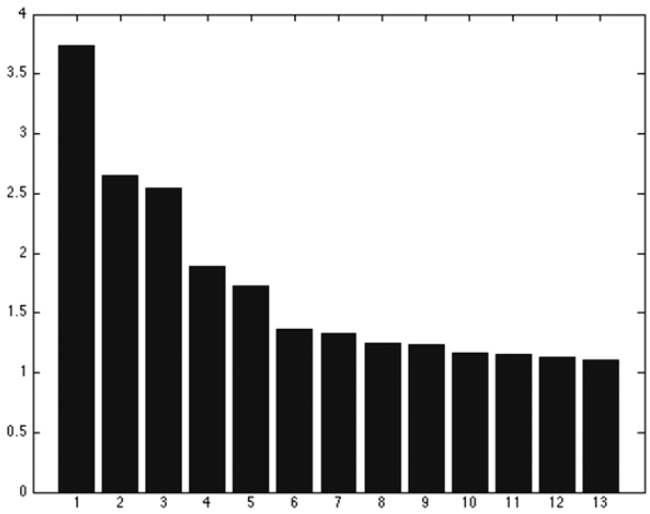

Fig. 3. Typical scree plot from PCA of 72 neurons in bins of 100ms. Note the long tail of PC variances. Unpublished data, courtesy of David Euston (same data set as Fig. 1).

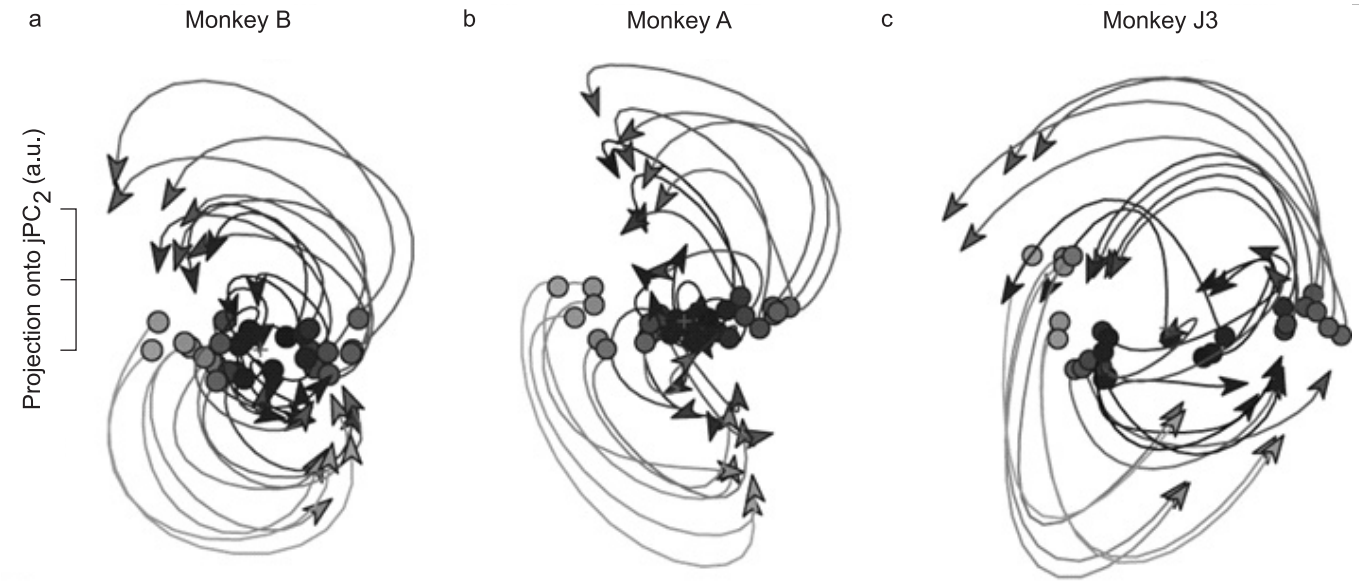

Fig. 4. Rotational structure in the dynamics of neural activity during reaching uncovered through jPCA method of (Churchland *et al.* 2012). Axes represent projection of first six PC axes onto a two-dimensional skew-symmetric subspace. Each trace represents on those axes the mean of many trajectories during reaching attempts under the same conditions; several conditions were tried. Green trajectories represent reaching to right and red reaching to left. [Copyright Nature; permission applied for]

expected, since there is much spontaneous activity in brain, unrelated to the task being measured. Nevertheless classical dimension reduction techniques often do capture some meaningful patterns in neuroscience data, and have been energetically applied to fMRI data. Variations of the classical techniques have been developed to try to adapt to characteristics of spiking data. Other multivariate exploratory techniques, such as Hidden Markov Models (HMMs) and cluster analysis, have been applied to capture important transitions in activity patterns.

