bioRxiv preprint doi: https://doi.org/10.1101/339564; this version posted June 5, 2018. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aManaYstoptisubiniteb

CalmAn: An open source tool for scalable Calcium Imaging data **Analysis**

- Andrea Giovannucci^{1*}, Johannes Friedrich^{1,2}, Pat Gunn¹, Jérémie Kalfon^{3†}, Sue
- Ann Koay⁴, Jiannis Taxidis⁵, Farzaneh Najafi⁶, Jeffrey L. Gauthier⁴, Pengcheng
- Zhou², David W. Tank⁴, Dmitri Chklovskii¹, Eftychios A. Pnevmatikakis^{1*} 6

*For correspondence:

agiovannucci@flatironinstitute.org 7 (AG); epnevmatikakis@ flatironinstitute.org (EAP)

[†]IK contributed to this work during₁₀ an internship at the Flatiron Institute 11

¹Center for Computational Biology, Flatiron Institute, New York, USA; ²Department of Statistics and Center for Theoretical Neuroscience, Columbia University, New York, USA: ³ECE Paris, Paris, France; ⁴Princeton Neuroscience Institute, Princeton University, New Jersey, USA; ⁵Department of Neurology, UCLA, California, USA; ⁶Cold Spring Harbor Laboratory, New York, USA

- Abstract Advances in fluorescence microscopy enable monitoring larger brain areas in-vivo with 13
- finer time resolution. The resulting data rates require reproducible analysis pipelines that are 14
- reliable, fully automated, and scalable to datasets generated over the course of months. Here we 15
- present CAIMAN, an open-source library for calcium imaging data analysis. CAIMAN provides 16
- automatic and scalable methods to address problems common to pre-processing, including motion 17
- correction, neural activity identification, and registration across different sessions of data collection. 18
- It does this while requiring minimal user intervention, with good performance on computers 19
- ranging from laptops to high-performance computing clusters. CAIMAN is suitable for two-photon 20
- and one-photon imaging, and also enables real-time analysis on streaming data. To benchmark the 21
- performance of CAIMAN we collected a corpus of ground truth annotations from multiple labelers 22
- on nine mouse two-photon datasets. We demonstrate that CAIMAN achieves near-human 23
- performance in detecting locations of active neurons. 24
- 25

8

12

Introduction 26

Understanding the function of neural circuits is contingent on the ability to accurately record and 27 modulate the activity of large neural populations. Optical methods based on the fluorescence 28 activity of genetically encoded calcium binding indicators (Chen et al., 2013) have become a standard 29 tool for this task, due to their ability to monitor *in vivo* targeted neural populations from many 30 different brain areas over extended periods of time (weeks or months). Advances in microscopy 31 techniques facilitate imaging larger brain areas with finer time resolution, producing an ever-32 increasing amount of data. A typical resonant scanning two-photon microscope produces data at a 33 rate greater than 50GB/Hour¹, a number that can be significantly higher (up to more than 1TB/Hour) 34 with other custom recording technologies (Sofroniew et al. (2016); Ahrens et al. (2013); Flusberg 35 et al. (2008); Cai et al. (2016); Prevedel et al. (2014); Grosenick et al. (2017); Bouchard et al. (2015)). 36 This increasing availability and volume of calcium imaging data calls for automated analysis 37 methods and reproducible pipelines to extract the relevant information from the recorded movies, 38 i.e., the locations of neurons in the imaged Field of View (FOV) and their activity in terms of raw 39

¹Calculation performed on a 512×512 FOV imaged at 30Hz producing an unsigned 16-bit integer for each measurement.

- ⁴⁰ fluorescence and/or neural activity (spikes). The typical steps arising in the processing pipelines are
- the following (Fig. 1a): i) Motion correction, where the FOV at each data frame (image or volume)
- ⁴² is registered against a template to correct for motion artifacts due to the finite scanning rate and
- 43 existing brain motion, ii) source extraction where the different active and possibly overlapping
- 44 sources are extracted and their signals are demixed from each other and from the background
- 45 neuropil signals (Fig. 1b), and iii) activity deconvolution, where the neural activity of each identified
- ⁴⁶ source is deconvolved from the dynamics of the calcium indicator.

47 Related work

48 Source extraction

- ⁴⁹ Some source extraction methods attempt the detection of neurons in static images using supervised
- ⁵⁰ or unsupervised learning methods. Examples of unsupervised methods on summary images include
- s1 graph-cut approaches applied to the correlation image (Kaifosh et al., 2014; Spaen et al., 2017),
- ⁵² and dictionary learning (*Pachitariu et al., 2013*). Supervised learning methods based on deep
- neural networks have also been applied to the problem of neuron detection (Apthorpe et al., 2016;
- 54 Klibisz et al., 2017). While these methods can be efficient in detecting the locations of neurons, they
- cannot infer the underlying activity nor do they readily offer ways to deal with the spatial overlap of
- ⁵⁶ different components.

To extract temporal traces together with the spatial footprints of the components one can use 57 methods that directly represent the full spatio-temporal data in a matrix factorization setup e.g. 58 independent component analysis (ICA) (Mukamel et al., 2009), constrained nonnegative matrix 59 factorization (CNMF) (Pneymatikakis et al., 2016) (and its adaptation to one-photon data (Zhou 60 et al., 2018)), clustering based approaches (Pachitariu et al., 2017), dictionary learning (Petersen 61 et al., 2017), or active contour models (Reynolds et al., 2017). Such spatio-temporal methods are 62 unsupervised, and focus on detecting active neurons by considering the spatio-temporal activity of 63 a component as a contiguous set of pixels within the FOV that are correlated in time. While such 64 methods tend to offer a direct decomposition of the data in a set of sources with activity traces 65 in an unsupervised way, in principle they require processing of the full dataset, and thus can be 66 rendered intractable very quickly. Possible approaches to deal with the data size include distributed 67 processing in High Performance Computing (HPC) clusters (Freeman et al., 2014), spatio-temporal 68 decimation (Friedrich et al., 2017a), and dimensionality reduction (Pachitariu et al., 2017). Recently, 69 Giovannucci et al. (2017) prototyped an online algorithm (ONACID), by adapting matrix factorization 70 setups (*Pnevmatikakis et al., 2016*: *Mairal et al., 2010*), to operate on calcium imaging streaming 71

⁷² data and thus natively deal with large data rates.

73 Deconvolution

- ⁷⁴ For the problem of predicting spikes from fluorescence traces, both supervised and unsupervised
- ⁷⁵ methods have been explored. Supervised methods rely on the use of ground truth data to train
- ⁷⁶ or fit biophysical or neural network models (*Theis et al., 2016; Speiser et al., 2017*). Unsupervised
- 77 methods can be either deterministic, such as sparse non-negative deconvolution (Vogelstein
- 78 et al., 2010; Pnevmatikakis et al., 2016) that give a single estimate of the deconvolved neural
- ⁷⁹ activity, or probabilistic, that aim to also characterize the uncertainty around these estimates
- e.g., (Pnevmatikakis et al., 2013; Deneux et al., 2016)). A recent community benchmarking effort

81 (Berens et al., 2017) characterizes the similarities and differences of various available methods.

82 CAIMAN

- ⁸³ Here we present CAIMAN, an open source suite for the analysis pipeline of both two-photon and one-
- ⁸⁴ photon calcium imaging data. CAIMAN includes frameworks for both offline analysis (CAIMAN BATCH)
- ⁸⁵ where all the data is processed at once at the end of experiment, and online analysis on streaming
- data (CAIMAN ONLINE). Moreover, CAIMAN requires very moderate computing infrastructure (e.g., a

- personal laptop or workstation), thus providing automated, efficient, and reproducible large-scale
- ⁸⁸ analysis on commodity hardware.

89 Contributions

⁹⁰ Our contributions can be roughly grouped in three different directions:

Methods: CAIMAN BATCH improves on the scalability of the source extraction problem by employing 91 a MapReduce framework for parallel processing and memory mapping which allows the 92 analysis of datasets larger than would fit in RAM on most computer systems. It also improves 93 on the qualitative performance by introducing automated routines for component evaluation 94 and classification, better handling of neuropil contamination, and better initialization methods. 95 While these benefits are here presented in the context of the widely used CNMF algorithm 96 of *Pnevmatikakis et al.* (2016), they are in principle applicable to any matrix factorization 97 approach. 98 CAIMAN ONLINE improves and extends the ONACID prototype algorithm (Giovannucci et al., 99 2017) by introducing, among other advances, new initialization methods and a convolutional 100 neural network (CNN) based approach for detecting new neurons on streaming data. Our 101 analysis on *in vivo* two-photon and light-sheet imaging datasets shows that CAIMAN ONLINE 102 approaches human-level performance and enables novel types of closed-loop experiments. 103 Apart from these significant algorithmic improvements CAIMAN includes several useful anal-104 vsis tools such as, a MapReduce and memory-mapping compatible implementation of the 105 CNMF-E algorithm for one-photon microendoscopic data (*Zhou et al., 2018*), a novel efficient 106 algorithm for registration of components across multiple days, and routines for segmentation 107

¹⁰⁸ of structural (static) channel information which can be used for component seeding.

Software: CAIMAN comes as a complete open source software suite implemented in Python, and 109 is already widely used by, and has received contributions from, the community. It contains 110 efficient implementations of the standard analysis pipeline steps (motion correction - source 111 extraction - deconvolution - registration across different sessions), as well as numerous other 112 features. Apart from Python, several of the tools presented here are also available in MATLAB®. 113 Data: We benchmark the performance of CAIMAN against a previously unreleased corpus of manu-114 ally annotated data. The corpus consists of 9 mouse *in vivo* two-photon datasets manually 115 annotated by 3-4 independent labelers that were instructed to select active neurons in a 116 principled and consistent way, and who subsequently combined their annotations to create 117 a "consensus" ground truth that is also used to quantify the limits of human performance. 118 The manual annotations are also released to the community providing a valuable tool for 119

¹²⁰ benchmarking and training purposes.

121 Paper organization

The paper is organized as follows: We first give a brief presentation of the analysis methods and 122 features provided by CAIMAN. In the Results section we benchmark CAIMAN ONLINE and CAIMAN 123 BATCH against a corpus of manually annotated data. We apply CAIMAN ONLINE to a zebrafish whole 124 brain lightsheet imaging recording, and demonstrate how such large datasets can be processed 125 efficiently in real time. We also present applications of CAIMAN BATCH to one-photon data, as well as 126 examples of component registration across multiple days. We conclude by discussing the utility of 127 our tools, the relationship between CAIMAN BATCH and CAIMAN ONLINE and outline future directions. 128 Detailed descriptions of the introduced methods are presented in *Methods and Materials*. 129

130 Methods

¹³¹ Before presenting the new analysis features introduced with this work, we overview the analysis

132 pipeline that CAIMAN uses and builds upon.

Overview of analysis pipeline

¹³⁴ The standard analysis pipeline for calcium imaging data used in CAIMAN is depicted in Fig. 1a. ¹³⁵ The data in movie format is first processed to remove motion artifacts. Subsequently the active

- components (neurons and background) are extracted as individual pairs of a spatial footprint that
- describes the shape of each component projected to the imaged FOV, and a temporal trace that
- captures its fluorescence activity (Fig. 1b-d). Finally, the neural activity of each fluorescence trace
- ¹³⁹ is deconvolved from the dynamics of the calcium indicator. These operations can be challenging
- ¹⁴⁰ because of limited axial resolution of 2-photon microscopy (or the much larger integration volume
- in one-photon imaging). This results in spatially overlapping fluorescence from different sources
- and neuropil activity. Before presenting the new features of CAIMAN in more detail, we briefly review
- ¹⁴³ how it incorporates existing tools in the pipeline.

144 Motion Correction

145 CAIMAN uses the NORMCORRE algorithm (Pnevmatikakis and Giovannucci, 2017) that corrects non-

- rigid motion artifacts by estimating motion vectors with subpixel resolution over a set of overlapping
- patches within the FOV. These estimates are used to infer a smooth motion field within the FOV
- ¹⁴⁸ for each frame. For two-photon imaging data this approach is directly applicable, whereas for
- one-photon micro-endoscopic data the motion is estimated on high pass spatially filtered data,
- a necessary operation to remove the smooth background signal and create enhanced spatial
- ¹⁵¹ landmarks. The inferred motion fields are then applied to the original data frames.

152 Source Extraction

Source extraction is performed using the constrained non-negative matrix factorization (CNMF) 153 framework of **Pneymatikakis et al.** (2016) which can extract components with spatial overlapping 154 projections (Fig. 1b). After motion correction the spatio-temporal activity of each source can be 155 expressed as a rank one matrix given by the outer product of two components; a component in 156 space that describes the spatial footprint (location and shape) of each source, and a component 157 in time that describes the activity trace of the source (Fig. 1c). The data can be described by the 158 sum of all the resulting rank one matrices together with an appropriate term for the background 159 and neuropil signal and a noise term (Fig. 1d). For two-photon data the neuropil signal can be 160 modeled as a low rank matrix (*Pnevmatikakis et al., 2016*). For microendoscopic data the larger 161 integration volume leads to more complex background contamination (Zhou et al., 2018). Therefore, 167 a more descriptive model is required (see Methods and Materials (Mathematical model of the CNMF 163 framework) for a mathematical description). CAIMAN BATCH embeds these approaches into a general 164 algorithmic framework that enables scalable automated processing with improved results in terms 165

¹⁶⁶ of quality and processing speed.

167 Deconvolution

Neural activity deconvolution is performed using sparse non-negative deconvolution (*Vogelstein et al., 2010; Pnevmatikakis et al., 2016*) and implemented with both the near-online OASIS algorithm (*Friedrich et al., 2017b*) and an efficient convex optimization framework (*Pnevmatikakis et al., 2016*). The algorithm is competitive to the state of the art according to recent benchmarking studies
 (*Berens et al., 2017*). Prior to deconvolution, the traces are detrended to remove non-stationary

¹⁷³ effects, e.g., photo-bleaching.

174 Online Processing

The three processing steps described above can be implemented in an online fashion on streaming data using the ONACID algorithm (*Giovannucci et al., 2017*). The method builds upon the online dictionary learning framework presented in *Mairal et al. (2010*) for source extraction, by adding the capability of finding new components as they appear and also incorporating the steps of motion

correction and deconvolution (Fig. 1e). CAIMAN ONLINE extends and improves the ONACID prototype

bioRxiv preprint doi: https://doi.org/10.1101/339564; this version posted June 5, 2018. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a Manuscriptisubmittedcense.

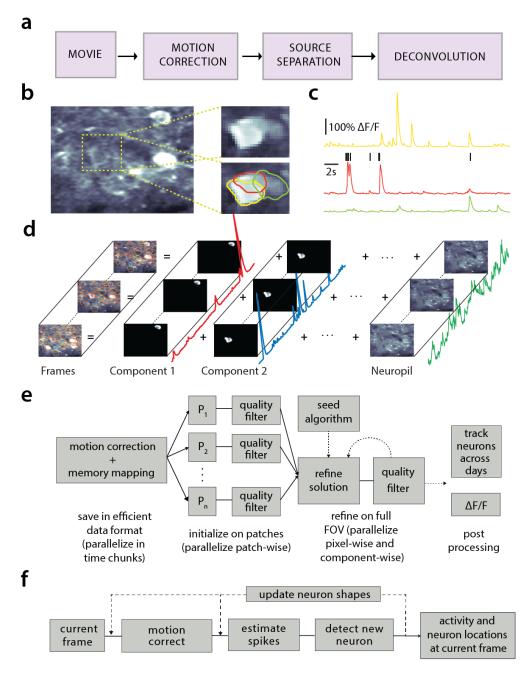


Figure 1. Processing pipeline of CAIMAN for calcium imaging data. (a) The typical pre-processing steps include (i) correction for motion artifacts, (ii) extraction of the spatial footprints and fluorescence traces of the imaged components, and (iii) deconvolution of the neural activity from the fluorescence traces. (b) Time average of 2000 frames from a two-photon microscopy dataset (left) and magnified illustration of three overlapping neurons (right), as detected by the CNMF algorithm. (c) Denoised temporal components of the three neurons in (b) as extracted by CNMF and matched by color (in relative fluorescence change, $\Delta F/F$). (d) Intuitive depiction of CNMF. The algorithm represents the movie as the sum of rank-one spatio-temporal components capturing either neurons and processes, plus additional non-sparse low-rank terms for the background fluorescence and neuropil activity. (e) Flow-chart of the CAIMAN BATCH processing pipeline. From left to right: Motion correction and generation of a memory efficient data format. Initial estimate of somatic locations in parallel over FOV patches using CNMF. Refinement and merging of extracted components via seeded CNMF. Removal of low quality components. Final domain dependent processing stages. (f) Flow-chart of the CAIMAN ONLINE algorithm. After a brief mini-batch initialization phase, each frame is processed in a streaming fashion as it becomes available. From left to right: Correction for motion artifacts. Estimate of activity from existing neurons, identification and incorporation of new neurons. Periodically, the spatial footprints of inferred neurons are updated (dashed lines).

algorithm by introducing a number of algorithmic features and a CNN based component detection
 approach, leading to a major performance improvement.

We now present the new methods introduced by CAIMAN. More details are given in *Methods and Materials* and pseudocode descriptions of the main routines are given in the *Appendix*.

