Analysis of functional imaging data at single-cellular resolution*

Eftychios A. Pnevmatikakis[1] and Liam Paninski[2]
[1] Center for Computational Biology, Flatiron Institute, New York, NY
[2] Departments of Statistics and Neuroscience,
Grossman Center for the Statistics of Mind,
Center for Theoretical Neuroscience, Columbia University

September 18, 2018

Abstract

Calcium imaging (CI) has become the dominant method for recording from large populations of neurons, due to several well-known advantages: CI offers sub-cellular spatial resolution with cell-type specificity and can be coupled easily with a variety of genetic tools; CI has proven scalability to record simultaneously from $O(10^4)$ neurons in vivo; finally, CI allows for longitudinal tracking of cellular activity across multiple days.

At the same time, CI presents some important analysis challenges: 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. In addition, CI produces datasets that can be quite large (hundreds of GB per hour in some cases); thus analysis methods must be scalable and as automated and reproducible as possible.

This chapter presents a (non-exhaustive) review of current state of the art analysis methods for CI data, with a focus on data recorded at single-neuron resolution. We will also touch briefly on the analysis of voltage imaging data and widefield CI data, both of which come with their own related but distinct analysis challenges.

Introduction

We begin with an overview: the typical steps in a modern calcium imaging analysis pipeline are shown in Fig. 1. After correction of low-level issues (e.g., different gains or noise levels on different pixels in CMOS cameras, or line-to-line phase errors in multiphoton scanning imaging), the dataset in movie format is motion corrected to remove artifacts arising from brain motion and slow imaging rate. Next the movie is denoised and compressed into a smaller format. Next the movie is *demixed* to extract a shape and temporal trace corresponding to each neural component. These temporal traces represent the average fluorescence within each spatial component in each temporal frame and are therefore an indirect measure of neural activity; these traces can be temporally *deconvolved* to estimate the underlying activity of each corresponding neuron. Finally, an imaging experiment can visit the same field of view (FOV) over the course of multiple sessions/days. To combine the results from multiple sessions, the components from the different sessions need to be registered. Below we discuss each of these steps in more detail.

*This chapter was written for the Society for Neuroscience 2018 short course on "Functional, structural, and molecular imaging, and big data analysis," and is partially adapted from "Analysis pipelines for calcium imaging data" by Eftychios Pnevmatikakis (under review) and (Paninski and Cunningham, 2018).

