

Neural data science: accelerating the experiment-analysis-theory cycle in large-scale neuroscience

Paninski, L. and Cunningham, J.P.
Departments of Statistics and Neuroscience
Grossman Center for the Statistics of Mind
Zuckerman Mind Brain Behavior Institute
Center for Theoretical Neuroscience
Columbia University

Abstract

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 twenty years 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, with 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; this review focuses on recent advances in methods for analyzing neural time-series data with single-neuronal precision.




Figure 1. The central role of data science in modern large-scale neuroscience. Topics reviewed herein are indicated in black.

I. High-throughput neural signal processing methods

Neuroscientists have long dreamed of recording from many thousands of 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 on the order of *terabytes/hour*, and data rates continue to increase. At these huge rates processing and even storing the data is challenging (Freeman et al 2014 Nature Methods), let alone optimally extracting all the useful information in these data streams; 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; imaging approaches 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; Andilla and Hamprecht 2014), (Pnevmatikakis et al 2016 *; Pachitariu et al 2017; Haeffele and Vidal 2017) present constrained nonnegative or penalized matrix factorization (NMF) approaches to simultaneously solve these demixing and deconvolution problems. (Zhou et al 2016 *) extend this approach to handle data from one-photon imaging approaches, 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, unpublished, 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).

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 computational analysis approach should 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. That is, 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 trade-off even further: once we know the locations and shapes of the cells in the field of view, 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 significant problem requires further development: 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; see also (Kim et al 2017 *) for impressive recent progress in larval zebrafish.

While we have focused on calcium imaging in this section, many similar themes will hold for 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 at subcellular resolution, 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 studies on spike sorting activity from large dense MEAs (Pachitariu et al 2016 *, Yger et al 2016; Jun et al 2017; Lee et al 2017; Dhawale et al 2017; Chung 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, as 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 quality 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 most easily-discriminable 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 also exploit information from the non-separable units and local field potential signals) 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 significant effort 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 multi-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: GLMs enable spike trains to be regressed against covariates such as behavioral parameters, hidden factors, and other spiking in the population; 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 Rahnama 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 large 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 a promising 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 prior models for networks.

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 2017b) propose a scalable adaptive closed-loop optimal experimental design 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; Chen et al 2017a).

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 visualized 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 signals 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; Goris et al 2014; Ecker et al 2014; Gao et al 2015; Gao et al 2016 *; Zhao and Park 2016; Zhao and Park 2017), systems with switching dynamical structure (Petreska et al 2012; Kato et al 2015; Wiltshcko et al 2015 *; Johnson et al 2016; Linderman et al 2017 *), and systems with recurrent neural network dynamical structure (Krishnan et al 2015; Sussillo et al 2016b). An alternative direct approach to dimensionality reduction is to stipulate an objective or loss function that encodes the features one would like to capture in the data and then optimize a map from data to low-dimensional latent factors (Cunningham and Ghahramani 2015 *).

Dimensionality reduction approaches have been widely used in neuroscience (Cunningham and Yu 2014). 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 (Kaufman et al 2014; Elsayed et al 2016 *; Stavisky et al 2017). Third, (Sadler et al 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). As more connectomic and cell type constraints become available for population activity recordings, we expect this literature to continue to mature and deepen methodologically, and to elucidate interactions between cell type-specific subpopulations in multiple brain areas (Semedo et. al. 2014; Bittner et. al. 2017).

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. One important line of research, which 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 concern that novel analyses of large-scale neural data may not be discovering new phenomena, but are rather rediscovering simpler, previously known features of the data that appear new 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 (Harrison et al 2015 *; Elsayed and Cunningham 2017 *; Laoiza-Ganem et al 2017, Savin and Tkacik 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. 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 2017 *).

Another key point of skepticism is whether recording larger and larger datasets will produce fundamentally new findings. The answer may depend on the complexity of the experimental paradigm: if the number of recorded neurons grows while the task the animal needs to solve is kept relatively simple, will new scientific insights follow, or must the complexity of the task grow in concordance with that of the data? 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 in others that it does not (Williamson et al 2017). Significant additional theoretical and experimental work is required to provide clearer conclusions here.

