



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

L Paninski^{1,2} and JP Cunningham¹

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.

Addresses

¹ Department of Statistics, Grossman Center for the Statistics of Mind, Zuckerman Mind Brain Behavior Institute, Center for Theoretical Neuroscience, Columbia University, United States

² Department of Neuroscience, Grossman Center for the Statistics of Mind, Zuckerman Mind Brain Behavior Institute, Center for Theoretical Neuroscience, Columbia University, United States

Corresponding author: .Paninski, L (liam@stat.columbia.edu)

Current Opinion in Neurobiology 2018, 50:232–241

This review comes from a themed issue on **Neurotechnologies**

Edited by **Polina Anikeeva** and **Liquan Luo**

For a complete overview see the [Issue](#) and the [Editorial](#)

<https://doi.org/10.1016/j.conb.2018.04.007>

0959-4388/© 2018 Elsevier Ltd. All rights reserved.

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

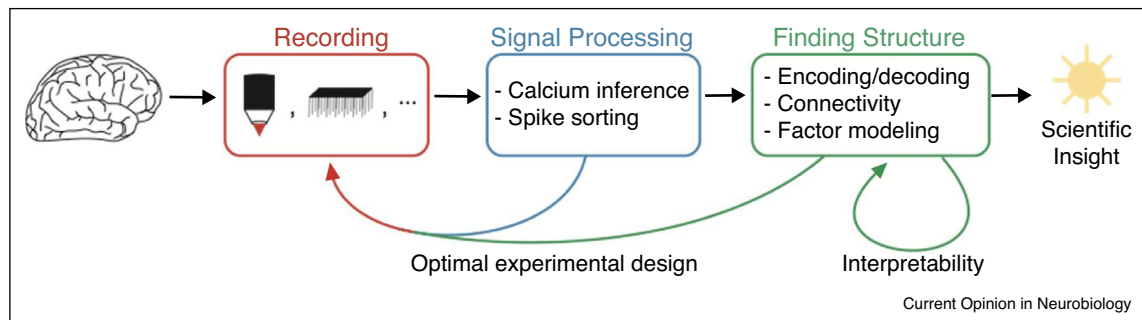
methods and MEAs output data at rates on the order of *terabytes/hour*, and data rates continue to increase. At these huge rates processing and even storing the data is challenging, let alone optimally extracting all the useful information in these data streams; without the right analytical technology, we will never unlock the true potential of these experimental advances (Figure 1).

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 (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 [1,2,3,4–6] present constrained and/or nonnegative matrix factorization (NMF) approaches to simultaneously solve these demixing and deconvolution problems. Ref. [7] 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 [8,9] will likely benefit from this approach or modifications thereof. Refs. [10,11] developed real-time implementations that process incoming data online, one imaging frame at a time, enabling closed-loop experiments. Ref. [12] developed an improved and more general implementation of the hidden Markov model deconvolution approach of [13]. Ref. [14] developed hierarchical Bayesian methods for sharing statistical information across behaviorally similar

Figure 1



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

trials to enable temporal super-resolution of estimated neural activity. Ref. [15] developed useful mathematical theory on exactly solving the sparse deconvolution problem addressed in [3*,10]. Ref. [16] 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); Ref. [17] presented a promising first step. Since NMF is a non-convex problem, accurate initialization of the estimates is critical; Refs. [18,11*] explore these issues further. 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 [19*]; see [20] 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 [21,22].

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. The suboptimality of, for example, optimizing an imaging apparatus in isolation is widely recognized; 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. Refs. [23–25,26*,27*] have all offered variations on a 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 more without negatively impacting the quality of the recovered temporal neural activity [27*]. Ref. [28] 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 [29] 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, for example, larval zebrafish [30], *Drosophila* [8], or hydra [31]. Good solutions have been developed in *Caenorhabditis elegans* [32*,33], though demixing of fast cytosolic (non-nuclear-localized) signals remains an unsolved problem. We expect non-rigid registration approaches similar to those developed by [34] to be helpful here; see also [35*] 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 [36], which is expected to be a major growth area over the next decade; see for example [37], 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; earlier works offer

algorithmic starting points for modeling voltage data with subcellular resolution [38–42].