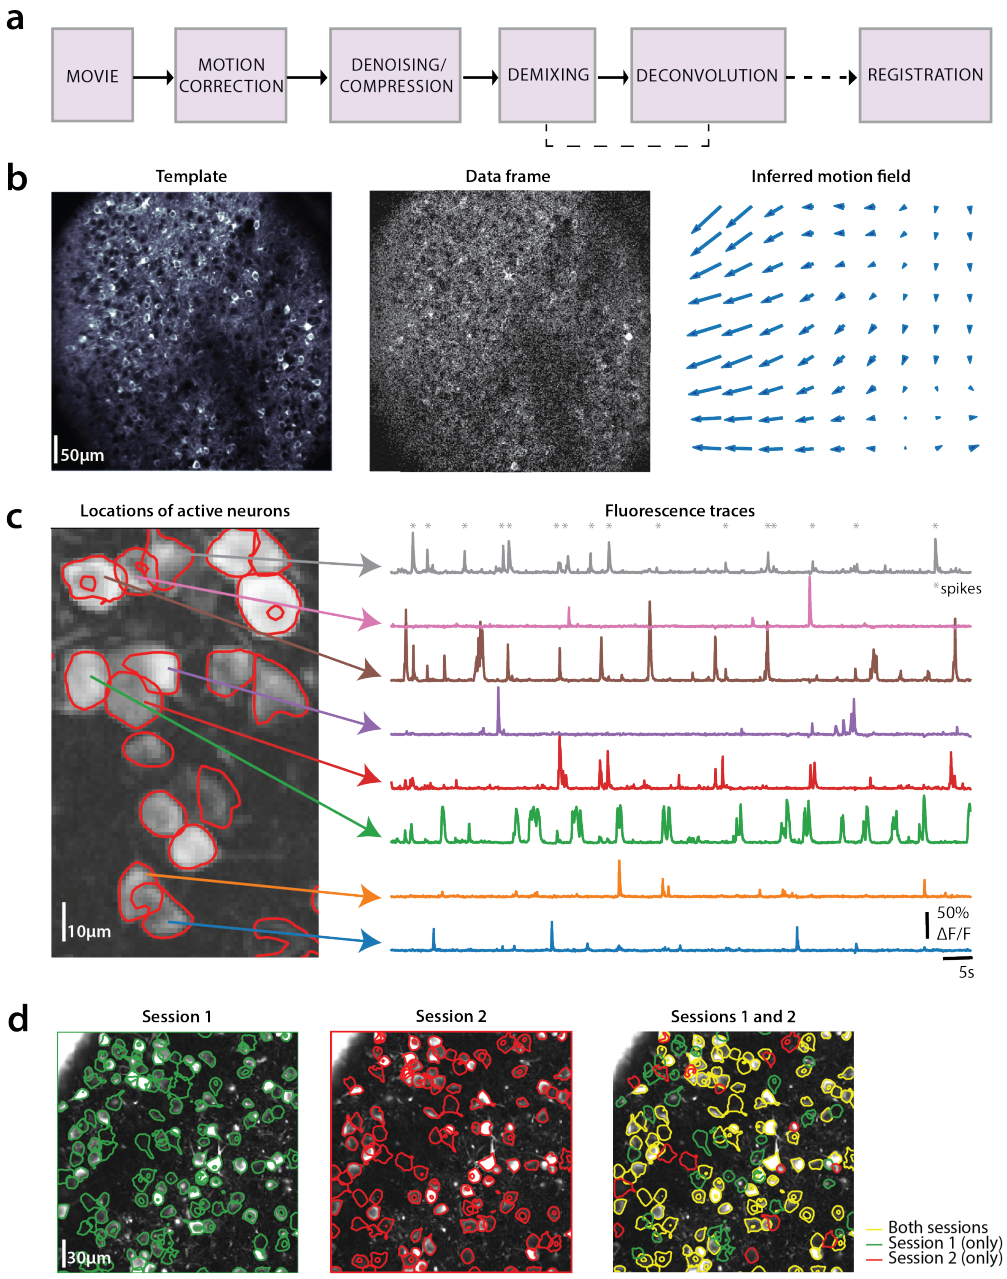

Figure 1: Typical analysis pipeline for calcium imaging data. (a) The data is first processed for removing motion artifacts (b). This can be done by estimating a motion field from aligning each data frame to a template. Subsequently the locations of the neurons in the imaged FOV and their activity are extracted (c). Neurons can appear as spatially overlapping due to limited axial resolution, and their activity needs to be demixed. The activity of each neuron (spikes - gray stars) can be estimated from its corresponding fluorescence trace. (d) Registration of components produced by imaging sessions with the same FOV over the course of different days (*left, middle*). Neurons that are active in all or only some of the imaged sessions are identified (*right*). (The different steps of the pipeline are displayed on mouse *in vivo* cortex data, courtesy of S.A. Koay and D. Tank (Princeton University). Results were obtained using the `CaImAn` package (Giovannucci et al., 2018).)

1 Motion correction

Motion artifacts in calcium imaging datasets can arise from natural brain movement. For a small FOV this motion can be approximated as rigid, and can usually be corrected[1] using standard template based registration methods (Thevenaz et al., 1998). However, in multi-photon imaging data, brain motion can be faster than the raster scanning imaging rate, resulting in non uniform motion artifacts within a data frame; non-rigid registration methods have been developed to handle these effects (Dombeck et al., 2007; Greenberg and Kerr, 2009). For larger FOVs, the rigid motion approximation is often insufficient; to handle this issue, we can split the FOV into smaller spatial patches, compute motion corrections within each patch, and then combine the results over patches (Pnevmatikakis and Giovannucci, 2017). This local "patchwise" processing approach helps parallelize computation and enables scalability to very large datasets, and will be a recurring theme in this chapter.

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 (Christensen et al., 2015; Venkatachalam et al., 2016; Nguyen et al., 2017), though demixing of fast cytosolic (non-nuclear-localized) signals in small flexible animals remains an unsolved problem. We expect non-rigid registration approaches similar to those developed by (Pnevmatikakis and Giovannucci, 2017) to be helpful here.

2 Denoising and compression

To facilitate visualization and further processing at this stage it is useful to *denoise* the data — i.e., to separate the signal from the noise, and discard the noise — and *compress* the signal into a format that can be stored and processed more efficiently. Denoising is particularly useful for fast or low-intensity imaging methods for which photon count noise may be relatively high compared to the signal. Principal component analysis (PCA) is often applied here (Mukamel et al., 2009; Pachitariu et al., 2017). Recently (Buchanan et al., 2018) pointed out that the statistical model underlying PCA is not very well-matched to functional imaging data, leading to relatively slow computation and suboptimal separation of signal from noise. By incorporating a more appropriate statistical model (exploiting the fact that signals are local in space, and the dominant noise sources are temporally and spatially uncorrelated), (Buchanan et al., 2018) develop a patchwise penalized matrix decomposition approach that achieves significantly faster computation and improved denoising and compression ($\sim$2-4x increases in SNR and compression rates of 20-300x, with minimal visible loss of signal). Notably, these methods are effective across a wide variety of functional imaging data, including single-photon, multi-photon, calcium, voltage, wide-field, and single-cellular-resolution data, with no manual parameter adjustment required.

In some datasets it is necessary to include a detrending step here to remove photo-bleaching effects (Buchanan et al., 2018). In addition, some approaches (e.g., widefield calcium imaging) require another step to remove contamination from hemodynamic signals (Ma et al., 2016). Due to space constraints we will not review either of these issues in depth here.

3 Demixing

The next task is to demix activity from the multiple spatially-overlapping neurons visible in the FOV into separate components. (This problem is analogous to the "spike-sorting" problem from classical extracellular electrophysiology.) A natural approach is to model the observed movie data

[1] In general motion within the imaging plane is easier to correct; motion out of the imaging plane can cause artifacts in which cells pop in and out of the plane and therefore appear to turn on and off, respectively. Extending the depth of field can ameliorate this issue to some degree.

as

$$Y(x,t) = \sum_{i=1}^{K} a_i(x)c_i(t) + B(x,t) + \varepsilon_{x,t}. \quad (1)$$

Here $Y(x,t)$ denotes the observed fluorescence at location x and time t, and $\boldsymbol{a_i}$ and $\boldsymbol{c_i}$ denote the spatial footprint and fluorescence trace, respectively, of the i-th neural component, with K denoting the number of neurons visible in the FOV[2]. $B(x,t)$ denotes the neuropil/background activity and $\varepsilon_{x,t}$ measurement noise, respectively; the background $B(x,t)$ represents the summed contributions from processes that can not be reliably separated into single-neuron components.

