

Draft in progress; title TBD.

Artificial meets Biological Intelligence: Machine Learning for Neuroscience, or
Statistical Neuroscience: Machine Learning for Advancing Study of the Brain, or
Neural data science: large-scale statistical neuroscience at single-neuron resolution

Paninski, L. & Cunningham, J.P.
Departments of Statistics and Neuroscience
Grossman Center for the Statistics of Mind
Columbia University

Introduction

Modern large-scale multineuronal recording methodologies, including multielectrode arrays, calcium imaging, and optogenetic techniques, produce single-neuron resolution data of a magnitude and precision that were the realm of science fiction a decade ago. The major bottlenecks in systems and circuit neuroscience no longer lie in simply collecting data from large neural populations, but also in *understanding* this data: developing novel scientific questions and corresponding analysis techniques and experimental designs to fully harness these new capabilities and meaningfully interrogate these questions. Advances in methods for signal processing, network analysis, dimensionality reduction, and optimal control – developed in lockstep with advances in experimental neurotechnology -- promise major breakthroughs in multiple fundamental neuroscience problems. These trends are clear in a broad array of subfields of modern neuroscience; here we focus on recent advances in methods for analyzing neural time-series data with single-neuronal precision.

Figure 1 as a signpost to what the paper is about and why it is a coherent set of topics

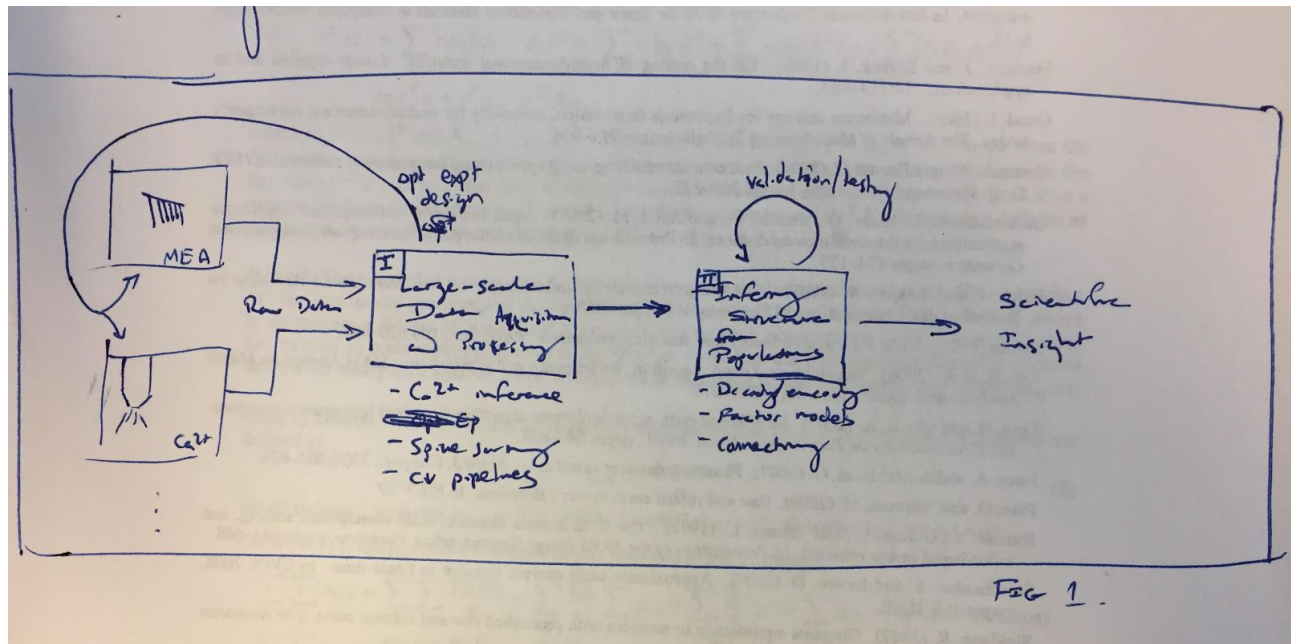


Fig 1

I. High-throughput neural signal acquisition methods

Neuroscientists have long dreamed of recording from many thousands (or more) neurons simultaneously. This goal is the major motivation of the BRAIN initiative and related efforts, and with new calcium imaging methods and large-scale multielectrode array (MEA) devices this dream is quickly becoming a reality. But now a major bottleneck exists. Cutting-edge calcium imaging methods and MEAs output data at rates measured in *terabytes/hour*, and data rates continue to increase. At these huge rates analyzing and even storing the data is challenging (Freeman et al 2014 Nature Methods, Cunningham 2014 Nature Methods), let alone optimally extracting all the useful information in these datastreams; without the right analytical technology, we will never unlock the true potential of these experimental advances.

Calcium imaging

Calcium imaging has become the dominant method for recording from large populations of neurons, due to several well-known advantages: calcium imaging offers cell-type specificity and can be coupled easily with a variety of genetic tools; calcium imaging can be less invasive and damaging to brain tissue than inserting an MEA; calcium imaging has proven scalability to record simultaneously from $O(10^4)$ neurons in vivo (over an order of magnitude larger than achieved by an MEA to date); and finally, imaging approaches enable significantly greater experimental design flexibility than MEAs in terms of which subsets of neurons in the imaging volume are interrogated at which times, and how many pixels are assigned to each neuron (we expand on the importance of this point below). At the same time, calcium imaging suffers from

some clear disadvantages: calcium signals represent a slow, nonlinear encoding of the underlying spike train signals of interest, and therefore it is necessary to denoise, temporally deconvolve, and spatially demix calcium video data into estimates of neural activity.