¹⁸⁴ Batch processing of large scale datasets on standalone machines

The batch processing pipeline mentioned above can become a computational bottleneck when 185 tackled without customized solutions. For instance, a naive approach to the problem might have as 186 a first step to load in-memory the full dataset; this approach is non-scalable as datasets typically 187 exceed available RAM (and extra memory is required by any analysis pipeline). To limit memory 188 usage, as well as computation time. CAIMAN BATCH relies on a MapReduce approach (Dean and 189 Ghemawat, 2008). Unlike previous work (Freeman et al., 2014). CAIMAN BATCH assumes minimal 100 computational infrastructure (up to a standard laptop computer), is not tied to a particular parallel 191 computation framework, and is compatible with HPC scheduling systems like SLURM (Yoo et al., 192 2003). 193

Naive implementations of motion correction algorithms need to either load in memory the full 194 dataset or are constrained to process one frame at a time, therefore preventing parallelization. Mo-195 tion correction is parallelized in CAIMAN BATCH without significant memory overhead by processing 196 several temporal chunks of a video data on different CPUs. CAIMAN BATCH broadcasts to each CPU a 197 meta-template, which is used to align all the frames in the chunk. Each process writes in parallel to 198 the target file containing motion-corrected data, which is stored in as a memory mapped array. This 190 allows arithmetic operations to be performed against data stored on the hard drive with minimal 200 memory use, and slices of data to be indexed and accessed without loading the full file in memory 201 More details are given in *Methods and Materials (Memory mapping*). 202

Similarly, the source extraction problem, especially in the case of detecting cell bodies, is 203 inherently local with a neuron typically appearing in a neighborhood within a small radius from its 204 center of mass (Fig. 2a). Exploiting this locality, CAIMAN BATCH splits the FOV into a set of spatially 205 overlapping patches which enables the parallelization of the CNMF (or any other) algorithm to 206 extract the corresponding set of local spatial and temporal components. The user specifies the size 207 of the patch, the amount of overlap between neighboring patches and the initialization parameters 208 for each patch (number of components and rank background for CNMF, stopping criteria for CNMF-209 E). Subsequently the patches are processed in parallel by the CNME/CNME-E algorithm to extract 210 the components and neuropil signals from each patch. 211

Apart from harnessing memory and computational benefits due to parallelization, processing in 212 patches acts indirectly as a dynamic range equalizer and enables CAIMAN BATCH to detect neurons 213 across the whole FOV, a feature absent in the original CNMF, where areas with high absolute 214 fluorescence variation tend to be favored. This results in better source extraction performance. 215 After all the patches have been processed, the results are embedded within the FOV (Fig. 2a). 216 and the overlapping regions between neighboring patches are processed so that components 217 corresponding to the same neuron are merged. The process is summarized in algorithmic format in 218 Alg. 1 and more details are given in Methods and Materials (Combining results from different patches). 219

220 Initialization Methods

221 Initialization methods for matrix factorization problems can impact results due to the non-convex

nature of their objective function. CAIMAN BATCH provides an extension of the GREEDYROI method

- used in *Pnevmatikakis et al. (2016*), that detects neurons based on localized spatiotemporal activity.
- ²²⁴ CAIMAN BATCH can also be seeded with binary masks that are obtained from different sources, e.g.,
- through manual annotation or segmentation of structural channel (SEEDEDINITIALIZATION, Alg. 2).
- ²²⁶ More details are given in *Methods and Materials (Initialization strategies)*.

bioRxiv preprint doi: https://doi.org/10.1101/339564; this version posted June 5, 2018. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under

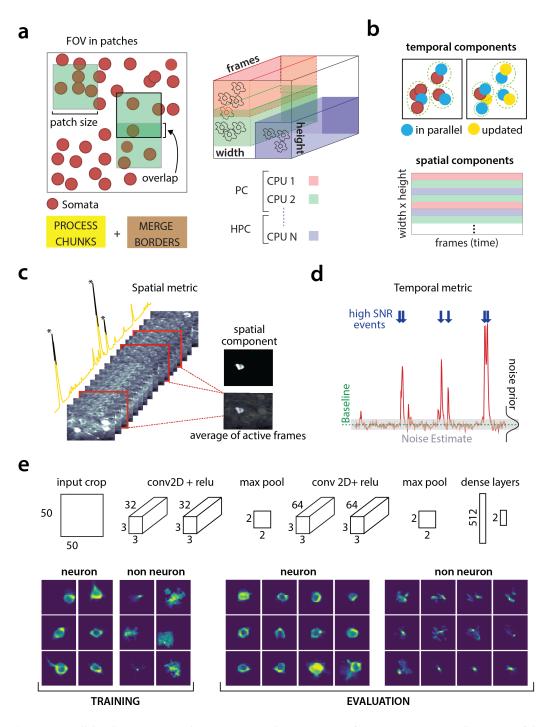


Figure 2. Parallelized processing and component quality assessment for CAIMAN BATCH. (a) Illustration of the parallelization approach used by CAIMAN BATCH for source extraction. The data movie is partitioned into overlapping sub-tensors, each of which is processed in an embarrassingly parallel fashion using CNMF, either on local cores or across several machines in a HPC. The results are then combined. (b) Refinement after combining the results can also be parallelized both in space and in time. Temporal traces of spatially non-overlapping components can be updated in parallel (top) and the contribution of the spatial footprints for each pixel can be computed in parallel (bottom). Parallelization in combination with memory mapping enable large scale processing with moderate computing infrastructure. (c) Quality assessment in space: The spatial footprint of each real component is correlated with the data averaged over time, after removal of all other activity. (d) Quality assessment in time: A high SNR is typically maintained over the course of a calcium transient. (e) CNN based assessment. *Top:* A 4-layer CNN based classifier is used to classify the spatial footprint of each component into neurons or not. *Bottom:* Positive and negative examples for the CNN classifier, during training (left) and evaluation (right) phase. The CNN classifier can accurately classify shapes and generalizes across datasets from different brain areas.

227 Automated component evaluation and classification

A common limitation of matrix factorization algorithms is that the number of components that 228 the algorithm seeks during its initialization must be pre-determined by the user. For example, 229 **Pneymotikakis et al.** (2016) suggest a large number of components which are then heuristically 230 ordered according to their size and activity pattern. When processing large datasets in patches 231 the target number of components is passed on to every patch implicitly assuming a uniform 232 density of (active) neurons within the entire FOV. In general this assumption does not hold and can 233 generate a large number of spurious components. CAIMAN introduces tests to assess the quality 234 of the detected components and eliminate possible false positives. These tests are based on the 235 observation that active components are bound to have a distinct localized spatio-temporal signature 236 within the FOV. We present below unsupervised and supervised tests employed by CAIMAN for 237 component classification. In CAIMAN BATCH, they are initially applied after the processing of each 238 patch is completed, and additionally as a post-processing step after the results from the patches 239 have been merged and refined, whereas in CAIMAN ONLINE they are used to screen new candidate 240 components. We briefly present these tests below and refer to Methods and Materials (Details of 241 auality assessment tests) for more details: 242 **Spatial footprint consistency:** To test whether a detected component is spurious, we correlate 243 the spatial footprint of this component with the average frame of the data, taken over the 244

- ²⁴⁵ interval when the component, with no other overlapping component, was active (Fig. 2c). ²⁴⁶ The component is rejected if the correlation coefficient is below a certain threshold θ_{sp} (e.g.,
- 247 $\theta_{\rm sp} < 0.5$).

Trace SNR: Similarly, for each component we computed the peak SNR of its temporal trace averaged over the duration of a typical transient (Fig. 2d). The component is rejected if the computed SNR is below a certain threshold θ_{SNR} (e.g., $\theta_{SNR} = 2$).

CNN based classification: We also trained a 4-layer convolutional neural network (CNN) to classify spatial footprints into true or false components (Fig. 2e), where a true component here corresponds to a spatial footprint that resembles the soma of a neuron. The classifier, named batch classifier, was trained on a small corpus of manually annotated datasets (full description given in section *Benchmarking against ground truth*) and exhibited similar high classification performance on test samples from different datasets.

²⁵⁷ While CAIMAN uses the CNMF algorithm, the tests described above can be applied to results obtained ²⁵⁸ from any source extraction algorithm, highlighting the modularity of our tools.

259 Online analysis with CAIMAN ONLINE

²⁶⁰ CAIMAN supports online analysis on streaming data building on the core of the prototype algorithm

²⁶¹ of *Giovannucci et al. (2017)*, and extending it in terms of qualitative performance and computational

²⁶² efficiency:

Initialization: Apart from initializing CAIMAN ONLINE with CAIMAN BATCH on a small time interval,
 CAIMAN ONLINE can also be initialized in a bare form over an even smaller time interval, where
 only the background components are estimated and all the components are determined during the online analysis. This process, named BAREINITIALIZATION, can be achieved by running

the CNMF algorithm (*Pnevmatikakis et al.*. 2016) over the small interval to estimate the back-

ground components and possibly a small number of components. The SEEDEDINITIALIZATION

of Alg. 2 can also be used.

270 **Deconvolution:** Instead of a separate step after demixing as in *Giovannucci et al.* (2017), decon-

- volution here can be performed simultaneously with the demixing online, leading to more stable traces especially in cases of low-SNR, as also observed in *Pnevmatikakis et al.* (2016).
- stable traces especially in cases of low-SNR, as also observed in *Pnevmatikakis et al.* (2016).
 Online deconvolution can also be performed for models that assume second order calcium
- dynamics, bringing the full power of *Friedrich et al.* (2017b) to processing of streaming data.

- **Epochs:** CAIMAN ONLINE supports multiple passes over the data, a process that can detect early activity of neurons that were not picked up during the initial pass, as well as smooth the
- activity of components that were detected at late stages during the first epoch.

New component detection using a CNN: To search for new components in a streaming setup. 278 ONACID keeps a buffer of the residual frames, computed by subtracting the activity of already 279 found components and background signals. Candidate components are determined by 280 looking for points of maximum energy in this residual signal, after some smoothing and 28 dynamic range equalization. For each such point identified, a candidate shape and trace are 282 constructed using a rank-1 NMF in a local neighborhood around this point. In its original 283 formulation (Giovannucci et al., 2017), the shape of the component was evaluated using the 284 space correlation test described above. Here, we introduce a CNN classifier approach that 285 tests candidate components by examining their spatial footprint as obtained by the average 286 of the residual buffer across time. This online classifier (different from the batch classifier 287 for quality assessment described above), is trained to be strict, minimizing the number of 288 false positive components that enter the online processing pipeline. It can test multiple 280 components in parallel, and it achieves better performance with no hyper-parameter tuning 290 compared to the previous approach. More details on the architecture and training procedure 291 are given in Methods and Materials (Classification through CNNs). The identification of candidate 292 components is further improved by performing spatial high pass filtering on the average 293 residual buffer to enhance its contrast. The new process for detecting neurons is described in 294 Algs. 3 and 4. See Supplemental Movies 1 and 2 on a detailed graphic description of the new 295 component detection step. 296

297 Component registration across multiple sessions

CAIMAN provides a method to register components from the same FOV across different sessions. 298 The method uses a simple intersection over union metric to calculate the distance between different 290 cells in different sessions and solving a linear assignment problem to perform the registration in 300 a fully automated way (REGISTERPAIR, Alg. 5). To register the components between more than 2 301 sessions (REGISTERMULTI, Alg. 6), we order the sessions chronologically and register the components 302 of the current session against the union of component of all the past sessions aligned to the current 303 FOV. This allows for the tracking of components across multiple sessions without the need of 30/ pairwise registration between each pair of sessions. More details as well as discussion of other 305 methods (Sheintuch et al., 2017) are given in Methods and Materials (Component registration). 306 Benchmarking against ground truth 307

³⁰⁸ To quantitatively evaluate CAIMAN we benchmarked its results against ground truth data.

³⁰⁹ Creating ground truth data through manual annotation

We collected manual annotations from multiple independent labelers who were instructed to find round or donut shaped² *active* neurons on 9 two-photon *in vivo* mouse brain datasets. The datasets were collected at various labs and from various brain areas (hippocampus, visual cortex, parietal cortex) using several GCaMP variants. A summary of the features of all the annotated datasets is given in Table 2. Details about the annotation procedure are given in *Methods and Materials*.

To address human variability in manual annotation each dataset was labeled by 3 or 4 independent labelers, and the final ground truth dataset was created by having the different labelers

reaching a *consensus* over their disagreements (Fig. 3a). The result of this process was defined as

318 ground truth for the evaluation of CAIMAN as well as each individual labeler against the consensus

³¹⁹ (Fig. 3b)³. More details are given in *Methods and Materials* (Collection of manual annotations and

²Since proteins expressing the calcium indicator are confined outside the cell nuclei, neurons will appear as ring shapes, with a dark disk in the center.

³It is possible that this process generated slightly biased results in favor of each individual annotators since the ground truth was always a subset of the union of the individual annotations.

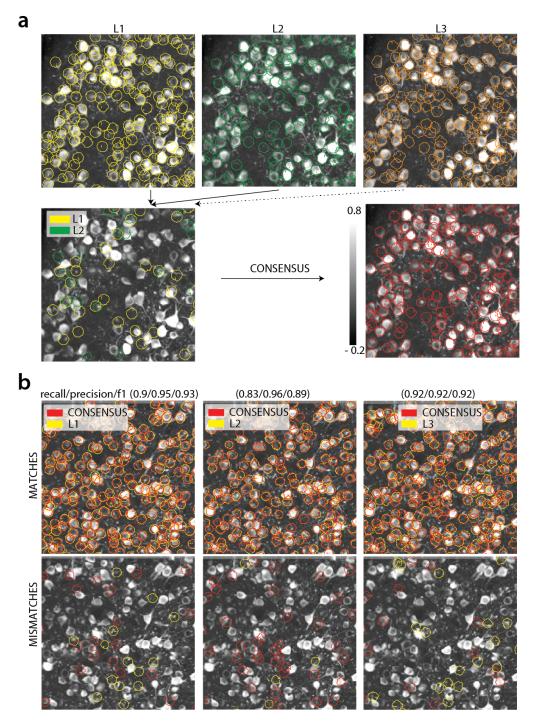


Figure 3. Ground truth generation. (a) *Top:* Individual manual annotations on the dataset N.04.00.t (only part of the FOV is shown) for labelers L1 (left), L2 (middle), L3(right). *Bottom:* Disagreements between L1 and L2 (left), and ground truth labels (right) after the consensus between all labelers has been reached. In this example, consensus considerably reduced the number of initially selected neurons. (b) Matches (top) and mismatches (bottom) between each individual labeler and consensus ground truth. Red contours on the mismatches panels denote false negative contours, i.e., components in the consensus not selected by the corresponding labeler, whereas yellow contours indicate false positive contours. Performance of each labeler is given in terms of precision/recall and F_1 score and indicates an unexpected level of variability between individual labelers.

- 320 ground truth). We believe that the current database, which will be made publicly available, presents
- an improvement over the existing neurofinder database (http://neurofinder.codeneuro.org/) in
- 322 several aspects:
- 323 **Consistency:** The datasets are annotated using exactly the same procedure (see *Methods and*
- Materials), and in all datasets the goal is to detect only active cells. In contrast, the annotation of
- the various neurofinder datasets is performed either manually or automatically by segmenting
- an image of a static (structural) indicator. Even though structural indicators could be used for
- ³²⁷ ground truth extraction, the segmentation of such images is not a straightforward problem ³²⁸ in the case of dense expression, and the stochastic expression of indicators can lead to
- in the case of dense expression, and the stochastic expression of indicators can lead to
- mismatches between functional and structural indicators.
- Uncertainty quantification: By employing more than one human labeler we discovered a surprising level of disagreement between different annotators (see Table 1, Fig. 3b for details),
 which renders individual annotations somewhat unreliable for benchmarking purposes, and
 non-reproducible. The combination of the various annotations leads to more reliable ground
- truth and also quantifies the limits of human performance.

335 Comparing CAIMAN against ground truth

To compare CAIMAN against the consensus ground truth, the manual annotations were used as

binary masks to construct the ground truth spatial and temporal components, using the SEEDEDINI-

TIALIZATION procedure (Alg. 2) of CAIMAN BATCH. The set of spatial footprints obtained from CAIMAN

is registered against the set of ground truth spatial footprints (derived as described above) using

the REGISTERPAIR algorithm (Alg. 5) for component registration described above. Performance is then quantified using a precision/recall framework similar to other studies (*Apthorpe et al.*, 2016:

then quantified using a precision/recall framework similar to o
 Giovannucci et al., 2017).

Giovannucci et al., 201

343 Software

³⁴⁴ CAIMAN is developed by and for the community. Python open source code for all the methods

described above is available at https://github.com/flatironinstitute/CalmAn. The repository contains documentation, numerous demos, and lupyter notebook tutorials, as well as visualization tools, and

documentation, numerous demos, and Jupyter notebook tutorials, as well as visualization tools, and an message/discussion board. The code, which is compatible with Python 2 and Python 3⁴, uses tools

from several open source libraries, such as OpenCV (*Bradski, 2000*), scikit-learn (*Pedregosa et al.*

³⁴⁹ 2011), and scikit-image (Van der Walt et al., 2014). Most routines are also available in MATLAB[®] at

³⁵⁰ https://github.com/flatironinstitute/CalmAn-MATLAB.

351 Results

352 Manual annotations show a high degree of variability

We compared the performance of each human annotator against a consensus ground truth. The 353 performance was quantified with a precision/recall framework and the results of the performance 354 of each individual labeler against the consensus ground truth for each dataset is given in Table 1. 355 The range of human performance in terms of F_1 score was 0.69-0.94, with average 0.83+ 0.07 (mean 356 + STD). All annotators performed similarly on average (0.83+0.05, 0.83+0.08, 0.84+0.06, 0.85+0.08) 357 We also ensured that the performance of labelers was stable across time (i.e. their learning curve 358 plateaued, data not shown). As shown in Table 1 (see also Fig 4b) the F_1 score was never 1, and in 359 most cases it was less or equal to 0.9, demonstrating significant variability between annotators. 360 Fig. 3 (bottom) shows an example of matches and mismatches between individual labelers and 361 consensus ground truth for dataset K53, where the level of agreement was relatively high. The high 362 degree of variability in human responses indicates the challenging nature of the source extraction problem and raises reproducibility concerns in studies relying heavily on manual ROI selection. 364

⁴All future development of CAIMAN will be in Python 3, eventually rendering it incompatible with Python 2.x.