Equation 1 is a reasonable starting point but it is not yet a fully-specified statistical model in which all the components are identifiable (since we could trivially set $B(x,t)$ to equal $Y(x,t)$ and fully explain the observed data). To make further progress we need to introduce statistical assumptions or constraints on the model components, and then develop efficient, scalable algorithms for inferring the components from data. Different choices for these constraints have led to different algorithms based on equation 1. For example, the independent components analysis (ICA) approach in (Mukamel et al., 2009) does not constrain the spatial components $\boldsymbol{a_i}$ (and discards $B(x,t)$ from the model) while searching for maximally independent temporal components $\boldsymbol{c_i}$. The constrained non-negative matrix factorization (CNMF) approach developed in (Pnevmatikakis et al., 2016) imposes non-negativity and sparsity constraints on $\boldsymbol{a_i}$ and $\boldsymbol{c_i}$ and models $B(x,t)$ as a low-rank matrix within small spatial patches; by incorporating a more appropriate statistical model, CNMF achieves better performance than the less-structured ICA approach. (Pachitariu et al., 2017) use a similar constrained NMF approach, but with a slightly different objective function and background model. (Giovannucci et al., 2017) developed a real-time CNMF implementation that processes incoming data online, one imaging frame at a time, enabling closed-loop experiments. (Zhou et al., 2018) introduce a more flexible model for $B(x,t)$ to handle data from one-photon imaging approaches, where background contributions from out-of-focus light are much more severe than in multi-photon data; if these background effects are not handled correctly, strong spurious correlations between neighboring neurons can corrupt downstream analyses. (Buchanan et al., 2018) introduce a new more robust method for initializing the CNMF model and demonstrate significantly improved performance on spatially-extended dendritic signals; this paper also shows that the CNMF approach can be extended to handle voltage imaging data. Finally[3], (Giovannucci et al., 2018) incorporate a convolutional artificial neural network to help initialize new neural components in online CNMF, building on earlier work from (Apthorpe et al., 2016; Klibisz et al., 2017). The approaches developed in (Zhou et al., 2018; Giovannucci et al., 2018; Buchanan et al., 2018) all incorporate patchwise processing to scale to very large movies.

3.1 Benchmarking

One major open issue 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, 2017). Some gold standard datasets have been developed for the *segmentation* problem of finding nuclei or somas in calcium imaging data, either via manual annotation or the segmentation of datasets in which neurons co-express a static structural indicator; see e.g. `http://neurofinder.codeneuro.org/` for some example datasets. However, these segmented datasets can be unreliable:

[2]It is useful to distinguish this demixing problem from simpler *segmentation* problems that are the subject of large computer vision and biological image processing literatures: in a segmentation problem we would want to assign each pixel to at most a single neuron, so for each location x, at most one $a_i(x)$ would be allowed to be nonzero. Instead, in equation 1 we allow multiple neurons i to contribute to a given pixel x, since even in the multi-photon imaging setting a single diffraction-limited excitation spot will often excite fluorophores from multiple neurons. Correctly assigning fluorescence signals to each neuron (i.e., solving the full demixing problem) is particularly critical for any downstream analyses of the correlations between neighboring neurons.

[3]A number of other papers have developed approaches based on equation 1, including (but not limited to) (Maruyama et al., 2014; Andilla and Hamprecht, 2014; Haeffele and Vidal, 2017; Inan et al., 2017; Takekawa et al., 2017; Nöbauer et al., 2017; Petersen et al., 2017), but we lack the space to review all of this work here.

expression of structural indicators does not discriminate between active and non-active neurons, and is not guaranteed to be constrained only to neurons where the functional indicator is expressed. Moreover, as shown in (Giovannucci et al., 2018), individual manual annotations can be highly variable, with different labelers disagreeing up to a 20% level on the same dataset. Finally, these segmented datasets only partially indicate some of the (somatic) pixels within the spatial components $\boldsymbol{a_i}$, and do not provide ground truth data for the primary objects of interest in the *demixing* problem (i.e, the full demixed spatial and temporal components $\boldsymbol{a_i}$ and $\boldsymbol{c_i}$). Ground truth for the temporal components would be important to assess the robustness of demixing methods to neuropil contamination or contamination from small spatially overlapping neurites (Gauthier et al., 2017).

Thus the curation of fully spatiotemporal gold standard demixing datasets remains a critical challenge; the IARPA MICRONS project (`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. Meanwhile, realistic generative models (Song et al., 2017a) can also generate useful simulated ground truth data.

3.2 Towards optimal computational imaging and extensions beyond calcium imaging

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. (Pnevmatikakis and Paninski, 2013; Yang et al., 2016; Prevedel et al., 2016; Friedrich et al., 2017a; Song et al., 2017b; Lu et al., 2017; Kazemipour et al., 2018) 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 and/or with higher temporal resolution 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 (Pnevmatikakis and Paninski, 2013; Yang et al., 2016; Friedrich et al., 2017a; Kazemipour et al., 2018). All of these approaches can be cast in the same mathematical framework: instead of observing Y in equation 1 directly, we observe WY instead, where W is some linear operator that depends on the details of the imaging technique. A number of generalized demixing approaches have been developed to handle these data types; we expect to see continued algorithmic development in this direction in the near future. Simulators such as those developed in (Song et al., 2017a) will again likely play a useful role in the ongoing joint optimization of demixing methods and hardware design, as we push the limits of labeling density, imaging speed, SNR, FOV size (and the number of neurons observed simultaneously) and other critical imaging system parameters.

