4-2014

Statistical Research and Training Under the Brain Initiative

Robert E. Kass
Carnegie Mellon University, kass@andrew.cmu.edu

Genevera Allen
Rice University

Brian S. Caffo
Johns Hopkins University

John Cunningham
Columbia University

Uri Eden
Boston University

See next page for additional authors

Authors
Robert E. Kass, Genevera Allen, Brian S. Caffo, John Cunningham, Uri Eden, Timothy D. Johnson, Martin A. Lindquist, Thomas A. Nichols, Hernando Ombao, Liam Paninski, Russell T. Shinohara, and Bin Yu

STATISTICAL RESEARCH AND TRAINING UNDER THE BRAIN INITIATIVE

A Working Group of the American Statistical Association*

April 2014

1 Introduction and Summary

The BRAIN (Brain Research through Advancing Innovative Neurotechnologies) Initiative aims to produce a sophisticated understanding of the link between brain and behavior and to uncover new ways to treat, prevent and cure brain disorders.[1] Success in meeting these multifaceted challenges will require scientific and technological paradigms that incorporate novel statistical methods for data acquisition and analysis. Our purpose here is to substantiate this proposition, and to identify implications for training.

Brain research relies on a wide variety of existing methods for collecting human and animal neural data, including neuroimaging (radiography, fMRI, MEG, PET), electrophysiology from multiple electrodes (EEG, ECoG, LFP, spike trains), calcium imaging, optical imaging, optogenetics, and anatomical methods (diffusion imaging, electron microscopy, fluorescent microscopy). Each of these modalities produces data with its own set of statistical and analytical challenges. As neuroscientists improve these techniques and develop new ones, data are being acquired at very large scales. For example, advances in multiple-electrode recording and two-photon calcium imaging have led to an exponential growth in the size of neural populations that can be observed simultaneously, at single-cell resolution (2; 3; 52). Similarly, new anatomical methods have led to a rapid rise in the size and the scale of data, and the resulting level of detail with which brain structure can be investigated (17; 10; 32; 39). Furthermore, both new and existing technologies are often used together, and are increasingly accompanied by rich characterizations of individuals and their behavior, ranging from genetic information to sensor-based monitors of activity.

*The working group included Robert E. Kass, Carnegie Mellon University (Chair); Genevera Allen, Rice University; Brian Caffo, Johns Hopkins University; John Cunningham, Columbia University; Uri Eden, Boston University; Timothy D. Johnson, University of Michigan; Martin A. Lindquist, Johns Hopkins University; Thomas A Nichols, University of Warwick; Hernando Ombao, University of California, Irvine; Liam Paninski, Columbia University; Russell T. Shinohara, University of Pennsylvania; Bin Yu, University of California, Berkeley. Affiliations are for identification purposes only and do not imply an institution's endorsement of this document.

[1] http://www.whitehouse.gov/share/brain-initiative

These advances have begun to produce exciting breakthroughs. But, to realize their potential, new analysis and computational techniques are needed to optimize data acquisition, manage acquired data on the fly, screen and segment the data, correct for artifacts, and align and register data across multiple time points, multiple experiments, multiple subjects, or different laboratories. In addition, as the data-generation process becomes more complex, and the data sets themselves get larger and more varied, it is crucial that reliability and scientific relevance of results be assessed against the backdrop of natural variation and measurement noise. This is the essential role of statistical analysis.

It is important to keep in mind the fundamental advantage of large and varied data sets: it is the ability to pose questions at a much finer level of detail than would otherwise be available. For instance, it is becoming possible to investigate the connectivity of neural networks, and the way connectivity changes across subjects or behavioral circumstances. Rodents, non-human primates, and, under special circumstances, even humans, can be monitored by many kinds of sensors, including video, while large numbers of neurons are recorded in relevant brain regions. As the questions being asked become increasingly specialized to particular circumstances, however, the number of possible configurations of outcomes grows rapidly. Major challenges in such Big Data applications, therefore, come from the resulting complexity of the statistical models used to describe the variation in the data: they must capture key features of the data, while allowing relevant assessments concerning scientific issues, and their implementation must be made efficient enough to run on available hardware and software in manageable time. This leads to a host of statistical machine learning problems. Some of these are computational, and algorithmic, while others are inferential. In many brain-related contexts, data scientists of various backgrounds have made great progress in developing algorithms to mine data in order to produce snapshots of network structure. As promising as such quantitative summaries and visualizations appear, however, they are useless without some notion of their reliability. Procedures are needed to conduct rigorous statistical inference in conjunction with statistical machine learning algorithms, and attention must be paid to reproducibility of results across relevant populations. Furthermore, consideration of data analytic techniques can produce new ways to design experiments or important improvements in recording technologies (34; 35; 47).