Name	L1	L2	L3	L4	CAIMAN BATCH	CAIMAN ONLINE
N.03.00.t	Х	0.90	0.85 (0.78,0.93)	0.78 (0.73,0.83)	0.78 (0.77,0.79)	0.76 (0.77,0.75)
N.04.00.t	Х	0.69	0.75 (0.61,0.97)	0.87 (0.78,0.98)	0.67 (0.62,0.72)	0.68 (0.65,0.71)
N.02.00	0.89 (0.86,0.93)	$\underset{(0.88,0.85)}{0.87}$	0.84 (0.92,0.77)	0.82	0.79 (0.8,0.77)	0.77 (0.79,0.76)
N.00.00	Х	$\underset{(0.93,0.91)}{0.92}$	0.83	$\underset{(0.96,0.80)}{0.87}$	0.72 (0.83,0.64)	0.72 (0.83,0.64)
N.01.01	0.80 (0.95,0.69)	0.89 (0.96,0.83)	0.78 (0.73,0.84)	0.75 (0.80,0.70)	0.77 (0.88,0.69)	0.73 (0.78,0.68)
YST	0.78 (0.76,0.81)	0.90 (0.85,0.97)	0.82	0.79 (0.96,0.67)	0.76 (0.9,0.66)	0.78 (0.76,0.81)
K53	0.89 (0.96,0.83)	0.92	0.93	0.83 (1.00,0.72)	0.77 (0.83,0.72)	0.82 (0.80,0.83)
J115	Х	0.93 (0.94,0.91)	0.94 (0.95,0.93)	0.83 (1.00,0.71)	0.77 (0.9,0.68)	0.81 (0.75,0.88)
J123	$\underset{(0.96,0.76)}{0.85}$	0.83 (0.73,0.96)	0.90	0.91 (0.92,0.89)	0.68 (0.94,0.51)	0.80 (0.82,0.79)

Table 1. Results of each labeler, CAIMAN BATCH and CAIMAN ONLINE algorithms against consensus ground truth. Results are given in the form F_1 score (precision, recall), and empty entries correspond to datasets not manually annotated by the specific labeler. In *italics* the datasets used to train the CNN classifiers.

CAIMAN BATCH and CAIMAN ONLINE detect neurons with near-human accuracy 365

We first benchmarked CAIMAN BATCH and CAIMAN ONLINE against consensus ground truth for the 366 task of identifying neurons locations and their spatial footprints, using the same precision recall 367 framework (Table 1). Fig. 4a shows an example dataset (K53) along with neuron-wise matches 368 and mismatches between CAIMAN BATCH and consensus ground truth (top) and CAIMAN ONLINE vs 369 consensus ground truth (bottom). 370

The results indicate a similar performance between CAIMAN BATCH and CAIMAN ONLINE; CAIMAN 371 BATCH has F, scores in the range 0.68-0.79 and average performance 0.75+0.04 (mean+STD). On the 372 other hand CAIMAN ONLINE had F_1 scores in the range 0.68-0.82 and average performance 0.76 \pm 0.04. 373 While the two algorithms performed similarly on average, CAIMAN BATCH tends to perform better for 374 shorter datasets whereas online processing tends to lead to better results for longer datasets (see 375 Table 2 for characteristics of the various datasets). CAIMAN approaches but is in most cases below 376 the accuracy levels of human annotators (Fig. 4b). This can be attributed to a number of reasons: 377 First, to demonstrate the generality and ease of use of our tools, the results presented here are 378 obtained by running CAIMAN BATCH and CAIMAN ONLINE with exactly the same parameters for each 379 dataset (see Methods and Materials (Implementation details)); fine-tuning to each individual dataset 380 can significantly increase performance. Second, CNMF detects active components regardless of 381 their shape, and can detect non-somatic structures with significant transients. While non-somatic 382 components can be filtered out to some extent using the CNN classifier, their existence degrades 383 performance compared to the ground truth that consists only of neurons. Lastly, the ground truth 384 is by construction a subset of the union of all individual annotations, which can bias upwards the 385 scores of individual labelers. 386

Neurons with higher SNR transients are detected more accurately 387

396

While CAIMAN ONLINE had balanced performance with respect to precision and recall (mean precision 388 0.77+0.05, mean recall 0.76+0.07). CAIMAN BATCH showed significantly higher precision than recall 389 (mean precision 0.83+0.09, mean recall 0.69+0.08). We looked into this behavior, by analyzing 390 CAIMAN BATCH performance as a function of the SNR of the inferred and ground truth traces 391 (Fig. 4c-d). The SNR measure of a trace corresponds to the peak-SNR averaged over the length of a 392 typical trace (see Methods and Materials (Detecting fluorescence traces with high SNR)). An example is 393 shown in Fig. 4c where the scatter plot of SNR between matched ground truth and inferred traces is 39/ shown (false negative/positive components are shown along the x- and y- axis, respectively). To 395 evaluate the performance we computed a precision metric as the fraction of inferred components

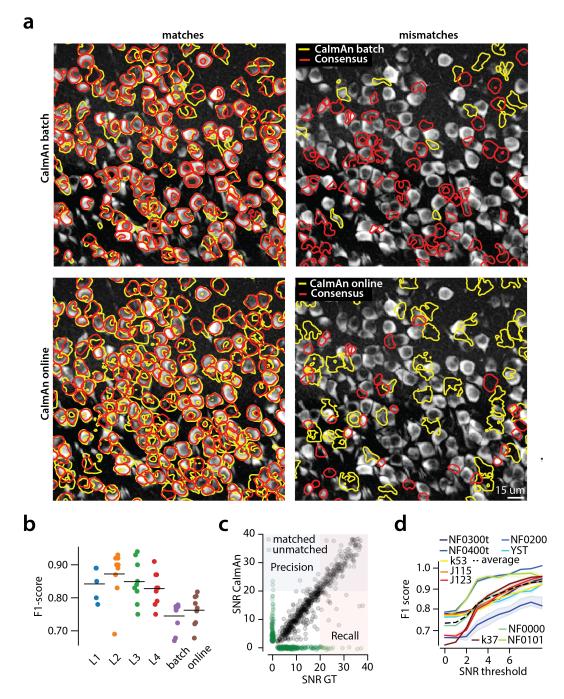


Figure 4. Evaluation of CAIMAN performance against manually annotated data. (a) Comparison of CAIMAN BATCH (top) and CAIMAN ONLINE (bottom) when benchmarked against consensus ground truth for dataset K53. For a portion of the FOV, correlation image overlaid with matches (left panels, true positives red for consensus ground truth, yellow for CAIMAN) and mismatches (right panels, red for false negatives, yellow for false positives). (b) Performance of CAIMAN BATCH, CAIMAN ONLINE and all labelers (L1, L2, L3, L4) for all 9 datasets in terms of F_1 score. CAIMAN BATCH and CAIMAN ONLINE reach near-human accuracy for neuron detection. Complete results with precision and recall for each dataset are given in Table 1. (c-d) Performance of CAIMAN BATCH increases with peak SNR. (c) Example of scatter plot between SNRs of matched traces between CAIMAN BATCH and ground truth for dataset K53. False negative/positive pairs are plotted in green along the x- and y-axes respectively, perturbed as a point cloud to illustrate the density. Most false positive/negative predictions occur at low SNR values. Shaded areas represent thresholds above which components are considered for matching (blue for CAIMAN BATCH selected components and red for GT selected components) (d) F_1 score and upper/lower bounds for all datasets as a function of various peak SNR thresholds. Performance increases significantly for neurons with high peak SNR traces (see text for definition of metrics and the bounds). bioRxiv preprint doi: https://doi.org/10.1101/339564; this version posted June 5, 2018. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a Manuscriptisubmittedcense.

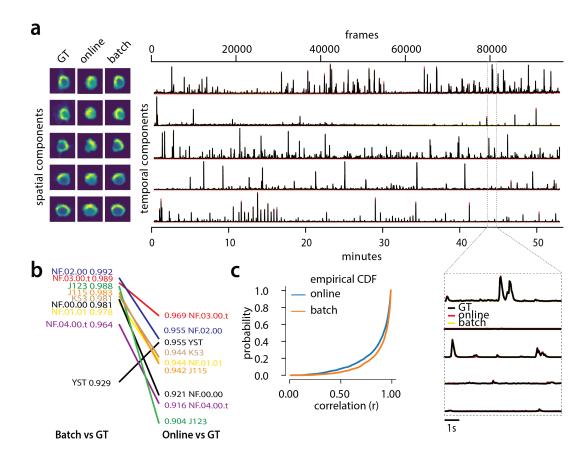


Figure 5. Evaluation of CAIMAN extracted traces against traces derived from ground truth. (a) Examples of shapes (left) and traces (right) are shown for five matched components extracted from dataset K53 for consensus ground truth (GT, black), CAIMAN BATCH (yellow) and CAIMAN ONLINE (red) algorithms. The dashed gray portion of the traces is also shown magnified (bottom-right). Spatial footprints and traces for ground truth are obtained by seeding CAIMAN with the consensus binary masks. The traces extracted from both versions of CAIMAN match closely the ground truth traces. (b) Slope graph for the average correlation coefficient for matches between ground truth and CAIMAN BATCH, and between ground truth and CAIMAN ONLINE. Batch processing produces traces that match more closely the traces extracted from the ground truth data. (c) Empirical cumulative distribution functions of correlation coefficients aggregated over all the tested datasets. Both distributions exhibit a sharp derivative close 1 (last bin), with the batch approach giving better results.

³⁹⁷ above a certain SNR threshold that are matched with a ground truth component (Fig. 4c, shaded

- ³⁹⁸ blue). Similarly we computed a recall metric as the fraction of ground truth components above
- $_{399}$ a SNR threshold that are detected by CAIMAN BATCH (Fig. 4c, shaded red), and an F_1 score as the
- ⁴⁰⁰ harmonic mean of the two (Fig. 4d). The results indicate that the performance significantly grows as
- a function of the SNR for all datasets considered, growing on average from 0.73 when all neurons

are considered to 0.92 when only neurons with traces having SNR ≥ 9 are considered (Fig. 4d)⁵.

403 CAIMAN reproduces the ground truth traces with high fidelity

404 Testing the quality of the inferred traces is a more challenging task due to the complete lack

- of ground truth data in the context of large scale *in vivo* recordings. As mentioned above, we
- ⁴⁰⁶ considered as ground truth the traces obtained by running the CNMF algorithm seeded with the

⁵These precision and recall metrics are computed on different sets of neurons, and therefore strictly speaking one cannot combine them to form an F_1 score. However, they can be bound from above by being evaluated on the set of matched and non-matched components where at least one trace is above the threshold (union of blue and pink zones in Fig. 4c) or below by considering only matched and non-matched components where both ground truth and inferred traces have SNR above the threshold (intersection of blue and pink zones in Fig. 4c). In practice these bounds were very tight for all but one dataset (Fig. 4d). More details can be found in *Methods and Materials (Performance quantification as a function of SNR)*.

binary masks obtained by consensus ground truth procedure. After alignment of the ground truth
 with the results of CAIMAN, the matched traces were compared both for CAIMAN BATCH and for

CAIMAN ONLINE. Fig. 5a, shows an example of 5 of these traces for the dataset K53, showing very
 similar behavior of the traces in these three different cases.

To quantify the similarity we computed the correlation coefficients of the traces (ground truth vs 41 CAIMAN BATCH, and ground truth vs CAIMAN ONLINE) for all the 9 datasets (Fig. 5b-c). Results indicated 412 that for all but one dataset (Fig. 5b) CAIMAN BATCH reproduced the traces with higher fidelity, and 413 in all cases the mean correlation coefficients was higher than 0.9, and the empirical histogram 414 of correlation coefficients peaked at the maximum bin 0.99-1 (Fig. 5c). The results indicate that 415 the batch approach extracts traces closer to the ground truth traces. This can be attributed to 416 a number of reasons: By processing all the time points simultaneously, the batch approach can 417 smooth the trace estimation over the entire time interval as opposed to the online approach where 418 at each timestep only the information up to that point is considered. Moreover, CAIMAN ONLINE 419 might not detect a neuron until it becomes strongly active. This neuron's activity before detection is 420 unknown and has a default value of zero, resulting in a lower correlation coefficient. While this can 421 be ameliorated to a great extent with additional passes over the data, the results indicate trade-offs 422 between using the online and offline versions of CAIMAN. 423

424 Online analysis of a whole brain zebrafish dataset

We tested CAIMAN ONLINE with a 380GB whole brain dataset of larval zebrafish (Danio rerio) acquired 425 with a light-sheet microscope (Kawashima et al., 2016). The imaged transgenic fish (Tg(elavl3:H2B-426 GCaMP6f)if7) expressed the genetically encoded calcium indicator GCaMP6f in almost all neuronal 427 nuclei. Data from 45 planes (FOV 820x410 μ m², spaced at 5.5 μ m intervals along the dorso-ventral 428 axis) was collected at 1Hz for 30 minutes (for details about preparation, equipment and experiment 429 refer to Kawashima et al. (2016)). With the goal of simulating real-time analysis of the data, we run 430 all the 45 planes in parallel on a computing cluster with 9 nodes (each node is equipped with 24 43 CPUs and 128-256 GB RAM). Data was not stored locally in each machine but directly accessed from 432 a network drive 433

The algorithm was initialized with CAIMAN BATCH run on 200 initial frames and looking for 500 434 components. The small number of frames (1885) and the large FOV size (2048×1188 pixels) for this 435 dataset motivated this choice of increased number of components during initialization. In Fig. 6 we 436 report the results of the analysis for plane number 11 of 45. For plane 11, CAIMAN ONLINE found 437 1524 neurons after processing 1685 frames. Since no ground truth was available for this dataset. 438 it was only possible to evaluate the performance of this algorithm by visual inspection. CAIMAN 430 ONLINE identified all the neurons with a clear footprint in the underlying correlation image (higher 440 SNR. Fig. 6a) and missed a small number of the fainter ones (low SNR). By visual inspection of 441 the components the authors could find very few false positives. Given that the parameters were 442 not tuned and that the classifier was not trained on zebrafish neurons, we hypothesize that the 443 algorithm is biased towards a high precision result. Spatial components displayed the expected 444 morphological features of neurons (Fig. 6b-c). Considering all the planes (Figs 6e and 11) CAIMAN 445 ONLINE was able to identify in a single pass of the data a total of 66108 neurons. See Supplemental 446 Movie 3 for a summary across all planes. The analysis was performed in 21 minutes, with the first 447 3 minutes allocated to the initialization and the remaining 18 to process the rest of the data in 448 streaming mode (and in parallel for each plane). This demonstrates the ability of CAIMAN ONLINE 449 to process large amounts of data in real-time (see also Fig. 8 for a discussion of computational 450 performance). 451

452 Analyzing 1p microendoscopic data using CAIMAN

⁴⁵³ We tested the CNMF-E implementation of CAIMAN BATCH on *in vivo* microendosopic data from ⁴⁵⁴ mouse dorsal striatum, with neurons expressing GCaMP6f. 6000 frames were acquired at 30 ⁴⁵⁵ frames per second while the mouse was freely moving in an open field arena (for further details bioRxiv preprint doi: https://doi.org/10.1101/339564; this version posted June 5, 2018. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aManuscriptisubmittedcense.

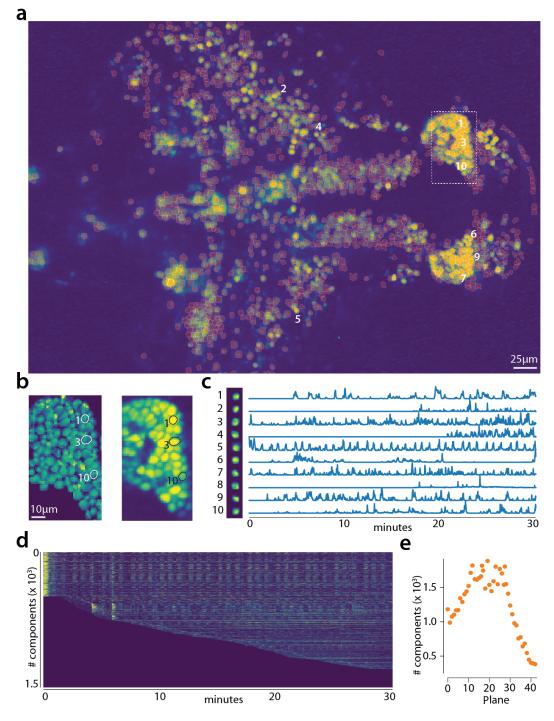


Figure 6. Online analysis of a 30 min long whole brain recording of the zebrafish brain. (a) Correlation image overlaid with the spatial components found by the algorithm (portion of plane 11 out of 45 planes in total). (b) Left: Spatial footprints found in the dashed region in (a), contours represent neurons displayed in (c). Right: Correlation image for the same region. (c) Spatial (left) and Temporal (right) components associated to the ten example neurons marked in panel (a). (d) Temporal traces for all the neurons found in the FOV in (a), the initialization on the first 200 frames contained 500 neurons (present since time 0). (e) Number of neurons found per plane (See also Supplementary Fig. 11 for a summary of the results from all planes).

refer to *Zhou et al. (2018)*). In Fig. 7 we report the results of the analysis using CAIMAN BATCH with patches and compare to the results of the MATLAB[®] implementation of *Zhou et al. (2018*). Both

implementations detect similar components (Fig. 7a) with an F_1 -score of 0.89. 573 neurons were

found by both implementations. 106 and 31 additional components were detected by **Zhou et al.**

(2018) and CAIMAN BATCH respectively. The median correlation between the temporal traces of

⁴⁶¹ neurons detected by both implementations was 0.86. Similar results were also obtained by running

⁴⁶² CAIMAN without patches. Ten example temporal traces are plotted in Fig. 7b.