There has recently been a flurry of research activity addressing these issues. Building on earlier work (Vogelstein et al 2010), (Pnevmatikakis et al 2016) and (Pachitariu et al 2017) present constrained nonnegative matrix factorization (NMF) approaches to simultaneously solve these demixing and deconvolution problems. (Zhou et al 2016) extended this approach to handle one-photon data, where large “background” contributions from out-of-focus light complicate the demixing problem; recent large-scale approaches to acquiring one-photon imaging data (Bouchard et al 2015; Kim et al 2016) will likely benefit from this approach or modifications thereof. (Friedrich et al 2017a; Giovanucci et al 2017) developed real-time implementations that process incoming data online, one imaging frame at a time, enabling closed-loop experiments. (Deneux et al 2016) developed an improved and more general implementation of the hidden Markov model deconvolution approach of (Vogelstein et al 2009). (Picardo et al 2016) developed hierarchical Bayesian methods for sharing statistical information across behaviorally-similar trials to enable temporal super-resolution of estimated neural activity. (Jewell and Witten 2017) developed useful mathematical theory on exactly solving the sparse deconvolution problem addressed in (Pnevmatikakis et al 2016; Friedrich et al 2017a). (Wei et al 2017) investigate the impact of calcium indicator nonlinearities on downstream analyses of neural population activity, concluding that some caution is warranted in interpreting neural dynamics inferred from calcium imaging data.

In the near future, we expect that modern computational vision approaches (e.g., based on artificial neural networks (ANNs)) can be incorporated into the NMF framework for further improvements (as we will see below, the incorporation of ANNs to replace modules in various analysis pipelines is a recurring theme in this review); (Apthorpe et al 2016) presented a promising first step. Since NMF is a non-convex problem, accurate initialization of the estimates is critical; the approach presented in (Petersen et al 2017) is a promising alternative approach here (see also Giovanucci et al 2017). One major issue that has slowed progress is the lack of “gold standard” datasets that can be used to objectively score algorithm performance. The iterative optimization of open-sourced algorithms on agreed-upon standard datasets has been a critical theme enabling progress in modern machine learning (Donoho 2015); see (Berens et al 2017) for a recent application of this general program to improve available calcium deconvolution methods. The curation of “gold standard” spatiotemporal calcium imaging datasets remains a critical challenge; the IARPA MICRONS project (<https://www.iarpa.gov/index.php/research-programs/microns>) will soon deliver public datasets that combine large-scale electron microscopy with calcium imaging in the same cortical volumes, and will therefore serve as a major step forward in this direction (see also Lee et al 2016; Vishwanathan et al 2017; Bae et al 2017).

Given the explosive growth of data acquisition, one major trend that we see guiding research in this area over the next several years involves the optimization of experimental design and

analysis methods jointly in order to image larger populations at higher temporal resolution. It is widely recognized that it does not make sense to e.g. optimize an imaging apparatus in isolation; instead, the full experimental preparation, imaging technology, and analysis approach should each be considered as parts of a pipeline that should be optimized as a whole. (Yang et al 2016; Prevedel et al 2016; Song et al 2017a; Lu et al 2017; Friedrich et al 2017b) have all offered variations on a similar theme: spatial resolution can be usefully traded off for temporal resolution -- we can record from more cells if we are willing to accept a lower ratio of pixels per cell; and moreover, prior information about cell shapes and locations can shift the favorable point of this tradeoff even further -- once we know where the cells are we can reduce our spatial resolution even further without negatively impacting the quality of the recovered temporal neural activity (Friedrich et al 2017b). (Noebauer et al 2017) presented another example of this theme in the context of a computationally-challenging light-field microscopy application; we expect that performance here can be improved significantly with stronger signal models. Simulators such as those developed in (Song et al 2017b) will likely play a useful role in the ongoing joint optimization of demixing methods and hardware design.

One major problem that remains open: tracking activity with single-neuron resolution in small moving animals with flexible nervous systems, e.g. larval zebrafish (Cong et al 2017), drosophila (Bouchard et al 2015), or hydra (Dupre and Yuste, 2017). Good solutions have been developed in *C. elegans* (Nguyen et al 2016, Venkatachalam et al 2016), though demixing of fast cytosolic (non-nuclear-localized) signals remains an unsolved problem. We expect non-rigid registration approaches similar to those developed by (Pnevmatikakis and Giovanucci 2017) to be helpful here.

While we have focused on calcium imaging in this section, many similar themes will hold for superthreshold voltage imaging at single-cell resolution, which is expected to be a major growth area over the next decade; see e.g. (Xu et al 2017) for a recent review. Of course voltage imaging also provides the opportunity to record from subthreshold signals at multiple points along the dendrite and axon; once these imaging methods become more mature we expect to see rapid growth in statistical methods for extracting information from this noisy data, with (Huys et al 2006; Huys and Paninski 2009; Paninski 2010; Paninski et al 2012; Pakman et al 2014) providing a foundation for further development.

Spike sorting data from large-scale MEAs.

Spike sorting has been a not-quite-solved problem for decades. In small-scale recordings, a large degree of manual supervision over the spike sorting process is viable; additionally, it is feasible to manually optimize the depth of a few electrodes or tetrodes to ensure high-SNR recordings. Neither strategy is possible with large-scale MEAs: recordings with hundreds of electrodes are routine now, and much larger MEAs are on the way (Tsai et al 2015; Rios et al 2016; www.darpa.mil/program/neural-engineering-system-design). This looming bottleneck has driven a recent uptick of papers on spike sorting papers on large dense MEAs (Pachitariu et al

2016; Yger et al 2016; Jun et al 2017; Lee et al 2017); see also (Barnett et al 2016) for an approach for validating spike sorting methods with no ground truth and (Farina et al 2017) for a discussion of similar issues in the context of EMG signals. This recent literature has emphasized computational scalability and proper handling of spike events that overlap across many electrodes; in particular, (Pachitariu et al 2016) introduced a fast implementation of a matching-pursuit algorithm to detect these spike overlaps, and (Lee et al 2017) built on this work with a more robust and efficient “triage-then-cluster” approach in which an ANN detects putative spike events and then “clean” spikes are clustered first, followed by more difficult overlapping spikes.