Part of this effort involves careful consideration of theoretical conceptions of brain function, phrased in the language of mathematics. By merging ideas from statistics together with those from applied mathematics and biophysics, productive new frameworks are likely to be spawned.

As the Interim Report of the NIH BRAIN initiative working group says,[2]

> All signs point to a major increase in the quantity of neuronal recordings. ... We will need new tools to analyze these complex datasets, as well as new tools and algorithms for data acquisition and interpretation. ... Data analytic and theoretical problems are likely to emerge that we cannot anticipate at the present time. Resources should be available for experts from essential disciplines such as statistics, optimization, signal processing and machine learning. ... To enable progress in theory and data analysis, we must foster collaborations between experimentalists and scientists from statistics, physics, mathematics, engineering, and computer science.

We elaborate briefly on these points below, considering major research areas and the need for more intensive cross-disciplinary training.

[2] http://www.nih.gov/science/brain/

2 Research

The methods of neuroscience involve behavioral studies, neuroimaging, electrophysiology, and anatomy, as well as cellular and molecular methods. All are crucial to understanding brain function. The approaches that have, so far, generated the most research in the mathematical sciences are neuroimaging and electrophysiology.

Neuroimaging is widely applied, familiar to non-specialists through the popular press, and likely to remain central to the BRAIN initiative. The field has gained enormously from statistical research. Over the past 20 years, for example, fMRI has become a standard methodology in cognitive neuroscience in large part due to a massive communal effort to develop useful statistical procedures and software (8; 14; 25; 41; 51; 56). Earlier, in tomography and nuclear medicine, mathematical scientists and statisticians played similar roles (30; 49; 50). In electrophysiology, statistical methods have also been crucial in advancing the field, e.g., (12; 13; 45), and see examples in (28). Historically, the key to progress has been insight in mathematics and statistics informed by knowledge of the underlying biophysics.

Under the BRAIN initiative, new analytical issues arise not only from improved recording technologies, but also from the availability of vast new repositories of data, some of which contain diverse data types, such as the 1,000 functional connectome, NIH connectome, ADHD 200, and ABIDE (4; 21; 22; 40; 54; 55). A central bottleneck is the need to combine different sets of data, often called data fusion. This includes combining data from different human subjects. While computational neuroanatomy and shape analysis have been successful in particular contexts (6; 26), essential anatomical distinctions between individuals make the problem of inter-subject comparison difficult. New statistical insights are needed that can incorporate biologically-informed, flexible templates in order to make effective use of these key data sources being driven by the BRAIN initiative and related activities.

Fundamental issues also arise in electrophysiology. New methods are needed to extract information from the large spatio-temporal data sets produced by calcium imaging, particularly in cases where many neuronal processes overlap spatially. These would include efficient methods for de-noising, de-convolving, and de-mixing of calcium-sensitive and voltage-sensitive movie data into estimates of single-neuronal (or sub-neuronal) activity, along with measures of uncertainty about these estimates (46). Similar “spike-sorting” problems arise in the context of large-scale electrical recordings from multielectrode arrays (15; 16).

These examples highlight the need for attention to statistical detail when attempts are made to tap into the great new collections of data that are becoming available. Studies having the even grander ambition of integrating knowledge across multiple levels of investigation by explicitly connecting cellular properties with anatomical constraints, physiology, and/or behavior could be especially valuable, but they are likely to produce varied data that introduce unfamiliar and potentially subtle complications. To be successful, research teams will have to be interdisciplinary, including members with cutting-edge statistical expertise. Below we discuss a few example problems where this kind of interdisciplinary approach promises to be especially fruitful.

2.1 Mapping and modeling of neuronal networks

One of the greatest challenges in neuroscience is to reconstruct and interpret dynamic connectivity of large, complex neuronal networks. The ability to understand the way such networks evolve under varying experimental and/or behavioral conditions would have a fundamental impact on understanding the functional organization of the nervous system. Indeed, electrophysiological and neuroimaging studies are undergoing a paradigm shift toward network-level hypotheses and analyses.