463 Computational performance of CAIMAN

We examined the performance of CAIMAN in terms of processing time for the various analyzed 464 datasets presented above (Fig. 8). The processing time discussed here excludes motion correction. 465 which is highly efficient and primarily depends on the level of the FOV discretization for non-rigid 166 motion correction (Pnevmatikakis and Giovannucci, 2017). For CAIMAN BATCH, each dataset was 467 analyzed using three different computing architectures: i) a single laptop (MacBook Pro) with 8 468 CPUs and 16GB of RAM (blue in Fig. 8a), ii) a linux-based workstation (CentOS) with 24 CPUs and 469 128GB of RAM (magenta), and iii) a linux-based HPC cluster (CentOS) where 112 CPUs (4 nodes, 28 470 CPUs each) were allocated for the processing task (vellow). Fig. 8a shows the processing of CAIMAN 471 BATCH as a function of dataset size on the 5 longest datasets, whose size exceeded 8GB, on log-log 472 plot. 473

Results show that, as expected, employing more processing power results in faster processing 474 CAIMAN BATCH on a HPC cluster processes data faster than acquisition time (Fig. 8a) even for very 475 large datasets. Processing of an hour long dataset was feasible within 3 hours on a single laptop. 476 even though the dataset has size multiple times the available RAM memory. Here, acquisition time 477 is defined as number of frames times imaging rate, computed based on the assumption of imaging 478 a FOV discretized over a 512×512 grid at a 30Hz rate (a typical two-photon imaging setup with 479 resonant scanning microscopes), and a representation of the measurements using single precision 480 arithmetic, which is the minimum precision required for standard algebraic processing. These 481 assumptions lead to a data rate of ~105GB/hour. In general the performance scales linearly with 482 the number of frames (and hence, the size of the dataset), but a dependence is also observed with 483 respect to the number of components. The majority of the time (Fig. 8b-left) the majority of the 484 time required for CAIMAN BATCH processing is taken by CNMF algorithmic processing either during 485 the initialization in patches (orange bar) or during merging and refining the results of the individual 486 patches (green bar). 487

Fig. 8a also shows the speed performance of CAIMAN ONLINE (red markers). Because of the 488 low memory requirements of the streaming algorithm, this performance only mildly depends on 489 the computing infrastructure allowing for near real-time processing speeds on a standard laptop 490 (Fig. 8a). As discussed in *Giovannucci et al.* (2017) processing time of CAIMAN ONLINE depends 491 primarily on i) the computational cost of tracking the temporal activity of discovered neurons, ii) 492 the cost of detecting and incorporating new neurons, and iii) the cost of periodic updates of spatial 493 footprints. Fig. 8b-right shows that the two first steps, which are required for each frame, can 494 be done in real-time. In Fig. 8c the cost per frame is plotted for the analysis of the whole brain 495 zebrafish recording. The lower imaging rate (1Hz) allows for the tracking of neural activity to be 496 done with computational cost significantly lower than the 1 second between volume imaging time 497 (Fig. 8c), even in the presence of a large number of components (typically more than 1000 per plane. 498 Fig. 6) and the significantly larger FOV (2048 × 1188 pixels). As expected the cost of updating spatial 499 footprints can be significantly larger if done simultaneously for all components (Fig. 8c, bottom). 500 However, the average cost of updating a single spatial footprint is roughly 8ms, enabling real-time 501 processing for each frame, when this step is evenly distributed among different frames/volumes, or 502 is performed by a parallel independent process (Giovannucci et al.. 2017). 503

⁵⁰⁴ The cost of processing 1p data in CAIMAN BATCH using the CNMF-E algorithm (*Zhou et al., 2018*) ⁵⁰⁵ is shown (Fig. 8d) for the workstation hardware. Splitting in patches and processing in parallel can

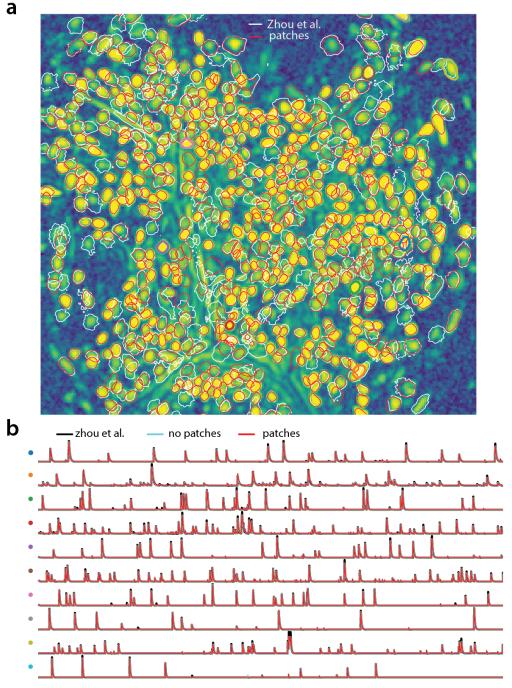


Figure 7. Analyzing microendoscopic 1p data with the CNMF-E algorithm using CAIMAN BATCH. (a) Contour plots of all neurons detected by the CNMF-E (white) implementation of *Zhou et al.* (*2018*) and CAIMAN BATCH (red) using patches. Colors match the example traces shown in (b), which illustrate the temporal components of 10 example neurons detected by both implementations. CAIMAN BATCH reproduces with reasonable fidelity the results of *Zhou et al.* (*2018*).

bioRxiv preprint doi: https://doi.org/10.1101/339564; this version posted June 5, 2018. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a Manuscriptisubmittedcense.

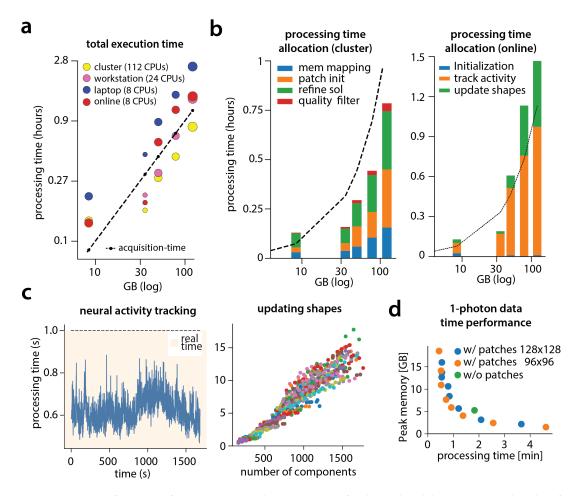


Figure 8. Time performance of CAIMAN BATCH and CAIMAN ONLINE for the analyzed datasets. (a) Log-log plot of total processing time as a function of data size for CAIMAN BATCH for the 5 largest two-photon datasets using three different processing infrastructures: i) a laptop with 8 CPUs (blue), ii) a desktop workstation with 24 CPUs (magenta), and iii) a HPC where 112 CPUs are allocated (yellow). The results indicate a near linear scaling of the processing time with the size of dataset, with additional dependence on the number of found neurons (size of each point). Even very large datasets (> 100GB) can be processed efficiently with a single laptop, whereas access to a HPC enables processing with speed faster than the acquisition time (considered 30Hz for a 512×512 FOV here). The results of CAIMAN ONLINE using the laptop are also plotted in red indicating near real-time processing speed. (b) Break down of processing time for CAIMAN BATCH (left) and CAIMAN ONLINE (right) (excluding motion correction). (Left) Processing with CNMF in patches and refinement takes most of the time for CAIMAN BATCH. Right: Tracking neural activity and new neuron detection can be done in real-time for CAIMAN ONLINE. (c) (Left) Cost of neural activity online tracking for the whole brain zebrafish dataset (maximum time over all planes per frame). Tracking can be done in real-time. (Right) The most expensive part during online processing occurs while updating the spatial footprints, a step that can be distributed or parallelized. Each color corresponds to the update cost for the various different planes. (d) Cost analysis of CNMF-E implementation for processing a 6000 frames long 1p dataset. Processing in patches in parallel induces a time/memory tradeoff and can lead to speed gains (patch size in legend).

lead to computational gains at the expense of increased memory usage. This is because the CNMF-E
 introduces a background term that has the size of the dataset and needs to be loaded and updated
 in memory in two copies. This leads to processing times that are slower compared to the standard
 processing of 2p datasets, and higher memory requirements. However, as Fig. 8d demonstrates,
 memory usage can be controlled enabling scalable inference at the expense of slower processing

511 speeds.

512 CAIMAN successfully tracks neurons across multiple days

Fig. 9 shows an example of tracking neurons across 6 different sessions corresponding to 6 different 513 days of mouse cortex in vivo data using our multi-day registration algorithm REGISTERMULTI (see 514 Methods, Alg. 6), 453, 393, 375, 378, 376, and 373 active components were found in the six sessions. 515 respectively. Our tracking method detected a total of 686 distinct active components. Of these, 172 516 108, 70, 92, 82, and 162 appeared in exactly 1, 2, 3, 4, 5, and all 6 sessions respectively. Contour 517 plots of the 162 components that appeared in all sessions are shown in Fig. 9a, and parts of the 518 FOV are highlighted in Fig. 9d showing that components can be tracked in the presence of non-rigid 519 deformations of the FOV between the different sessions. 520

To test the stability of REGISTERMULTI for each subset of sessions, we repeated the same 521 procedure running backwards in time starting from day 6 and ending at day 1, a process that 522 also generated a total of 686 distinct active components. We identified the components present 523 in at least a given subset of sessions when using the forward pass, and separately when using 524 the backwards pass, and compared them against each other (Fig. 9b) for all possible subsets. 525 Results indicate a very high level of agreement between the two approaches with many of the 526 disagreements arising near the boundaries (data not shown). Disagreements near the boundaries 527 can arise because the forward pass aligns the union with the FOV of the last session, whereas the 528 backwards pass with the FOV of the first session, potentially leading to loss of information near the 529 boundaries. 530

A step by step demonstration of the tracking algorithm for the first three sessions is shown in the appendix (Fig. 10). Our approach allows for the comparison of two non-consecutive sessions through the union of components without the need of a direct pairwise registration (Fig. 10f), where it is shown that registering sessions 1 and 3 directly and through the union leads to nearly identical results. Fig. 9c compares the registrations for all pairs of sessions using the forward (red) or the backward (blue) approach, with the direct pairwise registrations. Again, the results indicate a very high level of agreement, indicating the stability and effectiveness of the proposed approach.

538 Discussion

Reproducible and scalable analysis for the 99%

Significant advances in the reporting fidelity of fluorescent indicators, and the ability to simulta-540 neously record and modulate neurons granted by progress in optical technology, have propelled 541 calcium imaging to being the main experimental method in systems neuroscience alongside elec-542 trophysiology recordings. The resulting increased adoption rate has generated an unprecedented 543 wealth of imaging data which poses significant analysis challenges. The goal of CAIMAN is to provide 544 the experimentalist with a complete suite of tools for analyzing this data in a formal, scalable 545 and reproducible way. The goal of this paper is to present the features of CAIMAN and examine 546 its performance in detail. CAIMAN embeds existing methods for preprocessing calcium imaging 547 data into a MapReduce framework and augments them with supervised learning algorithms and 548 validation metrics. It builds on the CNMF algorithm of *Pnevmatikakis et al.* (2016) for source 5/19 extraction and deconvolution, extending it along the lines of i) reproducibility and performance 550 improvement, by automating quality assessment through the use of unsupervised and supervised 551 learning algorithms for component detection and classification, and ii) scalability, by enabling fast 552 large scale processing with standard computing infrastructure (e.g., a commodity laptop or worksta-553

bioRxiv preprint doi: https://doi.org/10.1101/339564; this version posted June 5, 2018. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a Manuscriptisubmitted ense.

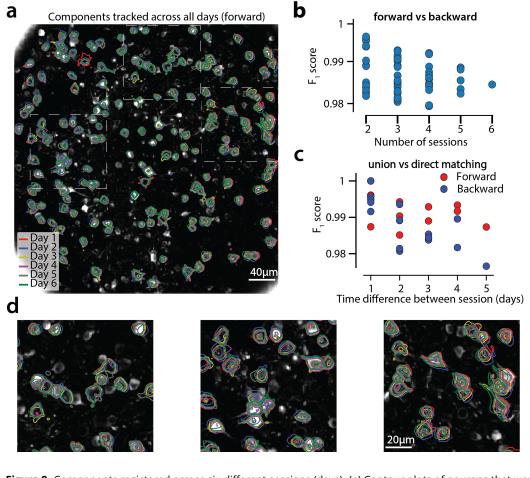


Figure 9. Components registered across six different sessions (days). (a) Contour plots of neurons that were detected to be active in all six imaging sessions overlaid on the correlation image of the sixth imaging session. Each color corresponds to a different session. (b) Stability of multiday registration method. Comparisons of forward and backward registrations in terms of F_1 scores for all possible subsets of sessions. The comparisons agree to a very high level indicating the stability of the proposed approach. (c) Comparison (in terms of F_1 score) of pair-wise alignments using readouts from the union vs direct alignment. The comparison is performed for both the forward and the backwards alignment. For all pairs of sessions the alignment using the proposed method gives very similar results compared to direct pairwise alignment. (d) Magnified version of the tracked neurons corresponding to the squares marked in panel (a). Neurons in different parts of the FOV exhibit different shift patterns over of the course of multiple days, but can nevertheless be tracked accurately by the proposed multiday registration method.

tion). Scalability is achieved by either using a MapReduce batch approach, which employs parallel
 processing of spatially overlapping, memory mapped, data patches; or by integrating the online
 processing framework of *Giovannucci et al.* (2017) within our pipeline. Apart from computational
 gains both approaches also result in improved performance. Towards our goal of providing a
 single package for dealing with standard problems arising in analysis of imaging data, CAIMAN
 also includes an implementation of the CNMF-E algorithm of *Zhou et al.* (2018) for the analysis
 of microendoscopic data, as well as with a novel method for registering analysis results across

561 multiple days.

562 Towards surpassing human neuron detection performance

To evaluate the performance of CAIMAN BATCH and CAIMAN ONLINE, we generated a corpus of 563 multiply apported two-photon imaging datasets. The results indicated a surprising level of dis-564 agreement between individual labelers, highlighting both the difficulty of the problem, and the 565 non-reproducibility of the laborious task of human annotation. CAIMAN reached near-human 566 performance with respect to this ground truth, by using the same parameters for all the datasets 567 without dataset dependent parameter tweaking. Such tweaking could for example include setting 568 the SNR threshold based on the noise level of the recording, the complexity of the neuropil signal 569 based on the level of background activity, or specialized treatment around the boundaries of the 570 FOV to compensate for eventual imaging artifacts. 571

Apart from being used as a benchmarking tool, the set of manual annotations can also be used as labeled data for supervised learning algorithms. CAIMAN uses two CNN based classifiers trained on (a subset of) this data, one for post processing component classification in CAIMAN BATCH, and the other for detecting new neurons in residual images in the CAIMAN ONLINE. The deployment of these classifiers resulted in significant gains in terms of performance, and we expect further advances in the future. The annotations will be made freely available to the community upon publication of the paper for benchmarking and training purposes.