Third, at the broadest level, we might ask if our current approaches will ever produce a coherent mechanistic understanding of the neural system. (Jonas and 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, then made simulated recordings, applied a battery of statistical analyses, and failed to arrive at a satisfactory understanding of the system's function or design. Thus their answer to their question, "could a neuroscientist understand a microprocessor," seems to be negative. While we don't share the pessimism implicit here, we do agree that despite rapid progress in our field over the last decade, neural data science remains in an early stage, largely because the curve of increasing neural data complexity that we have emphasized above has only recently begun to accelerate sharply upwards. Moreover, many of our theories of the brain have been allowed to flourish largely untethered from data that could constrain and winnow these theories. But now we have to grapple seriously with the question of what we will do when we have in hand, for example, a matrix of the spatially- and temporally-resolved activity of all neurons in an animal performing an interesting behavior. This remains a yet-distant dream in mammals but is close to reality in several invertebrate species, and our field needs to think critically and deeply about what to do now that this century-long goal is almost in our grasp. We believe the way forward is an acceleration of the experiment-analysis-theory cycle; there is a rapidly growing need for new theories to guide our exquisite new experimental tools, and as these theories develop we will continue to need well-matched scalable analysis methods that can connect experiment and theory in a tightly closed loop.

III. Future outlook

We close by summarizing 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 constraints (Macosko et al 2015; Lee et al 2014; Kebschull et al 2016) become available to inform functional models (Jonas and Koerding 2015 *).
- With this growth in the scale, quantity, and complexity of datasets and analysis methodologies, statistical techniques for validating and appropriately contextualizing resulting findings will become increasingly essential (Elsayed and Cunningham 2017 *).
- 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. Similarly, as noted above, the richness and complexity of available experimental data is beginning to outstrip the sophistication of the theory that we need to guide new experiments and the development of new analysis approaches. This is becoming a critical bottleneck (Akil et al 2016); increased investment in neural data science and neurotheory training will pay rich dividends in improving our understanding of neural systems over the next decade.

Acknowledgments

Thanks to Scott Linderman, E.J. Chichilnisky, and Sean Bittner for helpful conversations and suggestions, and Anthony Cruz and Meghan Kase for technical assistance. This work was funded by the DARPA NESD program, NSF BIGDATA IIS-1546296 (LP), IARPA MICRONS D16PC00003 (LP) and D16PC00008 (LP), NIH NINDS 1U01NS103489-01 (LP), NIBIB R01EB22913 (LP), and NEI R21EY02759201 (LP); NIH CRCNS R01 NS100066-01 (JPC), the Sloan Research Fellowship (JPC), the McKnight Scholar award (JPC), and the Simons Collaboration on the Global Brain (JPC and LP).

References

- Ahmadian Y, Packer A, Yuste R, Paninski, L: **Designing optimal stimuli to control neuronal spike timing.** *J. Neurophysiology* 2011, **106**: 1038-1053.
- Akil H, Balice-Gordon R, Cardozo DL, Koroshetz W, Norris SMP, Sherer T, Sherman SM, Thiels E: **Neuroscience Training for the 21st Century.** *Neuron* 2016, **90**: 917–926.
- Andilla FD, Hamprecht AF: **Sparse Space-Time Deconvolution for Calcium Image Analysis.** *NIPS* 2014, **27**.
- Antolik J, Hofer SB, Bednar JA, Mrsic-Flogel TD: **Model Constrained by Visual Hierarchy Improves Prediction of Neural Responses to Natural Scenes.** *PLOS Comput. Bio* 2016, **12**: e1004927.
- * A nice recent example of a hierarchical forward encoding model that shares information from simultaneously-recorded cells to estimate a hidden layer that better explains the observed responses.
- Apthorpe JN, Riordan AJ, Aguilar RE, Homann J, Gu Y, Tank DW, Seung HS: **Automatic Neuron Detection in Calcium Imaging Data Using Convolutional Networks.** *NIPS* 2016, **29**.
- Bae JA, Mu S, Kim JS, Turner NL, Tartavull I, Kemnitz N, Jordan CS, Norton AD, Silversmith WM, Prentki R, Sorek M, David C, Jones DL, Bland D, Sterling ALR, Park J, Briggman KL, Seung HS, the EyeWriters: **Structural and functional diversity of a dense sample of retinal ganglion cells.** *Biorxiv preprint* 2017, 182758.
- Bansal AJ, Truccolo W, Vargas-Irwin CE, Donoghue JP: **Decoding 3D reach and grasp from hybrid signals in motor and premotor cortices: spikes, multiunit activity, and local field potentials.** *J. Neurophysiology* 2012, **107**: 1337-1355.
- Barnett AH, Magland JF, Greengard LF: **Validation of neural spike sorting algorithms without ground-truth information.** *J Neuroscience Methods* 2016, **264**: 65-77.

Batty E, Merel J, Brackbill N, Heitman A, Sher A, Litke A, Chichilnisky EJ, Paninski L: **Multilayer Network Models of Primate Retinal Ganglion Cells**. *International Conference on Learning Representations* 2017.