It is important to distinguish anatomical from functional connectivity. Anatomical approaches aim to reconstruct the physical connectivity of the network. This could include viral tracing (23) and volumetric electron microscopy (i.e., slicing the brain very thin) and then tracking the axons and dendrites travelling from one cell body to another, typically using a combination of computer vision techniques along with substantial hands-on manpower. Recent genetic methods provide a complementary technique (57). Functional approaches aim to identify the influence that each neuron in a network has upon the other neurons in the network by, for example, stimulating one neuron while recording the strength of the synaptic response in another neuron, or by recording from many neurons simultaneously and then inferring network weights statistically. The burgeoning field of statistical network theory can potentially offer many useful tools and ideas for estimating and interpreting these neuronal networks.

Crucial problems involve assessment of uncertainty due to incomplete observation of the network as well as biological variation (across specimens or subjects). Again, proper experimental design can have an important impact (29). Statistical work should also involve the combination of different types of experimental data from two or more of these approaches, perhaps at different spatial scales: for example, information from functional recordings can constrain and guide efforts in anatomical reconstructions (5; 11; 53). An additional, major outstanding problem is that there does not exist a comprehensive statistical methodology for relating dynamic networks to experimental conditions, features of behavior, or characteristics of subjects. This should be a focus of research attention as BRAIN-related network data start to accrue.

2.2 Finding lower-dimensional dynamic structure

In the case of electrophysiology, analysis of networks of co-varying neurons aims to illuminate the remarkable computational abilities of the nervous system. A starting point is to acknowledge that neural responses are not independent degrees of freedom, but rather they are noisy observations of a scientifically meaningful hidden system of lower dimension. This system often has interesting spatiotemporal structure across that low-dimensional space. Accordingly, questions of neural population response dimensionality, time-series analysis, and dynamical systems are now critical to the neuroscience community. Basic methods of this sort have been fruitfully applied to population recordings in studies of the motor system (19), decision making (36), working memory (48), visual attention (20), the auditory system (33), the olfactory system (37), rule learning (24), speech (7), and more. These studies only represent a beginning. As a specific example, recent studies (19) have demonstrated the presence of non-trivial low-dimensional dynamics, but most key questions remain unanswered. Does a given brain area switch between different classes of dynamics? Given that many of these critical transitions have no externally observable behavior (e.g., decision making or attention) and can occur at arbitrary times even in well controlled experiments, how can these transitions be detected from neural population responses alone? New analytical tools will be needed to reveal the intrinsic dimensionality within covarying multidimensional patterns across a neuronal network that evolves across time.

Related issues arise in neuroimaging, where the physiological process underlying the measurements is expected to have lower-dimensional structure that produces large numbers of time-varying signals. In the case of EEG or MEG, thousands of brain-source time series are generated across millions of time points. In fMRI there are hundreds of thousands of brain sources across hundreds of time points. Furthermore, studies involve multiple subjects whose brain anatomy and function varies in subtle ways. Analyses must take account of both slowly-evolving mean effects and more rapidly-evolving variance and covariance effects, with the latter typically analyzed via frequency decompositions. In the frequency domain it is, in principle, possible to use well-established methods of dimension reduction by assuming the signals to be stationary linear processes. This requires segmentation of time into short blocks during which the signals are time-invariant, and methods exist for this purpose (43). These, however, are computationally intensive and are currently unable to handle large numbers of simultaneous signals. New methods that allow similar capabilities while scaling up to much larger

data sets would provide flexible identification of interactions among time-varying networks of brain activity.

2.3 Optimal spatiotemporal control of large neural populations

Most of the topics discussed above involve recording information from the brain. There are also a number of applications that involve stimulating the brain, using electrical, magnetic, or optical input signals. Indeed, one of the most exciting recent advances in neuroscience involves new optogenetic stimulation methods that allow scientists to excite or inhibit genetically-defined neural populations with fine temporal control (9). Most applications of these methods involve rather crude, purely temporal control stimuli, in which all of the targeted neurons in a given brain region are activated (or inactivated) simultaneously (i.e., the control signal has no particular spatial structure). An important and relatively open area for future research involves optimal spatiotemporal control of neural populations, in which different spatial locations are stimulated with different temporal control signals (1), potentially in a closed-loop paradigm, necessitating real-time statistical signal processing. Such applications may arise in the near future in the context of brain-machine interface systems involved in deep-brain stimulation, epilepsy control, or sensory prosthetics (42).