As in the calcium imaging context, it is clear that agreed-upon gold standard datasets will lead to accelerated progress here. Acquisition of ground truth data in this context is a notoriously challenging problem; for now we can only hope for partial gold standard datasets, in which e.g. ground truth spiking for single neurons is available (Neto et al 2016). Optogenetic tagging methods (in which a sparse subset of neurons is activated at known times) could play a very useful role here. In some brain areas we can exploit useful side information to provide a sanity check on the sorting results: for example, the mosaic tiling of receptive fields in the primate retina provides some partial validation. In parallel, similar to the imaging context, the iterative improvement of simulators of electrical activity (Hagen et al 2015) remains an important direction. Another useful approach is to create “hybrid” datasets combining simulated spiking signals with real noise signals (Pachitariu et al 2016; Jun et al 2017). The time seems ripe for a community-based collaborative approach to develop a battery of gold standard datasets and metrics and then iteratively improve each module of these pipelines towards more scalable and accurate solutions.

A separate track of work has taken as a starting point the realization that spike sorting selects the biggest units from the observed voltage traces, but leaves behind a large amount of information in the lower-SNR units that can not be separated cleanly from the noise floor. “Clusterless” decoding approaches (Kloosterman et al 2014; Deng et al 2015) have been developed to extract information from these unsorted spikes. A combined strategy (in which we sort the sortable units, but exploit information from the unseparable units and LFP signals too) works best in practice (Bansal et al 2012; Todorova et al 2014).

Finally, as interest grows in bidirectional electrical neural interfaces that stimulate and record simultaneously, the problem of stimulation artifact cancellation becomes critical; see (Mena et al 2016) for a recent step forward.

II. Understanding large-scale neural signals

As emphasized above, acquiring and processing large-scale neural data with single-neuron and high temporal resolution has represented a critical bottleneck that has attracted significant research effort over the last couple years. While significant challenges remain, these efforts

have established a clear path forward towards eliminating this bottleneck. The next frontier then is to extract understanding from the resulting high-dimensional neural activity data.

Historically, the analysis of spike train data has focused on three broad questions. (1) How is information encoded in spike trains, and how can we decode this information? (2) Can we infer network connectivity from multiple spike train data? (3) Can we model the activity of large neural populations in terms of a lower-dimensional set of factors? We will review recent progress on each of these three themes in turn below, but it is worth emphasizing up front that models developed to address any of these questions can be profitably combined: e.g., factor analysis models developed to address question (3) can lead to improved decoding of neural data (a la question 1).

Encoding and decoding.

How the brain *encodes* external variables into spike trains, and the converse problem of *decoding* external variables from spike trains, are classic problems in statistical neuroscience.

For the first problem, generalized linear models (GLMs) have for years provided the methodological foundation; see Paninski et al 2007 for a review. Some recent work has targeted the computational efficiency of GLM estimation methods (Mena and Paninski 2014; Ramirez and Paninski, 2014; Wu et al 2015). Of course, GLMs, being simply a generalization of linear regression methods, have effectiveness dependent entirely on the chosen “feature set” -- that is, the collection of variables against which we choose to regress neural activity. Chang and Tsao 2017 describe an exciting recent application showing that with a good choice of features it is possible to predict highly nonlinear responses. One major trend is to learn features adaptively using a hierarchical approach to combine information over many cells / experiments. This leads to significantly richer and more powerful models. Field et al 2010 and Antolik et al 2016 are two examples of this approach, in which we share information from simultaneously-recorded cells to estimate a hidden layer that better explains the observed responses. (See also Rahnema Rad et al 2017 for a different method for sharing statistical strength across cells.) Again, modern ANN methods are well-suited to this task of learning a useful shared feature representation from a big dataset of many neural responses: (McIntosh et al 2016, Batty et al 2017) provide two examples of this idea (see also Mineault et al 2012 for an earlier example), and we expect to see more applications of this approach in the near future. Alternatively, we can repurpose ANNs trained to perform computer vision tasks (eg object recognition) and use the resulting feature sets to predict responses; see (Kriegeskorte and Diedrichsen 2016; Yamins and DiCarlo 2016) for perspectives on this growing literature.

Regarding the converse problem of decoding, there is a large ongoing engineering and clinical literature on brain-machine interfaces (including not only systems to extract motor control information from the brain but also sensory devices such as cochlear and retinal prosthetics) that we will not attempt to review systematically here. The “ReFIT” decoding algorithm

proposed in (Gilja et al 2012) contributed a notable empirical advance in decoding motor signals; (Merel et al 2016) provides a rigorous theoretical foundation and generalization of this algorithm. Another thread of work has shown that stronger models of joint neural variability can be used to construct better decoders (Lawhern et al 2010, Kao et al 2015); see also (Burkhart et al 2016) for an interesting converse approach. Finally, ANNs have recently been applied to decoding problems (Sussillo et al 2016; Glaser et al 2017); (Parthasarathy et al 2017) notably developed a straightforward procedure for converting an encoding model (i.e., a probabilistic model of the neural responses to an arbitrary stimulus), plus samples from the prior stimulus distribution, into an easily-computed approximation of the optimal Bayesian decoder. We expect to see more applications of similar ideas in the near future.