While this chapter focuses on calcium imaging, many similar themes will hold for voltage imaging at single-cell resolution (Buchanan et al., 2018), 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 methods also provide 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 spatiotemporal data; earlier works offer algorithmic starting points for modeling voltage data with subcellular resolution (Huys et al., 2006; Huys and Paninski, 2009; Paninski, 2010; Pakman et al., 2014).

4 Deconvolution

The demixing methods discussed in the previous section output temporal components $\boldsymbol{c_i}$ for each neuron in the field of view; $c_i(t)$ is proportional to the average fluorescence in neuron i at time t. *Deconvolution* methods aim to estimate the activity of each neuron i given the extracted fluorescence trace $\boldsymbol{c_i}$. This problem can be challenging due to the unknown and often non-linear nature of the indicator dynamics, the presence of measurement noise, and the relatively low imaging rate of typical calcium recordings. A popular approach assumes linear and time invariant dynamics, expressing the fluorescence trace $\boldsymbol{c}$ as the sum of the calcium transients due to neural activity plus measurement noise:

$$c(t) = \sum_{j=1}^{N} s_j h(t - t_j) + b + \varepsilon_t. \tag{2}$$

Here t_j, s_j denote the time and amplitude of the j-th transient, $\boldsymbol{h}$ is a causal function denoting the shape of the transient, and b, ε_t denote the (possibly time varying) baseline and measurement noise at time t. Transients are typically characterized by fast rise followed by a slower decay. This behavior can be modeled by a simple single or double exponential model. Under these assumptions the deconvolution problem can be cast as a convex optimization problem that can be solved efficiently; see (Vogelstein et al., 2010) for early work in this direction, and (Pnevmatikakis et al., 2016; Tubiana et al., 2017; Jewell and Witten, 2017) for more recent advances. Fast online formulations are also available (Friedrich et al., 2017b; Jewell et al., 2018), and methods based on more detailed biophysical models including nonlinearities have also been developed (Vogelstein et al., 2009; Deneux et al., 2016; Greenberg et al., 2018). Finally, these models can be used to denoise the trace $\boldsymbol{c_i}$; this denoising step can be incorporated into the demixing step to improve performance (Pnevmatikakis et al., 2016).

Supervised learning methods have also been applied to the spike inference problem (Theis et al., 2016). Unlike the unsupervised methods reviewed above, supervised methods rely on training data, usually in the form of dual electrophysiological and imaging recordings in small neuron populations, which can be hard to obtain. If a good generative model of the data is available then the model can be used to generate an unlimited amount of training data (Berens et al., 2018). The recent community benchmark effort (Berens et al., 2018) found that current supervised and unsupervised learning algorithms perform similarly on the labeled datasets available at `http://spikefinder.codeneuro.org/`. On the other hand, (**?**) 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.

The accuracy and temporal resolution of the recovered activity depends on several factors, most critically the imaging frame rate, SNR, and indicator dynamics (particularly the rise time and degree of nonlinearity of the indicator). It is not possible to recover spike times at millisecond resolution given standard 30 Hz frame rates without side information. However, Bayesian methods that incorporate prior information about spike timing (using information from the activity of other cells, or from external covariates such as stimulus or movement timing) can improve inference accuracy and temporal resolution beyond the original frame rate (Vogelstein et al., 2009; Deneux et al., 2016; Picardo et al., 2016; Aitchison et al., 2017).

Voltage imaging data will present some new challenges. In this context we do not just care about recovering spike times; instead we wish to denoise and recover subthreshold voltage fluctuations as well. Thus models that combine both sparse spiking effects with constraints on the smoothness of the subthreshold voltage will likely be critical. In addition there are important tradeoffs to optimize between imaging framerate, SNR, and the speed and brightness of the indicator: for example, slower indicators may be brighter, and if deconvolution methods can improve the resulting temporal resolution then it might be preferable to use a slower indicator. As in the demixing problem discussed in the previous section, we expect that an integrated computational imaging approach (where we optimize jointly over the indicator, imaging approach, and computational deconvolution method) will lead to improved recovery of voltage signals.