2.4 Data assimilation for biophysical models

Even as new technologies allow us to record from increasingly large neural populations, only a small subset of the signals related to neural activity (e.g., a neuron's voltage, or calcium flux) are typically observed. A fundamental question in neuroscience is how to link observed neural activity to the unobserved biophysical mechanisms that generate this activity. As emerging technologies accelerate the number of neurons observed simultaneously, targeted experimental manipulations that permit detailed assessment of these biophysical mechanisms become infeasible or intractable. There is therefore a critical need for data assimilation methods that can incorporate the limited and imprecise data that we observe with detailed, mechanistic models of neural activity. Such methods are sometimes described as a “mathematical microscope" that would permit researchers to look more deeply into the mechanisms responsible for data structure beyond what ordinary modeling, visualization, or statistical inference can provide.

Several groups have demonstrated recently that data assimilation methods can successfully identify classes of biophysical models that are consistent with observed activity using spikes, intracellular voltage recordings, or local fields in small-scale neural systems (27; 38). One important feature of such methods is that they can often identify a full space of models, rather than a single model, consistent with the observed data. An important statistical challenge moving forward is to develop data assimilation methods that scale well as the complexity of the mechanistic models and the size of the data sets increase.

3 Training

To achieve the goals of the BRAIN initiative, a multifaceted approach to statistical training must be undertaken. First, the next generation of biologically-oriented scientific investigators will need a solid understanding of statistical ideas and a broad familiarity with available techniques. Furthermore, computationally-oriented neuroscientists need to be exposed to statistical thinking, so that they can avoid wasted effort and be fully productive. Finally, statisticians, computer scientists, and bioengineers must be recruited to the brain sciences, and educated as to where their efforts will be best directed.

3.1 Training for neurobiologists

There currently seem to be relatively few doctoral programs in neurobiology that provide statistical training aimed specifically at these students. General courses in biostatistics can be helpful, but are not ideal. A suitable textbook has recently become available (28), and there will likely be some movement toward increased exposure to fundamental statistical ideas, but a significant hurdle is the difficulty of providing conceptual foundations, together with exposure to advanced methods, without getting bogged down in details and/or intimidating students who have minimal mathematical backgrounds. Additional teaching materials would be very valuable. (One example course, focusing on the statistical analysis of spike train data, is here: http://www.stat.columbia.edu/∼liam/teaching/neurostat-fall13/.) In particular, modularized electronic open access tools for instruction would both assist the creation of these courses and eliminate the need to duplicate efforts across universities.

3.2 Training for those with strong mathematical backgrounds

The lion's share of advanced data analysis in the brain sciences is carried out by neuroscientists trained in bioengineering, biophysics, and computer science. Many of these people have had little or no systematic exposure to the statistical paradigm. A course based on material such as that covered in (18; 28; 31; 44) could go a long way toward giving them essential tips that can keep them on the path toward effective application, and development, of statistical methods. Modularized instructional tools would again assist efforts. In fact, there is considerable overlap in content between the materials needed for statistical training of neurobiologists and quantitative neuroscientists.

We need to attract a greater share of our most talented mathematical, computational, and statistical scientists into brain-related research. To be effective, they will need to learn how to think like a laboratory scientist, and to use that thinking in designing data collection and analysis strategies. Trainees should take neuroscience classes, embed themselves in neuroscience labs, and get repeated practice in data-related projects. They should be critiqued and challenged on their data work, and be upheld to a writing standard comparable to that for a neuroscience graduate student. The trainees should also be educated in principles of ethical and effective collaborative behavior. Support for BRAIN-relevant interdisciplinary training programs in the mathematical sciences are desperately needed: they would provide appropriate incentives for increased production of effective researchers in this domain.

3.3 Specialized workshops

In addition to multi-course programs aimed at producing outstanding cross-disciplinary researchers, short-term workshops can be very useful: they can give quantitatively-oriented graduate students, and postdoctoral fellows, essential vocabulary and concepts they need in order to get involved in BRAIN-related statistical work, even when their home institution does not have a specific program in this area. Good examples are the former NIH-funded neuroinformatics course at Woods Hole; the former NSF-funded course on Mining and Modeling of Neuroscience Data at the Redwood Center, Berkeley; and the current NIH-funded Multi-Modal Neuroimaging Training Program at Carnegie Mellon and the University of Pittsburgh. These kinds of workshops should continue.