Connectivity estimation

Another classic problem in statistical neuroscience is to infer neuronal network connectivity from correlated activity in the network, and then to use the inferred connectivity to understand the network function and predict its dynamics. The major roadblock here has been the “common input” problem: without strong prior information, it is not possible to reliably distinguish causal connections between pairs of observed neurons versus correlations induced by common input from unobserved neurons. (Soudry et al 2015) introduced a novel “shotgun” experimental design that exploits the flexibility of imaging approaches for recording from large populations of cells: the idea is to image different subsets of the network in a serial manner, then use statistical methods to estimate the full network connectivity (note that this approach is enabled by optical approaches to neural recording; this approach would not be feasible with current multi-electrode arrays). In simulations, this approach enables the accurate estimation of networks an order of magnitude larger than was previously possible. (See Turaga et al 2014 for a simplified implementation of this idea.) Experimental methods are now becoming sufficiently fast and scalable to put this method into practice. Other relevant advances include (Harrison et al 2015), who introduce new conditional inference methods to address hypotheses about the precision of multineuronal responses, and (Linderman et al 2016), who discuss methods for incorporating stronger prior knowledge into network estimates; see also (Jonas and Koerding 2015) for improved network priors.

Once we have estimated the network connectivity, we need a high-throughput method for verifying our estimates (e.g., the inferred synaptic weights). Optogenetic approaches are well-suited to this task; (Shababo et al 2013; Chen et al 2017) propose a scalable adaptive approach towards mapping and verifying the connectivity onto single postsynaptic cells.

Finally, once these networks are inferred a major goal is to study their dynamical properties. (Gerhard et al 2017; Hocker and Park, 2017) point out that standard GLM estimation approaches can lead to dynamically unstable estimated networks, and propose approaches to correct this deficit; some relevant asymptotic theory is developed in (Hall et al 2017).

Factor models

In the language of machine learning, the encoding and decoding problems are *supervised*, in the sense that one seeks a mapping between two known signals: measurable behavioral variables and populations of spike trains. The *unsupervised* analog is often approached using *factor models*: high-dimensional neural population activity is assumed to be a noisy, redundant observation of some hidden (latent), often low-dimensional, signal of interest. This signal can then be interrogated with respect to a scientific hypothesis, used as a denoised and simpler representation of the neural activity, or viewed for compact exploratory analysis of the data.

Following the pioneering work of (Smith and Brown 2003), much of this literature has followed the Bayesian paradigm, where a generative probabilistic model is stipulated to link low-dimensional latent signal to high-dimensional neural spike trains, and then a computational inference procedure recovers the posterior distribution of the latent variable from the observed data. Examples of this paradigm include systems with simple latent temporal structure (e.g., Smith and Brown 2003, Yu et al 2009, Macke et al 2012 NIPS, Gao et al 2015 NIPS, Gao et al 2016 NIPS, Zhao and Park 2016, Zhao and Park 2017, Goris et al Nature Neuro 2014, Ecker et al 2014), systems with switching dynamical structure (Petreska et al 2012 NIPS, Kato... Zimmer Cell 2015, Wiltschko et al 2015, Johnson et al 2016 NIPS, Linderman et al 2017), and systems with recurrent neural network dynamical structure (Sussillo et al 2016b LFADS, Krishnan... Sontag 2015 arxiv). Following a similar scientific motivation, a methodologically distinct approach to factor modeling is dimensionality reduction: one stipulates an objective or loss function and then optimizes a map from data to low-dimensional latent factors (Cunningham and Ghahramani 2015 JMLR). This approach has been widely used in neuroscience and was recently reviewed in Cunningham and Yu Nature Neuroscience 2014.

One key question about factor models is to ask what they have and might be able to teach us about the brain. Several exemplars of the scientific potential of these approaches are worth noting. First, some of the earliest applications of factor models were to understand mixed selectivity in prefrontal cortex (Machens et al 2010); this work showed that despite the apparent complex responses displayed by single neurons, at the population level simple behavioral correlates can be effectively read out from the brain. Second, one natural but significant extension of this “de-mixing” perspective was the finding that different computations in certain brain areas are carried out in different subspaces of neural population activity, thus providing an implicit gating mechanism for irrelevant activity (Sanger and Kalaska 2014, Kaufman et al 2014, Elsayed et al 2016, Stavisky et al Neuron 2017). Third, Sadtler et al Nature 2014 used this notion of subspaces of activity along with a brain-machine interface to discover constraints (in terms of dimensions in neural population space) on learning. Fourth, population activity along with factor models has been used to understand the dynamical structure of motor and prefrontal cortices (Churchland et al 2012, Mante et al 2013). We expect this literature to continue to mature and deepen methodologically, and to incorporate stronger constraints based on neural circuitry and cell types (Okun et al 2015) as more of this information becomes available. In the coming years we expect these models to be increasingly used to deepen our understanding of population-level neural activity.

Interpretability of Large-Scale Neural Data Analysis

Of course, the impetus behind the analysis of large scale neural data is the belief that these efforts will lead to deeper insights into principles of neural computation, biomedically relevant understanding of brain dysfunction, and more. One important line of research, that is in its earliest chapter, is to ask: to what extent is that belief well founded? There are three categories of approach to address this critical question.

First, there is the growing risk and concern that novel analyses of large-scale neural data are not discovering new phenomena, but are rather rediscovering simpler, previously known features of the data that appear novel given the novel class of data and algorithms used to investigate these points. Recent work has created statistical hypothesis testing frameworks to enable researchers to quantitatively address this rigorous skepticism in their own data [Harrison et al 2015, Elsayed and Cunningham Nature Neuro 2017, Laoiza-Ganem et al ICLR 2017], by developing methods to generate datasets that contain these simpler, previously known features but are otherwise random, thus creating a null distribution against which novel large-scale data claims can be tested. Early applications to test the presence of linear dynamics in motor cortex (Churchland et al 2012) and the presence of de-mixed readouts in prefrontal cortex (Machens et al 2010) have clarified these previous results [Elsayed and Cunningham Nature Neuro 2017].