Principal components analysis (PCA) has been applied several times over the past two decades to identify coordinated firing patterns in ensembles of cells over time, and to understand their possible relation to function (Churchland *et al.* 2007, Peyrache *et al.* 2010). Unlike most applications of PCA, in applications to spike train data the loadings are of little interest, since the neurons are not otherwise identifiable; rather the interest is in the dynamic changes over time in the PC scores – the projections onto the eigenvectors of the correlation matrix associated with the largest eigenvalues – and in how they co-vary with interesting aspects of animal behavior.

In order to make continuous data out of multiple spike trains, the spikes are usually aggregated into bins of 20ms – 100ms, depending on the typical firing rates. Neurons with very few recorded spikes are often discarded at this stage. Then PCA is applied to the spike counts, which may be further variance-stabilized by a square root transform. This processing step typically gives time-series of several tens of thousands of steps for each neuron. Under standard asymptotic theory, then PCA should accurately identify several principal components, provided enough spikes occur in each bin to make a Gaussian approximation to the counts (or square root-counts) realistic. Because neural firing rates are highly skewed (Fig. 1), usually scaled PCA is done.

Several variations on this procedure have been proposed. Since it is generally easier to interpret positive loadings than negative loadings, factor analysis techniques applied to the count data, especially sparse factor analysis approaches, seem promising (Reimers, in preparation). Yu and colleagues introduced a mathematically more sophisticated approach called Gaussian Process Factor Analysis (GPFA) (Yu *et al*. 2009). They consider each neuron being recorded as a noisy reflection of an underlying smooth process, then use Gaussian Process theory to extract many smooth trajectories from the noisy, high-dimensional recorded firings. It is not clear how reliable these extracted latent factors are, if one compares data from different sampling epochs.

Some of the same researchers focused on identifying subspaces with specific characteristics within the leading PCA subspaces (Churchland *et al*. 2012). Their method, which they call jPCA, aims to identify a two-dimensional rotational component within the first six PCs of binned data from spike trains. The jPCA approach seems to bring out some striking rotational structure in recordings from primary motor cortex (see Fig. 4), but again the validity of these kinds of latent factors has not been studied statistically.

Four kinds of problems with all such procedures remain to be addressed. One is that the spike counts are typically low, since only a few principal neurons from most recorded regions (*e.g*. cortex or hippocampus) will fire more than five times in a typical bin. Therefore the Gaussian model that underlies all these multivariate techniques seems unlikely to be even approximately valid in most cases. The appealing application of jPCA to motor cortex, where some neurons fire more often than typical cortical neurons, may not generalize across most other regions. Second there is considerable evidence that timing relations among neurons are important on a scale of a few milliseconds, and neuroscientists are therefore very interested in understanding these timing relations. However in order to capture enough spikes for even a minority of neurons researchers use coarse bins of roughly 50 or 100 ms, which largely obscure this crucial timing information. Thirdly there is evidence for the importance of changes in the ways neurons fire together over time, especially as an animal learns a task. The classical multivariate measures and their new variants may pick up shifts over time from activity in some ensembles to activity in other ensembles, but do not help us characterize changes in coordination among ensembles. Finally none of these methods can represent a majority of the variance in the recorded datain just a few time-varying factors, and it remains unclear what exactly is the significance of the latent factors being identified. Therefore there is considerable scope for inventive statisticians to improve even this most basic phase of exploratory analysis.