Spike sorting data from large-scale MEAs

Spike sorting has been a not-quite-completely-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 [43–45] (see also www.darpa.mil/program/neural-engineering-system-design). This looming bottleneck has driven a recent uptick in studies of spike sorting activity from large dense MEAs [46*,47–51] see also [52] 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, Ref. [46*] introduced a fast implementation of a matching-pursuit algorithm to detect these spike overlaps, and [48] 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 solutions, for example datasets in which ground truth spiking for single neurons is available [53]. 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 partial validation. In parallel, as in the imaging context, the iterative improvement of simulators of electrical activity [54] remains an important direction. Another useful approach is to create ‘hybrid’ datasets combining simulated spiking signals with real noise signals [46*,49]. 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 [55,56] have been developed to extract information from these unsorted spikes. A combined strategy (in which one sorts

the sortable units, but also exploits information from the non-separable units and local field potential signals) has been shown to have superior performance in offline movement decoding experiments [57,58].

Finally, as interest grows in bidirectional electrical neural interfaces that stimulate and record simultaneously, the problem of stimulation artifact cancellation becomes critical; see [59] for a scalable Bayesian artifact removal algorithm applied to large retinal MEA data.

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. Firstly, How is information encoded in spike trains, and how can we decode this information? Secondly, Can we infer network connectivity from multi-spike train data? Thirdly, 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: for example, factor analysis models developed to address question 3 can lead to improved decoding of neural data (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 [60] for a review. Some recent work has targeted the computational efficiency of GLM estimation methods [61–63]. 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. Ref. [64*] describe an exciting recent application showing that with a good choice of features it is possible for simple regression models 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. Refs. [65,66*] 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 [67] 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: Refs. [68,69*,70] provide three examples of this idea (see also [71] 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 machine learning tasks (e.g. object recognition) and use the resulting feature sets to predict responses; see [72*,73] 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 [74] contributed substantial empirical performance improvement in brain-machine interface decoding from the motor cortex; Ref. [75] provides a rigorous theoretical foundation and generalization of this algorithm. Another thread of work has shown that more constrained models of joint neural variability can be used to construct better decoders [76,77]; see also [78] for a promising converse approach using a discriminative (rather than the more typical generative) model. Finally, ANNs have recently been applied to decoding problems [79,80]; Ref. [81*] 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. Ref. [82*] 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, and 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 also [83] 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 [84*], who introduce new conditional inference methods to address hypotheses about the precision of multineuronal responses, and [85], who discuss methods for incorporating stronger prior knowledge into network estimates; see also [86*] 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; Refs. [87,88] 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. Refs. [89,90] 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 [91,92].

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 [93], 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. [93–98,99*,100,101]), systems with switching dynamical structure [102,103,104*,105,106*], and systems with recurrent neural network dynamical structure

[107,108]. 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 [109*].

Dimensionality reduction approaches have been widely used in neuroscience [110]. 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 [111]; 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 [112,113*,114]. Third, Ref. [115] 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 [116,117]. 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 [118,119].

Interpretability of large-scale neural data analysis

Of course, the impetus behind the development of new large scale neural recording and analysis methods 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 question [84*,120*,121,122], 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 [116] and the presence of de-mixed readouts in prefrontal cortex [111] have clarified these previous results [120*].

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 theoretical notion of inherent data complexity [123], 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 [124]. 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. Ref. [125*] recently presented an arguably 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, and many analyses have similarly flourished without appropriate statistical testing to constrain their interpretation. 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 and testable analysis methods that can connect experiment and theory in a tightly closed loop.

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 [91,126,127].
- Fusion of multimodal datasets will represent another critical research area, as large-scale connectomic and cell type constraints [128–130] become available to inform functional models [125*].
- 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 [120*].
- 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 [131].
- More broadly, sociological trends towards more open and large-scale data sharing and open-source collaborative projects, supported by stronger pipelines [132] 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 [133,134]; increased investment in neural data science and neurotheory training will pay rich dividends in improving our understanding of neural systems over the next decade.

Conflict of interest statement

Nothing declared.

Acknowledgements

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 and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
- of outstanding interest

1. 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 Neurophysiol* 2010, **104**:3691–3704.
2. Andilla FD, Hamprecht AF: **Sparse space-time deconvolution for calcium image analysis.** *NIPS* 2014:27.
3. 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.
4. 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.
5. Haeffele BD, Vidal R: **Structured low-rank matrix factorization: global optimality, algorithms, and applications.** *Arxiv preprint* 2017. 1708.07850.
6. Inan H, Erdogan M, Schnitzer M: **Robust estimation of neural signals in calcium imaging.** *NIPS* 2017.
7. 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.** *eLife* 2018, **7**: e28728.
8. 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.** *Nat Photon* 2015, **9**:113–119.
9. 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.
10. Friedrich J, Zhou P, Paninski L: **Fast active set method for online spike inference from calcium imaging.** *PLOS Comput. Biol.* 2016, **13**: e1005423.
11. Giovanucci A et al.: **OnACID: online analysis of calcium imaging data in real time.** *NIPS* 2017:30.
12. 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.** *Nat Commun* 2016, **7**:12190.
13. Vogelstein J, Watson B, Packer A, Yuste R, Jedynak B, Paninski L: **Spike inference from calcium imaging using sequential Monte Carlo methods.** *Biophys J* 2009, **97**:636–655.
14. 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.