Second, another key point of scepticism is to ask to what extent recording larger and larger datasets will produce fundamentally new findings. An often-discussed example logically supposes that, if the experimental tasks that researchers ask of their subjects do not appropriately increase in complexity to match the increase in data, limitations on resulting scientific insights will follow. Work to address this question has to date been preliminary. One group has discussed a notion of required task complexity [gao and ganguli 2015], and two others have attempted to measure the complexity of neural population activity in the face of larger and larger datasets, finding both that complexity (as measured by the apparent dimensionality of the data) grows seemingly without bound (Pachitariu et al, unpublished 2017) in some cases, and that it does not in others (Williamson et al 2017 PLOS CB). Significant additional theoretical and experimental work is required to provide more satisfactory answers here.

Third, at the broadest level, we might ask if the paradigm of measuring individual neuron activities -- no matter how many neurons we might be able to record -- will ever produce a coherent mechanistic understanding of the neural system. Jonas + Koerding 2017 recently presented a somewhat pessimistic answer to this question; these authors used a man-made computer as a proxy for a small nervous system, recorded units from that system, and failed to arrive at a satisfactory understanding of the system's function or design. While we don't share the pessimism implicit in the conclusion of this paper, we do agree that this is a critical line of questioning. Specifically, what will we do when we have in hand a matrix of the spatially- and temporally-resolved activity of all the cells in the brain of an animal performing an interesting behavior? This remains a yet-distant dream in mammals but is close to reality in several

invertebrate species. Our field needs to think critically and deeply about what to do now that this century-long goal is almost in our grasp.

Future outlook

We close by discussing several trends that will guide development in this field over the next several years.

- Datasets will continue to grow in size as recording modalities are optimized and new approaches are introduced; the scalability of processing pipelines will remain a critical design constraint.
- Closed loop, many-degree-of-freedom optimal control of neuronal populations will represent a critical research subfield as optogenetic spatiotemporal control methods continue to mature (Ahmadian et al 2011; Grosenick et al 2015).
- Fusion of multimodal datasets will represent another critical research area as large-scale connectomic and cell type information (Macosko et al 2015; Lee et al 2015; Kobschull et al 2016) becomes available to inform and constrain functional models (Jonas and Koerding 2015).
- We expect to see more fruitful marriages of “classical” computational neuroscience theories (e.g., network dynamics, reinforcement learning) with statistical models for network inference and dimensionality reduction; see (Linderman and Gershman 2017) for an interesting step in this direction.
- More broadly, “sociological” trends towards more open and large-scale data sharing and open-source collaborative projects, supported by stronger pipelines (Yatsenko et al 2015) and reproducibility tools (e.g., <http://mybinder.org/>), will enable richer and more ambitious multilevel models of neural function that are beyond the reach of single laboratories. <https://www.internationalbrainlab.com> represents an example of this; we expect to see more. Automated curation and compression of data into useful shareable form are important underexplored steps in the analysis pipeline here.
- Finally, from our vantage point the number of critical neural data science projects is currently growing significantly more quickly than the number of young scientists with the necessary interdisciplinary training in machine learning, statistics, and neuroscience. This is becoming a critical bottleneck (Akil et al 2016); increased investment in neural data science training will pay rich dividends in improving our understanding of neural systems over the next decade.

Acknowledgments

Thanks to Scott Linderman and E.J. Chichilnisky for helpful conversations. Funding...

References

Ahmadian, Y., Packer, A., Yuste, R. & Paninski, L. (2011). Designing optimal stimuli to control neuronal spike timing. *J. Neurophys.* 106: 1038-1053.

Akil et al 2016: [http://www.cell.com/neuron/fulltext/S0896-6273\(16\)30209-4](http://www.cell.com/neuron/fulltext/S0896-6273(16)30209-4)

Antolik et al 2016 <http://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1004927>

Apthorpe et al 2016 <https://arxiv.org/abs/1606.07372>

Bae et al 2017 <http://www.biorxiv.org/content/early/2017/08/31/182758>

Bansal et al 2012 <http://jn.physiology.org/content/107/5/1337>

Barnett et al 2016. <https://www.ncbi.nlm.nih.gov/pubmed/26930629>

Batty, E., Merel, J., Brackbill, N., Heitman, A., Sher, A., Litke, A., Chichilnisky, E.J. & Paninski, L. (2017). Multilayer Network Models of Primate Retinal Ganglion Cells. ICLR.

Berens, P. et al. (2017) Spikefinder paper.

Bittner et al 2017 PLOS One <https://users.ece.cmu.edu/~byronyu/papers/BittnerPONE2017.pdf>

Bouchard et al 2015. <http://www.nature.com/nphoton/journal/v9/n2/full/nphoton.2014.323.html>

Buesing, L., Calabrese, A., Cunningham, J., Woolley, S. & Paninski, L. (2017). A Statistical Model of Shared Variability in the Songbird Auditory System. bioRxiv 113670.

Burkhardt et al 2016 (<https://arxiv.org/abs/1608.06622>)

Chang and Tsao 2017, [http://www.cell.com/cell/abstract/S0092-8674\(17\)30538-X](http://www.cell.com/cell/abstract/S0092-8674(17)30538-X)

Chen et al 2017 <http://sand.stat.cmu.edu/all-yi/#chen>

Churchland et al 2012 Nature

<http://stat.columbia.edu/~cunningham/pdf/ChurchlandNature2012.pdf>

Cong et al 2017 <http://www.biorxiv.org/content/early/2017/04/28/131532>

Cunningham 2014 Nature Methods
<http://stat.columbia.edu/~cunningham/pdf/CunninghamNMETH2014.pdf>

Cunningham and Yu 2014 Nature Neuro
<http://stat.columbia.edu/~cunningham/pdf/CunninghamNN2014.pdf>