A natural complement to exploratory multivariate analysis of neural activity measures, and then seeking correlations with behavior, is multivariate classification of behavior from neural activity measures. Can we use measures of many neurons, or patterns of brain activity over many regions, to infer what the animal or person is doing, or even what they are thinking, at the time of measurement? The neuroscience jargon term for this is 'decoding'. Such 'decoding' has been a staple of science fiction for a century, and of neuroscientific studies of navigation and decision-making for over a decade. In the hippocampus, an area important for navigation, and also a subject of intense research, most cells fire at peak rates when an animal is within a specific location, and rarely at other places; often firing rate is strongly dependent on the direction of motion. Such sparseness makes Bayesian approaches to inferring location and trajectory from patterns of cell activity fairly successful (Brown *et al.* 1998). Similar approaches have been applied to infer future paths the animal may take at decision points (Johnson and Redish 2007, Pfeiffer and Foster 2013). Similar methods have been extended to humans using fMRI (Hassabis *et al*. 2009). In primary sensory regions of cortex, in which neurons are clustered to form functional maps of the body or of space, these approaches also successful; from fMRI measures in visual cortex (Nishimoto *et al*. 2011) were able to reconstruct some striking images of what the subject in the scanner was seeing at the time of acquisition. Others have been able to reconstruct which image from a finite set a person is imagining at any one time (Norman *et al.* 2006). These 'decoding' techniques are fairly straightforward conceptually though technically demanding, often requiring advanced machine learning techniques. Moreover most such studies have worked with whole brain imaging, or with regions of the brain, such as visual cortex, in which sparse representation is the norm. This is not the case for most regions of the brain, and more innovative statistical methods may be needed to infer behavior or imagined objects from activity patterns in those areas. Moreover these methods crucially depend on choosing among a well-characterized and practiced finite set of real-world activities or imagined objects as counterparts to neural activity. There is a strong need for methods to characterize how neurons may respond to novel situations in relation to their activity during previous experience.

Neuroscientists want to characterize changes in neural activity in relation to changes in behavior for insight on how the mind shifts attention, or even how new insights form. It is difficult in principle to distinguish changes in patterns of activity across neural ensembles from neural plasticity (of which more below); for now let us consider the former to be changes in firing rates correlated with changes in

circumstance or behavior, and the latter to be changes in firing rates in the same circumstance and behavior. Changes in firing rates of individual neurons in relation to stimuli have been studied for decades, but recordings from multiple neurons allow for a more powerful characterization of systematic changes in firing patterns over a brain region. Durstowitz and colleagues have explored several statistical methods to capture abrupt transitions and have used multivariate characterizations of neural activity to show abrupt changes when the animal switches from one choice to another (Durstewitz *et al.* 2010). This observation has motivated them further to use Hidden Markov Models (HMMs) in an attempt to identify change points (Caracheo *et al.* 2013), and others have observed similar changes (Karlsson *et al.* 2012).

However it remains unclear whether the abruptness of transitions is primarily attributable to the sharply distinct behavioral alternatives available in these experiments, or is a general feature of neural dynamics. There is a need to use statistical techniques to characterize the dynamics and rates of changes in neural activity in freely roaming animals, where most behavioral transitions are more gradual.

A question related to that of detecting changes in patterns of neural activity is whether the same patterns of neural activity are ever repeated in a stereotyped way. The idea that there might be repeating patterns suggests that relative timing of individual neuron firings might be crucial for their function; this idea is attractive if one thinks that such repeating motifs might be the basis of a kind of memory. There have been several reports of statistically significant recurrence of precise patterns (Ikegaya *et al*. 2004). However the statistical significance of these apparent recurrence is still very much debated (Mokeichev *et al*. 2007, Roxin *et al.* 2008), because the problem of specifying a null distribution is more sophisticated than simple Poisson statistics would lead one to expect.

4. COMPUTATIONAL MODELING AND DATA ANALYSIS

Many researchers have been fascinated by the prospect of building computer models of brain function. So-called 'neural network models' have even given rise to effective machine-learning approaches, although few neuroscientists think these machine learning tools are models for how the brain functions. Nevertheless there is an active community building more sophisticated computational models intending to represent and study some aspects of cognitive function. It would seem natural that they would want to connect their models with data on actual brain function. Nevertheless although many groups have tried to match certain qualitative features between their models and observed behavior, most have been diffident about quantitatively comparing models with activity measures from real brains during performance of tasks they are modeling.

This is changing. Recently (Hunt *et al*. 2012) simulated three popular models for decision-making, while storing summaries of aggregate simulated activity, and used these to predict the measured neural signal itself. That is they compared their model's total synaptic input to the high time resolution physiological MEG signal from specific brain regions recorded from individuals doing an analogous task. The recorded dynamics clearly favored one model of decision-making over the others.