579 CAIMAN BATCH VS CAIMAN ONLINE

Our results suggest similar performance between CAIMAN BATCH and CAIMAN ONLINE in terms of 580 processing speed and quality of results with CAIMAN ONLINE outperforming CAIMAN BATCH ON 58 longer datasets in terms of neuron detection, possibly due to its inherent ability to adapt to non-582 stationarities arising during the course of a large experiment, and underperforming on shorter 583 datasets potentially due to lack of enough information. By contrast, CAIMAN BATCH extracts better 584 traces compared to CAIMAN ONLINE with respect to "ground truth" traces. While multiple passes 585 over the data with CAIMAN ONLINE can mitigate these shortcomings, this still depends on good 586 initialization with CAIMAN BATCH, as the analysis of the whole brain zebrafish dataset indicates. In 587 offline setups. CAIMAN ONLINE could also benefit from the post processing component evaluation 588 tools used in batch mode, e.g., using the batch classifier for detecting false positive components at 580 the end of the experiment. 590

What sets the two algorithms apart is the streaming processing mode of CAIMAN ONLINE which, 591 besides lowering memory requirements, can be used to enable novel types of closed-loop all-592 optical experiments (Packer et al., 2015; Carrillo-Reid et al., 2017). As discussed in Giovannucci 593 et al. (2017), typical all-optical closed-loop experiments require the pre-determination of ROIs that 594 are monitored/modulated. Processing with CAIMAN ONLINE can improve upon this by allowing 595 identification and modulation of new neurons on the fly, greatly expanding the space of possible 596 experiments. Even though our simulated online processing setup is not integrated with hardware to 597 an optical experimental setup, our results indicate that CAIMAN ONLINE performed close to real-time 598 in most cases, without optimizing for speed. This suggest that large scale closed-loop experiments 599 with single cell resolution are feasible by combining existing all-optical technology and our proposed 600 analysis method. 60'

602 Future directions

While CAIMAN uses a highly scalable processing pipeline for two-photon datasets, processing of 603 one-photon microendoscopic imaging data is less scalable due to the more complex background 604 model that needs to be retained in memory during processing. Adapting CAIMAN ONLINE to the one-605 photon data processing algorithm of *Zhou et al.* (2018) is a promising way for scaling up efficient 606 processing in this case. The continuing development and quality improvement of neural activity 607 indicators has enabled direct imaging of neural processes (axons/dendrites), imaging of synaptic 608 activity (Xie et al., 2016), or direct imaging of voltage activity in vivo conditions (Piatkevich et al., 609 2018). While the approach presented here is tuned for somatic imaging through the use of various 610 assumptions (space localized activity, CNN classifiers trained on images of somatic activity), the 611 technology of CAIMAN is largely transferable to these domains as well. These extensions will be 612 pursued in future work. 613

614 Methods and Materials

615 Memory mapping

In order to efficiently access data in parallel, CAIMAN BATCH relies on memory mapping. With 616 memory mapped (mmap) arrays, arithmetic operations can be performed on data residing on the 617 hard drive without explicitly loading it to RAM, and slices of data can be indexed and accessed 618 without loading the full file in memory, enabling out-of-core processing (*Toledo, 1999*). The order 619 in which data in a memory mapped file is stored on the hard drive can dramatically affect the 620 read-write performance of out-of-core operations on spinning disks, and to a lesser degree on solid 621 state drives. On modern computers tensors are stored in linear format, no matter the number of 622 the array dimensions. Therefore, one has to decide which elements of an array are contiguous in 623 memory: in row-major order, consecutive elements of a row (first-dimension) are next to each other. 624 whereas in column-major order consecutive elements of a column (last dimension) are contiguous. 625 Such decisions significantly affect the speed at which data is read or written: in column-major order 626 reading a full column is fast because memory is read in a single sequential block, whereas reading a 627 row is inefficient since only one element can be read at a time and all the data needs to be accessed. 628 Therefore, the original dataset must be saved in the right order to avoid performance problems. 629 In the context of calcium imaging datasets, CAIMAN BATCH represents the datasets in a matrix 630 form Y, where each row corresponds to a different imaged pixel, and each column to a different 631 frame. As a result, a column-major order mmap file enables the fast access of individual frames at a 632 given time, whereas a row-major order files enables the fast access of an individual pixel at all times. 633 To facilitate processing in patches CAIMAN BATCH stores the data in row-major order. In practice, 634 this is opposite to the order with which the data appears, one frame at a time. In order to reduce 635

memory usage and speed up computation CAIMAN BATCH employs a MapReduce approach, where either multiple files or multiple chunks of a big file composing the original datasets are processed and saved in mmap format in parallel. This operation includes two phases, first the chunks/files are saved in multiple row-major mmap format, and then chunks are simultaneously combined into a single large row-major mmap file. In order to reduce preprocessing steps, if the file(s) need to be corrected for motion artifacts, chunks of the registered data can be stored on-the-fly during motion correction.

643 Mathematical model of the CNMF framework

The CNMF framework (Fig. 1d) for calcium imaging data representation can be expressed in mathe matical terms as (*Pnevmatikakis et al., 2016*)

$$Y = AC + B + E. \tag{1}$$

Here, $Y \in \mathbb{R}^{d \times T}$ denotes the observed data written in matrix form, where *d* is the total number of observed pixels/voxels, and *T* is the total number of observed timesteps (frames). $A \in \mathbb{R}^{d \times N}$ denotes the matrix of the *N* spatial footprints, $A = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_N]$, with $\mathbf{a}_i \in \mathbb{R}^{d \times 1}$ being the spatial footprint of component *i*. $C \in \mathbb{R}^{N \times T}$ denotes the matrix of temporal components. $C = [\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_n]^T$.

with $\mathbf{c}_i \in \mathbb{R}^{T \times 1}$ being the temporal trace of component *i*. *B* is the background/neuropil activity

matrix. For two-photon data it is modeled as a low rank matrix $B = \mathbf{bf}$, where $\mathbf{b} \in \mathbb{R}^{d \times n_b}$, $\mathbf{f} \in \mathbb{R}^{n_b \times T}$

 $_{652}$ correspond to the matrices of spatial and temporal background components, and n_b is the number

of background components. For the case of micro-endoscopic data the integration volume is much

⁶⁵⁴ larger and the low rank model is inadequate. A solution comes from the CNMF-E algorithm of **Zhou**

ŀ

et al. (2018) where the background is modeled as

$$B = W(Y - AC), \tag{2}$$

where $W \in \mathbb{R}^{d \times d}$ is an appropriate weight matrix, where the (i, j) entry models the influence of the neuropil signal of pixel *j* to the neuropil signal at pixel *i*.

Combining results from different patches

To combine results from the different patches we first need to account for the overlap at the bound-650 aries. Neurons lying close to the boundary between neighboring patches can appear multiple times 660 and must be merged. With this goal, we optimized the merging approach used in *Pnevmatikakis* 661 et al. (2016): Groups of components with spatially overlapping footprints whose temporal traces are 662 correlated above a threshold are replaced with a single component, that tries to explain as much of 663 the variance already explained by the "local" components (as opposed to the variance of the data 664 as performed in (*Pnevmatikakis et al., 2016*)). If A_{old} , C_{old} are the matrices of components to be 665 merged, then the merged component $\mathbf{a}_{m}, \mathbf{c}_{m}$ are given by the solution of the rank-1 NMF problem: 666

$$\min_{\mathbf{a}_m \ge 0, \mathbf{c}_m \ge 0} \| A_{\text{old}} C_{\text{old}} - \mathbf{a}_m \mathbf{c}_m^{\mathsf{T}} \|.$$
(3)

Prior to merging, the value of each component at each pixel is normalized by the number of patches
 that overlap in this pixel, to avoid counting the activity of each pixel multiple times.

We follow a similar procedure for the background/neuropil signals from the different patches. 669 For the case of two-photon data, the spatial background/neuropil components for each patch can 670 be updated by keeping their spatial extent intact to retain a local neuropil structure, or they can 671 be merged when they are sufficiently correlated in time as described above to promote a more 672 global structure. For the case of one-photon data, CNMF-E estimates the background using a 673 local auto-regressive process (see Eq. (2)) (Zhou et al., 2018), a setup that cannot be immediately 674 propagated when combining the different patches. To combine backgrounds from the different 675 patches, we first approximate the backgrounds B^{i} from all the patches i with a low rank matrix 676 using non-negative matrix factorization of rank g_b to obtain global spatial, and temporal background 677 components. 678

$$[\mathbf{b}^{i}, \mathbf{f}^{i}] = \mathsf{NNMF}(B^{i}, g_{b}).$$
(4)

The resulting components are embedded into a large matrix $B \in \mathbb{R}^{d \times T}$ that retains a low rank structure. After the components and backgrounds from all the patches have been combined, they are further refined by running CNMF iteration of updating spatial footprints, temporal traces, and neuropil activity. CAIMAN BATCH implements these steps in a highly parallel fashion (as also described in *Pnevmatikakis et al. (2016)*): Temporal traces whose corresponding spatial traces do not overlap can be updated in parallel. Similarly, the rows of the matrix of spatial footprints *A* can also be updated in parallel (2b). The process is summarized in algorithmic format in Alg. 1.

686 Initialization strategies

Source extraction using matrix factorization requires solving a bi-convex problem where initialization
 plays a critical role. The CNMF/CNMF-E algorithms use initialization methods that exploit the locality
 of the spatial footprints to efficiently identify the locations of candidate components (*Pnevmatikakis et al., 2016; Zhou et al., 2018*). CAIMAN incorporates these methods, extending them by using the

- temporal locality of the calcium transient events. The available initialization methods for CAIMAN
 BATCH include:
- GREEDYROI: This approach, introduced in *Pnevmatikakis et al.* (2016), first spatially smooths the
- data with a Gaussian kernel of size comparable to the average neuron radius, and then
- 695 initializes candidate components around locations where maximum variance (of the smoothed
- data) is explained. This initialization strategy is fast but requires specification of the number
- ⁶⁹⁷ of components by the user.
- ROLLINGGREEDYROI: The approach, introduced in this paper, operates like GREEDYROI by spatially
 smoothing the data and looking for points of maximum variance. Instead of working across
 all the data, ROLLINGGREEDYROI looks for points of maximum variance on a rolling window of
 a fixed duration, e.g., 3 seconds, and initializes components by performing a rank one NMF on
 a local spatial neighborhood. By focusing into smaller rolling windows, ROLLINGGREEDYROI
- can better isolate single transient events, and as a result detect better neurons with sparse
- activity. ROLLINGGREEDYROI is the default choice for processing of 2-photon data.
- **GREEDYCORR:** This approach, introduced in *Zhou et al.* (2018), initializes candidate components
 around locations that correspond to the local maxima of an image formed by the pointwise
 product between the correlation image and the peak signal-to-noise ratio image. By setting a
 threshold for acceptance, this approach does not require the prior specification of number of
 components. This comes at the expense of a higher computational cost. GREEDYCORR is the
- 710 default choice for processing of 1-photon data.
- SparseNMF: Sparse NMF approaches, when ran in small patches, can be effective for quickly uncov ering spatial structure in the imaging data, especially for neural processes (axons/dendrites)
- whose shape cannot be easily parametrized and/or localized.

714 Algorithm seeding with binary masks

Often locations of components are known either from manual annotation or from labeled data obtained in a different way, such as data from a static structural channel recorded concurrently

- vith the functional indicator. CAIMAN can be seeded with binary (or real valued) masks for the
- ⁷¹⁸ spatial footprints. Apart from *A*, these masks can be used to initialize all the other relevant matrices
- $_{719}$ C and B as well. This is performed by i) first estimating the temporal background components **f**
- using only data from parts of the FOV not covered by any masks and, ii) then estimating the spatial
- background components **b**, and then estimating A, C (with A restricted to be non-zero only at the
- locations of the binary masks), using a simple NMF approach. Details are given in Alg. 2.

723 Details of quality assessment tests

Here we present the unsupervised and supervised quality assessment tests in more detail (Fig. 2).

725 Matching spatial footprints to the raw data

Let $\mathbf{a}_i, \mathbf{c}_i$ denote the spatial footprint and temporal trace of component *i*, and the let A_{i}, C_{i} denote 726 the matrices A, C when the component i has been removed. Similarly, let $Y_i = Y - A_{ij}C_{ij} - B$ denote 727 the entire dataset when the background and the contribution of all components except i have been 728 removed. If component i is real then Y_i and $\mathbf{a}_i \mathbf{c}_i^{\mathsf{T}}$ will look similar during the time intervals when 729 the component *i* is active. As a first test CAIMAN finds the first N_n local peaks of c_i (e.g., $N_n = 5$), 730 constructs intervals around these peaks. (e.g., 50 ms in the past and 300ms in the future, to cover 731 the main part of a possible calcium transient around that point), and then averages Y_{i} across time 732 over the union of these intervals to obtain a spatial image $\langle Y_i \rangle$ (Fig. 2c). The Pearson's correlation 733 over space between $\langle Y \rangle$ and a (both restricted on a small neighborhood around the centroid of 734 \mathbf{a}_{i}) is then computed, and component *i* is rejected if the correlation coefficient is below a threshold 735

value θ_{sp} , (e.g., $\theta_{sp} < 0.5$). Note that a similar test is used in the online approach of *Giovannucci et al.*

737 (2017) to accept for possible new components.

738 Detecting fluorescence traces with high SNR

⁷³⁹ For a candidate component to correspond to an active neuron its trace must exhibit dynamics

740 reminiscent of the calcium indicator's transient. A criterion for this can be obtained by requiring

- the average SNR of trace \mathbf{c}_i over the course a transient to be above a certain threshold θ_{SNR} , e.g.,
- $\theta_{SNR} = 2$, (Fig. 2d). The average SNR is as a measure of how unlikely it is for the transients of \mathbf{c}_i (after
- ⁷⁴³ some appropriate z-scoring) to have been a result of a white noise process.
- To compute the SNR of a trace, let R = Y AC B be the residual spatiotemporal signal. We can obtain the residual signal for each component *i*, **r**_{*i*}, by projecting *R* into the spatial footprint **a**^{*i*}:

$$\mathbf{r}_i = \frac{1}{\|\mathbf{a}_i\|^2} R^{\mathsf{T}} \mathbf{a}_i \tag{5}$$

Then the trace $\mathbf{c}_i + \mathbf{r}_i$ corresponds to the non-denoised trace of component *i*. To calculate its SNR we first compute a type of z-score:

$$\mathbf{z}_{i} = \frac{\mathbf{c}_{i} + \mathbf{r}_{i} - \mathsf{BASELINE}(\mathbf{c}_{i} + \mathbf{r}_{i})}{\mathsf{NOISE}(\mathbf{c}_{i} + \mathbf{r}_{i})}.$$
(6)

The BASELINE(·) function determines the baseline of the trace, which can be varying in the case of

 $_{749}$ $\,$ long datasets exhibiting baseline trends, e.g., due to bleaching. The function ${\sf NOISE}(\cdot)$ estimates

the noise level of the trace. Since calcium transients around the baseline can only be positive, we

estimate the noise level by restricting our attention only to the points t_n where $\mathbf{c}_i + \mathbf{r}_i$ is below the

baseline value, i.e., $t_n = \{t : \mathbf{c}_i(t) + \mathbf{r}_i(t) \le \text{BASELINE}(\mathbf{c}_i + \mathbf{r}_i)\}$, and compute the noise level as the scale

⁷⁵³ parameter of a half-normal distribution (Fig. 2b):

NOISE
$$(\mathbf{c}_i + \mathbf{r}_i) = \operatorname{std}([\mathbf{c}_i + \mathbf{r}_i](t_n))/\sqrt{1 - \frac{2}{\pi}}.$$
 (7)

⁷⁵⁴ We then determine how likely is that the positive excursions of \mathbf{z}_i can be attributed just to noise. We

compute the probabilities $\mathbf{p}_i(t) = \Phi(-\mathbf{z}_i(t))$, where $\Phi(\cdot)$ denotes the cumulative distribution function

 $_{756}$ of a standard normal distribution, and compute the most unlikely excursion over a window of N_s

- timesteps that corresponds to the length of a typical transient, e.g., $N_s = [0.4s \times F]$, where 0.4s
- could correspond to the typical length of a GCaMP6f transient, and F is the imaging rate.

$$p_{\min}^{i} = \min_{t} \left(\prod_{j=0}^{N_{s}-1} \mathbf{p}_{i}(t+j) \right)^{1/N_{s}}.$$
(8)

⁷⁵⁹ The (averaged peak) SNR of component *i* can then be defined as

$$SNR_i = \Phi^{-1}(1 - p_{\min}^i) = -\Phi^{-1}(p_{\min}^i),$$
(9)

where Φ^{-1} is the quantile function for the standard normal distribution (logit function) and a

⁷⁶¹ component is accepted if SNR_i $\geq \theta_{SNR}$. Note that for numerical stability we compute p_{min}^i in the ⁷⁶² logarithmic domain and check the condition $p_{min}^i \leq \Phi(-\theta_{SNR})$.

⁷⁶³We can also use a similar test for the significance of the time traces in the spike domain after ⁷⁶⁴performing deconvolution. In this case, traces can be considered as spiking if the maximum height ⁷⁶⁵due to a spike transient is significantly larger than a threshold. If we assume that the shape of each ⁷⁶⁶calcium transient has been normalized to have maximum amplitude 1, then this corresponds to ⁷⁶⁷testing $\|\mathbf{s}_i\|_{\infty} \ge \theta_{\text{SNR}}\sigma_i$, where \mathbf{s}_i represents the deconvolved activity trace for component *i*, and θ_{SNR} ⁷⁶⁸is again an appropriate SNR threshold, e.g., $\theta_{\text{SNR}} = 2$, and σ_i is the noise level for trace *i*.

⁷⁶⁹ Classification through convolutional neural networks (CNNs)

The tests described above are unsupervised but require fine-tuning of two threshold parameters

 $(\theta_{sp}, \theta_{SNR})$ that might be dataset dependent and might be sensitive to strong non-stationarities. As a

third test we trained a 4-layer CNN to classify the spatial footprints into true or false components,

- ⁷⁷³ where a true component here corresponds to a spatial footprint that resembles the soma of a
- neuron (See Fig. 2e and section *Classification through convolutional networks* for details). A simple
- threshold θ_{CNN} can be used to tune the classifier (e.g., $\theta_{\text{CNN}} = 0.5$).

Name	Area brain	1 - 1-					
		Lab	Rate (Hz)	Size (T×X×Y)	Indicator	# labelers	# neurons GT
NF.03.00.t	Hippocampus	Losonczy	7	2250x498x467	GCaMP6f	3	178
NF.04.00.t	Cortex	Harvey	7	3000x512x512	GCaMP6s	3	257
NF.02.00	Cortex	Svoboda	30	8000x512x512	GCaMP6s	4	394
NF.00.00	Cortex	Svoboda	7	2936x512x512	GCaMP6s	3	425
NF.01.01 \	Visual Cortex	Hausser	7	1825x512x512	GCaMP6s	4	333
YST ۱	Visual Cortex	Yuste	10	3000x200x256	GCaMP3	4	405
K53 Pa	arietal Cortex	Tank	30	116043x512x512	GCaMP6f	4	920
J115 H	Hippocampus	Tank	30	90000x463x472	GCaMP5	3	891
J123 H	Hippocampus	Tank	30	41000x458x477	GCaMP5	4	183

Table 2. Properties of manually annotated datasets. For each dataset the duration, imaging rate and calcium indicator are given, as well as the number of active neurons selected after consensus of the manual annotations.