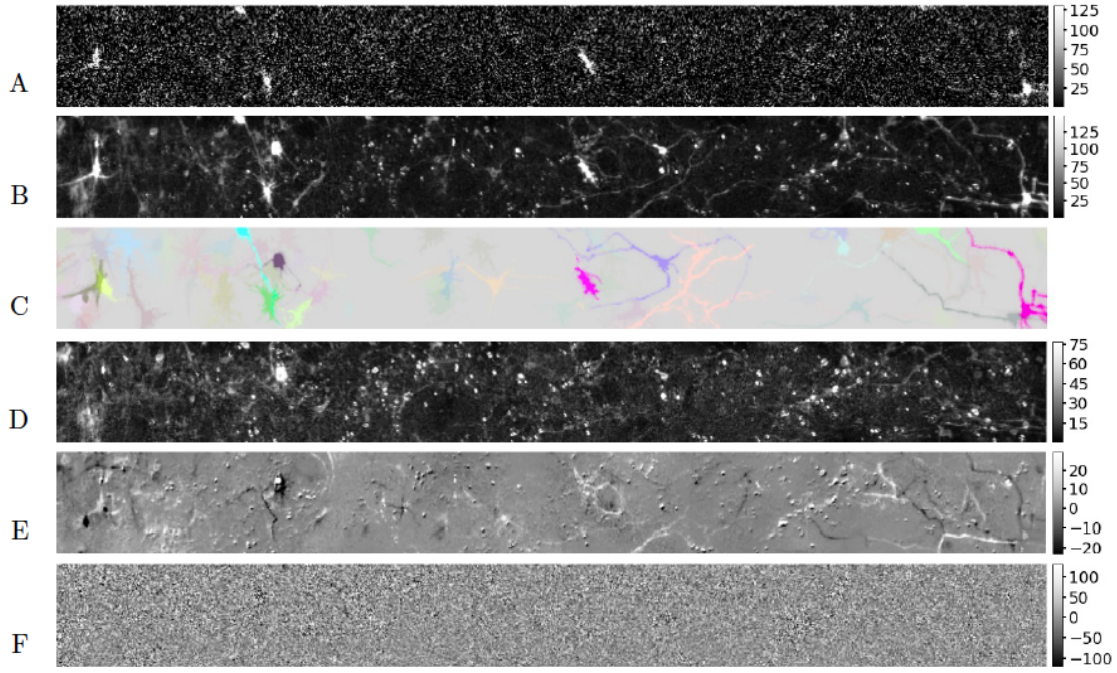

Figure 2: An example frame illustrating demixing on voltage imaging data. (From (Buchanan et al., 2018); see that paper for link to full diagnostic video.) (A) Detrended data $\mathbf{Y}$. (B) Denoised data $\hat{\mathbf{Y}}$. (C) Extracted signals $\mathbf{AC}$; each component i is assigned a unique color, and the intensity of each pixel at each time is determined by the corresponding value of $\mathbf{AC}$. (D) Estimated background $\mathbf{B}$ (constrained to be temporally constant in this case). (E) Residual $\hat{\mathbf{Y}} - \mathbf{AC} - \mathbf{B}$. Note the small scale of the residual compared to the original signal. (F) Noise removed in the denoising step.

5 Post-processing and visualization

A number of post-processing steps have been implemented to check for missing or clearly non-neuronal components $\boldsymbol{a_i}$ or $\boldsymbol{c_i}$ in the output of the demixing step; for example, in some cases it is useful to apply post hoc image processing methods to remove artifacts from spatial components. These approaches can also be incorporated within the demixing loop to improve overall demixing performance.

As in any complex statistical analysis, it is important to visualize the outputs of the pipelines described above. One basic but critical visualization is to simply view the raw data movie Y synchronized with and adjacent to corresponding movies of the components the pipeline has extracted from Y, along with the residual $Y - AC - B$. This visualization makes it easy to quickly identify poor motion correction, potential missing components, poor separation of background versus single-neuronal components, and other common artifacts. Figure 2 provides an illustration.

Another useful visualization involves sorting the components $\boldsymbol{a_i}$ and $\boldsymbol{c_i}$ by brightness and then viewing each component individually. Typically the brightest output components will be of high quality while the dimmest extracted components may be overly noisy or corrupted by artifacts. Quick visual inspection of the sorted components can determine a good value of the number of components to be retained (K in equation 1). Simple graphical user interfaces have been developed to aid in this procedure, but further effort in this direction would be useful.

In general, any software used for manual intervention and/or post-processing should include an automatic logger to ensure full reproducibility of the analysis.

6 Registration across multiple sessions

Calcium imaging enables the monitoring of large neural populations over many different sessions across multiple days. Several packages offer semi-automated methods for registering neurons across multiple sessions (Kaifosh et al., 2014; Pachitariu et al., 2017; Sheintuch et al., 2017; Giovannucci et al., 2018). As in the demixing problem, ground-truth data for validating these methods is difficult to obtain; thus the results of multi-session alignment from challenging datasets (particularly one-photon imaging datasets with limited SNR or very large background signals) should be interpreted with some caution (c.f. (Katlowitz et al., 2018)).

Software implementations

A critical requirement for the adoption of formal reproducible methods is the existence of reliable, well-documented software that is scalable to the size of modern datasets. Available packages for automated and/or interactive analysis include `SIMA` (Kaifosh et al., 2014, Python), `Suite2p` (Pachitariu et al., 2017, (MATLAB)), `ABLE` (Reynolds et al., 2017, (MATLAB)), `SCALPEL` (Petersen et al., 2017, (R)), `SamuROI` (Rueckl et al., 2017, (Python)), the toolbox of (Romano et al., 2017, (MATLAB)), `CaImAn` (Giovannucci et al., 2018, (Python and MATLAB)), `MIN1PIPE` (Lu et al., 2018, (MATLAB)), and `CNMF-E`[4] (Zhou et al., 2018, (MATLAB)). Of these, at least `CaImAn`, `CNMF-E`, and `Suite2p` have attracted a critical mass of users and a community of developers that are continuing to support and improve the software.

Open issues / future work

The methods described above are the first step in the analysis of datasets that are being acquired daily in hundreds of neuroscience labs; thus these pipelines represent a foundation upon which our understanding of the nervous system is currently being built. Therefore it is absolutely critical that this foundation be as sturdy as possible; while the state of the art in this field has progressed rapidly in the last several years, there is still significant work to be done.

We have already mentioned several directions for future work, including: better graphical user interfaces for visualization and analysis; methods for motion correction that can demix non-nuclear-localized signals in small flexible moving animals; better gold standard demixing data; improved generalized demixing methods to handle experiments in which we observe some linear projection WY of the data Y in equation 1; and development of methods optimized to process voltage imaging and neurotransmitter release imaging data.

A number of additional directions remain open. For example, we expect to see further development of on-line, real-time analysis approaches in the context of closed-loop experiments, building on the methods introduced in (Giovannucci et al., 2017); the extension of these methods to handle single-photon data with large background signals (as is typical in micro-endoscopic imaging) is an important next step. Scalable Bayesian methods for quantifying the reliability of each component output by the demixing pipeline would also be very valuable. There are also a number of open analysis challenges regarding multimodal data, i.e., functional imaging data collected in conjunction with e.g. electrophysiological data, spatial transcriptomics measurements, electron microscopy (or other very large-scale anatomical imaging approaches), or spatiotemporal patterned optogenetic stimulation.