* Another example of hierarchical forward encoding modeling. Provides a method to learn a shared feature space (sharing information over many cells that were not necessarily recorded at the same time) to build richer encoding models. See also (McIntosh et al 2016).

Berens P, Freeman J, Deneux T, Chenkov N, McColgan T, Speiser A, Macke JH, Turaga S, Mineault P, Rupprecht P, Gerhard S, Friedrich RW, Friedrich J, Paninski L, Pachitariu M, Harris KD, Bolte B, Machado TA, Ringach D, Stone J, Sofroniew NJ, Reimer J, Froudarakis E, Euler T, Roman-Roson M, Theis L, Tolias AS, Bethge M: **Community-based benchmarking improves spike inference from two-photon calcium imaging data**. *Biorxiv preprint* 2017, 177956.

Bittner SR, Williamson C, Snyder AC, Litwin-Kumar A, Doiron B, Chase SM, Smith MA, Yuet BM: **Population activity structure of excitatory and inhibitory neurons**. *PLOS One* 2017, **12**: e0181773.

Bouchard B, Voleti V, Mendes CS, Lacefield C, Grueber WB, Mann RS, Bruno RM, Hillman EMC: **Swept confocally-aligned planar excitation (SCAPE) microscopy for high-speed volumetric imaging of behaving organisms**. *Nature Photonics* 2015 **9**: 113–119.

Buesing L, Calabrese A, Cunningham J, Woolley S, Paninski L: **A Statistical Model of Shared Variability in the Songbird Auditory System**. *Biorxiv preprint* 2017, 113670.

Burkhart MC, Brandman DM, Vargas-Irwin CE, Harrison MT: **The discriminative Kalman filter for nonlinear and non-Gaussian sequential Bayesian filtering**. *Arxiv preprint* 2016, 1608.06622.

Chang L, Tsao DY: **The Code for Facial Identity in the Primate Brain**. *Cell* 2017, **169**: 1013-28.

* An impressive recent example that standard regression methods can be quite effective in predicting highly nonlinear neural responses if a good set of regression features are chosen.

Chen S, Shojaie A, Shea-Brown E, Witten D: **The Multivariate Hawkes Process in High Dimensions: Beyond Mutual Excitation**. *Arxiv preprint* 2017, 1707.04928.

Chen S, Shababo B, Deng X, Adesnik H, Paninski L: **Mapping Neural Microcircuits: Design and Inference**. *Statistical Analysis of Neural Data* 2017.

Chung J, Magland JF, Barnett A, Tolosa VM, Tooker AC, Lee KY, Shah KG, Felix SH, Frank LM, Greengard LF: **A Fully Automated Approach to Spike Sorting**. *Neuron* 2017, **95**:1381-1394.

Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV: **Neural population dynamics during reaching.** *Nature* 2012, **487**: 51-6.

Cong L, Wang Z, Chai Y, Hang W, Shang C, Yang W, Bai L, Du J, Wang K, Wen Q: **Rapid Whole Brain Imaging Of Neural Activities In Freely Behaving Larval Zebrafish.** *Biorxiv preprint* 2017, 131532.

Cunningham J, Yu BM: **Dimensionality reduction for large-scale neural recordings.** *Nature Neuroscience* 2014, **11**:1501-1509.

Cunningham J, Ghahramani Z: **Linear Dimensionality Reduction: Survey, Insights, and Generalizations.** *Journal of Machine Learning Research* 2015, **16**: 2859–2900.

* A unifying review of dimensionality reduction methods that surveys existing factor models of this type and provides algorithmic and software tools for designing bespoke methods for particular scientific problems of interest.

Deneux T, Kaszas A, Szalay G, Katona G, Lakner T, Grinvald A, Rózsa B, Vanzetta I: **Accurate spike estimation from noisy calcium signals for ultrafast three-dimensional imaging of large neuronal populations in vivo.** *Nature Communications* 2016, **7**: 12190.

Deng X, Liu DF, Kay K, Frank LM, Eden UT: **Clusterless Decoding of Position from Multiunit Activity Using a Marked Point Process Filter.** *Neural Computation* 2015, **7**:1438-60.

Dhawale A et al: **Automated long-term recording and analysis of neural activity in behaving animals.** *eLife* **6**: e27702.

Donoho D: **50 years of Data Science.** 2015, *Tukey Centennial workshop*.

* Reviews drivers of progress in machine learning, with important lessons for neural data science - particularly the importance of large public datasets (e.g., Imagenet) and clear open “gold standards” that the community can use to iteratively improve algorithms.