A significant advance in processing calcium imaging data, including open source code that is now widely applied.

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.

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

15. Jewell S, Witten D: **Exact spike train inference via ℓ_0 optimization**. *Arxiv preprint* 2017. 1703.08644.
16. Wei Z, Druckmann S *et al.*: **Different Recording Methods Offer Different Tradeoffs for Interrogating Neural Activity**. 2018 <http://www.im-phys.org>.
17. 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.
18. Petersen A, Simon N, Witten D: **SCALPEL: extracting neurons from calcium imaging data**. *Arxiv preprint* 2017. 1703.06946.
19. Donoho D: **50 years of data science**. *Tukey Centennial Workshop* 2015.
- 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.
20. 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.
21. Lee WCA, Bonin V, Reed M, Graham BJ, Hood G, Glatfelter K, Reid RC: **Anatomy and function of an excitatory network in the visual cortex**. *Nature* 2016, **532**:370-374.
22. Vishwanathan A, Daie K, Ramirez AD, Lichtman JW, Aksay ERF, Seung HS: **Electron microscopic reconstruction of functionally identified cells in a neural integrator**. *Curr Biol* 2017, **27**:2137-2147.
23. 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**:269-284.
24. 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**. *Nat Methods* 2016, **13**:1021-1028.
25. Song A, Charles AS, Koay SA, Gauthier JL, Thiberge SY, Pillow JW, Tank DW: **Volumetric two-photon imaging of neurons using stereoscopy (vTwINS)**. *Nat Methods* 2017, **14**:420-426.
26. 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**. *Nat Neurosci* 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.
27. 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 Biol* 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.
28. Noebauer T, Skocek O, Pernia-Andrade AJ, Weiglun L, Traub FM, Molodtsov MI, Vaziri A: **Video rate volumetric Ca^{2+} imaging across cortical layers using Seeded Iterative Demixing (SID) microscopy**. *Nat Methods* 2017, **14**:811-818.
29. Song A, Charles AS, Gauthier JL, Koay SA, Tank DW, Pillow JW: **Two-photon microscopy simulation for optics optimization and benchmarking**. *CoSyNe* 2017.
30. 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**. *eLife* 2017, **6**:e28158.
31. Dupre C, Yuste R: **Non-overlapping Neural Networks in Hydra vulgaris**. *Curr Biol* 2017, **8**:1085-1097.
32. Nguyen JP, Shipley FB, Linder AN, Plummer GS, Liu M, Setru SU, Shaevitz JW, Leifer AM: **Whole-brain calcium imaging with cellular resolution in freely behaving *Caenorhabditis elegans***. *Proc Natl Acad Sci U S A* 2015, **113**:E1074-E1081.
- Demonstrates calcium imaging of a large fraction of the worm brain during behavior. See also [33].
33. 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***. *Proc Natl Acad Sci U S A* 2017, **113**:E1082-E1088.
34. Pnevmatikakis E, Giovanucci A: **NoRMCorre: an online algorithm for piecewise rigid motion correction of calcium imaging data**. *J Neurosci Methods* 2017, **291**:83-94.
35. 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**. *Nat Methods* 2017 <http://dx.doi.org/10.1038/nmeth.4429>.
- An impressive recent paper demonstrating motion tracking and calcium imaging in freely behaving larval zebrafish. See also [30].
36. Buchanan EK, Friedrich J, Kinsella I, Stinson P, Zhou P, Gerhard F, Ferrante J, Dempsey G, Paninski L: **Constrained matrix factorization methods for denoising and demixing voltage imaging data**. *COSYNE* 2018. III-62.
37. Xu Y, Zou P, Cohen AE: **Voltage imaging with genetically encoded indicators**. *Curr Opin Chem Biol* 2017, **39**:1-10.
38. Huys Q, Ahrens M, Paninski L: **Efficient estimation of detailed single-neuron models**. *J Neurophysiol* 2006, **96**:872-890.
39. Huys Q, Paninski L: **Smoothing of, and parameter estimation from, noisy biophysical recordings**. *PLOS Comput Biol* 2009, **5**:e1000379.
40. Paninski L: **Fast Kalman filtering on quasilinear dendritic trees**. *J Comput Neurosci* 2010, **28**:211-228.
41. Paninski L, Vidne M, DePasquale B, Ferreira D: **Inferring synaptic inputs given a noisy voltage trace**. *J Comput Neurosci* 2012, **33**:1-19.
42. Pakman A, Huggins J, Smith C, Paninski L: **Fast penalized state-space methods for inferring dendritic synaptic connectivity**. *J Comput Neurosci* 2014, **36**:415-443.
43. Tsai D *et al.*: **High-channel-count, high-density micro-electrode array for closed-loop investigation of neuronal networks**. *International Conference of the IEEE EMBS*. 2015.
44. Rios G *et al.*: **Nanofabricated neural probes for dense 3-D recordings of brain activity**. *Nano Letters* 2016, **16**:6857-6862.
45. Jun J *et al.*: **Fully integrated silicon probes for high-density recording of neural activity**. *Nature* 2017, **551**:232-236.
46. Pachitariu M, Steinmetz N, Kadir S, Matteo Carandini M, Kenneth D, Harris KD: **Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels**. *NIPS* 2016.
- A significant step forward in scalable spike sorting on large multi-electrode arrays; see also [48].
47. 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.
48. 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**. *NIPS* 2017.
49. 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.
50. Dhawale A *et al.*: **Automated long-term recording and analysis of neural activity in behaving animals**. *eLife* 2017, **6**:e27702.