More broadly, we would advocate continued efforts towards full analysis standardization and automation, to enable widespread, routine data sharing and reproducibility. Large collaborative experimental efforts such as the International Brain Laboratory (International Brain Lab, 2017) depend on these efforts, and we expect to see a number of large-scale projects with similar requirements in the near future.

[4]Most of the packages listed here focus on two-photon data; `CNMF-E` and `MIN1PIPE` can handle multi-photon data but is designed specifically for one-photon data.

Acknowledgements

We thank S.A. Koay and D. Tank for the mouse *in vivo* dataset, and E. Buchanan, J. Friedrich, A. Giovannucci, I. Kinsella, D. Peterka, J. Vogelstein, D. Zhou, and P.C. Zhou for useful discussions. E.P. was internally funded by the Flatiron Institute and did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. L.P. was funded by Army Research Office W911NF-12-1-0594 (MURI), the Simons Foundation Collaboration on the Global Brain, National Institutes of Health R01EB22913, R21EY027592, 1U01NS103489-01, and U19NS104649-01, and by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DoI/IBC) contract number D16PC00003. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

References

Aitchison, L., Russell, L., Packer, A. M., Yan, J., Castonguay, P., Hausser, M., and Turaga, S. C. (2017). Model-based Bayesian inference of neural activity and connectivity from all-optical interrogation of a neural circuit. In *Advances in Neural Information Processing Systems 30*, pages 3486–3495.

Andilla, F. D. and Hamprecht, F. A. (2014). Sparse space-time deconvolution for calcium image analysis. In *Advances in Neural Information Processing Systems*, pages 64–72.

Apthorpe, N., Riordan, A., Aguilar, R., Homann, J., Gu, Y., Tank, D., and Seung, H. S. (2016). Automatic neuron detection in calcium imaging data using convolutional networks. In *Advances in Neural Information Processing Systems*, pages 3270–3278.

Berens, P., Freeman, J., Deneux, T., Chenkov, N., McColgan, T., Speiser, A., Macke, J. H., Turaga, S. C., Mineault, P., Rupprecht, P., et al. (2018). Community-based benchmarking improves spike rate inference from two-photon calcium imaging data. *PLoS computational biology*, 14(5):e1006157.

Bouchard, M. B., Voleti, V., Mendes, C. S., Lacefield, C., Grueber, W. B., Mann, R. S., Bruno, R. M., and Hillman, E. M. (2015). Swept confocally-aligned planar excitation (scape) microscopy for high-speed volumetric imaging of behaving organisms. *Nature photonics*, 9(2):113–119.

Buchanan, E. K., Kinsella, I., Zhou, D., Zhu, R., Zhou, P., Gerhard, F., Ferrante, J., Ma, Y., Kim, S., Shaik, M., Liang, Y., Lu, R., Reimer, J., Fahey, P., Muhammad, T., Dempsey, G., Hillman, E., Ji, N., Tolias, A., and Paninski, L. (2018). Penalized matrix decomposition for denoising, compression, and improved demixing of functional imaging data. *bioRxiv*, page 334706.

Christensen, R. P., Bokinsky, A., Santella, A., Wu, Y., Marquina-Solis, J., Guo, M., Kovacevic, I., Kumar, A., Winter, P. W., Tashakkori, N., et al. (2015). Untwisting the caenorhabditis elegans embryo. *Elife*, 4:e10070.

Cong, L., Wang, Z., Chai, Y., Hang, W., Shang, C., Yang, W., Bai, L., Du, J., Wang, K., and Wen, Q. (2017). Rapid whole brain imaging of neural activity in freely behaving larval zebrafish (Danio rerio). *eLife*, 6.

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

Dombeck, D. A., Khabbaz, A. N., Collman, F., Adelman, T. L., and Tank, D. W. (2007). Imaging large-scale neural activity with cellular resolution in awake, mobile mice. *Neuron*, 56(1):43–57.

Donoho, D. (2017). 50 years of data science. *Journal of Computational and Graphical Statistics*, 26(4):745–766.

Dupre, C. and Yuste, R. (2017). Non-overlapping neural networks in Hydra vulgaris. *Current biology*, 27.

Friedrich, J., Yang, W., Soudry, D., Mu, Y., Ahrens, M. B., Yuste, R., Peterka, D. S., and Paninski, L. (2017a). Multi-scale approaches for high-speed imaging and analysis of large neural populations. *PLoS computational biology*, 13(8):e1005685.

Friedrich, J., Zhou, P., and Paninski, L. (2017b). Fast online deconvolution of calcium imaging data. *PLoS computational biology*, 13(3):e1005423.

Gauthier, J., Charles, A., Pillow, J., and Tank, D. (2017). Robust estimation of calcium transients by modeling contamination. *Cosyne abstracts*.

Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Koay, S. A., Taxidis, J., Najafi, F., Gauthier, J. L., Zhou, P., Tank, D. W., Chklovskii, D. B., and Pnevmatikakis, E. A. (2018). CaImAn: An open source tool for scalable calcium imaging data analysis. *bioRxiv*, page 339564.

Giovannucci, A., Friedrich, J., Kaufman, J., Churchland, A., Chklovskii, D., Paninski, L., and Pnevmatikakis, E. A. (2017). OnACID: Online analysis of calcium imaging data in real time. In *Neural Information Processing Systems (NIPS)*.