Dupre C, Yuste R: **Non-overlapping Neural Networks in *Hydra vulgaris*.** *Current Biology* 2017, **8**: 1085–1097.

Ecker AS, Berens P, Cotton RJ, Subramaniam M, Denfield GH, Cadwell CR, Smirnakis SM, Bethge M, Tolias AS: **State Dependence of Noise Correlations in Macaque Primary Visual Cortex.** *Neuron* 2014, **82**: 235–248.

Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP: *Nature Communications* 2016, **7**:13239.

* Introduces useful methodology for analyzing and testing orthogonality of various signals in population-level data; uses these methods to advance understanding of preparation and execution of movement.

Elsayed GF, Cunningham JP: **Structure in neural population recordings: an expected byproduct of simpler phenomena?** 2017 *Nature Neuroscience* **20**, 1310–1318

* Addresses a major debate in the field about the legitimacy of population-level neuroscience, and specifically clarifies existing results in motor and prefrontal cortices. Introduces a framework to test population-level neuroscience results against simpler alternatives such as tuning, temporal smoothness, and signal correlations.

Farina D, Vujaklija I, Sartori M, Kapelner T, Negro F, Jiang N, Bergmeister K, Andalib A, Principe J, Aszmann OC: **Man/machine interface based on the discharge timings of spinal motor neurons after targeted muscle reinnervation.** *Nature Biomedical Engineering* 2017, **1**: 0025.

Field G, Gauthier J, Sher A et al: **Functional connectivity in the retina at the resolution of photoreceptors.** *Nature* 2014, **467**: 673-677.

Friedrich J, Zhou P, Paninski L: **Fast Active Set Method for Online Spike Inference from Calcium Imaging.** *PLOS Comput. Bio.* **13**: e1005423.

Friedrich J, Yang W, Soudry D, Mu Y, Ahrens M, Yuste R, Peterka D, Paninski L: **Multi-scale approaches for high-speed imaging and analysis of large neural populations.** *PLOS Comput. Bio.* 2017, **13**, e1005685.

* Introduces a two-phase experimental design for imaging of large neuronal populations: first, obtain images at standard spatial resolution and determine the locations and shapes of neurons in the field of view; then switch to low-spatial-resolution imaging (which can be done faster over larger fields of view) and computationally demix the resulting data, exploiting the prior information obtained in the first imaging phase.

Gao P, Ganguli S: **On simplicity and complexity in the brave new world of large-scale neuroscience.** *Current Opinion in Neurobiology* 2015, **32**:148-155.

Gao Y, Buesing L, Shenoy KV, Cunningham JP: **High-dimensional neural spike train analysis with generalized count linear dynamical systems.** *NIPS* 2015, **28**.

Gao Y, Archer, E, Paninski L, Cunningham, J: **Latent linear-dynamical neural population models through nonlinear embedding.** *NIPS* 2016, **29**.

* An advance for nonlinear latent factor modeling of point process data: the method developed here simultaneously performs nonlinear dimensionality reduction on multiple spike train data and estimates a linear dynamical system in the resulting low-dimensional space. See also (Krishnan et al 2015; Johnson et al 2016; Sussillo et al 2016) for related approaches.

Gerhard F, Deger M, Truccolo W: **On the stability and dynamics of stochastic spiking neuron models: nonlinear Hawkes process and point process GLMs.** *PLOS Comput Biol* 2017, **13**: e1005390.

Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, Fan JM, Churchland MM, Kaufman MT, Kao JC, Ryu SI, Shenoy KV: **A high-performance neural prosthesis enabled by control algorithm design.** *Nature Neuroscience* 2012, **12**:1752-1757.

Giovanucci A *et al*: **OnACID: Online Analysis of Calcium Imaging Data in Real Time.** *NIPS*. 2017, **30**.

* Enables closed-loop experiments and brain-computer interfaces based on real-time denoised calcium imaging of hundreds of neurons.

Glaser JI, Chowdhury RH, Perich MG, Miller LE, Kording KP: **Machine learning for neural decoding.** *Arxiv preprint* 2017, 1708.00909.

Golub MD, Chase SM, Batista AP, Yue BM: **Brain computer interfaces for dissecting cognitive processes underlying sensorimotor control.** *Curr Opinion In Neurobiology* 2016, **37**: 53–58.

RLT Goris, JA Movshon, EP Simoncelli: **Partitioning neuronal variability.** *Nature Neuroscience* 2014, **17**: 858-865.

Grosenick L, Marshel J, Deisseroth K: **Closed-loop and activity-guided optogenetic control.** *Neuron* 2015, **86**: 106–139.