776 Collection of manual annotations and ground truth

We collected manual annotations from four independent labelers who were instructed to find 777 round or donut shaped neurons of similar size using the Imagel Cell Magic Wand tool Walker (2014). 778 We focused on manually annotating only cells that were active within each dataset and for that 779 reason the labelers were provided with two summary statistics: i) A movie obtained by removing a 780 running 20th percentile (as a crude background approximation) and downsampling in time by a 781 factor of 10, and ii) the max-correlation image. The correlation image (CI) at every pixel is equal 782 to the average temporal correlation coefficient between that pixel and its neighbors *Smith and* 783 Häusser (2010) (8 neighbors were used for our analysis). The max-correlation image is obtained 784 by computing the CI for each batch of 33 seconds (1000 frames for a 30Hz acquisition rate), and 785 then taking the maximum over all these images. Neurons that are inactive during the course of the 786 dataset will be suppressed both from the baseline removed video (since their activity will always be 787 around their baseline), and from the max-correlation image since the variation around this baseline 788 will mostly be due to noise leading to practically uncorrelated neighboring pixels. 9 different mouse 789 in vivo datasets were used from various brain areas and labs. A description is given in Table 2. To 790 create the consensus ground truth, the labelers were asked to jointly resolve the inconsistencies 791 with each others annotations. 792

The annotation procedure provides a binary mask per selected component. On the other 793 hand, the output of CAIMAN for each component is a non-negatively valued vector over the FOV 794 (a real-valued mask). The two sets of masks differ not only in their variable type but also in their 795 general shape: Manual annotation through the Cell Magic Wand tool tends to produce circular 796 shapes, whereas the output of CAIMAN will try to accurately estimate the shape of each active 797 component. To construct ground truth that can be directly used for comparison, the binary masks 798 from the manual annotations were used to seed the CNMF algorithm (Alg. 2). This produced a set 790 of ground truth real valued components with spatial footprints restricted to the areas provided by 800 the annotations, and a corresponding set of temporal components that can be used to evaluate 801 the performance of CAIMAN (Fig. 4). Registration was performed using the REGISTERPAIR algorithm 802 (Alg. 5) and match was counted as a true positive when the (modified) laccard distance (Eq. 11) was 803 below 0.7. Details of the registration procedure are given below (see Component registration). 804

805 Classification through convolutional neural networks (CNNs)

CAIMAN uses two CNN classifiers; one for post processing component screening in CAIMAN BATCH,

- and a different one for screening candidate components in CAIMAN ONLINE. In both cases a 4 layer
- ⁸⁰⁸ CNN was used, with architecture as described in Fig. 2e.

- 809 CAIMAN BATCH classifier for post processing classification
- ⁸¹⁰ The purpose of the batch classifier is to classify the components detected by CAIMAN BATCH into
- neuron somas or other shapes, by examining their spatial footprints. Only three annotated datasets
- (.03.00.t, NF.04.00.t, NF.02.00) were used to train the batch classifier. The set of estimated
- ⁸¹³ footprints from running CAIMAN BATCH initialized with the consensus ground truth was matched
- to the set of ground truth footprints. Footprints matched to ground truth components were
- considered positive examples, whereas the remaining components were labeled as negatives. The
- two sets were enriched using data augmentation (rotations, reflections, contrast manipulation etc.) through the Keras library (keras.io) and the CNN was trained on 60% of the data, leaving 20% for
- through the Keras library (keras.io) and the CNN was trained on 60% of the data, leaving 20% for validation and 20% for testing. The CNN classifier reached an accuracy of 97% on test data: that
- validation and 20% for testing. The CNN classifier reached an accuracy of 97% on test data; that also generalized to the rest of the datasets (Fig. 2e) without any parameter change.
- 820 Online classifier for new component detection
- ⁸²¹ The purpose of the CAIMAN ONLINE classifier is to detect new components based on their spatial
- footprints by looking at the mean across time of the residual buffer. To construct the ground truth data for the online classifier, CAIMAN BATCH was run on the first five annotated datasets seeded with the masks obtained through the manual annotations. Subsequently the activity of random
- subsets of found components and the background was removed from contiguous frames of the
- raw datasets to construct residual buffers, which were averaged across time. From the resulting images patches were extracted corresponding to positive examples (patches around a neuron that
- ⁸²⁷ images patches were extracted corresponding to positive examples (patches around a neuron that ⁸²⁸ was active during the buffer) and negative examples (patches around other positions within the
- FOV). A neuron was considered active if its trace attained an average peak-SNR value of 4 or higher
- during the buffer interval. Similarly to the batch classifier, the two sets were augmented and split
- into training, validation and testing sets. The resulting classifier reached a 98% accuracy on the
- testing set, and also generalized well when applied to different datasets.
- ⁸³³ Differences between the two classifiers
- Although both classifiers examine the spatial footprints of candidate components, their required 834 performance characteristics are different which led us to train them separately. The batch classifier 835 examines each component as a post-processing step to determine whether its shape corresponds 836 to a neural cell body. As such, false positive and false negative examples are treated equally 837 and possible mis-classifications do not directly affect the traces of the other components. By 838 contrast, the online classifier operates as part of the online processing pipeline. In this case, a new 839 component that is not detected in a residual buffer is likely to be detected later should it become 840 more active. On the other hand, a component that is falsely detected and incorporated in the online 8/11 processing pipeline will continue to affect the future buffer residuals and the detection of future 847 components. As such the online algorithm is more sensitive to false positives than false negatives. 843 To ensure a small number of false positive examples under testing conditions, only components 844 with average peak-SNR value at least 4 were considered as positive examples during training of the 845 online classifier. 846
- 847 Component registration
- Fluorescence microscopy methods enable imaging the same part of the brain across different 848 sessions that can span multiple days or weeks. While the microscope can visit the same location 849 in the brain with reasonably high precision, the FOV might might not precisely match due to 850 misalignments or deformations in the brain medium. CAIMAN provides routines for FOV alignment 851 and component registration across multiple sessions/days. Let $\mathbf{a}_1^1, \mathbf{a}_2^1, \dots, \mathbf{a}_N^1$ and $\mathbf{a}_2^2, \mathbf{a}_2^2, \dots, \mathbf{a}_N^2$ the 852 sets of spatial components from sessions 1 and 2 respectively, where N_1 and N_2 denote the total 853 number of components from each session. We first compute the FOV displacement by aligning 854 some summary images from the two sessions (e.g., mean or correlation image), using some non-855 rigid registration method, e.g., NoRMCorre (*Pnevmatikakis and Giovannucci, 2017*). We apply the 856

estimated displacement field to the components of A_1 to align them with the FOV of session 2. To perform the registration, we construct a pairwise distance matrix $D \in \mathbb{R}^{N_1 \times N_2}$ with $D(i, j) = d(\mathbf{a}_i^1, \mathbf{a}_j^2)$, where $d(\cdot, \cdot)$ denotes a distance metric between two components. The chosen distance corresponds to the Jaccard distance between the binarized versions of the components. A real valued component **a** is converted into its binary version $m(\mathbf{x})$ by setting to 1 only the values of **a** that are above the

maximum value of **a** times a threshold θ_b , e.g., $\theta_b = 0.2$:

$$m(\mathbf{a})((x)) = \begin{cases} 1, & \mathbf{a}((x)) \ge \theta_b \|\mathbf{a}\|_{\infty} \\ 0, & \text{otherwise} \end{cases}$$
(10)

To compute the distance between two binary masks m_1, m_2 , we use the Jaccard index (intersection over union) which is defined as

$$J(m_1, m_2) = \frac{|m_1 \cap m_2|}{|m_1 \cup m_2|},$$
(11)

⁸⁶⁵ and use it to define the distance metric as

$$d(\mathbf{a}_{i}^{1}, \mathbf{a}_{j}^{2}) = \begin{cases} 1 - J(m(\mathbf{a}_{i}^{1}), m(\mathbf{a}_{j}^{2})), & 1 - J(m(\mathbf{a}_{i}^{1}), m(\mathbf{a}_{j}^{2})) \le \theta_{d} \\ 0, & (m(\mathbf{a}_{i}^{1}) \subseteq m(\mathbf{a}_{j}^{2})) \text{ OR } (m(\mathbf{a}_{j}^{2}) \subseteq m(\mathbf{a}_{i}^{1})) \\ \infty, & \text{otherwise.} \end{cases}$$
(12)

where θ_d is a distance threshold, e.g., 0.5 above which two components are considered nonmatching and their distance is set to infinity to prevent false assignments.

After the distance matrix *D* has been completed, an optimal matching between the components of the two sessions is computed using the Hungarian algorithm to solve the linear assignment problem. As infinite distances are allowed, it is possible to have components from both sessions that are not matched with any other component. This process of registering components across two sessions (REGISTERPAIR) is summarized in Alg. 5.

To register components across multiple sessions, we first order the sessions chronologically 873 and register session 1 against session 2. From this registration we construct the union of the 874 distinct components between the two sessions by keeping the matched components from session 875 2 as well as the non-matched components from both sessions aligned to the FOV of session 876 2. We then register this union of components to the components of session 3 and repeat the procedure until all sessions are have been registered. This process of multi session registration 878 (REGISTERMULTI) is summarized in Alg. 6. At the end of the process the algorithm produces a list of 879 matches between the components of each session and the union of all active distinct components. 880 allowing for efficient tracking of components across multiple days (Fig. 9), and the comparison 881 of non-consecutive sessions through the union without the need of direct pairwise registration 882 (Fig. 10)). An alternative approach to the problem of multiple session registration (CellReg) was 883 presented recently by Sheintuch et al. (2017) where the authors register neurons across multiple 884 days by first constructing a similar union set of all the components which is then refined using a 885 clustering procedure, REGISTERMULTI differs from the CELLREG method of *Sheintuch et al.* (2017) in 886 a few key ways, that highlight its simplicity and robustness: 887

 REGISTERMULTI uses a very simple intersection over union metric to estimate the distance 888 between two neighboring neurons after the FOV alignment. Cells that have a distance above 889 a given threshold are considered different by default and are not tested for matching. This 890 parameter is intuitive to set a priori for each dataset. In contrast CELLREG uses a probabilistic 891 framework based on the joint probability distribution between the distance of two cells 892 and the correlation of their shapes that makes specific parametric assumptions about the 893 distributions of centroid distances between the same and different cells, as well as their shape 894 correlations. This model needs to be re-evaluated for every different set of sessions to be 895 registered and potentially requires a lot of data to learn the appropriate distance metric. 896

- REGISTERMULTI uses the Hungarian algorithm to register two different set of components, a
- ⁸⁹⁸ practice that solves the linear assignment problem optimally under the assumed distance ⁸⁹⁹ function. In contrast CELLREG uses a greedy method for initializing the assignment of cells to
- tunction. In contrast CELLREG uses a greedy method for initializing the assignment of cells to the union superset relving on the following clustering step to refine these estimates, and thus
- adding extra computational burden to the registration procedure.
- adding extra computational burden to the registration procedure

902 Implementation details for CAIMAN BATCH

Each dataset was processed using the same set of parameters, excepting the expected size of 903 neurons (estimated by inspecting the correlation image), the size of patches and expected number 904 of neurons per patch (estimated by inspecting the correlation image). For the dataset N.01.01, 905 where optical modulation was induced, the threshold for merging neurons was slightly higher (the 906 stimulation caused clustered synchronous activity). For shorter datasets, rigid motion correction 907 was sufficient; for longer datasets K53, J115 we applied non-rigid motion correction. The parameters 908 for the automatic selection of components were optimized using only the first three datasets and 909 fixed for all the remaining files. For all datasets the background neuropil activity was modeled as a 910 rank 2 matrix, and calcium dynamics were modeled as a first order autoregressive process. The 911 remaining parameters were optimized so that all the datasets could be run on a machine with less 912 than 128GB RAM. 913

914 Implementation details for CAIMAN ONLINE

Datasets were processed for two epochs with the exception of the longer datasets K53, J115, J123 915 where only one pass of the data was performed to limit computational cost. For each dataset 916 the online CNN classifier was used to detect new neurons, and five candidate components were 917 considered for each frame. The online CNN classifier had the same threshold 0.5 for all datasets. 918 with the exception of the longest datasets J115, J123 where the threshold was set to 0.75. Setting 910 the threshold to 0.5 for these datasets led to slightly poorer performance. Large datasets were 920 spatially decimated by a factor of 2 to enhance processing speed, a step that did not lead to changes 921 in detection performance. For all datasets the background neuropil activity was modeled as a rank 922 2 matrix, and calcium dynamics were modeled as a first order autoregressive process. For each 923 dataset. CAIMAN ONLINE was initialized on the first 200 frames, using the BAREINITIALIZATION on the entire FOV with only 2 neurons, so in practice all the neurons were detected during the online 925 mode. To highlight the *truly* online processing mode, no post-processing of the results was used, a 926 step that can further enhance the performance of the algorithm. Similarly to batch processing, the 927 expected size of neurons was chosen separately for each dataset after inspecting the correlation 928 image. 929

For the analysis of the whole brain zebrafish dataset, CAIMAN ONLINE was run for 1 epoch with 930 the same parameters as above, with only differences appearing in the number of neurons during 931 initialization (600 vs 2), and the value of the threshold for the online CNN classifier (0.75 vs 0.5). 932 The former decision was motivated by the goal of retrieving with a single pass neurons from a 933 preparation with a denser level of activity over a larger FOV in this short dataset (1885 frames). 934 To this end, the number of candidate neurons at each timestep was set to 10 (per plane). The 935 threshold choice was motivated by the fact that the classifier was trained on mouse data only, and 936 thus a higher threshold choice would help diminish potential false positive components. Rigid 937 motion correction was applied online to each plane. 938

939 Performance quantification as a function of SNR

⁹⁴⁰ To quantify performance as a function of SNR we approximate the ground truth traces by running

- 941 CAIMAN BATCH on the datasets seeded with the "consensus" binary masks obtained from the manual
- ₉₄₂ annotators. After that the average peak-SNR of a trace \mathbf{c} with corresponding residual signal \mathbf{r} (5) is
- 943 obtained as

$$SNR(\mathbf{z}) = -\Phi^{-1}(p_{\min}), \tag{13}$$

- where $\Phi^{-1}(\cdot)$ denotes the probit function (quantile function for the standard Gaussian distribution), z is the z-scored version of $\mathbf{c} + \mathbf{r}$ (6) and p_{\min} is given by (8).
- Let $c_1^{\text{gt}}, c_2^{\text{gt}}, \dots, c_N^{\text{gt}}$ be the ground truth traces and $c_1^{\text{cm}}, c_2^{\text{cm}}, \dots, c_N^{\text{cm}}$ be their corresponding CAIMAN
- ⁹⁴⁷ inferred traces. Here we assume that false positive and false negative components are matched with
- trivial components that have 0 SNR. Let also $\text{SNR}^{\text{gt}}_{i} = \text{SNR}(c_i^{\text{gt}})$ and $\text{SNR}^{\text{cm}}_{i} = \text{SNR}(c_i^{\text{cm}})$, respectively.
- 49 After we compute the SNR for both ground truth and inferred traces the performance algorithm
- can be quantified in multiple ways as a function of a SNR thresholds θ_{SNR} :
 - **Precision:** Precision at level θ_{SNR} , can be computed as the fraction of detected components with $SNR^{cm} > \theta_{SNR}$ that are matched with ground truth components. It quantifies the certainty that a component detected with a given SNR or above corresponds to a true component.

$$PREC(\theta_{SNR}) = \frac{|\{i : (SNR^{cm}_i > \theta_{SNR}) \& (SNR^{gt}_i > 0)\}|}{|\{i : (SNR^{cm}_i > \theta_{SNR})\}|}$$

Recall: Recall at level θ_{SNR} , can be computed as the fraction of ground truth components with $SNR^{gt} > \theta_{SNR}$ that are detected by the algorithm. It quantifies the certainty that a ground truth component with a given SNR or above is detected.

$$\text{RECALL}(\theta_{\text{SNR}}) = \frac{|\{i : (\text{SNR}^{\text{gt}}_i > \theta_{\text{SNR}}) \& (\text{SNR}^{\text{cm}}_i > 0)\}|}{|\{i : (\text{SNR}^{\text{gt}}_i > \theta_{\text{SNR}})\}|}$$

 F_1 **score:** An overall F_1 score at level θ_{SNR} , can be obtained by computing the harmonic mean between precision and recall

$$F_{1}(\theta_{SNR}) = 2 \frac{PREC(\theta_{SNR}) \times RECALL(\theta_{SNR})}{PREC(\theta_{SNR}) + RECALL(\theta_{SNR})}$$

⁹⁵¹ The cautious reader will observe that the precision and recall quantities described above are

not computed in the same set of components. This can be remedied by recomputing the quantities

953 in two different ways:

AND framework: Here we consider a match only if both traces have SNR above the given threshold:

$$PREC_{AND}(\theta_{SNR}) = \frac{|\{i : (SNR^{cm}_{i} > \theta_{SNR}) \& (SNR^{et}_{i} > \theta_{SNR})\}|}{|\{i : (SNR^{cm}_{i} > \theta_{SNR})\}|}$$
$$RECALL_{AND}(\theta_{SNR}) = \frac{|\{i : (SNR^{gt}_{i} > \theta_{SNR}) \& (SNR^{cm}_{i} > \theta_{SNR})\}|}{|\{i : (SNR^{gt}_{i} > \theta_{SNR})\}|}$$

OR framework: Here we consider a match if *either* trace has SNR above the given threshold and its match has SNR above 0.

$$\operatorname{RECALL}_{OR}(\theta_{SNR}) = \frac{|\{i : (\max(SNR^{gt}_{i}, SNR^{cm}_{i}) > \theta_{SNR}) \& (\min(SNR^{gt}_{i}, SNR^{cm}_{i}) > 0)\}|}{|\{i : (SNR^{cm}_{i} > 0)\}|}$$
$$\operatorname{RECALL}_{OR}(\theta_{SNR}) = \frac{|\{i : (\max(SNR^{gt}_{i}, SNR^{cm}_{i}) > \theta_{SNR}) \& (\min(SNR^{gt}_{i}, SNR^{cm}_{i}) > 0)\}|}{|\{i : (SNR^{gt}_{i} > 0)\}|}$$

It is easy to show that

$$\begin{split} & \text{PREC}_{\text{AND}}(\theta_{\text{SNR}}) \leq \text{PREC}(\theta_{\text{SNR}}) \leq \text{PREC}_{\text{OR}}(\theta_{\text{SNR}}) \\ & \text{RECALL}_{\text{AND}}(\theta_{\text{SNR}}) \leq \text{RECALL}(\theta_{\text{SNR}}) \leq \text{RECALL}_{\text{OR}}(\theta_{\text{SNR}}) \\ & F_{1\text{AND}}(\theta_{\text{SNR}}) \leq F_{1}(\theta_{\text{SNR}}) \leq F_{1\text{OR}}(\theta_{\text{SNR}}), \end{split}$$

with equality holding for $\theta_{SNR} = 0$. As demonstrated in Fig. 4d, these bounds are tight.