Greenberg, D., Wallace, D., Voit, K.-M., Wuertenberger, S., Czubayko, U., Monsees, A., Vogelstein, J., Seifert, R., Groemping, Y., and Kerr, J. (2018). Spike inference for genetically encoded calcium indicators with models of multistep binding kinetics. *Cosyne abstracts*.

Greenberg, D. S. and Kerr, J. (2009). Automated correction of fast motion artifacts for two-photon imaging of awake animals. *Journal of neuroscience methods*, 176(1):1–15.

Haeffele, B. D. and Vidal, R. (2017). Structured low-rank matrix factorization: Global optimality, algorithms, and applications. *arXiv preprint arXiv:1708.07850*.

Huys, Q., Ahrens, M., and Paninski, L. (2006). Efficient estimation of detailed single-neuron models. *Journal of Neurophysiology*, 96:872–890.

Huys, Q. and Paninski, L. (2009). Model-based smoothing of, and parameter estimation from, noisy biophysical recordings. *PLoS computational biology*, 5:e1000379.

Inan, H., Erdogdu, M. A., and Schnitzer, M. (2017). Robust estimation of neural signals in calcium imaging. In *Advances in Neural Information Processing Systems*, pages 2905–2914.

International Brain Lab (2017). An international laboratory for systems and computational neuroscience. *Neuron*, 96(6):1213–1218.

Jewell, S., Hocking, T. D., Fearnhead, P., and Witten, D. (2018). Fast nonconvex deconvolution of calcium imaging data. *arXiv preprint arXiv:1802.07380*.

Jewell, S. and Witten, D. (2017). Exact spike train inference via ℓ_0 optimization. *arXiv preprint arXiv:1703.08644*.

Kaifosh, P., Zaremba, J. D., Danielson, N. B., and Losonczy, A. (2014). Sima: Python software for analysis of dynamic fluorescence imaging data. *Frontiers in neuroinformatics*, 8:80.

Katlowitz, K. A., Picardo, M. A., and Long, M. A. (2018). Stable sequential activity underlying the maintenance of a precisely executed skilled behavior. *Neuron*.

Kazemipour, A., Novak, O., Flickinger, D., Marvin, J. S., King, J., Borden, P., Druckmann, S., Svoboda, K., Looger, L. L., and Podgorski, K. (2018). Kilohertz frame-rate two-photon tomography. *bioRxiv*.

Klibisz, A., Rose, D., Eicholtz, M., Blundon, J., and Zakharenko, S. (2017). Fast, simple calcium imaging segmentation with fully convolutional networks. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pages 285–293. Springer.

Lu, J., Li, C., Singh-Alvarado, J., Zhou, Z. C., Fröhlich, F., Mooney, R., and Wang, F. (2018). Min1pipe: A miniscope 1-photon-based calcium imaging signal extraction pipeline. *Cell Reports*, 23(12):3673–3684.

Lu, R., Sun, W., Liang, Y., Kerlin, A., Bierfeld, J., Seelig, J. D., Wilson, D. E., Scholl, B., Mohar, B., Tanimoto, M., et al. (2017). Video-rate volumetric functional imaging of the brain at synaptic resolution. *Nature neuroscience*, 20(4):620.

Ma, Y., Shaik, M. A., Kim, S. H., Kozberg, M. G., Thibodeaux, D. N., Zhao, H. T., Yu, H., and Hillman, E. M. C. (2016). Wide-field optical mapping of neural activity and brain haemodynamics: considerations and novel approaches. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 371(1705).

Maruyama, R., Maeda, K., Moroda, H., Kato, I., Inoue, M., Miyakawa, H., and Aonishi, T. (2014). Detecting cells using non-negative matrix factorization on calcium imaging data. *Neural Networks*, 55:11–19.

Mukamel, E. A., Nimmerjahn, A., and Schnitzer, M. J. (2009). Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron*, 63(6):747–760.

Nguyen, J. P., Linder, A. N., Plummer, G. S., Shaevitz, J. W., and Leifer, A. M. (2017). Automatically tracking neurons in a moving and deforming brain. *PLoS computational biology*, 13(5):e1005517.

Nöbauer, T., Skocek, O., Pernía-Andrade, A. J., Weilguny, L., Traub, F. M., Molodtsov, M. I., and Vaziri, A. (2017). Video rate volumetric ca 2+ imaging across cortex using seeded iterative demixing (sid) microscopy. *Nature methods*, 14(8):811.

Pachitariu, M., Stringer, C., Dipoppa, M., Schröder, S., Rossi, L. F., Dalgleish, H., Carandini, M., and Harris, K. D. (2017). Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *BioRxiv*, page 061507.

Pakman, A., Huggins, J., Smith, C., and Paninski, L. (2014). Fast state-space methods for inferring dendritic synaptic connectivity. *Journal of computational neuroscience*, 36(3):415–443.

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

Paninski, L. and Cunningham, J. (2018). Neural data science: accelerating the experiment-analysis-theory cycle in large-scale neuroscience. *Current Opinion in Neurobiology*, 50:232 – 241. Neurotechnologies.

Petersen, A., Simon, N., and Witten, D. (2017). Scalpel: Extracting neurons from calcium imaging data. *arXiv preprint arXiv:1703.06946*.

Picardo, M. A., Merel, J., Katlowitz, K. A., Vallentin, D., Okobi, D. E., Benezra, S. E., Clary, R. C., Pnevmatikakis, E. A., Paninski, L., and Long, M. A. (2016). Population-level representation of a temporal sequence underlying song production in the zebra finch. *Neuron*, 90(4):866–876.