51. 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.
52. 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**. *Nat Biomed Eng* 2017, **1**:0025.
53. 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.
54. Hagen E, Ness TV, 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**. *J Neurosci Methods* 2015, **145**:182-204.
55. Kloosterman F, Layton SP, Chen Z, Wilson MA: **Bayesian decoding using unsorted spikes in the rat hippocampus**. *J Neurophysiol* 2014, **111**:217-227.
56. Deng X, Liu DF, Kay K, Frank LM, Eden UT: **Clusterless decoding of position from multiunit activity using a marked point process filter**. *Neural Comput* 2015, **7**:1438-1460.
57. 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 Neurophysiol* 2012, **107**:1337-1355.
58. 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**. *J Neural Eng* 2014, **11**:056005.
59. 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**. *PLoS Comput Biol* 2017, **13**:e1005842.
60. Paninski L, Pillow J, Lewi J: **Statistical models for neural encoding, decoding, and optimal stimulus design**. *Prog Brain Res* 2007, **165**:493-507.
61. Mena G, Paninski L: **On quadrature methods for refractory point process likelihoods**. *Neural Comput* 2014, **26**:2790-2797.
62. Ramirez A, Paninski L: **Fast generalized linear model estimation via expected log-likelihoods**. *J Comput Neurosci* 2014, **36**:215-234.
63. Wu A, Park IM, Pillow JW: **Convolutional spike-triggered covariance analysis for neural subunit models**. *NIPS* 2015.
64. Chang L, Tsao DY: **The code for facial identity in the primate brain**. *Cell* 2017, **169**:1013-1028.
65. Field G, Gauthier J, Sher A *et al.*: **Functional connectivity in the retina at the resolution of photoreceptors**. *Nature* 2014, **467**:673-677.
66. Antolik J, Hofer SB, Bednar JA, Msršic-Flogel TD: **Model constrained by visual hierarchy improves prediction of neural responses to natural scenes**. *PLoS Comput Biol* 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.
67. Rahnama RK, Machado T, Paninski L: **Robust and scalable Bayesian analysis of spatial neural tuning function data**. *Ann Appl Stat* 2017, **11**:598-637.
68. McIntosh L, Maheswaranathan N, Nayebi A, Ganguli S, Baccus S: **Deep learning models of the retinal response to natural scenes**. *NIPS* 2016.
69. Batty E, Merel J, Brackbill N, Heitman A, Sher A, Litke A, Chichilnisky EJ, Paninski L: **Multilayer network models of primate retinal ganglion cells**. *Int Conf Learn Represent* 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 [68].
70. Cadena S, Denfield G, Walker E, Gatys L, Tolias A, Bethge M, Ecker A: **Deep convolutional models improve predictions of macaque V1 responses to natural images**. *Biorxiv preprint* 2017:201764.
71. Mineault PJ, Khawaja FA, Butts D, Pack CC: **Hierarchical processing of complex motion along the primate dorsal visual pathway**. *Proc Natl Acad Sci U S A* 2012, **109**:E972-E980.
72. Yamins DLK, DiCarlo JJ: **Using goal-driven deep learning models to understand sensory cortex**. *Nat Neurosci* 2016, **19**:356-365.
- Careful modeling of a hierarchical artificial neural network shows remarkable similarities to the cortical visual processing stream.
73. Kriegeskorte N, Diedrichsen J: **Inferring brain computational mechanisms with models of activity measurements**. *Philos Trans R Soc Lond B Biol Sci* 2016:371.
74. 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**. *Nat Neurosci* 2012, **12**:1752-1757.
75. Merel J, Carlson D, Paninski L, Cunningham J: **Neuroprosthetic decoder training as imitation learning**. *PLoS Comput Biol* 2016, **12**:e1004948.
76. Lawhern V, Wu W, Hatsopoulos N, Paninski L: **Population neuronal decoding using a generalized linear model with hidden states**. *J Neurosci Methods* 2010, **189**:267-280.
77. 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**. *Nat Commun* 2015, **6**:7759.
78. 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.
79. Sussillo D, Stavisky SD, Kao JC, Ryu SI, Shenoy KV: **Making brain-machine interfaces robust to future neural variability**. *Nat Commun* 2016, **7**:13749.
80. Glaser JI, Chowdhury RH, Perich MG, Miller LE, Kording KP: **Machine learning for neural decoding**. *Arxiv preprint* 2017. 1708.00909.
81. 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.
82. 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 Biol* 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.
83. Turaga SC, Buesing L, Packer AM, Dalgleish H, Pettit N, Hausser M, Macke JH: **Inferring neural population dynamics from multiple partial recordings of the same neural circuit**. *NIPS* 2014.
84. Harrison MT, Amarasingham A, Truccolo W: **Spatio-temporal conditional inference and hypothesis tests for neural ensemble spiking precision**. *Neural Comput* 2015, **27**:104-150.
- A hypothesis testing framework to test the timing precision of ensembles of neural spike trains.
85. Linderman S, Adams RP, Pillow JW: **Bayesian latent structure discovery from multi-neuron recordings**. *NIPS* 2017:30.
86. 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.