References

[1] AHMADIAN, Y., PACKER, A. M., YUSTE, R., AND PANINSKI, L. Designing optimal stimuli to control neuronal spike timing. *Journal of Neurophysiology 106 DOI: 10.1152/jn.00427.2010* (2011), 1038–1053.

[2] AHRENS, M. B., AND KELLER, P. J. Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nature Methods 10* (2013), 413–420.

[3] ALIVISATOS, A. P., ANDREWS, A. M., BOYDEN, E. S., CHUN, M., CHURCH, G. M., DEISSEROTH, K., DONOGHUE, J. P., FRASER, S. E., LIPPINCOTT-SCHWARTZ, J., LOOGER, L. L., MASMANIDIS, S., MCEUEN, P. L., NURMIKKO, A. V., PARK, H., PETERKA, D. S., REID, C., ROUKES, M. L., SCHERER, A., SCHNITZER, M., SEJNOWSKI, T. J., SHEPARD, K. L., TSAO, D., TURRIGIANO, G., WEISS, P. S., XU, C., YUSTE, R., AND ZHUANG, X. Nanotools for neuroscience and brain activity mapping. *ACS Nano 7* (2013), 1850–1866.

[4] BISWAL, B. B., MENNES, M., ZUO, X.-N., GOHEL, S., KELLY, C., SMITH, S. M., BECKMANN, C. F., ADELSTEIN, J. S., BUCKNER, R. L., COLCOMBE, S., ET AL. Toward discovery science of human brain function. *Proceedings of the National Academy of Sciences 107* (2010), 4734–4739.

[5] BOCK, D. D., LEE, W.-C. A., KERLIN, A. M., ANDERMANN, M. L., HOOD, G., WETZEL, A. W., YURGENSON, S., SOUCY, E. R., KIM, H. S., AND REID, R. C. Network anatomy and in vivo physiology of visual cortical neurons. *Nature 471* (2011), 177–182.

[6] BOOKSTEIN, F. L. *Morphometric Tools for Landmark Data: Geometry and Biology*. Cambridge University Press, Cambridge (UK), 1991.

[7] BOUCHARD, K. E., MESGARANI, N., JOHNSON, K., AND CHANG, E. F. Functional organization of human sensorimotor cortex for speech articulation. *Nature 495* (2013), 327–32.

[8] BOWMAN, F. D., CAFFO, B., BASSETT, S. S., AND KILTS, C. A bayesian hierarchical framework for spatial modeling of fmri data. *NeuroImage 39* (2008), 146–156.

[9] BOYDEN, E. S., ZHANG, F., BAMBERG, E., NAGEL, G., AND DEISSEROTH, K. Millisecond-timescale, genetically targeted optical control of neural activity. *Nature Neuroscience 8* (2005), 1263–1268.

[10] BRIGGMAN, K. L., AND DENK, W. Towards neural circuit reconstruction with volume electron microscopy techniques. *Current Opinion in Neurobiology 16* (2006), 562–570.

[11] BRIGGMAN, K. L., HELMSTAEDTER, M., AND DENK, W. Wiring specificity in the direction-selectivity circuit of the retina. *Nature 471* (2011), 183–188.

[12] BROWN, E. N., FRANK, L. M., TANG, D., QUIRK, M. C., AND WILSON, M. A. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *Journal of Neuroscience 18* (1998), 7411–7425.

[13] BROWN, E. N., KASS, R. E., AND MITRA, P. P. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nature Neuroscience 7* (2004), 456–461.

[14] CALHOUN, V. D., LIU, J., AND ADALI, T. A review of group ica for fmri data and ica for joint inference of imaging, genetic, and erp data. *Neuroimage 45* (2009), S163–S172.

[15] CARLSON, D., RAO, V., VOGELSTEIN, J., AND CARIN, L. Real-time inference for a gamma process model of neural spiking. *Neural Information Processing Systems (NIPS)* (2013).

[16] CARLSON, D., VOGELSTEIN, Q. W., LIAN, W., ZHOU, M., STOETZNER, C. R., KIPKE, D., WEBER, D., DUNSON, D., AND CARIN, L. Sorting electrophysiological data via dictionary learning and mixture modeling. *IEEE Transactions on Biomedical Engineering* (2013).