Cunningham and Ghahramani 2015 JMLR
<http://stat.columbia.edu/~cunningham/pdf/CunninghamJMLR2015.pdf>

Deneux et al 2016 <https://www.nature.com/articles/ncomms12190>

Deng et al 2015. <https://www.ncbi.nlm.nih.gov/pubmed/25973549>

Donoho 2015. <http://courses.csail.mit.edu/18.337/2015/docs/50YearsDataScience.pdf>

Dupre and Yuste 2017. [http://www.cell.com/current-biology/abstract/S0960-9822\(17\)30220-8](http://www.cell.com/current-biology/abstract/S0960-9822(17)30220-8)

Ecker et al 2014 [http://www.cell.com/neuron/fulltext/S0896-6273\(14\)00104-4](http://www.cell.com/neuron/fulltext/S0896-6273(14)00104-4)

Elsayed Lara et a 2016 Nature Comms
<http://stat.columbia.edu/~cunningham/pdf/ElsayedNCOMMS2016.pdf>

Elsayed GF and Cunningham JP Structure in neural population recordings: an expected byproduct of simpler phenomena? Nature Neuroscience 2017
<http://www.nature.com/neuro/journal/vaop/ncurrent/full/nn.4617.html>

Farina et al 2017. <https://www.nature.com/articles/s41551-016-0025>

Field, G., Gauthier, J., Sher, A. et al. (2010). Functional connectivity in the retina at the resolution of photoreceptors. Nature 467, 673-677.

Friedrich, J., Zhou, P. & Paninski, L. (2017a). [Fast Active Set Method for Online Spike Inference from Calcium Imaging](#). NIPS; PLoS Comput. Bio. 13: e1005423.

Friedrich, J., Yang W., Soudry D., Mu Y., Ahrens M., Yuste R., Peterka D. & Paninski L. (2017b). Multi-scale approaches for high-speed imaging and analysis of large neural populations. bioRxiv 091132.

Gao P, and Ganguli S 2015 <https://ganguli-gang.stanford.edu/pdf/15.BraveNewWorld.pdf>

Gao Y, Buesing L, Shenoy KV, Cunningham JP (2015) High-dimensional neural spike train analysis with generalized count linear dynamical systems. NIPS 2015.

Gao, Y., Archer, E., Paninski, L. & Cunningham, J. (2016). Latent linear-dynamical neural population models through nonlinear embedding. NIPS; Arxiv 1605.08454.

Gerhard F, Deger M, Truccolo W (2017) On the stability and dynamics of stochastic spiking neuron models: nonlinear Hawkes process and point process GLMs. PLoS Comput Biol 13(2): e1005390. doi:10.1371/journal.pcbi.1005390.

Gilja et al (2012) Nature Neuroscience <https://www.ncbi.nlm.nih.gov/pubmed/23160043>

Giovanucci, A. et al (2017). OnACID: Online Analysis of Calcium Imaging Data in Real Time. NIPS.

Glaser et al 2017. <https://arxiv.org/abs/1708.00909>

Golub et al CONB <https://users.ece.cmu.edu/~byronyu/papers/GolubCONB2016.pdf>

RLT Goris, JA Movshon, EP Simoncelli; Partitioning neuronal variability; Nature neuroscience 17 (6), 858-865

Grosenick et al 2015. [Closed-loop and activity-guided optogenetic control](#). L. Grosenick, J. Marshel, and K. Deisseroth. Neuron (2015) 86(1) 106–139.

Hagen et al 2015. <http://www.sciencedirect.com/science/article/pii/S0165027015000369>

Hall et al 2017. <https://arxiv.org/abs/1605.02693>

Harrison et al 2015; <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4457305/>

David Hocker and Il Memming Park 2017. Multistep Inference for Generalized Linear Spiking Models Curbs Runaway Excitation Accepted to 8th International IEEE EMBS Conference on Neural Engineering (NER 2017)

Huys, Q., Ahrens, M. & Paninski, L. (2006). [Efficient estimation of detailed single-neuron models](#). Journal of Neurophysiology 96: 872-890.

Huys, Q. & Paninski, L. (2009). [Smoothing of, and parameter estimation from, noisy biophysical recordings](#). PLOS Computational Biology 5: e1000379.

Jewell and Witten 2017. <https://arxiv.org/abs/1703.08644>

Johnson et al NISP 2016

<https://papers.nips.cc/paper/6378-composing-graphical-models-with-neural-networks-for-structured-representations-and-fast-inference>

Jonas and Koering 2015. <https://elifesciences.org/articles/04250>

Jonas and Koering 2017.

<http://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1005268>

Jun et al 2017 <http://www.biorxiv.org/content/early/2017/01/30/101030>

Kato et al 2015 Cell [http://www.cell.com/cell/pdfExtended/S0092-8674\(15\)01196-4](http://www.cell.com/cell/pdfExtended/S0092-8674(15)01196-4)

Kao et al 2015. <https://www.nature.com/articles/ncomms8759>

Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2014). Cortical activity in the null space: permitting preparation without movement. Nature Neuroscience, 17(3):440-8.

Kebschull et al 2016. <https://www.ncbi.nlm.nih.gov/pubmed/27545715>

Kim et al 2016. [http://www.cell.com/cell-reports/abstract/S2211-1247\(16\)31676-X](http://www.cell.com/cell-reports/abstract/S2211-1247(16)31676-X)

Kloosterman et al 2014. <https://www.ncbi.nlm.nih.gov/pubmed/24089403>

Kriegeskorte and Diedrichsen 2016. <https://arxiv.org/pdf/1608.02027.pdf>

Krishnan et al 2015 arxiv <https://arxiv.org/abs/1511.05121>

Lawhern, V., Wu, W., Hatsopoulos, N. & Paninski, L. (2010). Population neuronal decoding using a generalized linear model with hidden states. J. Neurosci. Methods 189: 267-280.