Not many examples yet exist of such detailed comparisons, and in these cases reduced summary statistics are computed and compared with putative analogs derived from real data. However the increasing sophistication of both modeling and data acquisition suggest that such comparisons will be a major enterprise, and statistical issues of adequate correspondence between trajectories will come to the fore.

5. CHARACTERIZING PLASTICITY

Neural plasticity refers to the changing of connection strengths as a result of experience, which in turn changes dynamics of neural activity. Plasticity is one of the distinctive functions of the brain, central to much research, and yet it remains difficult to characterize its effects on activity patterns quantitatively. There are several distinct molecular mechanisms for plasticity, affecting firing patterns at several time scales, and so there may not be a single measure of plasticity that suits all purposes. Nevertheless this is an area that needs attention from statisticians as bigger data enables more information to be gathered.

McNaughton and colleagues introduced the currently most common measure of plasticity induced by a learning (or training) experience, based on

comparing correlation matrices of spike counts of many neurons, before and after a training episode (McNaughton 1998). Researchers first aggregate the spikes in short time bins as for PCA, then they compute correlations among spike counts both before and after the training episode. The difference between these two correlation matrices is then regressed element-wise on the correlation matrix of the same neurons' spike counts during the training episode. The measure of learning is the R^2 of this regression, and hence the measure is known as Explained Variance (EV). Typical EV values in learning trials are 10% - 20%; but this number does not identify which connections may have been modified nor characterize the change in neural dynamics. More sophisticated measures are sorely needed.

6. OUTLOOK FOR THE COMING DECADE

We already see publications based on imaging tens of thousands of individual neurons in behaving animals at a (relatively) slow time-scale of one second (Ahrens *et al*. 2012, Ahrens *et al.* 2013a, Ahrens *et al*. 2013b) using calcium imaging. They extracted relatively coarse activation statistics, but these technologies are advancing rapidly, and we may expect within a few years data sets at time resolutions of under 100 ms, which should expose the temporal structure of rapid network dynamics (Misha Ahrens, personal communication). Statisticians will be able to contribute to the analysis of this rich data resource.

We may also expect closer integration of computational neuroscience, already a very active research endeavor, with analysis of big data sets. At first this will likely be through derived summary statistics. It seems unlikely that any time soon scientists will be able to model details of specific real neural circuits, and expect to predict the time series of individual neural activity. Nevertheless the Human Brain Project in EU (www.humanbrainproject.eu) seems poised to attempt such predictions.

Who might want to participate in the big data revolution in neuroscience? Although neuroscience is a very specific application domain, requiring much domain-specific knowledge, it is also technically very demanding. It will appeal to those with broad interdisciplinary interests, who will enjoy learning the necessary neuroscience and relish the mathematical and computing challenges.

REFERENCES

Ahrens, M.B., Huang, K.H., Narayan, S., Mensh, B.D. and Engert, F. (2013a). Two-photon calcium imaging during fictive navigation in virtual environments. *Frontiers Neural Circuits*, **7**, 104.

Ahrens, M.B. *et al*. (2012). Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature*, **485**, 471-477.

Ahrens, M.B., Orger, M.B., Robson, D.N., Li, J.M. and Keller, P.J. (2013b). Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nature Methods*, **10,** 413-420.

Brown, E.N., Frank, L.M., Tang, D., Quirk, M.C. and Wilson, M.A. (1998). A statistical poaradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *The J. Neurosci.: the official journal of the Society for Neuroscience*, **18,** 7411-7425.

Brown, E.N., Kass, R.E. and Mitra, P.P. (2004). Multiple neural spike train data analysis: State-of-the-art and future challenges. *Nature Neurosci.*, **7**, 456-461.

Caracheo, B.F., Emberly, E., Hadizadeh, S., Hyman, J.M. and Seamans, J.K. (2013). Abrupt changes in the patterns and complexity of anterior cingulate cortex activity when food is introduced into an environment. *Frontiers Neurosci.*, **7**, 74.

Churchland, M.M. *et al*. (2012). Neural population dynamics during reaching. *Nature*, **487**, 51-56.