[17] CHUNG, K., AND DEISSEROTH, K. CLARITY for mapping the nervous system. *Nature Methods 10* (2013), 508–513.

[18] CHUNG, M. K. *Statistical and Computational Methods in Brain Image Analysis*. World Scientific Publishing Company, Singapore, 2013.

[19] CHURCHLAND, M. M., CUNNINGHAM, J. P., KAUFMAN, M. T., FOSTER, J. D., NUYUJUKIAN, P., RYU, S. I., AND SHENOY, K. V. Neural population dynamics during reaching. *Nature 487* (2012), 51–56.

[20] COHEN, M. R., AND MAUNSELL, J. H. R. A neuronal population measure of attention predicts behavioral performance on individual trials. *The Journal of Neuroscience 30* (2010), 15241–53.

[21] CONSORTIUM, A.-., ET AL. The adhd-200 consortium: A model to advance the translational potential of neuroimaging in clinical neuroscience. *Frontiers in systems neuroscience 6* (2012).

[22] DI MARTINO, A., YAN, C. G., LI, Q., DENIO, E., CASTELLANOS, F. X., ALAERTS, K., ANDERSON, J. S., ASSAF, M., BOOKHEIMER, S. Y., DAPRETTO, M., ET AL. The autism brain imaging data exchange: Towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular psychiatry* (2013).

[23] DUM R P, S. P. L. Transneuronal tracing with neurotropic viruses reveals network macroarchitecture. *Curr Opin Neurobiol 23 doi: 10.1016/j.conb.2012.12.002* (2013), 245–249.

[24] DURSTEWITZ, D., VITTOZ, N. M., FLORESCO, S. B., AND SEAMANS, J. K. Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron 66* (2010), 438–448.

[25] FRISTON, K. J., HOLMES, A. P., WORSLEY, K. J., POLINE, J.-P., FRITH, C. D., AND FRACKOWIAK, R. S. Statistical parametric maps in functional imaging: A general linear approach. *Human brain mapping 2* (1994), 189–210.

[26] GRENANDER, U., AND MILLER, M. I. Representations of knowledge in complex systems. *Journal of the Royal Statistical Society. Series B (Methodological)* (1994), 549–603.

[27] HUYS, Q. J. M., AND PANINSKI, L. Smoothing of, and parameter estimation from, noisy biophysical recordings. *PLoS Computational Biology 5* (2009), e1000379.

[28] KASS, R. E., EDEN, U. T., AND BROWN, E. N. *Analysis of Neural Data*. Springer, New York, 2014.

[29] KESHRI, S., PNEVMATIKAKIS, E., PAKMAN, A., SHABABO, B., AND PANINSKI, L. A shotgun sampling solution for the common input problem in neural connectivity inference. *arXiv preprint arXiv:1309.3724* (2013).

[30] LANGE, K., CARSON, R., ET AL. EM reconstruction algorithms for emission and transmission tomography. *Journal of Computer Assisted Tomography 8* (1984), 306–316.

[31] LAZAR, N. A. *The Statistical Analysis of Functional MRI Data*. Springer, New York, 2008.

[32] LICHTMAN, J., LIVET, J., AND SANES, J. A technicolour approach to the connectome. *Nature Reviews Neuroscience 9* (2008), 417–422.

[33] LUCZAK, A., BARTHÓ, P., AND HARRIS, K. D. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron 62* (2009), 413–425.

[34] LUSTIG, M., DONOHO, D., AND PAULY, J. M. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic resonance in medicine 58* (2007), 1182–1195.

[35] LUSTIG, M., DONOHO, D. L., SANTOS, J. M., AND PAULY, J. M. Compressed sensing mri. *Signal Processing Magazine, IEEE 25* (2008), 72–82.

[36] MANTE, V., SUSSILLO, D., SHENOY, K. V., AND NEWSOME, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature 503* (2013), 78–84.

[37] MAZOR, O., AND LAURENT, G. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron 48* (2005), 661–73.

[38] MENG, L., KRAMER, M. A., MIDDLETON, S. J., WHITTINGTON, M. A., AND EDEN, U. T. A unified approach to linking experimental, statistical and computational analysis of spike train data. *PLoS ONE 9: e85269. doi:10.1371/journal.pone.0085269* (2014).