Haeffele BD, Vidal R: **Structured Low-Rank Matrix Factorization: Global Optimality, Algorithms, and Applications.** *Arxiv preprint* 2017, 1708.07850

Hagen E, TV Ness, Khosrowshahi A, Sorensen C, Fyhn M, Hafting T, Franke F, Einevoll, GT: **ViSAPy: A Python tool for biophysics-based generation of virtual spiking activity for evaluation of spike-sorting algorithms.** *Journal of Neuroscience Methods* 2015, **145**: 182-204.

Hall EC, Raskutti G, Willett R: **Inference of High-dimensional Autoregressive Generalized Linear Models.** *Arxiv preprint* 2017, 1605.02693.

Harrison MT, Amarasingham A, Truccolo W: **Spatio-temporal conditional inference and hypothesis tests for neural ensemble spiking precision** *Neural Computation* 2015, **27**: 104-50.

* A hypothesis testing framework to test the timing precision of ensembles of neural spike trains

D Hocker, Il M, Park: **Multistep Inference for Generalized Linear Spiking Models Curbs Runaway Excitation** *IEEE EMBS Conference on Neural Engineering* 2017.

Huys Q, Ahrens M, Paninski, L: **Efficient estimation of detailed single-neuron models** *J. Neurophysiology* **96**: 872-890.

Huys Q, Paninski L: **Smoothing of, and parameter estimation from, noisy biophysical recordings.** *PLOS Comput. Bio* 2016, **5**: e1000379.

Jewell S, Witten D: **Exact Spike Train Inference Via ℓ_0 Optimization** Arxiv preprint 2017, 1703.08644.

Johnson M, Duvenaud DK, Wiltchko A, Adams RP, Datta SR: **Composing graphical models with neural networks for structured representations and fast inference.** *NIPS* 2016, **29**.

Jonas E, Koerding K: **Automatic discovery of cell types and microcircuitry from neural connectomics.** *eLife* 2015, e04250.

* A nice first step towards automated clustering of cell types based on connectivity features.

Jonas E, Koerding K: **Could a Neuroscientist Understand a Microprocessor?** *PLoS Comput. Bio.* 2017, **13**: e1005268.

* A strong argument that our current suite of basic analysis approaches will likely fall short in achieving a deep understanding of complex nervous systems.

Jun JJ, Mitelut C, Lai C, Gratiy S, Anastassiou C, Harris TD: **Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction.** *Biorxiv preprint* 2017, 101030.

Kato S, Kaplan HS, Schrodell T, Yemini E, Lockery S, Zimmer M: **Global Brain Dynamics Embed the Motor Command Sequence of *Caenorhabditis elegans*.** *Cell* 2015, **163**: 656-69.

Kao JC, Nuyujukian P, Ryu SI, Churchland MM, Cunningham JP, Shenoy KV: **Single-trial dynamics of motor cortex and their applications to brain-machine interfaces.** *Nature Communications* 2015, **6**: 7759.

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

Kebschull JM, Garcia da Silva P, Reid AP, Peikon ID, Albeanu DF, Zador AM: **High-Throughput Mapping of Single-Neuron Projections by Sequencing of Barcoded RNA.** *Neuron* 2016, **91**: 975-987.

Kim TH, Zhang Y, Lecoq J, Jung J, Li J, Zeng H, Niell CM, Schnitzer MJ: **Long-Term Optical Access to an Estimated One Million Neurons in the Live Mouse Cortex.** *Cell* 2016, **17**: 3385-3394.

Kim DH, Kim J, Marques JC, Grama A, Hildebrand DGC, Gu W, Li JM, Robson DN: **Pan-neuronal calcium imaging with cellular resolution in freely swimming zebrafish.** *Nature Methods* 2017. doi:10.1038/nmeth.4429.

* An impressive recent paper demonstrating motion tracking and calcium imaging in freely behaving larval zebrafish. See also Cong et al 2017.

Kloosterman F, Layton SP, Chen Z, Wilson MA: **Bayesian decoding using unsorted spikes in the rat hippocampus.** *J. Neurophysiology* 2014, **111**: 217-27

Kriegeskorte N, Diedrichsen J: **Inferring brain computational mechanisms with models of activity measurements.** *Philos Trans R Soc Lond B Biol Sci.* 2016: **371**.

Krishnan RG, Shalit U, Sontag D: **Deep Kalman Filters.** *Arxiv preprint* 2015, 1511.05121.

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

Lee JH et al: **Highly multiplexed subcellular RNA sequencing in situ.** *Science* 2014, **343**: 1360-3.

Lee WCA, Bonin V, Reed M, Graham BJ, Hood G, Glattfelder K, Reid RC: **Anatomy and function of an excitatory network in the visual cortex.** *Nature* 2016, **532**: 370–374.