Lee et al 2015. <https://www.ncbi.nlm.nih.gov/pubmed/25675209>

Lee et al 2016. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4844839/>

Lee, J., Carlson, D., Shokri, H., Yao, W., Goetz, G., Hagen, E., Batty, E., Chichilnisky, E.J., Einevoll, G. & Paninski, L. (2017). YASS: Yet another spike sorter. Biorxiv 151928 / NIPS

Lewi, J., Butera, R. & Paninski, L. (2009). Sequential optimal design of neurophysiology experiments. Neural Computation 21: 619-687.

Scott W. Linderman, Ryan P. Adams, and Jonathan W. Pillow (2016). Bayesian latent structure discovery from multi-neuron recordings. In Advances in Neural Information Processing Systems (NIPS).

Linderman and Gershman 2017. <http://biorxiv.org/content/early/2017/01/31/104737>

Linderman, S., Miller, A., Adams, R., Blei, D., Paninski, L., Johnson, M. (2017). [Recurrent Switching Linear Dynamical Systems](#). AISTATS17 / ArXiv 1610.08466.

Loaiza-Ganem G*, Gao Y*, Cunningham JP (2017) "Maximum Entropy Flow Networks." ICLR 2017.

Lu et al 2017. <http://www.nature.com/neuro/journal/v20/n4/full/nn.4516.html>

Machens et al 2010 <https://www.ncbi.nlm.nih.gov/pubmed/20053916>

Macosko et al 2015. [http://www.cell.com/abstract/S0092-8674\(15\)00549-8](http://www.cell.com/abstract/S0092-8674(15)00549-8)

Macke et al 2012 NIPS <http://stat.columbia.edu/~cunningham/pdf/MackeNIPS2012.pdf>

Mante et al 2013 Nature
<http://www.nature.com/nature/journal/v503/n7474/abs/nature12742.html>

McIntosh et al 2016,
<http://papers.nips.cc/paper/6388-deep-learning-models-of-the-retinal-response-to-natural-scenes>

Mena, G., Grosberg, L., Madugula, S., Hottowy, P., Litke, A., Cunningham, J., Chichilnisky, E.J. & Paninski, L. (2016+). Removing Stimulation Artifacts From Neural Recordings Using Structured Gaussian Processes. bioRxiv 089912.

Mena, G. & Paninski, L. (2014). On quadrature methods for refractory point process likelihoods. Neural Computation 26: 2790-7.

Merel, J., Carlson, D., Paninski, L. & Cunningham, J. (2016). [Neuroprosthetic decoder training as imitation learning](#). PLoS Comp. Bio 12: e1004948.

Mineault et al 2012. <http://www.pnas.org/content/109/16/E972.full>

Neto et al 2016. <https://www.ncbi.nlm.nih.gov/pubmed/27306671>

Nguyen et al 2016. <http://www.pnas.org/content/113/8/E1074>

Noebauer et al 2017. <http://www.biorxiv.org/content/early/2017/06/27/155572>

Okun et al 2015; <http://www.nature.com/nature/journal/v521/n7553/abs/nature14273.html>

Pachitariu et al 2016 <http://www.biorxiv.org/content/early/2016/06/30/061481>

Pachitariu et al 2017 <http://www.biorxiv.org/content/early/2017/07/20/061507>

Pakman, A., Huggins, J., Smith, C. & Paninski, L. (2014). [Fast penalized state-space methods for inferring dendritic synaptic connectivity](#). J. Comput. Neurosci. 36: 415-43

Paninski, L. (2010). [Fast Kalman filtering on quasilinear dendritic trees](#). Journal of Computational Neuroscience 28: 211-28.

Paninski, L., Pillow, J. & Lewi, J. (2007). [Statistical models for neural encoding, decoding, and optimal stimulus design](#). Computational Neuroscience: Progress in Brain Research, eds. Cisek, P., Drew, T. & Kalaska, J.

Paninski, L., Vidne, M., DePasquale, B., & Ferreira, D. (2012). [Inferring synaptic inputs given a noisy voltage trace](#). J. Comput. Neuro. 33: 1-19.

Parthasarathy, N., Batty, E., Falcon, W., Rutten, T., Rajpal, M., Chichilnisky, E.J. & Paninski, L. (2017). Neural Networks for Efficient Bayesian Decoding of Natural Images from Retinal Neurons. Biorxiv 153759 / NIPS.

Petersen et al 2017; <https://arxiv.org/abs/1703.06946>

Petreska et al 2012 NIPS <http://stat.columbia.edu/~cunningham/pdf/PetreskaNIPS2012.pdf>

Picardo, M., Merel, J., Katlowitz, K., Vallentin, D., Okobi, D, Benezra, S., Clary, R., Pnevmatikakis, E., Paninski, L., and Long, M. (2016). Population-level representation of a temporal sequence underlying skilled behavior. Neuron 90: 866-876.

Pillow and Aoi 2017 Nature Neuroscience

<http://stat.columbia.edu/~cunningham/pdf/PillowNN2017.pdf>

Pnevmatikakis, E., Soudry, D., Gao, Y., Machado, T., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., Ahrens, M., Bruno, R., Jessell, T., Yuste, R., Peterka, D. & Paninski, L. (2016). Simultaneous denoising, deconvolution, and demixing of calcium imaging data. Neuron 89: 285-299.

Pnevmatikakis and Giovanucci 2017, <http://www.biorxiv.org/content/early/2017/02/17/108514>

Prevedel et al 2016. <https://www.nature.com/nmeth/journal/v13/n12/full/nmeth.4040.html>

Rahnama Rad, K., Machado, T. & Paninski, L. (2016+). [Robust and scalable Bayesian analysis of spatial neural tuning function data](#). In press, Ann. Applied Stat., Arxiv 1606.07845

Ramirez, A. & Paninski, L. (2014). Fast generalized linear model estimation via expected log-likelihoods. J. Comput. Neurosci. 36: 215-34.