[39] MICHEVA, K. D., BUSSE, B. L., WEILER, N. C., O'ROURKE, N., AND SMITH, S. J. Single-synapse analysis of a diverse synapse population: Proteomic imaging methods and markers. *Neuron 68* (2010), 639–653.

[40] MILHAM, M. P. Open neuroscience solutions for the connectome-wide association era. *Neuron 73* (2012), 214–218.

[41] NICHOLS, T. E., AND HOLMES, A. P. Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Human Brain Mapping 15* (2002), 1–25.

[42] NIRENBERG, S., AND PANDARINATH, C. Retinal prosthetic strategy with the capacity to restore normal vision. *Proceedings of the National Academy of Sciences 109* (2012), 15012–15017.

[43] OMBAO, H., VON SACHS, R., AND GUO, W. SLEX analysis of multivariate nonstationary time series. *Journal of the American Statistical Association 100* (2005), 519–531.

[44] OZAKI, T. *Time Series Modeling of Neuroscience Data*. CRC Press, Boca Raton, Florida, 2012.

[45] PILLOW, J. W., SHLENS, J., PANINSKI, L., SHER, A., LITKE, A. M., CHICHILNISKY, E. J., AND SIMONCELLI, E. P. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature 454* (2008), 995–999.

[46] PNEVMATIKAKIS, E. A., MEREL, J., PAKMAN, A., AND PANINSKI, L. Bayesian spike inference from calcium imaging data. *Arxiv 1311.6864*.

[47] PNEVMATIKAKIS, E. A., AND PANINSKI, L. Sparse nonnegative deconvolution for compressive calcium imaging: Algorithms and phase transitions. In *Advances in Neural Information Processing Systems 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2013, pp. 1250–1258.

[48] RIGOTTI, M., BARAK, O., WARDEN, M. R., WANG, X.-J., DAW, N. D., MILLER, E. K., AND FUSI, S. The importance of mixed selectivity in complex cognitive tasks. *Nature 497* (2013), 585–590.

[49] SHEPP, L. A., AND LOGAN, B. F. The fourier reconstruction of a head section. *IEEE Transactions on Nuclear Science 21* (1974), 21–43.

[50] SHEPP, L. A., AND VARDI, Y. Maximum likelihood reconstruction for emission tomography. *IEEE Transactions on Medical Imaging 1* (1982), 113–122.

[51] SMITH, S. M., JENKINSON, M., WOOLRICH, M. W., BECKMANN, C. F., BEHRENS, T. E. J., JOHANSEN-BERG, H., BANNISTER, P. R., DE LUCA, M., DROBNJAK, I., FLITNEY, D. E., ET AL. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage 23* (2004), S208–S219.

[52] STEVENSON, I. H., AND KORDING, K. P. How advances in neural recording affect data analysis. *Nature Neuroscience 14* (2011), 139–142.

[53] TAKEMURA, S.-Y., BHARIOKE, A., LU, Z., NERN, A., VITALADEVUNI, S., RIVLIN, P. K., KATZ, W. T., OLBRIS, D. J., PLAZA, S. M., WINSTON, P., ET AL. A visual motion detection circuit suggested by drosophila connectomics. *Nature 500* (2013), 175–181.

[54] VAN ESSEN, D. C., SMITH, S. M., BARCH, D. M., BEHRENS, T. E., YACOUB, E., AND UGURBIL, K. The WU-Minn human connectome project: An overview. *NeuroImage 80* (2013), 62–79.

[55] VAN ESSEN, D. C., UGURBIL, K., AUERBACH, E., BARCH, D., BEHRENS, T. E. J., BUCHOLZ, R., CHANG, A., CHEN, L., CORBETTA, M., CURTISS, S. W., ET AL. The human connectome project: A data acquisition perspective. *Neuroimage 62* (2012), 2222–2231.

[56] WAGER, T. D., LINDQUIST, M., AND KAPLAN, L. Meta-analysis of functional neuroimaging data: current and future directions. *Social Cognitive and Affective Neuroscience 2* (2007), 150–158.

[57] ZADOR, A., DUBNAU, J., OYIBO, H., ZHAN, H., CAO, G., AND PEIKON, I. Sequencing the connectome. *PLoS Biology 10:e1001411. doi: 10.1371/journal.pbio.1001411* (2012).