87. Shababo B, Paige B, Pakman A, Paninski L: **Bayesian inference and online experimental design for mapping neural microcircuits**. *NIPS* 2013.
 88. Chen S, Shababo B, Deng X, Adesnik H, Paninski L: **Mapping neural microcircuits: design and inference**. *Stat Anal Neural Data* 2017.
 89. 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.
 90. Hocker D, Park IM: **Multistep inference for generalized linear spiking models curbs runaway excitation**. *IEEE EMBS Conf Neural Eng* 2017.
 91. Chen S, Shojaie A, Shea-Brown E, Witten D: **The multivariate Hawkes process in high dimensions: beyond mutual excitation**. *Arxiv preprint* 2017. 1707.04928.
 92. Hall EC, Raskutti G, Willett R: **Inference of high-dimensional autoregressive generalized linear models**. *Arxiv preprint* 2017. 1605.02693.
 93. Smith AC, Brown EN: **Estimating a state-space model from point process observations**. *Neural Comput* 2003, **15**:965-991.
 94. 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 Neurophysiol* 2009, **102**:614-635.
 95. Macke JH, Busing GL, Cunningham JP, Yu BM, Shenoy KV, Sahani M: **Empirical models of spiking in neural populations**. *NIPS* 2012:25.
 96. Goris RLT, Movshon JA, Simoncelli EP: **Partitioning neuronal variability**. *Nat Neurosci* 2014, **17**:858-865.
 97. Ecker AS, Berens P, Cotton RJ, Subramaniyan 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.
 98. Gao Y, Buesing L, Shenoy KV, Cunningham JP: **High-dimensional neural spike train analysis with generalized count linear dynamical systems**. *NIPS* 2015:28.
 99. 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.* [107]; Johnson *et al.* [105]; Sussillo *et al.* [79]) for related approaches.
100. Zhao Y, Park IM: **Variational latent gaussian process for recovering single-trial dynamics from population spike trains**. *Neural Comput* 2017, **29**:1293-1316.
 101. Zhao Y, Park IM: **Interpretable nonlinear dynamic modeling of neural trajectories**. *NIPS* 2016:29.
 102. Petreska B, Yu B, Cunningham J, Santhanam G, Ryu S, Shenoy K: **Dynamical segmentation of single trials from population neural data**. *NIPS* 2012.
 103. 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**:656-669.
 104. Wiltschko 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.* [99], Krishnan *et al.* [107], Johnson *et al.* [105], and Linderman *et al.* [85].
105. Johnson M, Duvenaud DK, Wiltschko A, Adams RP, Datta SR: **Composing graphical models with neural networks for structured representations and fast inference**. *NIPS* 2016:29.
 106. Linderman S, Johnson M, Miller A, Adams R, Blei D, Paninski L: **Bayesian learning and inference in recurrent switching linear dynamical systems**. *Artif Intell Stat* 2017:914-922.
- A 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.
107. Krishnan RG, Shalit U, Sontag D: **Deep Kalman filters**. *Arxiv preprint* 2015. 1511.05121.
 108. Sussillo D, Jozefowicz R, Abbott LF, Pandarinath C: **LFADS – latent factor analysis via dynamical systems**. *Arxiv preprint* 2016. 1608.06622.
 109. Cunningham J, Ghahramani Z: **Linear dimensionality reduction: survey, insights, and generalizations**. *J Mach Learn Res* 2015, **16**:2859-2900.
- A unifying treatment of linear dimensionality reduction methods that surveys existing factor models of this type and provides algorithmic and software tools for designing methods for particular scientific problems of interest.