Gustavo Rios, Evgeniy V. Lubenov, Derrick Chi, Michael L. Roukes, and Athanassios G. Siapas, [Nanofabricated Neural Probes for Dense 3-D Recordings of Brain Activity](#). *Nano Letters* **16** (11), pp 6857–6862
DOI: [10.1021/acs.nanolett.6b02673](https://doi.org/10.1021/acs.nanolett.6b02673).

Sadtler et al 2014 <https://users.ece.cmu.edu/~byronyu/papers/SadtlerNature2014.pdf>

Sanger and Kalaska 2014 <http://www.nature.com/neuro/journal/v17/n3/full/nn.3663.html>

Shababo, B., Paige, B., Pakman, A. & Paninski, L. (2013). Bayesian inference and online experimental design for mapping neural microcircuits. NIPS.

Song et al 2017a. <https://www.ncbi.nlm.nih.gov/pubmed/28319111>

Smith AC and Brown EN 2003 Neural Computation

<http://www.stat.columbia.edu/~liam/teaching/neurostat-spr11/papers/brown-et-al/smith2003.pdf>

Song et al 2017b. A. Song, A.S. Charles, J.L. Gauthier, S.A. Koay, D.W. Tank, and J.W. Pillow. Two-photon microscopy simulation for optics optimization and benchmarking. CoSyNe Salt Lake City, Utah, February 2017.

Soudry, D., Keshri, S., Stinson, P., Oh, M.-W., Iyengar, G. & Paninski, L. (2015). Efficient “shotgun” inference of neural connectivity from highly sub-sampled activity data. PLoS Comp. Bio. 11: e1004464.

Stavisky SD, Kao JC, Ryu SI, Shenoy KV (2017) Motor cortical visuomotor feedback activity is initially isolated from downstream targets in output-null neural state space dimensions. Neuron. 95:195-208.

Sussillo D*, Stavisky SD*, Kao JC*, Ryu SI, Shenoy KV (2016) Making brain-machine interfaces robust to future neural variability. Nature Communications. 7:13749.

Sussillo D, Jozefowicz R, Abbott LF, Pandarinath C. (2016) LFADS – Latent Factor Analysis via Dynamical Systems. arXiv:1608.06315, 22 Aug. [arXiv] [PDF]

Todorova et al 2014. <http://iopscience.iop.org/article/10.1088/1741-2560/11/5/056005/meta>

D Tsai, E John, T Chari, R Yuste, K L Shepard, "[High-channel-count, high-density micro-electrode array for closed-loop investigation of neuronal networks](#)," Proceedings of the 37th Annual International Conference of the IEEE EMBS, 2015

Turaga SC, Buesing L, Packer AM, Dalgleish H, Pettit N, Hausser M and Macke JH (2014) Inferring neural population dynamics from multiple partial recordings of the same neural circuit In: Advances in Neural Information Processing Systems 26, , Twenty-Seventh Annual Conference on Neural Information Processing Systems (NIPS 2013), Curran, Red Hook, NY, USA, 539-547.

Venkatachalam, V. et al 2016. <http://www.pnas.org/content/113/8/E1082>

Vishwanathan et al 2017 [http://www.cell.com/current-biology/fulltext/S0960-9822\(17\)30730-3](http://www.cell.com/current-biology/fulltext/S0960-9822(17)30730-3)

Vogelstein, J., Packer, A., Machado, T., Sippy, T., Babadi, B., Yuste, R. & Paninski, L. (2010). [Fast non-negative deconvolution for spike train inference from calcium imaging](#). J. Neurophys. 104: 3691-3704

Wei, z et al 2017 (+ druckmann + svoboda). TK

Williamson et al 2017 PLOS CB

<https://users.ece.cmu.edu/~byronyu/papers/WilliamsonPLCB2016.pdf>

Wiltchko et al 2015. [http://www.cell.com/neuron/abstract/S0896-6273\(15\)01037-5](http://www.cell.com/neuron/abstract/S0896-6273(15)01037-5)

Wu et al 2015. http://pillowlab.princeton.edu/pubs/abs_Wu_NIPS15.html

Xu et al 2017, <http://www.sciencedirect.com/science/article/pii/S1367593117300212>

Yamins and DiCarlo 2016. <http://www.nature.com/neuro/journal/v19/n3/abs/nn.4244.html>

Yang, W., Miller, J., Carillo-Reid, L. Pnevmatikakis, E., Paninski, L., Yuste, R., & Peterka, D. (2016). Simultaneous multi-plane imaging of neural circuits. Neuron 89: 269-84.

Yatsenko et al (2015). <http://www.biorxiv.org/content/early/2015/11/14/031658>

Yger et al 2016. <http://www.biorxiv.org/content/early/2016/08/04/067843>

Yu Cunningham et al 2009 <http://stat.columbia.edu/~cunningham/pdf/YuJNP2009.pdf>

Zhou, P., Resendez, S., Rodriguez-Romaguera, J., Jimenez, J., Neufeld, S., Stuber, G., Hen, R., Kheirbek, M., Sabatini, B., Kass, R. & Paninski, L. (2016+). Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. Arxiv 1605.07266.

Zhao, Y. and Park, I.M.	Variational Latent Gaussian Process for Recovering Single-Trial Dynamics from Population Spike Trains	2017	Neural Computation Vol. 29(5)
-------------------------	---	------	----------------------------------

Zhao, Y. and Park, I.M.	Interpretable Nonlinear Dynamic Modeling of Neural Trajectories	2016	Advances in Neural Information Processing Systems (NIPS)
-------------------------	---	------	--