Lee J, Carlson D, Shokri H, Yao W, Goetz G, Hagen E, Batty E, Chichilnisky EJ, Einevoll G, Paninski L: **YASS: Yet another spike sorter.** *Biorxiv preprint* 2017, 151928.

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

Linderman S, Adams RP, Pillow JW: **Bayesian latent structure discovery from multi-neuron recordings.** *NIPS* 2017, **30**.

Linderman SW, Gershman SJ: **Using computational theory to constrain statistical models of neural data.** *Biorxiv preprint* 2017, 104737.

Linderman, S, Johnson M, Miller A, Adams R, Blei D, Paninski L: **Bayesian Learning and Inference in Recurrent Switching Linear Dynamical Systems.** *Artificial Intelligence and Statistics* 2017, 914-922.

* A significant methodological advance in estimating state-space models in which the latent dynamics can switch between one of a few discrete states; this significantly extends the range and applicability of these models.

Loaiza-Ganem G, Gao Y, Cunningham JP: **Maximum Entropy Flow Networks**. *International Conference on Learning Representations 2017*.

Lu R, Sun W, Liang Y, Kerlin A, Bierfeld J, Seelig JD, Wilson DE, Scholl B, Mohar B, Tanimoto M, Koyama M, Fitzpatrick D, Orger MB, Ji N: **Video-rate volumetric functional imaging of the brain at synaptic resolution**. *Nature Neuroscience* 2017, **20**: 620–628.

* An exciting method for imaging large neuronal populations; the success of the method is strongly dependent on computational demixing methods to extract the resulting data.

Machens CK, Romo R, Brody CD: **Functional, but not anatomical, separation of "what" and "when" in prefrontal cortex**. *J. Neuroscience*. 2010, **30**: 350-60.

Macosko EZ, Basu A, Satija R, Nemesh J, Shekhar K, Goldman M, Tirosh I, Bialas AR, Kamitaki N, Martersteck EM, Trombetta JJ, Weitz DA, Sanes JR, Shalek AK, Regev A, McCarroll SA: **Highly Parallel Genome-wide Expression Profiling of Individual Cells Using Nanoliter Droplets**. *Cell* 2015, **5**:1202–1214.

Macke JH, Busing GL, Cunningham JP, Yu BM, Shenoy KV, Sahani M: **Empirical models of spiking in neural populations**. *NIPS* 2012, **25**.

Mante V, Sussillo D, Shenoy KV, Newsome WT: **Context-dependent computation by recurrent dynamics in prefrontal cortex**. *Nature* 2013, **503**, 78-84

McIntosh L, Maheswaranathan N, Nayebi A, Ganguli S, Baccus S: **Deep Learning Models of the Retinal Response to Natural Scenes**. *NIPS* 2016, **29**.

Mena G, Grosberg L, Madugula S, Hottowy P, Litke A, Cunningham J, Chichilnisky EJ, Paninski L: **Removing Stimulation Artifacts From Neural Recordings Using Structured Gaussian Processes**. *Biorxiv preprint* 2016, 089912.

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

Merel J, Carlson D, Paninski L, Cunningham J: **Neuroprosthetic decoder training as imitation learning**. *PLOS Comput. Bio* 2016, **12**: e1004948.

Mineault PJ, Khawajaa FA, DA Butts, Packa CC: **Hierarchical processing of complex motion along the primate dorsal visual pathway**. *Proceedings of the National Academy of Sciences* 2012, **109**: E972-E980.

Neto JP, Lopes G, Frazão J, Nogueira J, Lacerda P, Baião P, Aarts A, Andrei A, Musa S, Fortunato E, Barquinha P, Kampff AR: **Validating silicon polytrodes with paired juxtacellular recordings: method and dataset.** *J. Neurophysiology* 2016, **116**: 892-903,

Nguyena JP, Shipleya FB, Linder AN, Plummera GS, Liua M, Setrua SU, Shaevitz JW, Leifera AM: **Whole-brain calcium imaging with cellular resolution in freely behaving *Caenorhabditis elegans*.** *Proceedings of the National Academy of Sciences* 2015, **113**: E1074-E1081.

* Demonstrates calcium imaging of a large fraction of the worm brain during behavior. See also Venkatachalam et al 2016.

Noebauer T, Skocek O, Pernia-Andrade AJ, Weilguny L, Traub FM, Molodtsov MI, Vaziri A: **Video rate volumetric Ca²⁺ imaging across cortical layers using Seeded Iterative Demixing (SID) microscopy.** *Nat Methods* 2017, **14**: 811-818.