Pnevmatikakis, E. A. and Giovannucci, A. (2017). Normcorre: An online algorithm for piecewise rigid motion correction of calcium imaging data. *Journal of Neuroscience Methods*, 291:83–94.

Pnevmatikakis, E. A. and Paninski, L. (2013). Sparse nonnegative deconvolution for compressive calcium imaging: algorithms and phase transitions. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems 26*, pages 1250–1258. Curran Associates, Inc.

Pnevmatikakis, E. A., Soudry, D., Gao, Y., Machado, T. A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., Ahrens, M., Bruno, R., Jessell, T. M., Peterka, D. S., Yuste, R., and Paninski, L. (2016). Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. *Neuron*, 89(2):285–299.

Prevedel, R., Verhoef, A. J., Pernia-Andrade, A. J., Weisenburger, S., Huang, B. S., Nöbauer, T., Fernández, A., Delcour, J. E., Golshani, P., Baltuska, A., et al. (2016). Fast volumetric calcium imaging across multiple cortical layers using sculpted light. *Nature methods*, 13(12):1021.

Reynolds, S., Abrahamsson, T., Schuck, R., Sjöström, P. J., Schultz, S. R., and Dragotti, P. L. (2017). Able: An activity-based level set segmentation algorithm for two-photon calcium imaging data. *eNeuro*, 4(5):ENEURO–0012.

Romano, S. A., Pérez-Schuster, V., Jouary, A., Boulanger-Weill, J., Candeo, A., Pietri, T., and Sumbre, G. (2017). An integrated calcium imaging processing toolbox for the analysis of neuronal population dynamics. *PLoS computational biology*, 13(6):e1005526.

Rueckl, M., Lenzi, S. C., Moreno-Velasquez, L., Parthier, D., Schmitz, D., Ruediger, S., and Johenning, F. W. (2017). Samuroi, a python-based software tool for visualization and analysis of dynamic time series imaging at multiple spatial scales. *Frontiers in neuroinformatics*, 11:44.

Sheintuch, L., Rubin, A., Brande-Eilat, N., Geva, N., Sadeh, N., Pinchasof, O., and Ziv, Y. (2017). Tracking the same neurons across multiple days in ca2+ imaging data. *Cell reports*, 21(4):1102–1115.

Song, A., Charles, A., Gauthier, J., Koay, S. A., Tank, D., and Pillow, J. (2017a). Two-photon microscopy simulation for optics optimization and benchmarking. *Cosyne abstracts*.

Song, A., Charles, A. S., Koay, S. A., Gauthier, J. L., Thiberge, S. Y., Pillow, J. W., and Tank, D. W. (2017b). Volumetric two-photon imaging of neurons using stereoscopy (vtwins). *Nature methods*, 14(4):420.

Takekawa, T., Asai, H., Ohkawa, N., Nomoto, M., Okubo-Suzuki, R., Ghandour, K., Sato, M., Hayashi, Y., Inokuchi, K., and Fukai, T. (2017). Automatic sorting system for large calcium imaging data. *bioRxiv*, page 215145.

Theis, L., Berens, P., Froudarakis, E., Reimer, J., Rosón, M. R., Baden, T., Euler, T., Tolias, A. S., and Bethge, M. (2016). Benchmarking spike rate inference in population calcium imaging. *Neuron*, 90(3):471–482.

Thevenaz, P., Ruttimann, U. E., and Unser, M. (1998). A pyramid approach to subpixel registration based on intensity. *IEEE transactions on image processing*, 7(1):27–41.

Tubiana, J., Wolf, S., and Debregeas, G. (2017). Blind sparse deconvolution for inferring spike trains from fluorescence recordings. *bioRxiv*, page 156364.

Venkatachalam, V., Ji, N., Wang, X., Clark, C., Mitchell, J. K., Klein, M., Tabone, C. J., Florman, J., Ji, H., Greenwood, J., Chisholm, A. D., Srinivasan, J., Alkema, M., Zhen, M., and Samuel, A. D. T. (2016). Pan-neuronal imaging in roaming caenorhabditis elegans. *Proceedings of the National Academy of Sciences*, 113(8):E1082–E1088.

Vogelstein, J., Watson, B., Packer, A., Jedynak, B., Yuste, R., and Paninski, L. (2009). Model-based optimal inference of spike times and calcium dynamics given noisy and intermittent calcium-fluorescence imaging. *Biophysical Journal*, 97:636–655.

Vogelstein, J. T., Packer, A. M., Machado, T. A., Sippy, T., Babadi, B., Yuste, R., and Paninski, L. (2010). Fast nonnegative deconvolution for spike train inference from population calcium imaging. *Journal of neurophysiology*, 104(6):3691–3704.

Xu, Y., Zou, P., and Cohen, A. E. (2017). Voltage imaging with genetically encoded indicators. *Current opinion in chemical biology*, 39:1–10.

Yang, W., Miller, J.-e., Carrillo-Reid, L., Pnevmatikakis, E., Paninski, L., Yuste, R., and Peterka, D. S. (2016). Simultaneous multi-plane imaging of neural circuits. *Neuron*, 89(2):269–284.

Zhou, P., Resendez, S. L., Rodriguez-Romaguera, J., Jimenez, J. C., Neufeld, S. Q., Giovannucci, A., Friedrich, J., Pnevmatikakis, E. A., Stuber, G. D., Hen, R., Kheirbek, M. A., Sabatini, B. L., Kass, R. E., and Paninski, L. (2018). Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. *eLife*, 7:e28728.