110. Cunningham J, Yu BM: **Dimensionality reduction for large-scale neural recordings**. *Nat Neurosci* 2014, **11**:1501-1509.
 111. Machens CK, Romo R, Brody CD: **Functional, but not anatomical, separation of “what” and “when” in prefrontal cortex**. *J Neurosci* 2010, **30**:350-360.
 112. Kaufman MT, Churchland MM, Ryu SI, Shenoy KV: **Cortical activity in the null space: permitting preparation without movement**. *Nat Neurosci* 2014, **17**:440-448.
 113. Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP: *Nat Commun* 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.
114. 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.
 115. 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.
 116. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV: **Neural population dynamics during reaching**. *Nature* 2012, **487**:51-56.
 117. Mante V, Sussillo D, Shenoy KV, Newsome WT: **Context-dependent computation by recurrent dynamics in prefrontal cortex**. *Nature* 2013, **503**:78-84.
 118. Semedo, Joao *et al.*: **Extracting latent structure from multiple interacting neural populations**. *NIPS* 2014.
 119. 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.
 120. Elsayed GF, Cunningham JP: **Structure in neural population recordings: an expected byproduct of simpler phenomena?** *Nat Neurosci* 2017, **20**:1310-1318.
- Addresses a major debate in the field about the interpretability of population-level neural data analysis approaches, and specifically clarifies existing results in motor and prefrontal cortices. Introduces a framework to test population-level analysis methods against simpler alternative models.
121. Loaiza-Ganem G, Gao Y, Cunningham JP: **Maximum entropy flow networks**. *Int Conf Learn Represent* 2017.
 122. Savin C, Tkacik G: **Maximum entropy models as a tool for building precise neural controls**. *Curr Opin Neurobiol* 2017, **46**:120-126.
 123. Gao P, Ganguli S: **On simplicity and complexity in the brave new world of large-scale neuroscience**. *Curr Opin Neurobiol* 2015, **32**:148-155.
 124. 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 Biol* 2017:e1005141.
125. Jonas E, Koerding K: **Could a neuroscientist understand a microprocessor?** *PLoS Comput Biol* 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
126. Ahmadian Y, Packer A, Yuste R, Paninski L: **Designing optimal stimuli to control neuronal spike timing.** *J Neurophysiol* 2011, **106**:1038-1053.
127. Grosenick L, Marshel J, Deisseroth K: **Closed-loop and activity-guided optogenetic control.** *Neuron* 2015, **86**:106-139.
128. Lee JH *et al.*: **Highly multiplexed subcellular RNA sequencing in situ.** *Science* 2014, **343**:1360-1363.
129. Macosko EZ, Basu A, Satija R, Nemes 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.
130. 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.
131. Linderman SW, Gershman SJ: **Using computational theory to constrain statistical models of neural data.** *Curr Opin Neurobiol* 2017, **46**:14-24.
132. 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.
133. 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.
134. Barone L, Williams J, Micklos D: **Unmet needs for analyzing biological big data: a survey of 704 NSF principal investigators.** *PLoS Comput Biol* 2017, **13**:e1005755.