Moore T, Hofer SB, Mrsic-Flogel TD, Carandini M, Harris KD: **Diverse coupling of neurons to populations in sensory cortex.** *Nature* 2012, **521**: 511–515.

Pachitariu M, Steinmetz N, Kadir S, Matteo Carandini M, Kenneth D Harris KD: **Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels.** *Biorxiv preprint* 2016, 061481.

* A significant step forward in scalable spike sorting on large multi-electrode arrays. See also Lee et al 2017.

Pachitariu M, Stringer C, Dipoppa M, Schröder S, Rossi LF, Dalgleish H, Carandini M, Harris KD: **Suite2p: beyond 10,000 neurons with standard two-photon microscopy.** *Biorxiv preprint* 2017, 061507.

Pakman A, Huggins J, Smith C, Paninski L: **Fast penalized state-space methods for inferring dendritic synaptic connectivity.** *Journal of Computational Neuroscience* 2014, **36**: 415-43

Paninski L: **Fast Kalman filtering on quasilinear dendritic trees.** *Journal of Computational Neuroscience* 2010, **28**: 211-28.

Paninski L, Pillow J, Lewi J: **Statistical models for neural encoding, decoding, and optimal stimulus design.** *Progress in Brain Research* 2007, **165**: 493-507.

Paninski L, Vidne M, DePasquale B, Ferreira D: **Inferring synaptic inputs given a noisy voltage trace.** *Journal of Computational Neuroscience* 2012, **33**: 1-19.

Parthasarathy N, Batty E, Falcon W, Rutten T, Rajpal M, Chichilnisky EJ, Paninski L: **Neural Networks for Efficient Bayesian Decoding of Natural Images from Retinal Neurons**. *NIPS 2017*.

* Develops a method for converting a forward encoding model into an approximation of the corresponding optimal Bayesian decoder, thus sidestepping what was previously a challenging computational problem.

Petersen A, Simon N, Witten D: **SCALPEL: Extracting Neurons from Calcium Imaging Data**. *Arxiv preprint 2017*. 1703.06946

Petreska B, Yu B, Cunningham J, Santhanam G, Ryu S, Shenoy K, **Dynamical segmentation of single trials from population neural data**. *NIPS 2012*.

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

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: **Simultaneous denoising, deconvolution, and demixing of calcium imaging data**. *Neuron 2016* **89**: 285-299.

* A significant advance in processing calcium imaging data, including critically the development of open source code that is now widely applied.

Pnevmatikakis E, Giovanucci A: **NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data**. *J Neurosci Methods*. 2017, **291**: 83-94.

Prevedel R, Verhoef AJ, Pernia-Andrade AJ, Weisenberger S, Huang B, Nobauer T, Fernandez A, Delcour JE, Golshani P, Baltuska A, Vaziri A: **Fast volumetric calcium imaging across multiple cortical layers using sculpted light**. *Nature Methods* 2016, **13** 1021-1028.

Rahnama RK, Machado T, Paninski L: **Robust and scalable Bayesian analysis of spatial neural tuning function data**. *Ann. Applied Stat.* 2017, Arxiv 1606.07845

Ramirez A, Paninski L: **Fast generalized linear model estimation via expected log-likelihoods**. *Journal of Computational Neuroscience*. 2014, **36**: 215-34.

Rios, G et al: **Nanofabricated Neural Probes for Dense 3-D Recordings of Brain Activity**. *Nano Letters* 2016, **16**: 6857–6862.

Sadtler PT, Kristin MQ, Matthew DG, Steven MC, Stephen IR, Elizabeth CT, Yu BM, Batista AP: **Neural constraints on learning**. *Nature* 2014, **512**: 423-426.

Savin C, Tkacik G: **Maximum entropy models as a tool for building precise neural controls.** *Current Opinion in Neurobiology* 2017, **46**: 120-126.

Semedo, Joao, *et al.*: **Extracting latent structure from multiple interacting neural populations.** *NIPS* 2014.

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

Song A, Charles AS, Koay SA, Gauthier JL, Thiberge SY, Pillow JW, Tank DW: **Volumetric two-photon imaging of neurons using stereoscopy (vTwINS).** *Nature Methods* 2017, **14**: 420-426.

Smith AC, Brown EN: **Estimating a State-Space Model from Point Process Observations.** *Neural Computation* 2003, **15**: 965-991.

Song A, Charles AS, Gauthier JL, Koay SA, Tank DW, Pillow JW: **Two-photon microscopy simulation for optics optimization and benchmarking.** *CoSyNe* 2017.