955 Additional features of CAIMAN

- 956 CAIMAN contains a number of additional features that are not presented in the results section for
- 957 reasons of brevity. These include:

- 958 Volumetric data processing
- 959 Apart from planar 2D data, CAIMAN BATCH is also applicable to 3D volumetric data arising either
- ⁹⁶⁰ from dense raster scanning methods, or from direct volume imaging methods such as light field
- ⁹⁶¹ microscopy (*Prevedel et al., 2014; Grosenick et al., 2017*).
- 962 Segmentation of structural indicator data
- 963 Structural indicators expressed in the nucleus and functional indicators expressed in the cytoplasm
- ₉₆₄ can facilitate source extraction and help identify silent or specific subpopulations of neurons
- 965 (e.g., inhibitory). CAIMAN provides a simple adaptive thresholding filtering method for segmenting
- ⁹⁶⁶ summary images of the structural channel (e.g., mean image). The obtained results can be used
- ⁹⁶⁷ for seeding source extraction from the functional channel in CAIMAN BATCH or CAIMAN ONLINE as
- ⁹⁶⁸ already discussed.
- 969 Duplicate Detection
- ⁹⁷⁰ The ground truth obtained through the consensus process was screened for possible duplicate
- selections. To detect for duplicate components we define the degree of spatial overlap matrix O as

$$O_{ij} = \begin{cases} 0, & i = j \\ \frac{|m(\mathbf{a}^i) \cap m(\mathbf{a}^j)|}{|m(\mathbf{a}^j)|}, & i \neq j \end{cases},$$
(14)

- that defines the fraction of component *i* that overlap with component *j*, where $m(\cdot)$ is the thresh-
- olding function defined in (10). Any entry of O that is above a threshold θ_o (e.g., $\theta_o = 0.7$ used
- ⁹⁷⁴ here) indicates a pair of duplicate components. To decide which of the two components should be
- 975 removed, we use predictions of the CAIMAN BATCH CNN classifier, removing the component with the
- 976 lowest score.
- 977 Extraction of $\Delta F/F$
- ⁹⁷⁸ The fluorescence trace \mathbf{f}_i of component *i* can be written as

$$\mathbf{f}_i = \|\mathbf{a}_i\|^2 (\mathbf{c}_i + \mathbf{r}_i). \tag{15}$$

The fluorescence due to the component's transients overlaps with a background fluorescence due to baseline fluorescence of the component and neuropil activity, that can be expressed as

$$\mathbf{f}_{0,i} = \mathsf{BASELINE}(\mathbf{f}_i + \boldsymbol{B}^{\mathsf{T}} \mathbf{a}_i), \tag{16}$$

where BASELINE : $\mathbb{R}^T \mapsto \mathbb{R}^T$ is a baseline extraction function, and *B* is the estimated background

⁹⁸² signal. Examples of the baseline extraction function are a percentile function (e.g., 10th percentile),

or a for longer traces, a running percentile function, e.g., 10th percentile over a window of a hundred

⁹⁸⁴ seconds⁶. To determine the optimal percentile level an empirical histogram of the trace (or parts of

it in case of long traces) is computed using a diffusion kernel density estimator (*Botev et al., 2010*),

⁹⁸⁶ and the mode of this density is used to define the baseline and its corresponding percentile level.

⁹⁸⁷ The $\Delta F/F$ activity of component *i* can then be written as

$$\mathbf{c}_{i}^{\Delta F/F} = \frac{\mathbf{f}_{i} - \mathsf{BASELINE}(\mathbf{f}_{i})}{\mathbf{f}_{0,i}}$$
(17)

The approach we propose here is conceptually similar to practical approaches where the $\Delta F/F$ is

⁹⁸⁹ computed by averaging over the spatial extent of an ROI (*Jia et al., 2011*) with some differences:

- $_{_{990}}$ i) instead of averaging with a binary mask we use the a weighed average with the shape of each
- ⁹⁹¹ component, ii) signal due to overlapping components is removed from the calculation of the

⁶Computing the exact running percentile function can be computationally intensive. To reduce the complexity we compute the running percentile with a stride of W, where W is equal or smaller to the length of the window, and then linearly interpolate the values.

- ⁹⁹² background fluorescence, and iii) the traces have been extracted through the CNMF process prior
- ₉₉₃ to the $\Delta F/F$ extraction. Note that the same approach can also be performed to the trace $\|\mathbf{a}_i\|^2 \mathbf{c}_i$
- ⁹⁹⁴ that does not include the residual traces for each component. In practice it can be beneficial to
- 995 extract $\Delta F/F$ traces prior to deconvolution, since the $\Delta F/F$ transformation can alleviate the effects
- of drifting baselines, e.g., due to bleaching. For the non-deconvolved traces \mathbf{f}_i some temporal
- ⁹⁹⁷ smoothing can also be applied to obtain more smooth $\Delta F/F$ traces.

998 Acknowledgments

We thank B. Cohen, L. Myers, N. Roumelioti, and S. Villani for providing us with manual annotations. 999 We thank V. Staneya and B. Deverett for contributing to the early stages of CAIMAN, and L. Paninski 1000 for numerous useful discussions. We thank N. Carriero, J. Fisk, and D. Simon from the Flatiron 1001 Institute (Simons Foundation) for useful discussions and suggestions to optimize High Performance 1002 Computing code. We thank T. Kawashima and M. Ahrens for sharing the whole brain zebrafish 1003 dataset. Last but not least, we thank the active community of CAIMAN users for their great help in 1004 terms of code/method contributions, bug reporting, code testing and suggestions that have led 1005 to the growth of CAIMAN into a widely used open source package. A partial list of contributors 1006 (in the form of GitHub usernames) can be found in https://github.com/flatironinstitute/CalmAn/ 1007 graphs/contributors (Python) and https://github.com/flatironinstitute/CalmAn-MATLAB/graphs/ 1008 contributors (MATLAB[®]). The authors acknowledge support from following funding sources: A.G., 1009 E.A.P., I.F., P.G. (Simons Foundation), I.G., S.A.K., D.W.T. (NIH NRSA F32NS077840-01.5U01NS090541, 1010 1U19NS104648 and Simons Foundation SCGB), P.Z. (NIH NIBIB R01EB022913, NSF NeuroNex DBI-1011 1707398, Gatsby Foundation), I.T. (NIH R01-MH101198), F.N. (MURI, Simons Collaboration on the 1012 Global Brain and Pew Foundation). 1013

1014 **References**

Ahrens MB, Orger MB, Robson DN, Li JM, Keller PJ. Whole-brain functional imaging at cellular resolution using
 light-sheet microscopy. Nature methods. 2013; 10(5):413–420.

Apthorpe N, Riordan A, Aguilar R, Homann J, Gu Y, Tank D, Seung HS. Automatic Neuron Detection in Calcium
 Imaging Data Using Convolutional Networks. In: *Advances in Neural Information Processing Systems*; 2016. p.
 3270–3278.

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, Reimer J, et al. Community-based benchmarking improves spike inference from two-photon calcium
 imaging data. bioRxiv. 2017 Aug; p. 177956. https://www.biorxiv.org/content/early/2017/08/18/177956, doi:
 10.1101/177956.

Botev ZI, Grotowski JF, Kroese DP. Kernel density estimation via diffusion. The Annals of Statistics. 2010; 38(5):2916–2957.

Bouchard MB, Voleti V, Mendes CS, Lacefield C, Grueber WB, Mann RS, Bruno RM, Hillman EM. Swept confocally aligned planar excitation (SCAPE) microscopy for high-speed volumetric imaging of behaving organisms.
 Nature photonics. 2015; 9(2):113–119.

Bradski G. The OpenCV Library. Dr Dobb's Journal: Software Tools for the Professional Programmer. 2000;
 25(11):120–123.

Cai DJ, Aharoni D, Shuman T, Shobe J, Biane J, Song W, Wei B, Veshkini M, La-Vu M, Lou J, et al. A shared neural
 ensemble links distinct contextual memories encoded close in time. Nature. 2016; 534(7605):115.

Carrillo-Reid L, Yang W, Kang Miller Je, Peterka DS, Yuste R. Imaging and Optically Manipulating Neuronal
 Ensembles. Annual review of biophysics. 2017; 46:271–293.

Chen TW, Wardill TJ, Sun Y, Pulver SR, Renninger SL, Baohan A, Schreiter ER, Kerr RA, Orger MB, Jayaraman V,
 Looger LL, Svoboda K, Kim DS. Ultrasensitive fluorescent proteins for imaging neuronal activity. Nature. 2013
 lul; 499(7458). https://www.nature.com/articles/nature12354, doi: 10.1038/nature12354.

- Cichocki A, Zdunek R, Amari Si. Hierarchical ALS algorithms for nonnegative matrix and 3D tensor factorization.
 In: International Conference on Independent Component Analysis and Signal Separation Springer; 2007. p. 169–176.
- Dean J, Ghemawat S. MapReduce: simplified data processing on large clusters. Communications of the ACM.
 2008; 51(1):107–113.

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

- Flusberg BA, Nimmerjahn A, Cocker ED, Mukamel EA, Barretto RP, Ko TH, Burns LD, Jung JC, Schnitzer MJ.
 High-speed, miniaturized fluorescence microscopy in freely moving mice. Nature methods. 2008; 5(11):935.
- Freeman J, Vladimirov N, Kawashima T, Mu Y, Sofroniew NJ, Bennett DV, Rosen J, Yang CT, Looger LL, Ahrens MB.
 Mapping brain activity at scale with cluster computing. Nature methods. 2014; 11(9):941–950.
- **Friedrich J**, Yang W, Soudry D, Mu Y, Ahrens MB, Yuste R, Peterka DS, Paninski L. Multi-scale approaches for highspeed imaging and analysis of large neural populations. PLoS computational biology. 2017; 13(8):e1005685.
- Friedrich J, Zhou P, Paninski L. Fast online deconvolution of calcium imaging data. PLOS Computational Biology.
 2017; 13(3):e1005423.
- **Giovannucci A**, Friedrich J, Kaufman J, Churchland A, Chklovskii D, Paninski L, Pnevmatikakis EA. OnACID: Online analysis of calcium imaging data in real time. In: *Neural Information Processing Systems (NIPS)*; 2017.
- 1057 Grosenick LM, Broxton M, Kim CK, Liston C, Poole B, Yang S, Andalman AS, Scharff E, Cohen N, Yizhar O,
- Ramakrishnan C, Ganguli S, Suppes P, Levoy M, Deisseroth K. Identification Of Cellular-Activity Dynamics
- Across Large Tissue Volumes In The Mammalian Brain. bioRxiv. 2017 May; p. 132688. https://www.biorxiv.
 org/content/early/2017/05/01/132688, doi: 10.1101/132688.
- Jia H, Rochefort NL, Chen X, Konnerth A. In vivo two-photon imaging of sensory-evoked dendritic calcium signals in cortical neurons. Nature protocols. 2011; 6(1):28.
- Kaifosh P, Zaremba JD, Danielson NB, Losonczy A. SIMA: Python software for analysis of dynamic fluorescence
 imaging data. Frontiers in neuroinformatics. 2014; 8:80.
- Kawashima T, Zwart MF, Yang CT, Mensh BD, Ahrens MB. The serotonergic system tracks the outcomes of actions to mediate short-term motor learning. Cell. 2016; 167(4):933–946.
- Klibisz A, Rose D, Eicholtz M, Blundon J, Zakharenko S. Fast, Simple Calcium Imaging Segmentation with Fully
 Convolutional Networks. In: *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support* Springer; 2017.p. 285–293.
- Mairal J, Bach F, Ponce J, Sapiro G. Online learning for matrix factorization and sparse coding. Journal of Machine Learning Research. 2010; 11(Jan):19–60.
- 1072 **Mukamel EA**, Nimmerjahn A, Schnitzer MJ. Automated analysis of cellular signals from large-scale calcium 1073 imaging data. Neuron. 2009; 63(6):747–760.
- Pachitariu M, Packer AM, Pettit N, Dalgleish H, Hausser M, Sahani M. Extracting regions of interest from
 biological images with convolutional sparse block coding. In: *Advances in Neural Information Processing Systems*; 2013. p. 1745–1753.
- 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. 2017; p. 061507.
- Packer AM, Russell LE, Dalgleish HW, Häusser M. Simultaneous all-optical manipulation and recording of neural
 circuit activity with cellular resolution in vivo. Nature Methods. 2015; 12(2):140–146.
- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg
 V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay É. Scikit-learn: Machine Learning in
- Python. Journal of Machine Learning Research. 2011; 12(Oct):2825–2830. http://www.jmlr.org/papers/v12/ pedregosa11a.html.
- Petersen A, Simon N, Witten D. SCALPEL: Extracting Neurons from Calcium Imaging Data. arXiv preprint
 arXiv:170306946. 2017; .

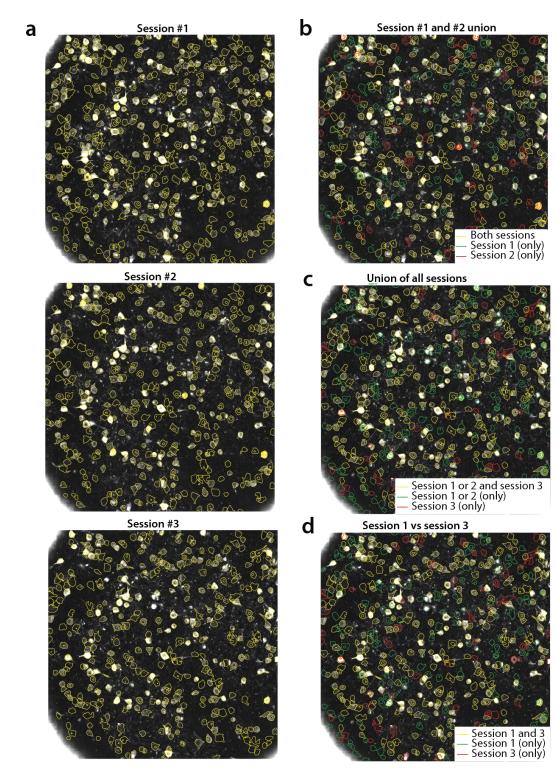
- 1087 Piatkevich KD, Jung EE, Straub C, Linghu C, Park D, Suk HJ, Hochbaum DR, Goodwin D, Pnevmatikakis E, Pak N,
- Kawashima T, Yang CT, Rhoades JL, Shemesh O, Asano S, Yoon YG, Freifeld L, Saulnier JL, Riegler C, Engert
 F, et al. A robotic multidimensional directed evolution approach applied to fluorescent voltage reporters.
- Nature Chemical Biology. 2018 Apr; 14(4):352–360. https://www.nature.com/articles/s41589-018-0004-9, doi:
 1091 10.1038/s41589-018-0004-9.
- Pnevmatikakis EA, Giovannucci A. NoRMCorre: An online algorithm for piecewise rigid motion correction of
 calcium imaging data. Journal of Neuroscience Methods. 2017; 291:83–94.
- Pnevmatikakis EA, Merel J, Pakman A, Paninski L. Bayesian spike inference from calcium imaging data. In:
 Signals, Systems and Computers, 2013 Asilomar Conference on IEEE; 2013. p. 349–353.
- Pnevmatikakis EA, Soudry D, Gao Y, Machado TA, Merel J, Pfau D, Reardon T, Mu Y, Lacefield C, Yang W,
 Ahrens M, Bruno R, Jessell TM, Peterka DS, Yuste R, Paninski L. Simultaneous Denoising, Deconvolution, and
 Demixing of Calcium Imaging Data. Neuron. 2016 Jan; 89(2):285–299. https://www.cell.com/neuron/abstract/