Churchland, M.M., Yu, B.M., Sahani, M. and Shenoy, K.V. (2007). Techniques for extracting single-trial activity patterns from large-scale neural recordings. *Current Opinion Neurobiol.*, **17,** 609-618.

Durstewitz, D., Vittoz, N.M., Floresco, S.B. and Seamans, J.K. (2010). Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron*, **66,** 438-448.

Grienberger, C. and Konnerth, A. (2012). Imaging calcium in neurons. *Neuron*, **73,** 862-885.

Hassabis, D. *et al*. (2009). Decoding neuronal ensembles in the human hippocampus. *Current Biol.*, **19,** 546-554.

Hill, D.N., Mehta, S.B. and Kleinfeld, D. (2011). Quality metrics to accompany spike sorting of extracellular signals. *The J. Neurosci.: the official journal of the Society for Neuroscience*, **31**, 8699-8705.

Hunt, L.T. *et al*. (2012). Mechanisms underlying cortical activity during value-guided choice. *Nature Neurosci.*, **15**, 470-476, S471-473.

Ikegaya, Y. *et al*. (2004). Synfire chains and cortical songs: Temporal modules of cortical activity. *Science*, **304,** 559-564.

Johnson, A. and Redish, A. D. (2007). Neural ensembles in Ca3 transiently encode paths forward of the animal at a decision point. *The J. Neurosci. : the official journal of the Society for Neuroscience*, **27,** 12176-12189.

Karlsson, M.P., Tervo, D.G. and Karpova, A.Y. (2012). Network resets in medial prefrontal cortex mark the onset of behavioral uncertainty. *Science*, **338,** 135-139.

McNaughton, B.L. (1998). The neurophysiology of reminiscence. *Neurobiol. Learning Memory*, **70,** 252-267.

Mohajerani, M.H. *et al*. (2013). Spontaneous cortical activity alternates between motifs defined by regional axonal projections. *Nature Neurosci.*, **16**, 1426-1435.

Mokeichev, A. *et al*. (2007). Stochastic emergence of repeating cortical motifs in spontaneous membrane potential fluctuations in vivo. *Neuron*, **53,** 413-425.

Mukamel, E.A., Nimmerjahn, A. and Schnitzer, M.J. (2009). Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron*, **63,** 747-760.

Neymotin, S.A., Lytton, W.W., Olypher, A.V. and Fenton, A.A. (2011). Measuring the quality of neuronal identification in ensemble recordings. *The J. Neurosci.: the official journal of the Society for Neuroscience*, **31,** 16398-16409.

Nishimoto, S. *et al*. (2011). Reconstructing visual experiences from brain activity evoked by natural movies. *Current Biol.*, **21,** 1641-1646.

Norman, K.A., Polyn, S.M., Detre, G.J. and Haxby, J.V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fmri data. *Trends Cognitive Scis.*, **10,** 424-430.

Petersen, C.C., Grinvald, A. and Sakmann, B. (2003). Spatiotemporal dynamics of sensory responses in layer 2/3 of rat barrel cortex measured in vivo by voltage-sensitive dye imaging combined with whole-cell voltage recordings and neuron reconstructions. *The J. Neurosci.: the official journal of the Society for Neuroscience*, **23,** 1298-1309.

Peyrache, A., Benchenane, K., Khamassi, M., Wiener, S.I. and Battaglia, F.P. (2010). Principal component analysis of ensemble recordings reveals cell assemblies at high temporal resolution. *J. Comput. Neurosci.*, **29**, 309-325.

Pfeiffer, B.E. and Foster, D.J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, **497**, 74-79.

Roxin, A., Hakim, V. and Brunel, N. (2008). The statistics of repeating patterns of cortical activity can be reproduced by a model network of stochastic binary neurons. *The J. Neurosci.: the official journal of the Society for Neuroscience*, **28,** 10734-10745.

Shoham, D. *et al*. (1999). Imaging cortical dynamics at high spatial and temporal resolution with novel blue voltage-sensitive dyes. *Neuron*, **24,** 791-802.

Vogelstein, J.T. *et al*. (2010). Fast nonnegative deconvolution for spike train inference from population calcium imaging. *J. Neurophysiol.*, **104,** 3691-3704.

Yu, B. M. *et al*. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.*, **102**, 614-635.