Soudry D, Keshri, S, Stinson P, Oh MW, Iyengar G, Paninski L: **Efficient “shotgun” inference of neural connectivity from highly sub-sampled activity data.** *PLOS Comput. Bio.* 2015, **11**: e1004464.

* Proposes a novel experimental design for solving the “common input” problem that has long stymied progress towards estimating connectivity from multiple spike train data

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

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

Sussillo D, Jozefowicz R, Abbott LF, Pandarinath C: **LFADS – Latent Factor Analysis via Dynamical Systems.** *Arxiv preprint* 2016, 1608.06315

Todorova S, Sadtler P, Batista A, Chase S, Ventura V: **To sort or not to sort: the impact of spike sorting on neural decoding performance.** *Journal of Neural Engineering* 2014, **11**, 056005.

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.** *International Conference of the IEEE EMBS* 2015.

Turaga SC, Buesing L, Packer AM, Dagleish H, Pettit N, Hausser M, Macke JH: **Inferring neural population dynamics from multiple partial recordings of the same neural circuit** *NIPS*, 2014.

Venkatachalam V, Ji N, Wang X, Clark C, Mitchell JK, Mason K, Tabone CJ, Florman J, Hongfei J, Greenwood J, Chisholm AD, Srinivasan J, Alkema M, Zhen M, Samuel ADT: **Pan-neuronal imaging in roaming *Caenorhabditis elegans***. *Proceedings of the National Academy of Sciences* 2017, **113**: E1082-E1088.

Vishwanathan A, Daie K, Ramirez AD, Lichtman JW, Aksay ERF, Seung HS: **Electron Microscopic Reconstruction of Functionally Identified Cells in a Neural Integrator**. *Current Biology* 2017, **27**: 2137-2147.

Vogelstein J, Packer A, Machado T, Sippy T, Babadi B, Yuste R, Paninski L: **Fast non-negative deconvolution for spike train inference from calcium imaging**. *J. Neurophysiology* 2010, **104**: 3691-3704.

Williamson RC, Cowley BR, Litwin-Kumar A, Doiron B, Kohn A, Smith MA, Yu BM: **Scaling Properties of Dimensionality Reduction for Neural Populations and Network Models**. *PLOS Comput. Bio* 2017, e1005141.

Wiltchko AB, Johnson JJ, Iurilli G, Peterson RE, Katon JM, Stan LP, Abaira VE, Adams RP, Datta SR: **Mapping Sub-Second Structure in Mouse Behavior**. *Neuron* 2015, **88**: 1121-1135.
* Introduces an elegant and effective approach for automated parsing of behavioral data, using factor models in the vein of Gao et al (2016), Krishnan et al (2015), Johnson et al (2016), and Linderman et al (2017).

Wu A, Park Im, Pillow JW: **Convolutional spike-triggered covariance analysis for neural subunit models**. *NIPS* 2015.

Xu Y, Zou P, Cohen, AE: **Voltage imaging with genetically encoded indicators**. *Current Opinion in Chemical Biology* 2017, **39**: 1-10.

Yamins DLK, DiCarlo JJ: **Using goal-driven deep learning models to understand sensory cortex**. *Nature Neuroscience* 2016, **19**: 356-365.

* Careful modeling of a hierarchical artificial neural network shows remarkable similarities to the cortical visual processing stream.

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

Yatsenko D, Reimer J, Ecker AS, Walker EY, Sinz F, Berens P, Hoenselaar A, Cotton R, Siapias AS, Tolias AT: **DataJoint: managing big scientific data using MATLAB or Python**.

Biorxiv preprint 2015, 031658.

Yger P, Spampinato GLB, Esposito E, LeFebvre B, Deny S, Gardella C, Stimberg M, Jetter F, Zeck G, Picaud S, Duebel J, Marre O: **Fast and accurate spike sorting in vitro and in vivo for up to thousands of electrodes**. *Biorxiv preprint 2016*, 067843.

Yu BM, Cunningham J, Gopal Santhanam, Ryu S, Shenoy KV, Sahani M: **Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity**. *J. Neurophysiology* 2009, **102**: 614–635.

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

* A significant step forward in processing microendoscopic calcium imaging data; sharply reduces the effect of background fluctuations due to out-of-focus fluorescence, which had previously posed a major problem in these recordings.

Zhao Y, Park IM: **Variational Latent Gaussian Process for Recovering Single-Trial Dynamics from Population Spike Trains**. *Neural Computation* 2017. **29**: 1293-1316.

Zhao Y, Park IM: **Interpretable Nonlinear Dynamic Modeling of Neural Trajectories**. *NIPS* 2016, **29**.