1099 S0896-6273(15)01084-3, doi: 10.1016/j.neuron.2015.11.037.

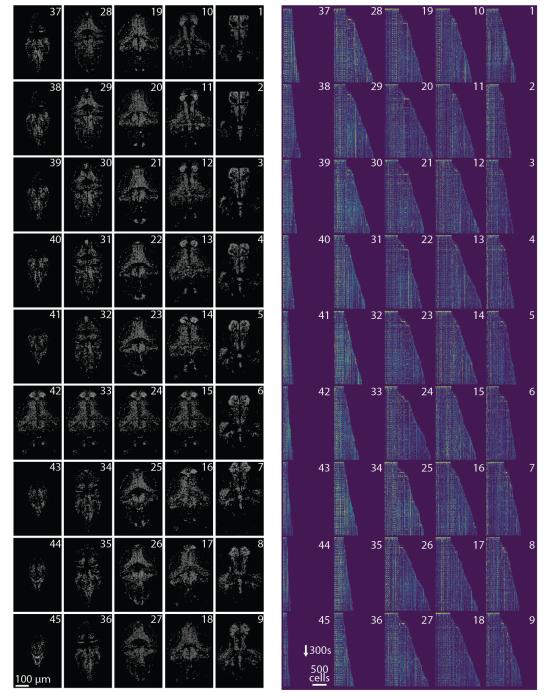
- Prevedel R, Yoon YG, Hoffmann M, Pak N, Wetzstein G, Kato S, Schrödel T, Raskar R, Zimmer M, Boyden ES, Vaziri
 A. Simultaneous whole-animal 3*D*-imaging of neuronal activity using light field microscopy. Nat Methods.
 2014; 11:727–730.
- **Reynolds S**, Abrahamsson T, Schuck R, Sjöström PJ, Schultz SR, Dragotti PL. ABLE: An Activity-Based Level Set Segmentation Algorithm for Two-Photon Calcium Imaging Data. eNeuro. 2017; 4(5):ENEURO-0012.
- Sheintuch L, Rubin A, Brande-Eilat N, Geva N, Sadeh N, Pinchasof O, Ziv Y. Tracking the Same Neurons across
 Multiple Days in Ca2+ Imaging Data. Cell reports. 2017; 21(4):1102–1115.
- Smith SL, Häusser M. Parallel processing of visual space by neighboring neurons in mouse visual cortex. Nature
 neuroscience. 2010; 13(9):1144–1149.
- Sofroniew NJ, Flickinger D, King J, Svoboda K. A large field of view two-photon mesoscope with subcellular
 resolution for in vivo imaging. Elife. 2016; 5.
- **Spaen Q**, Hochbaum DS, Asín-Achá R. HNCcorr: A Novel Combinatorial Approach for Cell Identification in Calcium-Imaging Movies. arXiv preprint arXiv:170301999. 2017; .
- Speiser A, Yan J, Archer EW, Buesing L, Turaga SC, Macke JH. Fast amortized inference of neural activity from
 calcium imaging data with variational autoencoders. In: *Advances in Neural Information Processing Systems*;
 2017. p. 4027–4037.
- **Theis L**, Berens P, Froudarakis E, Reimer J, Rosón MR, Baden T, Euler T, Tolias AS, Bethge M. Benchmarking spike rate inference in population calcium imaging. Neuron. 2016; 90(3):471–482.
- **Toledo S.** A survey of out-of-core algorithms in numerical linear algebra. External Memory Algorithms and
 Visualization. 1999; 50:161–179.
- Vogelstein JT, Packer AM, Machado TA, Sippy T, Babadi B, Yuste R, Paninski L. Fast nonnegative deconvolution for
 spike train inference from population calcium imaging. Journal of neurophysiology. 2010; 104(6):3691–3704.
- 1122
 Walker T, Cell magic wand tool; 2014.
 https://www.maxplanckflorida.org/fitzpatricklab/software/

 1123
 cellMagicWand/.
- Van der Walt S, Schönberger JL, Nunez-Iglesias J, Boulogne F, Warner JD, Yager N, Gouillart E, Yu T. scikit-image:
 image processing in Python. PeerJ. 2014; 2:e453.
- Xie Y, Chan AW, McGirr A, Xue S, Xiao D, Zeng H, Murphy TH. Resolution of high-frequency mesoscale intracortical maps using the genetically encoded glutamate sensor iGluSnFR. Journal of Neuroscience. 2016; 36(4):1261– 1272.
- 1129 **Yoo AB**, Jette MA, Grondona M. Slurm: Simple linux utility for resource management. In: *Workshop on Job* 1130 *Scheduling Strategies for Parallel Processing* Springer; 2003. p. 44–60.
- Zhou P, Resendez SL, Rodriguez-Romaguera J, Jimenez JC, Neufeld SQ, Giovannucci A, Friedrich J, Pnevmatikakis
 EA, Stuber GD, Hen R, Kheirbek MA, Sabatini BL, Kass RE, Paninski L. Efficient and accurate extraction of in
 vivo calcium signals from microendoscopic video data. eLife. 2018 Feb; 7:e28728. https://elifesciences.org/
 articles/28728, doi: 10.7554/eLife.28728.

bioRxiv preprint doi: https://doi.org/10.1101/339564; this version posted June 5, 2018. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a Manuscriptisubmitted ense.



Appendix 0 Figure 10. Tracking neurons across days, step-by-step description of multi session registration (Fig. 9). (a) Correlation image overlaid to contour plots of the neurons identified by CAIMAN BATCH in day 1 (top, 453 neurons), 2 (middle, 393 neurons) and 3 (bottom, 375 neurons). (b) Result of the pairwise registration between session 1 and 2. The union of distinct active components consists of the components that were active in i) both sessions (yellow - where only the components of session 2 are displayed), ii) only in session 2 (green), and iii) only in session 1, aligned to the FOV of session 2 (red). (c) At the next step the union of sessions 1 and 2 is registered with the results of session 3 to produce the union of all distinct components aligned to the FOV of session 3. (d) Comparison of non-consecutive sessions without pairwise registration. Keeping track of which session each component was active in, enables efficient and stable comparisons.



Appendix 0 Figure 11. Profile of spatial (left) and temporal (right) components found in each plane of the whole brain zebrafish recording. (Left) Components are extracted with CAIMAN ONLINE and then max-thresholded. (Right) See Results section for a complete discussion.

1135 Supplemental Data

- **1136 Description of Supplemental Movies**
- 1137 **Movie 1:** Depiction of CAIMAN ONLINE on a small patch of *in vivo* cortex data. Top left: Raw
- 1138 data. Bottom left: Footprints of identified components. Top right: Mean residual buffer and
- ¹¹³⁹ proposed regions for new components (in white squares). Enclosings of accepted regions are
- shown in magenta. Several regions are proposed multiple times before getting accepted. This is
- due to the strict behavior of the classifier to ensure a low number of false positives. Bottom right:
 Reconstructed activity.
- **Movie 2:** Depiction of CAIMAN ONLINE on a single plane of mesoscope data courtesy of E. Froudarakis,
- 1144 J. Reimers and A. Tolias (Baylor College of Medicine). Top left: Raw data. Top right: Inferred activity
- (without neuropil). Bottom left: Mean residual buffer and accepted regions for new components
- (magenta squares). Bottom right: Reconstructed activity.
- 1147 Movie 3: Results of CAIMAN ONLINE initialized by CAIMAN BATCH on a whole brain zebrafish dataset.
- Each panel shows the active neurons in a given plane (top-to-bottom) without any background
- activity. See the text for more details.

1150 Algorithmic Details

- In the following we present in pseudocode form several of routines introduced and used by CAIMAN.
- Note that the pseudocode descriptions do not aim to present a complete picture and may refer to
- 1153 other work for some of the steps.

Algorithm 1 PROCESSINPATCHES

Require: Input data matrix <i>Y</i> , patch size , overlap s	size, initialization method, rest of parameters.			
1: $Y^{(1)}, \dots, Y^{(N_p)} = \text{CONSTRUCTPATCHES}(Y, p_s, o_s)$	Break data into memory mapped patches.			
2: for $i = 1,, N_p$ do	Process each patch			
3: $[A^{(i)}, C^{(i)}, \mathbf{b}^{(i)}, \mathbf{f}^{(i)}] = CNMF(Y^{(i)}, options)$	Run CNMF on each patch			
4: end for				
5: $[A, C] = MERGECOMPONENTS[{A^{(i)}, C^{(i)}}_{i=1,,N}]$	Merge Components			
6: $[\mathbf{b}, \mathbf{f}] = MERGEBACKGROUNDS[\{\mathbf{b}^{(i)}, \mathbf{f}^{(i)}\}_{i=1,\dots,N}]$	Merge background components			
7: $M \leftarrow (A > 0)$.	Find masks of spatial footprints.			
8: repeat ▷ Optionally keep upda	ting <i>A</i> , <i>C</i> , b , f using HALS (<i>Cichocki et al., 2007</i>).			
9: $[\mathbf{b}, \mathbf{f}] \leftarrow NNMF(Y - AC, n_b)$				
10: $C \leftarrow \arg\min_{C \ge 0} \ Y - \mathbf{bf} - AC\ $				
11: $A \leftarrow \arg \min_{A \ge 0, A(\sim M) = =0} \ Y - \mathbf{bf} - AC\ $				
12: until Convergence				
13: return <i>A</i> , <i>C</i> , b , f				

Algorithm 2 SEEDEDINITIALIZATION

Require: Input data matrix Y, matrix of binary masks M, number of background components n_b . 1: $\mathbf{p} = \text{find}(A\mathbf{1} == 0)$ ▷ Find the pixels not covered by any component. 2: $[\sim, \mathbf{f}] \leftarrow \mathsf{NNMF}(Y[\mathbf{p}, :], n_b)$ ▷ Run NMF on these pixels just to get temporal backgrounds f 3: $\mathbf{b} \leftarrow \arg\min_{\mathbf{b}>0} \|Y - \mathbf{bf}\|$ ⊳ Obtain spatial background **b**. 4: $C \leftarrow \max\left((M^{\top}M)^{-1}M^{\top}(Y - \mathbf{bf}), 0\right)$ ▷ Initialize temporal traces. > Initialize spatial footprints constrained within the 5: $A \leftarrow \arg \min_{A \ge 0, A(\sim M) = =0} \|Y - \mathbf{bf} - AC\|$. masks. ▷ Optionally keep updating *A*, *C*, **b**, **f** using HALS. 6: repeat $[\mathbf{b}, \mathbf{f}] \leftarrow \mathsf{NNMF}(Y - AC, n_b)$ 7: $C \leftarrow \arg\min_{C>0} \|Y - \mathbf{bf} - AC\|$ 8: $A \leftarrow \arg\min_{A \geq 0, A(\sim M) = = 0} \|Y - \mathbf{bf} - AC\|$ 9:

- 10: **until** Convergence
- 11: **return** *A*, *C*, **b**, **f**

Algorithm 3 CAIMAN ONLINE (See Giovannucci et al. (2017) for explanation of routines)

Require: Data matrix *Y*, initial estimates *A*, **b**, *C*, **f**, *S*, current number of components *K*, current timestep *t*', rest of parameters.

1: $W = Y[:, 1:t']C^{\top}/t'$					
2: $M = CC^{T}/t'$ > Initialize sufficient statistics (<i>Giovannucci et al., 201</i>	7)				
3: $\mathcal{G} = DETERMINEGROUPS([A, \mathbf{b}], K)$ \triangleright Giovannucci et al. (2017) , Alg. S1-S	52				
4: $R_{\text{buf}} = [Y - [A, \mathbf{b}][C; \mathbf{f}]][:, t' - l_b + 1 : t']$ > Initialize residual buff	er				
5: $t = t'$					
6: for $i = 1, \dots, N_{\text{epochs}}$ do					
7: while there is more data do					
8: $t \leftarrow t+1$					
9: $\mathbf{y}_t \leftarrow MOTIONCORRECT(\mathbf{y}_t, \mathbf{bf}_{t-1})$ \triangleright (<i>Pnevmatikakis and Giovannucci, 201</i>	7)				
10: $[\mathbf{c}_t; \mathbf{f}_t] \leftarrow UPDATeTRACES([A, \mathbf{b}], [\mathbf{c}_{t-1}; \mathbf{f}_{t-1}], \mathbf{y}_t, \mathcal{G}) \succ Giovannucci et al. (2017), Alg. States of the $	33				
11: $C, S \leftarrow OASIS(C, \gamma, s_{\min}, \lambda)$ \triangleright Friedrich et al. (2017)	b)				
12: $A_{\text{new}}, C_{\text{new}} \leftarrow FINDNEwCOMPONENTS(R_{\text{buf}}, N_{\text{comp}})$ \triangleright Alg.	4				
13: $[A, \mathbf{b}], [C, \mathbf{f}], K, \mathcal{G}, R_{\text{buf}}, W, M \leftarrow \text{INTEGRATENEwCOMPONENTS}($					
14: $[A, \mathbf{b}], [C, \mathbf{f}], K, \mathcal{G}, A_{\text{new}}, C_{\text{new}}, R_{\text{buf}}, \mathbf{y}_t, W, M) \qquad \triangleright Giovannucci et al. (2017), Alg. Statement of the statement of the$	54				
15: $R_{\text{buf}} \leftarrow [R_{\text{buf}}[:, 2: l_b], \mathbf{y}_t - A\mathbf{c}_t - \mathbf{bf}_t]$ \triangleright Update residual buff	er				
16: $W, M \leftarrow UPDATeSuffStatistics(W, M, \mathbf{y}_t, [\mathbf{c}_t; \mathbf{f}_t])$					
17: if mod $(t - t', l_b) = 0$ then \triangleright Update $W, M, [A, b]$ every l_b timestep)S				
18: $[A, \mathbf{b}] \leftarrow UPDATeSHAPES[W, M, [A, \mathbf{b}]] \triangleright Giovannucci et al. (2017), Alg. S$	55				
19: end if					
20: end while					
21: $t \leftarrow 0$					
22: end for					
23: return <i>A</i> , b, <i>C</i> , f, <i>S</i>					

Algorithm 4 FINDNEWCOMPONENTS

Require: Residual buffer R_{buf} , number of new candidate components N_{comp} , neuron radius r.

1: $E \leftarrow \sum_{t} \max(R_{\operatorname{buf}(t)}, 0)^2$	
2: $E \leftarrow \text{HighPassFilter}(E)$	▷ Spatial high pass filtering for contrast enhancement.
3: $P = FINDLOCALPEAKS(E, N_{comp}, r)$	▷ Find local maxima at least 2r apart.
4: $A_{\text{test}} \leftarrow \emptyset$	
5: for $p \in P$ do	
6: $N_p = \{(x, y) : x - p_x \le r, y - p_y \le r\}$	} Define a neighborhood around p
7: $A_{\text{test}} \leftarrow A_{\text{test}} \cup MEAN(R_{\text{buf}})$	
8: end for	
9: $I_{\text{accept}} \leftarrow \text{ONLINECNNCLASSIFIER}(A_{\text{test}})$	Find indeces of accepted components
10: $A_{\text{new}} \leftarrow \emptyset, C_{\text{new}} \leftarrow \emptyset$	
11: for $i \in I_{accept}$ do	
12: $[\mathbf{a}, \mathbf{c}] \leftarrow NNMF(R_{buf}[N_{p^i}, :], 1)$	
13: $A_{\text{new}} \leftarrow A_{\text{new}} \cup \mathbf{a}$	
14: $C_{\text{new}} \leftarrow C_{\text{new}} \cup \mathbf{c}$	
15: end for	
16: return $A_{\text{new}}, C_{\text{new}}$	

Algorithm 5 REGISTERPAIR

Require: Spatial footprint matrices A_1, A_2 , field of view templates I_1, I_2 , thresholds for binarization θ_h and matching θ_m . 1: $S = \text{COMPUTEMOTIONFIELD}(I_1, I_2)$ ▷ Compute motion field between the templates. 2: $A_1 \leftarrow \mathsf{APPLYMOTIOnFIELD}(A_1, S)$ \triangleright Align A_1 to the template I_2 ▷ Turn components into binary masks. 3: $[M_1, M_2] = \text{BINARIZE}([A_1, A_2], \theta_b)$ ⊳ Compute distance matrix. 4: $D = \text{COMPUTEDISTANCEMATRIX}(M_1, M_2, \theta_D)$ 5: $P_1, P_2, L_1, L_2 = \text{HUNGARIAN}(D)$ ▷ Match using the Hungarian algorithm. 6: return Matched components P_1, P_2 , non-matched components L_1, L_2 and aligned components from first session A_1 .

Algorithm 6 REGISTERMULTI

Require: List of spatial footprint matrices A_1, A_2, \ldots, A_N , list of FOV templates I_1, I_2, \ldots, I_N , thresholds for binarization θ_b and matching θ_m .

1:	for $i = 1,, N$ do	
2:	$K_i = SIZE(A_i, 2)$	Number of components in each session.
3:	end for	
4:	$A_u \leftarrow A_1$	\triangleright Initialize A_u matrix
5:	$m[1] = [1, 2, \dots, K_1]$	Initialize matchings list
6:	$K_{tot} \leftarrow K_1$	Total # of distinct components so far.
7:	for $i = 2,, N$ do	
8:	$P_u, P_i, L_u, L_i, A_u = REGISTERPAIR(A_u, A_i, I_i)$	$(-1, I_i, \theta_b, \theta_m)$ \triangleright Register A_u to session i .
9:	$A_u[:, P_u] \leftarrow A_i[:, P_i]$	\triangleright Keep the matched components from session <i>i</i> .
10:	$A_u \leftarrow [A_u, A_i[:, L_i]] \qquad \qquad \triangleright \; Ir$	clude the non-matched components from session <i>i</i> .
11:	$m[i][P_i] = P_u \triangleright m[i][j] = k \text{ if components}$	ent <i>j</i> from session <i>i</i> is mapped to component k in A_u .
12:	$m[i][L_i] = [K_{tot} + 1, K_{tot} + 2, \dots, K_{tot} + L_i $] > Include newly added components.
13:	$K_{tot} \leftarrow K_{tot} + L_i $	Update total number of distinct components.
14:	end for	
15:	return Union of all distinct components A	m_{u} , and list of matchings <i>m</i> .